We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

6,300

Open access books available

172,000

International authors and editors

190M

Downloads

Our authors are among the

154

Countries delivered to

TOP 1%

most cited scientists

12.2%

Contributors from top 500 universities

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Building Clinical Trust in Automated Knowledge Acquisition

Anna Shillabeer
Carnegie Mellon University
Australia

1. Overview

The aim of this chapter is to describe the process of medical knowledge acquisition from a historical context and to define the requirements for producing knowledge which is able to be trusted and applied in a clinical setting. This is related to modern data mining approaches which do not yet adequately address these requirements. This is believed to be the most critical issue in the acceptance of data mining in the medical domain. The chapter will discuss how data mining *can* address these needs and will provide discussion of a technical solution to the stated issues. Overall this chapter aims to demonstrate that the individual needs of all medical professionals can be addressed and that data mining can be a valuable tool in the diagnostic and decision making toolkit. It will also empower medical professionals to take a greater role in the development of such systems by providing a checklist of features to include and pitfalls to avoid, thus ensuring greater success in future systems.

2. Introduction

2.1 Clinical data mining context

While acceptance of data mining technologies is growing, progress has been slow and it is not yet an integrated part of the medical data analysis toolkit. Many reasons have been documented for this but the primary issues are two fold; the decision making and knowledge acquisition processes of the medical domain are not adequately reflected in the technologies available and the systems are too often built to suit the specific analytical needs of an individual user. Whilst this has enabled the application of the technology in specific scenarios, it has resulted in the development of tools which cannot be utilised outside of the specific purpose for which they were built. These issues serve to limit the exposure, applicability and trust of data mining systems to the medical domain.

Data mining researchers have long been concerned with the application of tools to facilitate and improve data analysis on large, complex data sets for the purpose of knowledge acquisition. The current challenge is to make data mining and knowledge discovery systems applicable to a wider range of domains, among them medicine. Early work was performed over transactional, retail based data sets, but the attraction of finding previously unknown knowledge from the increasing volume of data collected from the medical domain is an emerging area of interest and specialisation. This chapter is primarily concerned with

Source: Data Mining in Medical and Biological Research, Book edited by: Eugenia G. Giannopoulou,
ISBN 978-953-7619-30-5, pp. 320, December 2008, I-Tech, Vienna, Austria

defining the manner in which new knowledge is acquired, what constitutes acceptability in new knowledge for the medical domain and how this can be measured and automated through the application of data mining technologies. There is a growing body of work which aims to qualify and define a process for the discovery of new knowledge across a range of domains, however this work has not focused on the unique needs of the medical domain and hence the domain has remained relatively untouched by the advances in data mining and knowledge discovery technology.

The primary challenge presented by medicine is to develop a technology that can apply a trusted knowledge acquisition process to reveal data patterns in the form of hypotheses which are based on measures that can be relied upon in medical health research and tested in a clinical environment (Ashby & Smith 2002). Due to the broad nature of medical informatics and the diversity of professional roles in medicine, the requirement is to find a solution that is both flexible enough to address the unique and varied data management and analysis requirements within the domain, whilst being specific enough to address the individual needs of an equally broad range of users. To date, automated data analysis systems have been developed for a particular role or field of investigation using a specific and relatively homogenous data set and are not transferable to other professional roles, fields or data sets within the domain. The primary requirements are therefore the provision of a methodology and system to facilitate the acquisition of knowledge which conforms to the requirements of the domain and, the production of statistically valid hypotheses which are appropriately targeted to the role of an individual user and which can provide a sound foundation for further research or clinical trials.

Data mining in medicine is most often used to complement and expand the work of the clinician and researcher by qualifying or expanding knowledge rather than providing new knowledge as is the trend in other domains. Very little health data mining is purely exploratory and hence the technology is generally not applied to provide novel knowledge i.e. to identify previously unknown patterns hidden within the data. One of the difficulties in providing new knowledge in the health domain is the need to sufficiently cross reference and validate the results. It is not sufficient to provide a standard rule in the form of A gives B in the presence of C without substantiating the information held therein. This information could already be known, may be contrary to known medical facts, incomplete due to missing attributes, may not be statistically valid by trusted measures or may simply not relate to the specialisation of the user and is therefore contextually irrelevant.

The barriers to the application of data mining to medical data can be generalised as follows:

1. The low level of flexibility in data mining systems and the need for medical analytical processes to adapt to data mining methodologies rather than data mining adapting to the needs of medicine,
2. The lack of opportunity for incorporating subjectivity when mining medical data,
3. The production of patterns in a technical language and format that are often not understandable or applicable in a clinical setting,
4. The broad range of users and analytical variance in medicine,
5. The production of too many irrelevant results, requiring a high level of user interpretation to discriminate those that are truly useful.

This chapter presents a flexible solution to these issues through discussion of a novel process to more closely reflect clinical medical processes in the technical data mining process. This is achieved through the development of two complementary systems; a

hypothesis engine and an Automated Data Pattern Translator (ADAPT) that provide a mechanism for the translation of technical data mining outputs into a language which can be better understood and accepted by medical professionals. As an integrated unit, the hypothesis engine and ADAPT are able to facilitate greater access to mining technologies, and the ability to apply some of the more complex mining technologies by all medical users without the risk of producing irrelevant or incomprehensible outputs.

The health domain is a myriad of complexity and standardised data mining techniques are often not applicable (Cios, 2002, Imberman & Domanski 2002; Hagland, 2004) hence the need for a deeper analysis of the potential for data mining in the medical domain which in turn requires knowledge of the process of medical knowledge acquisition and how the data mining technologies can facilitate this process. The remainder of this chapter aims to reduce this knowledge gap through a discussion of the processes of knowledge acquisition and diagnostic decision making in the medical domain and of the potential for novel data pattern evaluation methods to augment and automate these processes. Discussion of a collaborative two part solution developed through a merging of theories from data mining, medicine and information retrieval theory will be presented, together with experimental results to demonstrate the application of these solutions to the issues identified in this section.

2.2 The history of medical knowledge acquisition

History has taught us that standardised scientific processes have been followed for centuries to ensure that only trusted, proven knowledge is applied in a clinical setting. The process for defining what and how new medical knowledge is trusted does not readily correlate with current data mining processes. However where the two are combined the rate of acceptance of outputs is higher than in those where data mining process alone is enforced upon the medical knowledge acquisition process. To understand the similarities and differences between the two processes both processes must first be defined.

Throughout recorded history there has been debate over what constitutes knowledge and therefore what constitutes proof of knowledge. Early practitioners of medical science, such as Hippocrates, based their knowledge development in philosophy and their ability “to see with the eye of the mind what was hidden from their eyes” (Hanson, 2006) . By the first century A.D. physicians, such as Galen, were beginning to question the validity and contradictions of Hippocrates work which had stood mostly unopposed since the 5th Century B.C. It is not clear if there was any agreement or understanding of the methods applied by physicians to develop their knowledge base at this time as there was no empirical proof or scientific process documented. Galen was one of the first to suggest that there should be a process for the provision of substantiated evidence to convince others of the value of long held medical beliefs and hence raised the notion of a practical clinical method of knowledge acquisition which combined the Hippocratic concept of hypothesis development through considered thought and a priori knowledge, with clinical observation to evaluate and hence provide proof or otherwise of the hypothesis. This general methodology has survived to the present day and is reflected not only in the provision and acceptance of new knowledge but also in the process of clinical diagnosis.

The historical debate on knowledge acquisition methodologies has primarily focused on three philosophical groups; Methodists, Empiricists and Rationalists. Whilst these three groups are most frequently discussed in a Graeco-Roman context, they were either being

applied or paralleled in various other cultural contexts including India and Islam. All three of these cultural contexts are discussed here briefly to demonstrate the extent and foundations of medical knowledge acquisition debate in the ancient world.

2.2.1 The Graeco-Roman context

- Methodists

The first prominent physician practicing according to the Methodist philosophy was Hippocrates of Cos (460-380 B.C.) who is still referred to as the “Father of Medicine” (Hanson, 2006). It is believed by many that he initiated the production of over 60 medical treatises known as the Hippocratic Corpus. The corpus was written over a period of 200 years and hence had more than one author which is reflected in the sometimes contradictory material contained therein. The body of work was however consistent in its reliance on defining a natural basis for the treatment of illnesses without the incorporation or attribution of magic or other spiritual or supernatural means as had occurred previously. Methodists were defined as those whom attributed disease etiology and treatment primarily to an imbalance in bodily discharges. Illnesses were categorised by whether they represented a withholding of fluids, for example a blister holds water, or an excessive releasing of fluids, for example a weeping eye. This group founded its knowledge on an understanding of the nature of bodily fluids and developed methods for the restoration of fluid levels. They were not concerned with the cause of the imbalance or the effect on the body of the imbalance, only in recognising whether it was an excess or lack of fluid and the method for treating that observation.

- Rationalists

Rationalists believed that to understand the workings of the human body it was necessary to understand the mechanism of illness in terms of where and how it affected the body's functioning (Brieger, 1978). They were not interested in the treatment or diagnosis of illness but focused on understanding and recording the functioning of the living system. Two works are of prominence in this group (Corsans, 1997); the *Timaeus* by Plato which systematically described the anatomical organisation of the human body and; *Historia animalium* by Aristotle which discussed further both human and animal anatomy and the links between such entities as the heart and blood circulation. This method of knowledge acquisition was criticised as it effectively removed medicine from the grasp of the average man and moved it into a more knowledge based field where philosophical debate or an observational experiential approach was not deemed sufficient (Brieger, 1978). Essentially Rationalists did not believe in a theory unless it was accompanied by reason. They espoused the requirement for knowledge to be founded on understanding both cause and effect of physical change in the body (Horton, 2000).

- Empiricists

The Empiricists believed that it was not enough to understand how the body works and reacts to illness. They pursued a philosophy which stated that it was necessary to demonstrate the efficacy of treatments and provide proof that a treatment is directly responsible for the recovery of a patient rather than providing academic argument regarding why it should result in recovery. Galen is considered to be one of the earliest empiricists (Brieger, 1978). He was both a medical practitioner and a prolific scholarly writer and is certainly one of the best known and more frequently quoted empiricists. He was particularly interested in testing the theories proposed in the Hippocratic Corpus, especially

given its frequent contradictions. His work was also produced at a time when medicine as a science was evolving from its previous status as a branch of philosophy. In his work Galen argues that “medicine, understood correctly, can have the same epistemological certainty, linguistic clarity, and intellectual status that philosophy enjoyed” (Pearcy, 1985). Empiricists were the first to concentrate on the acquisition of knowledge through demonstrated clinical proof developed through scientific methodologies which provided conclusive statements of cause and effect.

2.2.2 The islamic context

- The Empiricists (Ashab al-Tajarib).

Dr. Mahdi Muhaqqiq was an early 20th century Iranian scholar who wrote texts on many subjects including medical knowledge acquisition throughout the history of Iran. He recorded that the early Empiricists believed that medical knowledge was derived from experience obtained through the use of the senses and that the knowledge is comprised of four types; “incident (ittifaq), intention (iradah), comparison (tashbih) and the adoption of a treatment that was used in another similar case (naql min shay' iki shabihhi)” (Muhaqqiq, 2007).

- Incident - this can either describe a natural event such as a sweat or headache, or an accidental event such as a cut or a broken limb.
- Intention - denotes an event experienced by choice for example taking a cool bath to reduce a fever.
- Comparison - A technique employed by a practitioner whereby he notes that one of the above techniques results in a useful effect which can be applied to other similar presentations. For example applying cold water to reduce localised burning of the skin following the observation that a cool bath can reduce generalised fever or body heat.
- Naql - A technique whereby the physician applies a treatment for a similar presentation in the instance of a presentation which has not been encountered before. An example might be the prescribing of a medication for a previously unencountered infected tooth where that medication had only previously been used for an infection elsewhere in the body.

The empiricists treated a patient through knowledge of that patient's demographics and therefore all patients of a certain age and sex with some similar complaint were treated the same whereas patients of the opposite sex may have been treated differently even though the condition was the same. Their knowledge was based on patient characteristics rather than a specific condition or set of symptoms. Whilst this seems to differ from the Graeco-Roman definition of empiricism, both groups believed that knowledge acquisition occurred through observing or testing the effect of a treatment and producing rules based on what is considered reliable empirical proof rather than conjecture and debate.

- The Dogmatists (Ashab al-Qiyas).

The Dogmatists believed that while scientific belief and knowledge should be derived from experience and observation this should be tempered by the use of thought and considered evaluation (Mohaghegh, 1988). They believed that changes in the bodily functions must be precipitated by some event and that it is necessary to not only understand what these changes are but also what the specific causes of those changes are in order to correctly diagnose and treat any condition. Changes are defined as being of two types (Muhaqqiq, 2007):

- Necessary change - drink reducing thirst. This is a change which is required for normal bodily functioning.
- Unnecessary change - dog bite causing bleeding. This change is not a requirement to aid or enhance bodily wellbeing.

Dogmatists based their treatments upon the nature of the condition rather than the type of patient as seen with the empiricists. The treatments were therefore selected through knowledge of the causes of illness and the effects of those treatments upon the illness or symptoms. This required an understanding of the physical body and the changes that result from illness in a similar manner to that of the Graeco-Roman Rationalists.

- The Methodists (Ashab al-Hiyal).

This group believed in a generalist view of illness and treatment and categorised conditions in terms of the extent to which bodily fluids and wastes are either retained and/or expelled. Treatments were generally natural remedies based upon adjusting the balance between such aspects of life as food and drink, rest and activity, etc. Methodists were not interested in the type of patient or cause and effect of illness and were hence considered to be more prone to error (Muhaqqiq, 2007). This is in direct parallel to the Methodist philosophy discussed in Section 2.2.1.

It has been suggested that in general, Islamic physicians relied primarily upon analogy which reflects their focus on logic in other scholarly areas (Mohaghegh 1988). This has resulted in widespread support for the Dogmatist methods of knowledge acquisition through research and understanding of cause and effect in the human system. However there is still debate between scholars with some believing that Dogmatism alone is the only method of ensuring progress in medical diagnosis and treatment as it is the only method which tries to seek new understanding rather than relying upon past experience or a closed assumption that there is a single cause for all illness (Mohaghegh, 1988). Others prefer to adhere to the Graeco-Roman perspective (developed by Plato) that a combination of experience and analogy is required if a holistic, 'correct' practice of medicine is to be achieved (Muhaqqiq, 2007).

2.2.3 The Indian context

India is not well known for its scientific contributions or texts, however it has a long history in the development of medical knowledge. In the 11th century a Spanish scholar, Said Al-Andalusi, stated that he believed that the Indian people were "the most learned in the science of medicine and thoroughly informed about the properties of drugs, the nature of composite elements and the peculiarities of the existing things" (al-Andalusí, 1991). The reasons for this apparent invisibility of Indian scientific progress may be due to religious debate in India which has frequently negated the influence of scientific explanation instead preferring to rely upon mystical or spiritual beliefs. There are however documented scientific approaches to the development of a body of knowledge regarding medicine from centuries before the texts of Hippocrates and which, although often earlier, discuss similar theories to those presented in the Graeco-Roman texts.

- The Rationalist schools

One of the earliest groups to produce texts concerning the acquisition of knowledge regarding the human state were the Upanishads which were believed to have been written between 1500 and 600B.C. and were concerned with knowledge regarding the spirit, soul and god (Tripod, 2002, Kaul & Thadani 2000). Although these texts were embedded in

mysticism and spirituality they used natural analogy to explain the notion of the soul and god and allowed the expression of scientific and mathematical thought and argument which formed the basis for the emergence of the rationalist period. Early rationalists included the Lokyata, Vaisheshika school and the Nyaya school. These groups espoused a scientific basis for human existence and a non-mystical relationship between the human body and mind. They also developed primitive scientific methodologies to provide "valid knowledge" (Tripod, 2002, Kaul & Thadani 2000).

The Lokyata were widely maligned by Buddhist and Hindu evangelicals as being heretics and unbelievers due to their refusal to "make artificial distinctions between body and soul" (Kaul & Thadani 2000). They saw all things in terms of their physical properties and reactions and gave little attention to metaphysical or philosophical argument, preferring to believe only what could be seen and understood. They developed a detailed understanding of chemistry, chemical interactions and relationships between entities. They are also believed to be the first group to document the properties of plants and their uses, this provided an elementary foundation for all pharmaceutical knowledge which followed.

The primary input of the Vaisheshika school toward the progression of human knowledge was their development of a process for classification of entities in the natural world and in their hypothesis that all matter is composed of very small particles with differing characteristics (Tripod, 2002). Their theory stated that particles, when combined, give rise to the wide variety of compounds found upon the earth and allowed them to be classified by the nature of the particles from which they were formed. This school also introduced the notion of cause and effect through monitoring and understanding temporal changes in entities. The importance of this work lay in the application of a methodology for identification and classification of relationships between previously unconnected entities. This early recognition of the need for a documented scientific process provided a mechanism for the schools which followed to present substantiated proof of evidence for theories in the sciences including physics, chemistry and medicine.

The Nyaya school further developed the work of the Vaisheshika school by continuing to document and elaborate a process for acquiring valid scientific knowledge and determining what is true. They documented a methodology consisting of four steps (Tripod, 2002):

- Uddesa was a process of defining a hypothesis.
- Laksan was the determination of required facts "through perception, inference or deduction".
- Pariksa detailed the scientific examination of facts.
- Nirnaya was the final step which involved verification of the facts.

This process would result in a conclusive finding which would either support or refute the original hypothesis.

The Nyaya school also developed definitions for three non scientific pursuits or arguments which were contrary to the determination of scientific truth but which were often applied by others to provide apparent evidence for theories or knowledge (Tripod, 2002, Kaul & Thadani, 2000). These included jalpa to describe an argument which contained exaggerated or rhetorical statements or truths aimed at proving a point rather than seeking evidence for or against a point; vitanda which aimed to lower the credibility of another person and their theories generally through specious arguments; and finally chal, the use of language to confuse or divert the argument.

Further to this again a set of five 'logical fallacies' were developed:

- savyabhichara - denotes the situation where a single conclusion is drawn where there could be several possible conclusions,
- viruddha - where contradictory reasoning was applied to produce proof of the hypothesis,
- kalatita - where the result was not presented in a timely manner and could therefore be invalidated,
- sadhyasama - where proof of a hypothesis was based upon the application of another unproven theory, and
- prakaranasama - where the process simply leads to a restating of the question.

These concepts were unique in their time but many remain applicable in modern scientific research.

This section has demonstrated that the quest for new medical knowledge and a deeper understanding of the human system is not a recent initiative but in fact one which has its foundations up to four centuries B.C. While there were several distinct cultural groups all were primarily concerned with defining the most reliable methodology for evaluating what knowledge could be trusted and applied clinically. The Graeco-Roman and Islamic practitioners were concerned with the means by which evidence was obtained and the Indians were more concerned with methods for proving the validity of knowledge after it had been discovered. Both of these foci remain topics of debate in the 21st century and as late as 1997 a report was published by the International Humanist and Ethical union regarding trusted versus untrusted clinical practices and the requirement for proof of the benefits of medical treatments. The opening of a Mantra Healing Centre at the Maulana Azad Medical College in New Delhi was described as “ridiculing the spirit of inquiry and science” through its application of “sorcery and superstition in their rudest form” (Gopal, 1997). The report did not however argue that there was no worth in mantra healing but that there was no proof of worth as per the requirements of the still flourishing rationalist opinion. The debate on what is trusted and clinically applicable knowledge forms the focus of this chapter as it investigates the application of new knowledge acquisition tools and aims to identify best practice procedures for automating the acquisition of new medical knowledge.

2.3 Non-scientific knowledge acquisition

History has shown that the acquisition of much currently accepted medical knowledge was based on serendipity or chance accompanied by a strong personal belief in an unproven hypothesis. Further to this, much knowledge was acquired through a process which directly contradicts accepted scientific practice. Whilst there was usually a scientific basis to the subsequent development of proof this was often produced through a non traditional or untrusted application of scientific processes. Unfortunately this frequently resulted in lengthy delays in acceptance of the work. The following list provides a range of such breakthroughs over the past 250 years which can be attributed to little more than chance, tenaciousness and the application of often radical methods to obtain proof.

1. James Lind (1716-1794) (Katch, 1997). Based upon an unsubstantiated personal belief that diet played a role in the development of scurvy on naval vessels, Dr Lind performed limited randomised trials to provide proof and then published his Treatise on the Scurvy which is still relevant to this day.
2. Edward Jenner (1749-1823) (Sprang, 2002). During his apprenticeship Dr Jenner overheard a milkmaid suggest that those who have had cowpox can not contract

smallpox. He then tested the theory by infecting a young boy sequentially with each pathogen and as a result created the concept of a vaccine and initiated the global eradication of smallpox.

3. John Snow (1813-1854) (Ucla, 2002, BBC, 2004). Dr Snow believed, without any direct evidence, that the transmission of viral agents was possible through contaminated water. In 1854 he applied the theory and provided an answer to the cholera epidemic.
4. Alexander Fleming (1881-1955) (Page, 2002). Dr Fleming stumbled upon a discarded culture plate containing a mould which was demonstrated to destroy staphylococcus. The mould was isolated and became the active ingredient in penicillin based antibiotics.
5. Henri Laborit (1914-1995) (Pollard, 2006). During his ward visits Dr Laborit noticed that patients given an antihistamine named promethazine to treat shock not only slept but reported pain relief and displayed a calm and relaxed disposition leading to the development of medications to treat mental disorders including schizophrenia.
6. Robert Edwards and Patrick Steptoe (1925 - , 1913-1988) (Swan, 2005, Fauser & Edwards, 2005). These doctors were the first men to deliver a baby through in-vitro fertilisation after 20 failed attempts and great ethical debate following a lack of proof in animal subjects.
7. Barry J Marshall (1951-) (Marshall, 1998). Dr Marshall worked against accepted medical knowledge to provide proof of the bacterial agent, *Helicobacter Pylori*, as the cause of stomach and duodenal ulcers. So strong was the opposition to initial clinical testing of the theory he resorted to using himself as the test subject.

Whilst each of these examples provided wide reaching benefits to human health and contributed significantly to the body of medical knowledge in some cases, they would not have been possible if only standardised scientific methodologies had been applied using only trusted traditional processes. This demonstrates that there is often a need to do things differently and not only apply what is comfortable and safe to enable the acquisition of knowledge, although there is always a requirement to provide substantiated proof and an argument based upon scientific principles. The applicability of this notion is particularly relevant to this chapter which focuses on the application of new techniques and technologies which have demonstrated an ability to provide an important impetus to the acquisition of knowledge in other domains and which have not been demonstrated to be detrimental to the process in the medical domain. However, the same proof of hypothesis hurdles must be overcome and an equally strong argument and testing methodology must be provided for the resulting knowledge to be accepted. Throughout history the same quality of evidence has been required and the omission of this evidence has often resulted in decades of latency between hypothesis statement and the generation of conclusive evidence in support (or otherwise) of that hypothesis.

Regardless of the methodology for producing the evidence required for knowledge acquisition, the above examples all had to fulfil a number of further requirements prior to the acquired knowledge being accepted. These requirements are summarised following:

1. Replication of results.
2. Non contradictory results.
3. Scientifically justified theories and hypotheses.
4. Ethical methodologies and measures.
5. Results demonstrated to be representative of the population.
6. Results derived from sufficient numbers of cases.
7. Publicly documented processes and results.

2.4 The application of data mining

Medical history has recorded many instances of the manual use of data mining techniques resulting in medical breakthroughs crucial to the preservation of thousands of human lives if not entire populations. Over centuries medical professionals have (often unknowingly) employed the same scientific analytical methods to data as are applied during data mining in order to develop hypotheses or to validate beliefs. Whilst these techniques have been applied in a simplistic form they clearly demonstrate the applicability of the founding principles of data mining to medical inquiry and knowledge acquisition. A number of the examples discussed in Section 2.3 are used to demonstrate this.

- Data sampling - James Lind (Katch, 1997). Dr Lind performed small randomised trials to provide proof of the cause of scurvy. In his position as Naval doctor he was able to test his theories on the crews of the vessels he sailed on, however without documented proof it was not possible to test the entire navy en mass. Developing sufficient proof in this manner was a lengthy process and it was 50 years before the British Admiralty accepted and applied his theories, a delay which cost the lives of many sailors. Lind's process shows the use of examining subsets of the population, being able to clearly identify the variant in the knowledge gained and then substantiating that knowledge by testing on similar populations to ensure the finding is representative is a suitable technique for hypothesis testing and knowledge substantiation.
- Association rules and support and confidence heuristics - Edward Jenner (Sprang, 2002). Following development of a hypothesis from the knowledge that milkmaids were less likely than members of the general population to develop small pox due to their increased contact with cow pox, Jenner conducted further tests over a period of 25 years to validate the relationship and publish his findings. This work demonstrates the use of the concept of support through Jenner's realisation that there was a frequently occurring and previously unknown pattern in a data set or population. That pattern was subsequently tested to provide confidence levels by showing that contracting cowpox almost always results in an inability to contract smallpox.
- Clustering - John Snow (Ucla, 2002, BBC, 2004). In his investigation of the cholera outbreak of 1854 Dr John Snow applied a meticulous process of interviewing to collect data. He used the information collected to develop a statistical map which clustered interview responses based upon the water pump which supplied water to the individual. This revealed that every victim had used a single supply of water and no non-sufferers had used that supply. Further investigation showed that this pump was in fact contaminated by a nearby cracked sewage pipe. This shows not only the power of the use of medical data for statistical purposes, but the benefits that can result from applying clustering techniques to that data.
- Association rules and classification - Henri Laboit (Pollard, 2006). Laboit extended the use of promethazine to treat mental disorders including schizophrenia by realising patterns in side effects from administering the drug during surgery on non mentally ill patients. This was achieved through identifying association rule style patterns to describe associations between focal and non focal attributes, for example combinations of relationships between diagnosis, treatment, symptoms, side effects and medications. Analytical techniques were employed to classify conditions exhibiting similar patterns of presentation and clinical testing was utilised to demonstrate the effect of applying an identified drug to control those classes of symptoms.

These techniques until recently were employed manually and hence were on a much smaller scale than we see today through the application of automated data mining systems, however they demonstrate the impressive potential for automated data analysis techniques to be applied with greater benefits and applicability than ever thought possible. There is a belief by some that the rate of medical breakthroughs of the calibre of those listed above has slowed dramatically since the 1970s (Horton, 2000). This could be attributed to the inability of the human mind to manage the volume of data available (Biedl et al., 2001, Lavrak et al., 2000) and that most if not all patterns in data which may reveal knowledge and which occur frequently enough to be noticed by the human analyst are now known. This adds significant weight to the argument for the application of more effective and efficient automated technologies to uncover the less visible knowledge or less frequent but equally important patterns in the data. We must however learn from history and ensure that the validation requirements for knowledge acquisition, as discussed previously, are adhered to by any automated process as for all other methods of knowledge acquisition even though this has been described as “the hardest part of the expert system development task” (Lavrak et al., 2000). Too often in recent work there has been a focus on developing new methods for determining the quality and value of outputs which does not take into consideration the many lessons we can learn from history, and is often little more than a process of reinventing an already rolling wheel (Ordonez et al., 2001).

3. The automation of medical knowledge acquisition

To automate it is not sufficient to simply understand how the process has occurred manually or how that process developed, although this is important in ensuring the results can be trusted. There is also a requirement to develop a seamless transfer from the clinical application of the process to the technical automated application of the same process and to provide accountability so that the results are trusted, justified and actionable in a clinical environment. In data mining systems, these accountability values are often measured through a concept termed ‘interestingness’ which essentially aims to measure the level of interest a user may have in the outputs provided by a system

3.1 What is interesting?

The term ‘interest’ is one which is widely accepted and used within data mining to denote output which has value on some level. This section is concerned with how we might define interest or value in clinical terms. The work discussed in this chapter began with the naive idea that interest can be measured and quantified for medicine as a homogenous entity as occurs in many other domains. This notion is now considered laughable at best. Medicine is not a homogenous entity and is not even a single entity in any contextual argument. Whilst it is defined generally (Oxford, 2007) as the science of studying, diagnosing, treating, or preventing disease and other damage to the body or mind, or treatment of illness by the ingestion of drugs, modification of diet or exercise, or other non-surgical means. The means by which those engaged in the practice of medicine achieve this outcome or engage in this activity varies widely depending upon individual needs, experience and data. If we are developing an automated system to assist or guide in this activity then the system must also be able to perform according to individual needs, experience and data. As discussed in the chapter introduction, data mining systems are able to manage a broad range of data types and analytical processes, but they are usually tailored to a user and hence are designed to

work with their individual technical proficiency and analytical requirements. Whilst work is ongoing in this area, uncertainty remains regarding the ability of automated systems to address the varied and fluid requirements of the user population. There are therefore two questions posed:

1. How can we identify the varied individual requirements of the user body?
 2. Can we automate these requirements into a system which caters to a range of users?
- To overcome this issue it is necessary to define what is interesting and what makes one thing more or less interesting than another and further to develop an algorithm to measure the extent to which the outputs of data mining conform to the user definition of interesting. As each of us finds different things interesting it is not possible to define a single specification for what is interesting to any group of people. Within the context of the focus of this chapter however, we can say that the degree to which something is deemed interesting can be quantified through the application of statistical methods. This method of value measurement is one which is applied in clinical testing where such measures as those shown in Table 1 are frequently used as a basis for determining which trial or trial arm has provided evidence that is clinically applicable or worthy of progression to the next level of testing (Gebeski & Keech 2003, Moser et al., 1999, Hilderman & Hamilton 2001, Geng & Hamilton 2006). Many of the same statistical methods are also applied within data mining systems but the provision of individual methods or combinations of methods are fixed and provide an inflexible analytical toolbox unlike in the clinical setting where any method can be chosen and applied. The requirement is therefore to provide a more flexible approach and not 'invent' another formula to determine the level of interest but to allow the user to determine how to define interest for each run or data set as occurs in clinical analysis. An added benefit in this is the ability to reinforce and support medical professionals control of their domain and the processes they apply to an analysis task. It should not be acceptable for technology to dictate how a medical professional (or any other) should practice.

Measure	Type	Application	Domain
P	Statistical	Determine degree of difference in results	Medicine
chi2	Statistical	Determine degree of difference in results	Medicine
chi2	Statistical	Result comparison	Medicine
Pearson's Correlation	Comparative	Measure of difference	Bio-inf
Euclidian Distance	Comparative	Measure of difference	Bio-inf
Cosine Similarity	Comparative	Measure of similarity of text	Linguistics
Support	Statistical	Probability, Frequency	Retail
Confidence	Statistical	Probability, Frequency	Retail
Accuracy	Domain	Determination of class membership	Medicine
Sensitivity	Domain	Measure of ability to find true positives	Medicine
Specificity	Domain	Measure of ability to reject true negatives	Medicine

Table 1. Some of the more commonly applied statistical tests.

In essence if the results of analysis are to be deemed interesting they must fall within defined thresholds as measured by one of more statistical method. The selection of appropriate methods is both objective and subjective. Objectively, certain qualities must be present in an interesting result and methods will be selected to provide evidence of this quality. For example if the result must be indicative of an accurate prediction of disease classification then such measures as sensitivity and specificity could be chosen as they are designed to measure this quality. Subjectively an analyst may place greater trust in one method over another due to experience or availability even though they both provide a method for measuring a particular quality. An example here may be the choice between using a p-value or χ^2 which can both provide similar quantified evidence. There is a long list of statistical methods which can be applied and they are selected based upon a combination of user objectivity and subjectivity which will potentially be different for each user. Although the provision of a fixed subset of available statistical methods represents the commonly applied data mining technique for evaluating interest, for the reasons discussed herein, it is not an appropriate methodology for the medical domain where there is no fixed notion of what is interesting or of how to measure the qualities which define interest. Whilst the process would be vastly simplified by ignoring the concept of individual interest it is necessary to overcome a number of issues with the application of data mining to medical data. The most important issue is that of non acceptance of the technology even though its benefits are many. Non acceptance is due to a defined set of factors:

- The complexity of medical data frequently results in a huge number of results; too many to be evaluated by the human user and a method for reducing the results to only those of greatest interest is necessary (Roddick et al., 2003; Ordonez et al., 2001)
- Each user has a trusted set of methods which are applied during clinical analysis and which are rarely seen in a system that is not purpose built for that user or their analytical requirements; this increases the cost of providing the technology, the frustration in trying to use the technology and allows the technology to dictate the analytical process rather than the other way around. Being able to facilitate and apply a range of interest definitions in a single system would open up the technology to a greater audience.
- Many users are not technically adept and do not trust something they do not understand; the provision of a recognised process which offers a personalised perspective within a generalised framework provides comfort and security.

To overcome these barriers to the technology it is necessary to move away from a user focussed approach, which has in fact created many of the issues presented here, and towards an interest based approach which is guided by individual needs as it is the level of interest to each user that primarily determines the acceptability of results. If we can develop a generic approach to interest that can be individually adapted then we can apply this to develop data mining systems which can be similarly founded upon a generic principle that can be personalised for individual use. The provision of such a solution is the focus of the remainder of this chapter.

3.2 A role based approach

The issue of developing a flexible data mining system with the intention of enabling the generic production of results with an acceptable level of interest for each individual user has been the focus of recent work by the author. The approach has been to investigate the

concept of a role in various forms and its relationship to the concept of interest as defined above.

Each of us plays many roles in our daily lives as a student, doctor, nurse, teacher, parent, guitar player, amateur photographer etc. etc. Each of these roles will define a level of interest in the world around us, which will be determined by how relevant to each role the world is at any specific time. As each of us has a unique set of roles then each individual will have a corresponding unique set of interests or interest triggers, and different information will appeal to us at different times. Our role is therefore defined by how we measure the value of our interest in information that is presented to us. For example an Oncologist would most likely have been interested in an article on new cancer treatments in a Weekend Australian newspaper entitled "Hype or Hope?" (Cornwell, 2005), as it relates to their professional role and, even an evaluation of keywords contained in the article, would have revealed many matches to a similar evaluation of keywords in their set of interests. In contrast an Electrical Engineer would most likely not have had the same level of interest unless they also held the role of cancer sufferer or carer of a cancer sufferer for example. Therefore we can define a role as being a collection of quantifiable interests. The focus of work presented here has been to develop a system that will allow the identification of these interest sets and develop a method for determining how to measure and evaluate how strongly the information is able to trigger interest. To achieve this it was necessary to provide a quantitative evaluation of interest by evaluating the requirements of an interest set and measuring the applicability of the information or data mining results to that unique set.

3.2.1 The application of role

The application of a role in determining which data mining outputs are of relevance is more complex than simply looking for the presence of keywords. In the field of epidemiology for example the simple presence of the word 'flu' is not sufficient to trigger interest, there needs to be statistical augmentation to the information. In particular it needs to be shown that the incidence of the condition is sufficiently different to that expected for a population at a defined time. The difficulty is in determining which heuristics will give an acceptable measure by which we can include or exclude results for each role. A role based result evaluation engine has been designed in preference to other options including new heuristics and new heuristic combinations as these fixed solutions can not provide a generic answer to the issue of evaluating the level of interest for the health domain as a whole given the complexities noted above. It was necessary for an evolution in current thinking in the area and for single-user solutions to be discarded. A generalised solution to the issue of result reduction for this domain cannot be achieved due to the broad range of roles and requirements to be addressed. Early work has since evolved into developing a system that can incorporate the range of roles without the need for a separate system for each. As it is the role that determines how the strength of general interest is measured, it was a natural step to discriminate the analysis by the role of the current user as defined by a set of measurable interests. Users should be able to analyse and focus their data mining outputs using a single system regardless of their speciality or analytical requirements. Whilst the needs of each role are unique there is also considerable overlap and the heuristics required to determine interest strength varies from role to role and also within each user role depending upon the nature of the analysis being undertaken (Kuonen, 2003; Bresnahan,

1997; Tan et al., 2002). A system with a high level of flexibility and methods to facilitate user definition was deemed to provide the best use of resources to accommodate this.

Role based access models have been successfully implemented in a wide range of domains and have demonstrated an ability to overcome issues such as those seen in health including a need for careful management of sensitive information and the need to provide enterprise level security policies which discriminate on a local level based on the role of the user (Cabri et al., 2004; Ferrailo & Kuhn, 1995). Role based access models have provided a fixed framework from which to apply highly flexible system definition and this concept was the major attraction in the creation of a role based results evaluation process for data mining applications in the health domain.

The following features of role based systems have been adapted and incorporated into a hypothesis engine as discussed in Section 4.

- The accommodation of roles that allow for overlapping requirements and measures.
- The ability for a user to have more than one role at a time.
- The ability to enforce constraints on data access where required to accommodate ethical sensitivities.
- The ease of modifying the role of a group of users to accommodate new technologies or methodologies.
- The ability to constrain at a global level and provide flexibility at a local level.

Whilst the choice of interest strength heuristics will often fluctuate little across mining runs, some vary greatly, become redundant or require supplementation by new or existing measures. This high level of flexibility is not currently available in documented data mining systems. Algorithms and selected heuristics are applied singly or in a fixed combination with others within a specific system designed for a specific use. The ability to utilise the interest role as a means of selecting and applying a significant number of the range of measures in a unique combination as required has not been documented and is believed to be a novel approach to the issues presented here. Support for such an approach has been provided by health domain professionals and domain based publications including the Medical Journal of Australia (MJA) which stated that appropriate statistical methods for analysing trial data are critical and suggested that the statistical methods used in each trial should be specifically tailored to each analysis (Gebiski & Keech, 2003). Each specialist field and role has its own requirements and hence the level of flexibility in data mining software packages must be equally flexible, open to adaptation and tailored for the user role at run time.

As discussed in the previous section, there are many methods which can be applied to measure the strength of interest a user may feel towards any information. An initial aim was to group these measures based on the role that uses them, this was rejected due to the overlaps discussed earlier. The aim was thus modified to group the measures into classes based upon their type and the characteristic of interest they are able to quantify, thus allowing each role to select and apply them as required. While there is no fixed notion of what defines interest strength for a particular role in each instance, there is agreement on the characteristics that indicate strength and these can be grouped and measured to quantify their level of expression in results presented. These classes were verified in discussion with a range of medical professionals during a work in progress seminar (Workshop in population health, 2005) presented to by the author. In attendance were medical specialists from several fields including cardiology, epidemiology, biostatistics, nursing, clinical

research and government and all agreed that the proposed classes defined the qualities they looked for during hypothesis development and results testing. It was noted that whilst most of those present could not adequately describe what determined a strong interest in a result for the domain generally, it was felt that the classes presented would provide an acceptable quantification for any role within the domain if applied uniquely for each role or field. It was also proffered that each test result was considered individually depending upon their needs at the time often the heuristics employed were often not selected until the time of evaluation thus strengthening the argument away from a generalised approach. This, in fact, emphasised the need to utilise traditional measures but in a flexible combination for each evaluation.

By allowing a subjective selection of heuristics and evaluating their application objectively it is possible to take the outputs of data mining and measure their value uniquely and flexibly for each role rather than utilising a unique but fixed sub set of heuristics for each system. Based upon the values achieved by each heuristic, unqualified mining outputs can be eliminated from presentation thus providing only those outputs which meet the requirements of the role and adequately contain the desired characteristics to be of interest.

Six classes or criteria are provided for interest strength measurement and each of these may require a number of statistical, comparative or other tests to determine the overall strength for each criterion. The individual values are then combined to provide a comprehensive measure of interest strength for a mining output based on the total requirements for the user role. A greater strength suggests an output that is more likely to be of value to the role that defined the heuristics and their scopes. The criteria for measuring interest and hence strength of new information or knowledge patterns produced through data mining are discussed following.

- **Novelty** - Is it unknown in the body of domain knowledge? This is more complex than simply not duplicating existing knowledge or presenting expected patterns. New patterns based upon existing knowledge may still be of interest if the strength or content of the new knowledge differs sufficiently from that which is expected. For example, medical professionals would reject as new a pattern which states that 3.6% of pregnant women develop gestational diabetes mellitus (GDM) as this is known and expected knowledge even though it would have sufficient strength by some traditional measures to warrant further investigation (Stone et al., 2002). However if a pattern were to report that the incidence rate of GDM in a data set primarily for a North Asian population was 3.6% then the interest in this may be greater as the rate would be expected to be higher. Hence it is the pattern novelty as a whole which is being evaluated and which thus determines the strength of interest. There are a number of measures that can be applied to quantify the expectedness or similarity of hypotheses to existing knowledge and it may be necessary to test this criterion using several classes of tests to adequately assess the novelty of a pattern.
- **Applicability** - Is it relevant to the current user? This infers that either some contextual information is required, or that previous patterns are tagged as interesting (or not) so that the system can learn and reference. The definition of applicability (or relevance) is context based and should be maintained on an individual level. An outlier that is strong in every aspect except for prevalence may not be relevant to an epidemiologist as it is not representative of the population but still may potentially be of interest to a clinical specialist or medical researcher and should be tagged for reference by that role. The implication is that any derived pattern produced from a medical data store is

potentially valuable to some role in the medical domain. If accepted, this suggests the importance of strength determination at a role based level to ensure that each role sees only patterns they are most likely to have an interest in and be able to act upon but that no strong pattern is omitted completely from consideration.

- **Relativity** – Is it valid relative to the data from which it originated or a class of object that it describes? Once again the applicability of this criterion is determined by the context within which it is measured. Within epidemiology it is important for pattern to be shown to be applicable for a generalised population. Results therefore need be demonstrated to apply across the human race or a definable sub section of it. A recent study published in the MJA discussed a potential but low correlation between passive smoking and breast cancer (Smith & Ebrahim, 2004). Whilst the link was biologically plausible in 1999 it was not deemed to be representative of the female population in an epidemiological sense and hence was not deemed interesting. Further work was done which focussed on the effect of environmental tobacco smoke across the age variable specifically. It is now accepted that there is enough evidence to suggest that passive smoking specifically in the early years of a females' life has a measurable impact upon the incidence of breast cancer later in life. Investigation at a finer granularity resulted in a hypothesis that is accepted as representative of a defined sub section of the population. This suggests that strength should be measured for all applicable classes, not only the most obvious or highest ranking. Patterns also need to be shown to be representative of the data set from which it came and there are standardised checklists such as that provided by CONSORT (Consolidated Standards of Reporting Trials) (Lord et al., 2004, Gebeski & Keech 2003, Altman et al., 2001) which are widely used within medical research and should be incorporated into the planning of data mining systems.
- **Provability** – Can it be proven through clinical testing? This reflects the actionability of the outcomes of data mining and incorporates the need to adhere to guidelines such as CONSORT discussed earlier. Whilst there are perceived difficulties in automatically determining what could be tested clinically, there are several requirements which define what the foundations of a clinical hypothesis should be and these should be present in hypotheses in the form of patterns derived through data mining also (Lord et al., 2004). For organisations that adhere to research guidelines, it is important that the pre-requisites are met for further work so that the potential for follow up clinical testing is not prevented. This criterion aims to ensure that potential hypotheses are not rendered inactionable due to the methodology employed for their derivation rather than trying to determine what will be actionable.
- **Understandability** – Can it be understood through appropriate presentation? New knowledge that cannot be described easily or accurately is of little use. The inability for the human brain to assimilate and perform functions upon large amounts of complex data is the very foundation upon which the field of data mining was based. When presenting patterns, this must be given due consideration. An overly complex or lengthy pattern may be overlooked in favour of those that can be read and understood quickly. Consideration must also be given to domain specific terminology and semantic hierarchies (Ashby & Smith, 2002). This will ensure that patterns are presented using uniform, accurate and appropriate terminology (Bojarczuk et al., 2001, Lavrak et al., 2000). Results should also be presented via a medium that is accepted as standard by each role or domain and there is a body of work in the fields of visualisation and linguistics that is attempting to address some of these issues.

- Validity - Is it statistically valid according to trusted domain measures? There are a wide range of statistical measures available to test these classes and each user role should be able to apply measures to each analysis based on the nature of the analysis and personal experience. The authors work argues for the use of role based metrics which are manipulated and utilised according to the individual needs of each user and suggests that pre-defining specific heuristics for statistical validity is redundant and archaic.

3.2.2 Concept formalisation

Whilst it is not possible to give a single definition to what is interesting it is possible to formalise the nature of interest as described in this chapter. As discussed, interest can be determined by a variety of objective and subjective criteria which in combination can provide an indication of the degree to which this output can be trusted. We can therefore formalise interest as following:

$I = \{ m1 \dots mn \}$ where I is Interest which is defined by a set of statistical methods or other heuristics (m).

Furthermore the applied heuristic or method (m) can be denoted in the following form:

$m = \{ \text{metric}, T_{\min}, T_{\max}, \text{var1} \dots \text{varn} \}$ where $T_{\min}$ is the minimum acceptable threshold for the metric and $T_{\max}$ is the maximum acceptable threshold for the metric and var is any other variable which may affect the application of the metric. A var example may be 'sex' and the metric 'weight', therefore there would be different $T_{\min}$ and $T_{\max}$ for weight thresholds for men and women. Thus the method is described as an object which has a number of qualities through which it can be defined or represented. All qualities except for metric name would be optional.

By applying the definition of interest (I) we can determine that the outcomes of data mining are likely to be trusted as they apply trusted metrics and are likely to be understood as again they are qualified using known metrics whose outputs are in a standardised form. The selection of these metrics, thresholds and variables should be done by each user at run time to allow for a measure of subjectivity to be applied to the definition of interest.

Further to this, it is necessary to determine how interesting an outcome might be, as one pattern may not adhere to all parts of the definition but should not necessarily be excluded on that basis. Many things are interesting not because they adhere to our schema but because they 'almost' adhere to the schema or in contrast because they do not at all adhere to the schema (RoddickRice 2000). It was therefore deemed useful to give some indication of how close to the interest schema the pattern is as defined by the metric set. This binary classification is formalised as:

$PI = \{ m1r \dots mnr \}$ where r denotes that this is the result of applying metric m to a pattern P . $m1r$ to mnr were coded as either 1 or 0 depending upon whether or not they fell within the thresholds with (1) denoting within and (0) denoting the result was outside of the threshold for that method. PI is an expression of the likely interest strength in a pattern based upon consideration of all applied metrics.

Whilst this was a satisfactory representation for results which had to absolutely comply with the stated thresholds and are hence critical to acceptance, it provided no indication of where results were outside of, but close enough to, the required thresholds as determined again by a user measure of acceptable flexibility. There are numerous statistical measures and methods available to test a pattern (Imberman & Domanski, 2002; Beals et al., 1999), however they may not be critical measures of pattern interest strength to a particular user.

Hence there is a need to define both needs and wants and be able to discriminate on that basis. For example, a confidence of between 80 and 95% may be wanted, however patterns with a lower confidence may still be of interest if other heuristics achieve acceptable levels, hence confidence would be considered a flexible measure. Conversely in an epidemiological context the incidence of a condition would need to be greater than background levels to be of interest, and anything less than or equal to background levels will not be interesting regardless of other factors. These heuristics are deemed to be critical measures as the scope for acceptance is inflexible.

For those methods that could be applied flexibly a fuzzy logic was applied to allow for a more considered pattern evaluation to occur. This fuzziness allowed for a result to be described in terms of four classes; far from interesting (f); close but lower (cl); close but higher (ch); and within thresholds (wt) thereby providing for a range of interest outcomes as follows:

$$f \dots cl \dots wt \dots ch \dots f$$

f appears at both ends of the scale as it can be uninteresting because it is too far above the stated thresholds or because it is too far below the stated threshold, however in some cases for example in pathogen monitoring it may be that both f's and/or values for cl and ch may not be applicable as a user may not have defined both a T_{min} and T_{max}.

By applying all of these constructs it may be possible to provide a meta description of a pattern as an array in the form of A and B appear in the presence of C {(M1 (metric, T_{min}, T_{max}, var1 ... varn)(wt)), (M2(metric, T_{min}, T_{max}, var1 ... varn)(1)), Mn(metric, T_{min}, T_{max}, var1 ... varn)(result))}

4. A flexible solution

4.1 The hypothesis engine

Human intuition often plays a strong role in determining what is interesting in a health context and many breakthroughs are born out of a serendipitous discovery that is subsequently validated through further research not by statistical validity (Bresnahan, 1997). A final decision on what is interesting may be based on little more than gut feeling and hence the need for flexibility in determining the metrics for inclusion and exclusion of a pattern that becomes a hypothesis for clinical testing is required to mirror the natural process. The foundation for applying the theories presented here is the development of a hypothesis engine. The engine provides a flexible means to discriminate data mining patterns and therefore hypotheses based upon individual role based requirements. The engine provides a system that allows the following functionality:

1. The flexible application of a wide range of heuristics through the provision of a wide range of heuristics for user selection,
2. The run time selection and scoping of heuristics provided through an interface which allows the selection of any combination of heuristics and thresholds for each,
3. A role based default heuristic selection developed through analysis of the most commonly applied heuristics and thresholds for the user,
4. A means of discriminating hypotheses based on critical and non-critical requirements by applying the selected heuristical methods and developing an interest array as described in Section 3.2.2
5. A measurement of information strength based on trusted classes of heuristics through combining heuristic values in hypotheses of interest.

This functionality provides for a two phase hypothesis culling process as shown in Figure 1. The hypotheses are systematically culled as they attempt to propagate upwards through the evaluation phases from individual heuristics to a quantified information strength for the pattern. Any individual heuristic achieving within the role defined thresholds would automatically propagate to the next level. A metric that does not achieve a level within the thresholds would be filtered through a switch. If the metric is critical (needed) to determine strength then the switch would cull that hypothesis. If the metric is not critical (only wanted) then the switch would not be activated resulting in that heuristic value and the hypothesis being included in the next level providing no other critical heuristics fail for that hypothesis. The switches have the dual purpose of reducing the numbers and increasing the validity of hypotheses presented.

The hypothesis engine allows for patterns to be produced at the broadest level and then evaluated to allow the provision of only those which match a pre-defined interest role as defined by a user designed schema at run time. This addresses a number of the major barriers to the acceptance and application of data mining in the medical domain described earlier in this chapter;

- 1. The low level of flexibility in data mining systems requiring medical analytical processes to adapt to data mining methodologies rather than vice versa.
- 2. The lack of opportunity for incorporating subjectivity when mining medical data,
- 3. The broad range of users and analytical variance in medicine,
- 4. The production of too many irrelevant results, requiring a high level of user interpretation to discriminate those that are truly useful.

Fig. 1. Hypothesis Engine Overview

4.2 Medical diagnostic decision making

The process of interdisciplinary research and analysis discussed in this chapter has resulted in a methodology for identifying knowledge which is interesting in a clinical context and has facilitated the definition of a mechanism for evaluating the level of interest an individual user may have in the knowledge or rules produced as a result of mining medical data. This mechanism was built into a hypothesis engine as described in the previous section. However, whilst useful as a standalone adjunct to the data mining process, the engine does

not address all of the issues identified in Section 2. To achieve this end a demonstration system named ADAPT (Automated DATA Pattern Translator) has been built. This system aims to facilitate greater access to mining technologies for all medical users, and to provide an ability to apply some of the more complex mining technologies to acquire new medical knowledge without the risk of producing incomprehensible outputs. It also adds weight to the solutions provided by the hypothesis engine, and addresses the issues not able to be solved by the engine. The development of ADAPT grew primarily out of an understanding of the issues with applying data mining technologies in the medical domain and an aim to create a solution to those issues. It required not only an understanding of the medical knowledge acquisition process but also an understanding of the more common clinical processes. One of the most frequently performed processes of a practicing clinician is that of diagnosis. This process is briefly described here and forms the basis for gaining a deeper understanding of the requirements in an automated knowledge application and decision making system. The decision making process of medical practitioners usually if not always occurs at a subjective level following an objective information gathering process. This information is collected and interrogated during the consultation phase of the patient episode and generally includes all or some of the following artefacts;

- The symptoms as described by the patient
- The findings of the physical examination
- The results of tests
- Diagnostic imaging
- Clinical referrals
- Historical reports.

Once collated this information is generally evaluated through either of two common diagnostic methodologies; by exclusion or by pattern categorisation (Frenster, 1989; Merck, 2003; Elstein & Schwartz, 2002):

1. The exclusion method is considered the safer but more expensive of the two options. Here a doctor will reach a decision on the most likely diagnosis and treatment pattern for a set of stated symptoms exclusively through a series of objective and often invasive tests. These tests allow the doctor to reach a decision based upon an ability to exclude potential diagnoses based upon the results of the tests. A final diagnosis is usually not made until the results of all required tests are known. The potential disadvantages of this option are that it results in a lengthy and costly process which can yield an unacceptable level of false negatives or positives and cause psychological stress for the patient during the wait; especially if the tests are for a condition for which the outcome is potentially life threatening, for example cancer, meningococcal infection or HIV.
2. Pattern categorisation is the preferred option in the medical community as it is more cost and time effective than diagnosis by exclusion, however it is generally only applied by more experienced doctors and specialists. In this scenario a doctor will reach a diagnosis by comparing the presenting patient's pattern of symptoms and test results to known patterns of condition diagnosis, progression and treatment. It is often the case that rather than making a single diagnosis the practitioner will develop a cluster of potential diagnoses based upon the similarity of the potential diagnoses to the pattern of the patient's condition that initiated the medical consultation as shown in Figure 2. Probabilistic pattern completion is often required and the diagnosis and treatment patterns are derived from a comparison to all others in the practitioner's body of knowledge and those sufficiently similar to the

patients’ pattern are selected. The diagnosis is therefore more holistically targeted to the patient and case rather than the cause or effect exclusively. A drawback of this is that a rare case or combination may mean the most suitable pattern is missed or incorrectly identified and errors can be made.

With experience, doctors become efficient at recognising what’s expected in a pattern and which patient or condition characteristics are most influential or critical when forming a diagnosis. The speed of pattern recognition may be increased by fuzzy matching which takes into account only matches between critical attributes or attribute values which fall within a range. Attributes or values which occur frequently cannot be applied to discriminate between patterns, but knowing which attributes to focus on and which to eliminate from the comparison is a skill developed through years of experience. Given the importance of the ability to recognise which pattern elements or attributes are important in the clinical diagnostic process, it is a logical step to move towards an automated process of identifying the important elements of patterns derived from the application of data mining technologies.

Fig. 2. Diagnostic process model. Note: Darker shading denotes higher level of subjective input.

For automated processes to be applied and accepted it is important that they are able to mirror or facilitate current practices. This requires the provision of knowledge which is able to be trusted, descriptive from an objective viewpoint and informative in a way which assists in the subjective appraisal of a patient. Unfortunately many data mining outputs (patterns) are presented in a non-intuitive and frequently opaquely coded form. The above

diagnostic methods suggest that there is a requirement for heuristic triggers and substantiated hypotheses which can suggest rather than dictate a course of action based upon pattern similarity and a need to be able to deconstruct the pattern into more and less influencing pattern elements upon which to base the diagnosis or treatment. If we accept these requirements, it is a logical step to move towards an automated process of identifying and describing the important elements of patterns derived from the application of data mining technologies and it is this precept that has defined the development of ADAPT.

For brevity and to ensure understanding of the system at an algorithmic level several definitions are provided here:

- Data mining is essentially an automated data analysis process, however there are almost as many formal definitions of data mining as there are data miners. One of the more commonly quoted definitions is that it is a “non-trivial process of identifying valid, novel, potentially useful and ultimately understandable patterns in data” (Fayyad et al., 1996). These patterns are often described through the standard heuristics of support and confidence.
- Support essentially quantifies the frequency with which a pattern occurs within the original data source, expressed as a percentage. The minimum support threshold is generally set by the user at runtime. A pattern is defined as a set of attribute values produced through a process of association rule mining, which occur equal to, or greater than, the minimum support required by the user. Patterns take a general format as shown in the following example; Condition ‘A’ + Treatment ‘B’ lead to an outcome of ‘Recovery Time C’ with support of 20%.
- Confidence quantifies as a percentage the degree to which we can rely on the presence of part of a pattern given another part of the same pattern. For example, given Condition ‘A’, the confidence states how often we could expect Treatment ‘B’ to also occur in the same pattern.
- A pattern element is a part of a pattern. In the example above, Condition ‘A’, Treatment ‘B’ and Recovery Time ‘C’ are the three elements of the pattern. Treatment ‘B’ leads to a Recovery Time ‘C’ is also an element, as it represents only a part of the pattern.
- Elemental support is the support for a pattern element. This concept is explained in further detail in foundation work by the author (Shillabeer & Roddick 2006).

4.3 ADAPT

Data mining outputs are currently deficient in their application to the medical domain as they are not able to identify which pattern elements directly affect the medical outcome or which have no effect and are therefore little more than confounders, although as explained earlier, this is necessary if patterns derived from medical data are to assist in processes such as diagnosis.

The ADAPT system aims to provide guidance and structure to the application of translating technical outputs into clinically relevant patterns of diagnosis or treatment, and presents a novel approach to post-mining pattern evaluation. There are six steps identified in the process:

1. Identify a set of interesting patterns,
2. Calculate the representative pattern element weightings,
3. Deconstruct the support for each pattern and determine its elemental supports,

4. Determine the positive, negative and inert elements,
5. Order patterns by their degree of representation,
6. Present patterns with their associated heuristics.

Data mining processes that have attempted to evaluate the patterns presented have focussed primarily on step 1, however approaches taken have often been unsuitable for the medical domain due to the wide range of user types and definitions of interest. Medically focussed solutions have been developed but have generally provided a user specific solution. This chapter has described a hypothesis engine developed specifically to address the requirement for flexibility in interest evaluation and quantification. Unfortunately this step alone still presents patterns which are potentially unsuitably formatted, and/or do not inform sufficiently well for direct clinical application. The need to supplement information content to aid understanding of the patterns produced as a result of data mining is the focus of steps 2 to 4. Presentation of patterns in a logical and value driven order is the focus of steps 5 and 6. Steps 2 to 4 are described following.

4.3.1 The argument for elemental support

The concept of elemental support arose out of an understanding of the need to determine which pattern elements are important in medical decision making as raised in the previous section, and the realisation that not all elements of a pattern necessarily carry the same importance. Two examples are provided to demonstrate the relevance of the argument.

1. Taking dymadon and lemsip together are unnecessary to control the symptoms of a cold, as they both contain the same active ingredient, but some patients may choose to do so from habit or ignorance. In a mining sense the pattern of dymadon + lemsip = reduction in symptoms, would be viable however we can see that in reality the removal of any one of the treatments would probably not affect the outcome but this information would not be available in traditional reporting of the presented pattern.
2. In the treatment of AIDS many medications were prescribed singularly and in combination before the specific combinations were found that made a genuine contribution to the well being and longevity of the patient. The individual medications may have demonstrated sufficient mining support for their ability to associate with a positive outcome which would have therefore been suggestive of an ability to facilitate a positive outcome alone, as demonstrated in Figure 3. However, if in fact their efficacy were true only when combined with specific other medications this should be evidenced through the metadata or heuristics provided to describe the pattern. It should be that as the single medication does not achieve sufficient support in isolation it is therefore not reported unless combined with the other medications in a pattern that can be substantiated statistically and medically.

Through considering these issues it was identified that data mining may not be able to ensure completeness, soundness and medical accuracy of results in the medical domain and a method for determining the importance of pattern elements rather than whole patterns is deemed necessary. As shown in figure 3, traditional support does not provide sufficiently granular information and may in fact be misleading although it is a frequently applied determinate. In the example, treatment A and procedure B are both quantified as strong using the support metric and this could suggest that their individual implementation would lead to a positive outcome for the patient. However as the value for the pattern element includes instances of that element in all other patterns also, and does not isolate to the

prevalence of that specific pattern no conclusions should be drawn. It is necessary to be able to discriminate and differentiate by knowing the support for treatment A or procedure B in isolation rather than only as an element in a longer pattern as shown in Figure 4. This more clearly denotes that applying either element alone will not necessarily lead to a positive outcome but when applied together there is a far greater chance of a positive outcome for the patient.

This understanding resulted in the following requirement list:

- The requirement for mining outcomes to be evaluated with a non-traditional application of the support measure.
- The requirement to discriminate between the overall incidence of an element in any pattern and the incidence of that element in isolation.
- The requirement to provide a clear description of the information held in a pattern to aid subjective judgement.

Fig. 3. Traditional support values for a sample data pattern.

Fig. 4. Elemental support values for the sample data pattern.

Two methods were developed to address these requirements. These methods have been termed support deconstruction and element weighting.

4.3.2 Support deconstruction

Support deconstruction is a novel method which aims to determine the value of each element to a pattern. If 3 element patterns have the same support as 2 element patterns then we can logically deduce that the third element does not add any extra value as the likelihood of achieving the same outcome is equal. In our earlier example of cold

medications, the addition of lemsip to the pattern ‘dymadon = relief of symptoms’ would probably not significantly increase the support for a relief in symptoms. To determine the effect of each element we need to calculate the support for the pattern before and after adding each extra element to the pattern. In effect we need to take the incidence of BC from BCD to determine the effect of D. In a pattern ABCD = a positive outcome, each element appears as often as each other; therefore by applying traditional support metrics, $suppA = suppB = suppC = suppD$ and each element could be considered equally strong. To determine the effect of element ‘A’ on pattern ‘ABCD’, its participation in all of the following relationships must be considered:

Level 1		A	
Level 2	AB	AC	AD
Level 3	ABC	ABD	ACD
Level 4		ABCD	

From being a single element (level 1) to being an element of the 4 element pattern (level 4), element ‘A’ participates in 7 elements or relationships and in this example it would lend part of its traditional support to each of these. Support sharing on any level can be denoted as the sum of supports for all elements on lower levels occurring alone. Whilst this is logical there is further complexity, as traditionally the support for AB has included the instances of ABC, ABD and ABCD and so on, and it is not clear how many times AB occurs in seclusion. To determine this there is a need to reverse engineer or deconstruct the support as demonstrated in Table 2. To reveal the deconstructed elemental support for each element the participation of each element at all levels must be calculated get an accurate value for each element in isolation not simply in those patterns with a direct linear relationship. This can be denoted as:

$Esupport(focus\ element) = Tsupport(focus\ element) - Esupport\ for\ each\ lower\ level\ element\ in\ which\ the\ focus\ element\ occurs.$ (E denotes elemental and T traditional).

Elements	A	AB	ABC	ABCD	Total
Traditional support %	16	10	5	4	35
Elemental support %	6	5	1	4	16

Table 2. Comparison of traditional and elemental supports.

In the simplified example in Table 2, we can see that the effect of adding element C to element AB is negative in that it has a lower elemental support than AB alone. In contrast adding element D to element ABC increases the elemental support of ABC, and hence denotes an increase of a positive outcome overall using this example. The obvious other comparisons needed here are the elemental supports for ABD, ACD, AC and AD to determine whether C actually adds value in the presence of D or if it is redundant and D is really the differentiator. This method has demonstrated a potential to show how much each element affects the outcome described in the pattern. In a real world example this would, for example, show the most or least effective combinations of medications which would allow for more accurate targeting of medications and potentially a reduction in the number of medications taken.

Through developing a greater knowledge regarding the importance of an element to an overall pattern three types of element can be identified; positive, negative and inert.

- Positive elements are those which increase the elemental support of an element it joins with in a manner which would influence the subjective interest in the resultant pattern.

- Negative elements are those which decrease the elemental support of an element it joins with and again would influence the subjective interest in the resultant pattern.
- Inert elements are those which do not affect the elemental support of an element it joins with. The subjective interest in this element cannot be determined. It could be that the lack of effect would in itself be valuable knowledge and would therefore affect a decision based upon it.

Elemental support can report on how important each element is to each pattern it participates in. A remaining issue is how to determine those pattern elements which have been recorded frequently enough to be eliminated due to the probability that they would already be known. Essentially the more relationships an element is involved in the more likely it is to be uninteresting. By evaluating the overall frequency of the pattern we can determine which are more unique or important and this would also assist in reducing the overall numbers of patterns by facilitating the removal of patterns or pattern elements which contain knowledge encountered frequently. This issue can be addressed by a determination of element weighting.

4.3.3 Element weighting

A frequent criticism of automated data analysis in medicine is that too many rules are produced and they often represent knowledge which is commonly known. For example if an element 'A' participates in only one pattern whereas 'C' is involved in many patterns it could be assumed that the effect of 'C' is more likely to be known, as it has been recorded more frequently. Solutions have been documented which involve the development and referencing of a knowledge base to hold 'known' patterns (Kononenko, 1993; Lucas, 1996; Perner, 1996; Moser, 1999). Whilst this has been demonstrated to be a valuable tool for reducing the numbers of known patterns presented, there are potential issues with this approach:

- The human resource time required to build and maintain a sufficiently complete knowledge base for practical use;
- The need to apply the knowledge subjectively to prevent exclusion of unusual or marginally different patterns and inclusion of different but non-informative patterns, and;
- The need to be sure that the knowledge base was developed from a sufficiently similar data set so that a comparison can be confidently made

This section presents a solution that begins to address these issues through a method of evaluating patterns based on pattern element weightings. ADAPT uses element weightings to compare the importance of elements within patterns, and combined element weights to compare patterns within levels. In this way element weights allow for subjective judgements to be made from original data rather than from comparison to external sources which may have been developed using different heuristics and/or from data with different characteristics, and/or origins. Pattern element weightings also facilitate subjective pattern evaluation through an understanding that the patterns and elements are representative of the data set from which they were derived and the rate of occurrence in the patterns can be compared with the occurrence in the data source to ensure that what seems frequent or rare is actually so. As suggested above, an element which participates in many relationships is more likely to be known and cannot effectively be used as a discriminator, but elements which participate in few relationships can be considered a discriminating participant and the application of element weightings will allow for accurate quantification of this quality.

As a first step toward automating the determination of element representation, a technique commonly used in the field of Information Retrieval (IR) that finds its origins in Information Theory (IT) has been utilised. IR is concerned with the classification of text and text based documents and offers a valuable perspective on the evaluation of worth for individual text elements of a set in the context of document classification, indexing and searching. Modern applications have included the evaluation of internet search results to quantify the relevance of a document to the user supplied search terms. Results containing a higher frequency of terms or a greater number of terms would be classified as more relevant. Researchers in the field suggest that rather than having a binary weighting, a non-binary measure should be used to determine the degree of similarity between sets of terms in a document or query (Baeza-Yates & Ribeiro-Neto, 1999). This would be applied to rank results even if they only partially matched the requirements. Fuzzy set theory is applied to allow for consideration of partial membership of a set and gives a measure of closeness to ideal membership i.e. how similar each set of terms is to each other. Seminal work in this area was published in 1988 by Salton and Buckley (Salton & Buckley, 1988) which described the use of a formula named *tf.idf* (TFIDF) to quantify the similarity between the contents of a document and user requirements as defined through a set of query terms.

TFIDF's ability to represent the amount of information value contained in a word in a sequence or document suggests a similar approach may be useful in determining element and/or pattern relevance. TFIDF is a weighting comprised of the two functions Term Frequency (TF) & Inverse Document Frequency (IDF). Term frequency is the frequency of a specific term in a document which indicates the importance of a term t_i to that document. It is formally described as;

$$TF = \frac{n_i}{\sum_k n_k} \quad (1)$$

where n_i is the number of occurrences of the specific term in a document and $\sum_k n_k$ the number of occurrences of all terms in the document.

Inverse document frequency indicates a terms importance with regard to the corpus. Based in part on information theory it is the logarithm of all the documents divided by the number of documents containing a specific term. Formally it is;

$$IDF = \log\left(\frac{|D|}{|(d_i \supset t_i)|}\right) \quad (2)$$

where $|D|$ is the total number of documents in the corpus and $|(d_i \supset t_i)|$ is the number of document within which term t_i occurs assuming $n_i \neq 0$.

TFIDF quantifies word frequency in a document such that the more frequently a word appears in a document (e.g., its TF, term frequency is high) the more it is estimated to be significant in that document while IDF measures how infrequent a word is in the collection and accordingly if a word is very frequent in the corpus (collection of documents), it is considered to be non-representative of this document (since it occurs in most documents). In contrast, if the word is infrequent in the text collection, it is suggested to be very relevant for the document. This characteristic facilitates the filtering out of common terms due to their lower ranking scores/weights.

This approach can be translated to the evaluation of medical patterns by simply re-stating the concepts and substituting comparable terms as follows:

TF can be applied to determine the frequency with which a pattern element occurs within a particular pattern sub-set at the same level or cardinality for a given support e.g. all 4 element patterns with a minimum support of 40%. This is given the assumption that patterns of the same cardinality and support are contextually similar enough to be associated/chunked much like that of the sentences of a document and thus form a pseudo document. IDF can be applied to determine the frequency with which a pattern element occurs within the original data set. When brought together in the TFIDF calculation, the information represented by patterns at an equal level can be compared to the total set to derive an entropic measure at the element, pattern and sub-set levels. This is useful for the subjective analysis of medical patterns as it allows for the recognition of the amount of intra and inter-pattern importance (information value) of individual elements.

In data-mining terms TF can be seen as comparable to the traditional support of the pattern as it could be applied to determine the frequency of the pattern within the total pattern set or data. The novelty of the solution becomes apparent through the application of IDF. In pattern analysis this could be defined as either the frequency of the pattern against all patterns in the present set or data source or the frequency of the element in all patterns in the set or data source. As the focus is on the value of an element within a pattern, the second definition was applied. By combining both of these into a value for TFIDF the frequency of element ABC within the pattern set is calculated and multiplied by the log of inverse frequency of element ABC in the data source. This will show us if that element participates in many or few other patterns and whether the element frequency is representative of the data source. If we achieve a low value this will show that the occurrence of the element in the pattern set is dissimilar to that in data source. If we have a high value then this is indicative of an inert element for example the gender element in the gender = female and condition = pregnancy pattern which would occur with high frequency in both the data source and the pattern set.

Traditional metrics have not been documented in such an application, but for the reasons discussed above it is an important issue in medical pattern analysis.

5. Experimental results

5.1 Methodology

ADAPT was tested using data collected from the Royal Australasian College of Surgeons National Breast Cancer Audit^{1,2}.

Association rule patterns were produced from these data to demonstrate that ADAPT is able to perform well on the most problematic mining algorithm for medical data due to the potential for a large number of patterns to be presented. It is hence the area which stands to benefit most from ADAPT.

¹ The supplied data has since been superseded and hence results detailed herein may or may not accurately reflect current data or practices within the management of breast cancer in Australia or New Zealand and hence should not be relied upon in a clinical setting.

² Acknowledgement and thanks are hereby given to the ASERNIP-S Adelaide team for their permission to use this data.

The original source data held 23,259 patient records, 53,950 separate surgical procedure records for those patients and 82 attributes with an average of 5 potential values per attribute across three tables. After combining patient records with their surgical procedures, allowing for multiple procedures performed together, then removing duplicate columns and extraneous data including dates, 42,705 records with 39 columns were mined using a simple association rule mining tool with support set at 1%.

ADAPT was applied to test the validity of four theories, and provide information relating to one area of interest described in the data dictionary provided with the data. The 5 cases analysed were:

1. "Some treatment recommendations differ by menopausal status, e.g. Ovarian ablation is not indicated after the menopause."
2. "The omission of axillary dissection should be considered only in the case of small primary tumours and is least desirable in pre-menopausal women but is standard in cases of DCIS"
3. "The presence of a significant amount of DCIS in or adjacent to an invasive tumour is a predictor of high relapse rates following wide local excision and radiotherapy."
4. "For DCIS, 'clear margins' as defined by no tumour within 1mm of any inked margins are associated with a high rate of recurrence."
5. It was also suggested that "correlations between excision technique and cosmetic result may be useful."

Interesting patterns were defined as those which contained the particular attribute values required to test the cases - as described in the results section.

5.2 Results

Association rules were developed to provide patterns relating to the five cases listed above using the attributes and values detailed in Table 5. For some cases it was necessary to apply different sets of attributes and develop more than one set of patterns for analysis to test each part of a more complex theory in isolation.

The results of the five test cases are detailed in Table 6 which shows elemental supports and TFIDF values. TFIDF was categorised on whether it fell in the upper or lower values for representativeness. Positive elements are shown in bold, negative in italic and inert in standard text and were classified as described earlier. One is denoted as both positive and negative based upon the elements which preceded it. Table 7 shows the descriptive potential of ADAPT based on these heuristics. Overall, only one of the theories was fully supported by the data and two were fully discounted with one gaining partial support. The TFIDF value was able to indicate if this result is generally representative. Where the theory pattern was in the lower level of representativeness it suggests that the theory should not be depended upon and its foundation should be investigated further.

5.3 Discussion

Data mining has many documented benefits for the medical domain and much work is being done to provide medicine specific solutions to data analysis problems. However this work has focussed on the adaptation of technologies to manage the complexities and non-uniform nature of medical data and while the management of input is being addressed, research by the author has been unable to find evidence of a similar effort in the management and translation of outputs. This is an issue when non statisticians or computer scientists are

attempting to make sense of the technology and its products. It has been reported that medical practitioners in general have little experience in statistical analysis and as a result are unable to apply or comprehend results which do not match their sphere of knowledge and understanding and as such they need assistance in translating or deciphering the meaning or content of statistical analysis which is beyond the complexity of a 0.05 p-value or confidence interval of 95% (Shillabeer & Pfitzner, 2007, Link & Marz, 2006). This has a

Case	Focal Element A	Element B	Element C	Element D
1	Menopause Status 1	Ovarian ablation 1	NA	NA
2.1	Ax Dis. 3	DCIS 2	NA	NA
2.2	Ax Dis. 3	Menopause Status 1	NA	NA
2.3	Size > 15	Aux Dis. 0	NA	NA
3.1	Relapse 2	EIC 1	Surgery 2	Radio 1
3.2	Relapse 1	EIC 1	Surgery 2	Radio 1
4	Relapse 2	DCIS 2	Clear mar.	NA
5.1	Surgery	Symmetry3	NA	NA
5.2	Surgery	Symmetry1	NA	NA

Table 5. Pattern attributes and values for each case

Case	A	AB	AC	AD	ABC	ACD	ABD	ABCD	TF/IDF	Theory proven/ Representative
1	59.85	0							NA	Yes/NA
2.1	10.01	4.61							0.1975	No/upper - generally applicable
2.2	12.88	13.04							0.1421	Yes/lower - conditionally applicable
2.3	18.04	0							0.1837	Yes/upper - generally applicable
3.1	2.27	0	0	0					NA	No/NA
3.2	12.62	2.27	2.44	4.18	0.53	2.04	9.91	1.46	0.0704	No/lower - conditionally applicable
4	0.22	1.03	0						NA	No/NA
5.1	30.98	1.02							0.1488	NA/lower - conditionally applicable
5.2	13.59	7.2							0.1510	NA/lower - conditionally applicable

Table 6. Elemental supports and TFIDF heuristics

Case	Traditionally formatted data mining patterns for each case	ADAPT description
1	"Meno_Status":2.00, "Ovarian_Ablation":2.00 (49.48%) "Meno_Status":2.00, "Ovarian_Ablation":9.00 (9.83%)	No post-menopausal patients received an ovarian ablation.
2.1	"Insitu_Necrosed":2.00 (14.62%) "Insitu_Necrosed":2.00, "Ax_Type":3.00 (4.61%)	Axillary Dissection is equally applied and omitted in DCIS cases.
2.2	"Meno_Status":1.00 (25.92%) "Meno_Status":1.00, "Ax_Type":3.00 (13.04%)	Half of all pre-menopausal women have axillary dissection and there was no level of omission.
2.3	"Tumor_Size":20.00, "Ax_Type":3.00 (3.10%) "Tumor_Size":0.00, "Ax_Type":0.00 (10.35%)	No cases of larger tumours avoid axillary dissection.
3.1	None	Relapse is not paired with relevant attributes.
3.2	"Surgical_Event":2.00, "Status":1.00, "EIC_Status":1.00, "Radiotherapy":1.00 (1.46%)	Non relapse is paired with the theoretical pattern.
4	None	Clear margins with DCIS are not an indicator of relapse.
5.1 5.2	"Surgical_Event":5.00, "Symmetry":1.00 (1.49%) "Surgical_Event":4.00, "Symmetry":3.00 (1.02%) "Surgical_Event":4.00, "Symmetry":1.00 (15.44%) "Surgical_Event":3.00, "Symmetry":1.00 (3.53%) "Surgical_Event":2.00, "Symmetry":1.00 (24.44%) "Surgical_Event":1.00, "Symmetry":1.00 (7.20%)	Only total mastectomy was associated with a poor result. Open biopsy has ranked highest for good results >50%

Table 7. Comparison of data mining pattern output and ADAPT description output

potential two pronged effect, it puts data mining technologies and the like out of the reach of many medical practitioners and, it necessitates the application of only the most basic analytical tools thus negating the potential of the more powerful tools.

The contribution of the work presented here is in its ability to broaden the applicability and marketability for the technology by making outputs approachable to all user types and levels without minimising the complexity or range of processes available. It does not attempt to modify the data mining process as the algorithms and statistical methods available currently are generally applicable and effective for the medical domain and have been used successfully by many medical teams. However these projects are often undertaken by specialists who understand the tools and technologies or they have required

manual post mining translation or interpretation of outputs by domain specialists (Imberman & Domanski 2002; Moser et al., 1999). The results presented here show that ADAPT is able to facilitate the ordering and presentation of complex outputs in a more intuitive language and provide the knowledge items required for decision making as described earlier. It is also able to overcome some of the more potent criticisms of the technology, for example the belief that the results of data mining are not always representative of the data set from which they were created and they are often misreported as a result (Milloy, 1995; Raju, 2003; Smith & Ebrahim, 2002). Whilst ADAPT pattern interpretations are currently manually created, a future challenge will be to create an automated natural language description of patterns based on the heuristics provided by the process.

5.4 Conclusion

Results have been developed from patterns which are far simpler than the technology is capable of providing; however the theories were often binary and did not require the production of multi-faceted patterns. Data mining would not generally be applied to develop such simple patterns but it has proved to be a suitable demonstration of the process in a realistic application. Also, the focus here was on the application of the ADAPT process rather than on the application of data mining technologies and the simplistic patterns provided an unambiguous and non-complex platform from which to work. ADAPT has demonstrated an ability to mirror the earlier identified trusted process of medical knowledge acquisition thereby facilitating the application of data mining technologies into clinical decision making, theory validation and hypothesis generation. Of particular relevance is the ability to apply both objective and subjective measures of interest to guide the knowledge acquisition process. The novelty and power of the solution is demonstrated through its ability to provide individualised case based knowledge reflecting the growth in personalised, evidence based clinical care. The logical next step is to process the results of more complex association rule mining through ADAPT and validate the output through empirical testing. This work tested and demonstrated the potential for ADAPT on essentially one tailed tests, but the real power of the methodology and its parts will become evident when applied to purely exploratory data mining.

The primary areas of contribution are in its ability to respond to the following issues;

1. The low level of flexibility in data mining systems and the need for medical analytical processes to adapt to data mining methodologies rather than data mining adapting to the needs of medicine,
2. The production of patterns in a technical language and format that are often not understandable or applicable in a clinical setting,
3. The production of results that require a high level of user interpretation to discriminate those that are truly useful.

The functionality and process flow of the system was designed to reflect the process of medical diagnostic decision making and is demonstrated through its ability to discriminate and translate the technical outputs of data mining into a clinically understandable and applicable form. Both the hypothesis engine and ADAPT are able to manage any data mining outputs but testing has been performed on the most problematic form of mining output; association rules. It is in this use of the technology that we see the greatest problem

with large numbers of uninteresting or irrelevant results that require intensive user evaluation and interpretation. It is believed that by addressing the area of greatest issue the power of the solution has most clearly been demonstrated.

There are many applications of data mining technologies in medicine with some being more successful than others as discussed in this chapter.

Specific applications and benefits of the work presented in this chapter include:

- providing a more approachable interface to data mining technologies
- provision of a more medicine specific format for outputs both in terms of user needs and technical knowledge
- ability to reduce pharmaceutical costs through the ability to identify most beneficial or effective combinations of medications for particular conditions or sub sets of the population
- Ability to enhance the knowledge base of the practicing clinician as required and in a real time context
- Ability to provide an automated solution which can mirror trusted methodologies
- provision of a system which can apply trusted metrics for the measurement of validity and applicability of outcomes
- provision of a system which can provide statistically valid and trustworthy hypothesis
- provision of an automated system that can assimilate the subjective and objective needs of any medical professional
- provision of a system which incorporates current technologies and provides knowledge in an understandable format and language
- provision of a system which is developed through an understanding of the unique needs of the medical profession and is based upon the application of methodologies developed over centuries of medical research
- development of solutions to the documented impediments to the acceptance and utilisation of data mining in the medical field
- development of solutions to address many of the documented issues in medical data mining

5.7 Future research

The need now is to build systems which allow for changes and improvements without the need to rebuild systems and carry the burden of lost productivity due to development time and cost, neither of which are insubstantial. Also it will become increasingly important in the future to consider a wider range of users as patients themselves become more empowered and willing to collaborate with their treating professionals in regard to their own treatment. There is an increasing thirst for personalised knowledge and information and an increasing expectation that this thirst will be quenched in real time thus necessitating a tool such as that discussed in this chapter. It was estimated in 2001 that up to one third of all Internet surfing was related to the search for health information (Kapur, 2001). Increasing use of the Internet has empowered the general public but this empowerment has not always been embraced by medical professionals. The environment in which medicine is practiced has changed dramatically through history although the founding principles have remained relatively static. The biggest change has been in the informedness and hence power of the people to influence and engage in their diagnosis and treatment, thus requiring a greater breadth and depth of knowledge in the treating specialist.

Whilst the discussion presented in this chapter has clearly demonstrated the potential benefits of the application of automated data analysis techniques such as data mining in the search for new, clinically applicable knowledge in the medical domain, the following issues remain;

- The lack of perceived trust in automated systems
- Lack of computer literacy across the medical profession making the application of technical solutions to data analysis needs currently unfeasible.
- The lack of uniform data standards within and across jurisdictions.
- The difficulty in sourcing suitable datasets for the development and testing of automated systems
- Lack of data sharing protocols for data sharing between health providers.

Although this is a leading edge area of research and the building blocks of solutions to address the changing needs of medicine have already been built and tested, work must now continue to perfect solutions to these issues before automated knowledge acquisition becomes integrated as a natural part of standard medical data analysis practice.

6. References

- Said al-Andalusí, Science in the Medieval World: "Book of the Categories of Nations" (trans. Sema'an I. Salem and Alok Kumar, 1991, University of Texas Press.
- Ashby, D., Smith, A.: The best medicine? Plus magazine - living mathematics (2002)
- BBC. :John Snow (1813-1854). Accessed 15/05/04 at www.ph.ucla.edu/epi/snow.html
- Beals, M., Gross, L., Harrell, S.: Diversity indices: Shannon's H and E (1999)
- Beals, M., Gross, L., Harrell, S.: Diversity indices: Simpsons (1999)
- Baeza-Yates, R. & Ribeiro-Neto, B., (1999), "Modern Information Retrieval", Addison-Wesley, Sydney
- Biedl, T., Brejova, B., Demaine, E. D., Hamel, A. M., Vinar, T. Optimal Arrangement of Leaves in the Tree Representing Heirarchical Clustering of Gene Expression Data. Technical Report 2001-14. University of Waterloo, ON, Canada.
- Bojarczuk C. C., Lopes H. S., Freitas A. A., Data mining with constrained-syntax genetic programming: applications in medical data sets. In proc. Intelligent Data Analysis in Medicine and Pharmacology (IDAMAP-2001)
- Bresnahan, J.: Data mining - a delicate operation. CIO Magazine (1997)
- Brieger, G. H. (1977), 'H 1 coulter. divided legacy: A history of the schism in medical thought.', Isis 69(1), 103--105.
- Cabri, G., Ferrari, L., Leonardi, L., Zambonelli, F. Role based approaches for Agent Development. AAMAS'04. July 19-23 2004. New York U.S.A.
- Cios, K.J., Moore, G.W. Uniqueness of Medical Data Mining. Artificial Intelligence in Medicine Journal. Vol26. No. 1-2, September-October 2002 1-24.
- Cornwell, J. Hype or Hope? The Weekend Australian Magazine. Aug 27-28 2005. 36-41.
- Cosans, C. E. (1997), 'Galen's critique of rationalist and empiricist anatomy', Journal of the History of Biology 30, 35--54.
- Elstein AS, and Schwartz A. Evidence Base of Clinical Diagnosis. BMJ March 2002: Vol. 324: 729-732.

- Fausser, B. C. & Edwards, R. G. (2005), 'The early days of IVF', *Human Reproduction Update* 11(5), 437--438.
- Fayyad U.M, Piatetsky-Shapiro G, and Smyth P. From Data Mining to Knowledge Discovery in Databases. *AI Magazine*. 1996: Vol 7 Issue 3: 37-54.
- Ferrailo, D., Kuhn, R. An Introduction to Role-Based Access Control. *NIST/ITL Bulletin* December 1995.
- Frenster JH. Matrix Cognition in Medical Decision Making. In: *Proceedings of the AAMSI Congress on Medical Informatics*. San Fransisco: Volume 7 1989; pp 131-134.
- Gebski V, and Keech A. Statistical methods in clinical trials. *Medical Journal of Australia*. 2003: 178: 182-184.
- Geng L, and Hamilton HJ. Interestingness Measures for Data Mining: A Survey. *ACM Computing Surveys* 2006; Vol.38(3): Article 9.
- Gopal, K. (1997), Rationalist victory, Technical report, Online at <http://www.iheu.org/node/668>.
- Hagland M. Data mining. *Healthcare Informatics Online* 2004: 21:33-6. Accessed 28/3/2005.
- Hanson, A. E. (2006), Hippocrates: The "greek miracle" in medicine, Technical report, Online at www.medicineantiqua.org.uk/sa_hippint.html.
- Hilderman, R.J., Hamilton, H.J.: Evaluation of interestingness measures for ranking discovered knowledge. In *5th PAKDD 2001*. Volume 2035 of LNCS., Hong Kong, China, Springer (2001) 247-259
- Horton, R. (2000), 'How sick is modern medicine?', *The New York Review of Books* 47(17).
- Imberman SP, and Domanski B. Using dependancy/association rules to find indications for computerised tomography in a head trauma dataset. *Artificial Intelligence in Medicine*. September 2002: Vol. 26: 55-68.
- Katch, F. I. (1997), History makers, Technical report, Online at www.sportsci.org/news/history/lind/lind_sp.html.
- Kaul, M. & Thadani, S. Development of philosophical thought and scientific method in ancient India, Technical report, Online at http://members.tripod.com/INDIA_RESOURCE/scienceh.htm.
- Kononenko I. Inductive and Bayesian Learning in Medical datasets. *Journal of Applied Artificial Intelligence*. (1993) 317-337.
- Kuonen, D.: Challenges in bioinformatics for statistical data miners. *Bulletin of the Swiss Statistical Society* (2003) 10-17
- Lavrac N., Keravnou E., Zupan B. Intelligent data analysis in medicine. *Encyclopaedia of Library and Information Science*. Vol. 68 Supplement 31. (2000) 209-254.
- Link TM, and Marz R. Computer literacy and attitudes towards e-learning among first year medical students. *BMC Medical Education*. June 2006: 6:34.
- Lord, S., Gebski, V., Keech, A.: Multiple analyses in clinical trials: sound science or data dredging? *MJA* 181 (2004) 452-454
- Lucas P. Enhancement of Learning by Declarative Expert-based Models. In *proc. Intelligent data analysis in biomedicine and pharmacology*. Budapest, Hungary. (1996)
- Marshall, B. J. (1998), Peptic ulcers, stomach cancer and the bacteria which are responsible, Accessed 15/03/06 from www.knaw.nl/heinekenprizes/pdf/37.pdf.

- Merck & Co. Medical Decision Tests. www.merck.com/mmhe/print/sec25/ch300/ch300.html Accessed 14/09/2006.
- Milloy S. Science Without Sense. The Risky Business of public health, Cato Institute, Washington DC.
- Mohaghegh, M. (1988), 'Miftah al-tibb wa minhaj altullab (a summary translation)', Medical Journal of the Islamic Republic of Iran 2(1), 61--63.
- Moser SA, Jones WT, and Brossette SE. Application of Data Mining to Intensive Care Unit Microbiological Data. *Emerging Infectious Diseases*. 1999;5:3: 454-457.
- Muhaqqiq, M. (n.d.), 'Medical sects in islam', al-Tawhid Islamic Journal 8(2).
- Ordonez, C., Omiecinski, E., deBaal, L., Santana, C., Ezquerro, N., Taboada, J., Cooke, C., Krawczynska, E., Garcia, E.: Mining constrained association rules to predict heart disease. In: *ICDM'01*, San Jose. California. (2001) 433-440
- Page Wise. The History of Penicillin. Accessed 15/05/04 from http://oh.essortment.com/historyofpen_pnd.htm
- Pearcy, L. (1985), 'Galen: a biographical sketch', *Archaeology* 38(6 (Nov/Dec)), 33--39.
- Perner P. Mining Knowledge in X-Ray Images for Lung Cancer Diagnosis. In *proc. Intelligent data analysis in biomedicine and pharmacology*. Budapest, Hungary. (1996).
- Pollard, R. (2006), Fortuitous discovery led to a revolution in treatment, Technical report, Online at <http://www.smh.com.au/news/science/fortuitous-discovery-led-to-a-revolution-in-treatment/2006/10/11/1160246197925.html>.
- Raju S. Data Flaws. American Council on Science & Health. www.healthfactsandfears.com/high_priorities/vs/2003/data072303.html Accessed 23/9/2003.
- Roddick, J.F., Fule, P., Graco, W.J.: Exploratory medical knowledge discovery : Experiences and issues. *SigKDD Explorations* 5 (2003) 94-99
- Roddick, J.F., Rice, S.P.: What's interesting about cricket? - on thresholds and anticipation in discovered rules. *SIGKDD Explorations* 3 (2001) 1-5
- Salton, G. and Buckley, C., (1988), "Term-Weighting Approaches in Automatic Text Retrieval", *Information Processing Management*, Vol. 24, No. 5, pp.513-523.
- Shillabeer A, and Pfitzner D. Determining pattern element contribution in medical datasets. In: *ACSC*, Vol.29. No.7. *ACSW Frontiers* 2007; pp 209-216.
- Shillabeer A, and Roddick JF. Towards Role Based Hypothesis Evaluation for Health Data Mining. *Electronic Journal of Health Informatics* 2006;1(1):1-8.
- Smith GD, and Ebrahim S. Data dredging - bias or confounding. *BMJ* Dec. 2002; Vol.325:1437-1438.
- Sprang, K. (2002), Dr. Edward Jenner and the smallpox vaccination, Technical report, Online at http://scsc.essortment.com/edwardjennersm_rmfk.htm.
- Stone, C., McLachlan, K., Halliday, J., Wein, P., Tippet, C.: Gestational diabetes in victoria in 1996: incidence, risk factors and outcomes. *MJA* 177 (2002) 486-491
- Swan N. IVF Pioneers. Radio National Australia. Accessed 25/04/2005 from www.abc.net.au/rn/talks/8.30/helthrpt/stories/s1349685.htm
- Tan P. N., Kumar V., Srivastava J. Selecting the right interestingness measure for association patterns. *SIGKDD*. Alberta, Canada. (2002)

Tripod South Asian History Project (2002), Philosophical development from Upanishadic metaphysics to scientific realism, Technical report, Online at http://india_resource.tripod.com/upanishad.html.

Work In progress workshop in population health. May 19th 2005. Flinders Medical Centre.

IntechOpen

IntechOpen

Data Mining in Medical and Biological Research

Edited by Eugenia G. Giannopoulou

ISBN 978-953-7619-30-5

Hard cover, 320 pages

Publisher InTech

Published online 01, November, 2008

Published in print edition November, 2008

This book intends to bring together the most recent advances and applications of data mining research in the promising areas of medicine and biology from around the world. It consists of seventeen chapters, twelve related to medical research and five focused on the biological domain, which describe interesting applications, motivating progress and worthwhile results. We hope that the readers will benefit from this book and consider it as an excellent way to keep pace with the vast and diverse advances of new research efforts.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Anna Shillabeer (2008). Building Clinical Trust in Automated Knowledge Acquisition, Data Mining in Medical and Biological Research, Eugenia G. Giannopoulou (Ed.), ISBN: 978-953-7619-30-5, InTech, Available from: http://www.intechopen.com/books/data_mining_in_medical_and_biological_research/building_clinical_trust_in_automated_knowledge_acquisition

INTECH
open science | open minds

InTech Europe

University Campus STeP Ri
Slavka Krautzeka 83/A
51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

InTech China

Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China
中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821

© 2008 The Author(s). Licensee IntechOpen. This chapter is distributed under the terms of the [Creative Commons Attribution-NonCommercial-ShareAlike-3.0 License](https://creativecommons.org/licenses/by-nc-sa/3.0/), which permits use, distribution and reproduction for non-commercial purposes, provided the original is properly cited and derivative works building on this content are distributed under the same license.

IntechOpen

IntechOpen