© 2004 IMIA. All rights reserved

Clinical Quality Needs Complex Adaptive Systems and Machine Learning

Stephen Marsland^a, Iain Buchan^{a,b}

^aImaging Science and Biomedical Engineering, University of Manchester, Manchester UK ^bEvidence for Population Health Unit, School of Epidemiology and Health Sciences

Abstract

The vast increase in clinical data has the potential to bring about large improvements in clinical quality and other aspects of healthcare delivery. However, such benefits do not come without cost. The analysis of such large datasets, particularly where the data may have to be merged from several sources and may be noisy and incomplete, is a challenging task. Furthermore, the introduction of clinical changes is a cyclical task, meaning that the processes under examination operate in an environment that is not static.

We suggest that traditional methods of analysis are unsuitable for the task, and identify complexity theory and machine learning as areas that have the potential to facilitate the examination of clinical quality. By its nature the field of complex adaptive systems deals with environments that change because of the interactions that have occurred in the past. We draw parallels between health informatics and bioinformatics, which has already started to successfully use machine learning methods.

Keywords:

Clinical quality, complex adaptive systems, complexity, machine learning, clinical data.

Introduction

Motivation

The drive to electronically archive and store progressively more healthcare data is both a boon and a potential danger. Data as varied as individual health records, hospital episode statistics, general practitioner death rates and levels of diet and health in the population as a whole are now routinely recorded [1,2]. Dealing with large datasets is not easy, and the challenge of finding ways in which to automatically analyse this data should be a major area of research in health informatics.

However, the benefits of having such data available are obvious – successful analysis can inform future decisions. In this paper we focus on the area of clinical quality, which is of fundamental importance in healthcare delivery. The improvement in clinical quality follows a development cycle; current practice is reviewed using hard and soft methods of data analysis, and various improvements are considered and, if appropriate, implemented. This amended practice is then reviewed and the cycle iterates.

The analysis of data is key in this cycle. Some of the difficulties lie in the heterogeneous nature of the problem – the culture and

politics of the institution play a large role, as does the question of whose quality is being considered; an individual patient does not want to be a waiting list statistic, while a hospital manager has to trade this off against having empty beds. This question of whose quality also highlights the difference between patient-centred clinical quality and clinical policy, which is targeted on the needs of the population as a whole. This difference has analogies with the difference between microeconomics and macroeconomics. Furthermore, the feedback inherent in the cycle, as failures in current practice drive the adoption of new procedures means that the environment that is being analysed is not static.

However, practical data considerations are also a considerable difficulty. Data held by different sources may well have been coded differently, and even different data recorded. In addition, the stored data is liable to be noisy (that is, to contain inaccuracies), contains a mixture of correlated and uncorrelated variables with the correlation not being known *a priori*, and has missing data. The sheer quantity of data means that any methods that are used must be entirely automatic, as any human intervention will take an incredibly long time.

Current Methods are Not Enough

Current statistical approaches to clinical quality and performance are under-developed [3]. Statistical process control methods from operational research have been applied to the surveillance of clinical quality [4]; we have suggested that this approach might be valuable in some situations but will often be an over-simplification of clinical systems [5]. Noting the absence of research in this area, we suggest that adaptive approaches should be sought for the surveillance of problematic patterns in clinical systems, mainly because there is implicit and facilitated (e.g., audit) feedback between observed clinical outcomes, interventions and the organisation and delivery of healthcare.

As substantial change over time is inevitable in clinical environments, unlike in highly controlled industrial processes, we doubt the appropriateness of models of analysis of clinical quality that are fixed with respect to time. Some time-referenced analyses, usually of policy changes such as the introduction of new therapeutic guidance [6], have used linear methods such as interrupted time-series, which may be appropriate for simple overview questions. If the unit of surveillance is to be the healthcare worker, the clinical team or some other set thereof, then the process of surveillance will usually involve many changes, some interacting, over time. For example, the introduction of incentives for

primary care physicians to treat a greater proportion of depressed patients via behavioural rather than drug therapies might affect the classification of patients as depressed, the proportion of patients treated from a certain socio-economic group, the amount of information recorded in clinical records, and the waiting time for access to behavioural therapy; all of which may interact with one another differently over time.

It is therefore reasonable to assume that there are non-linear relationships between these commonplace factors of healthcare and health outcomes. Fixed models of surveillance and linear reduction of complex clinical environments are therefore unlikely to be adequate methods for measuring clinical quality and performance.

Complex Problems Need Complex Solutions

Other areas, most notably in recent years bioinformatics, have faced the same difficulties of dealing with very large quantities of noisy data, for example from gene expression microarrays or protein sequence analysis. Rather than simply looking at common statistical methods bioinformatics has formed a symbiotic relationship with machine learning and artificial intelligence – information analysis needs in bioinformatics can drive theoretical development in machine learning [7,8,9]. We believe that a similar relationship should be fostered between health informatics and machine learning, and we describe how such methods can be of benefit for clinical quality.

We describe two areas of study that seem particularly suitable. The first is that of complex adaptive systems. Complexity theory, the theory of such systems, is relatively new, although it has its origins in the work of Prigogine on non-equilibrium thermodynamics [10,11]. Complex adaptive systems are comprised of groups of interacting agents. The agents themselves can learn and adapt to changes in their environment, and their interactions with the other agents in the model are controlled by rules. The key feature of such systems is emergence - high-level organisation arising from the relatively simple interactions of the agents.

Our second area is machine learning, a subset of artificial intelligence that is concerned with adaptive learning. While there has been some interest in machine learning for various aspects of medical data, particularly the analysis of medical images and signals [12,13] and the automatic diagnosis of patients [14,15,16], we are proposing that machine learning methods should also be considered for the analysis of many types of health data.

Complex Adaptive Systems

Introduction to Complex Adaptive Systems

Complexity theory, the study of complex adaptive systems, argues that it is the interactions between 'agents' that is important. This differs from the standard reductionist methods, where the component parts are each examined in isolation, so that the interactions are not considered at all.

The agents in the complex adaptive systems can be as simple as individual neurons within a brain and as complex as companies competing for business, but it is the underlying interactions that produce overall behaviour that is far more complicated than that

of any individual agent. For example, in essence a neuron simply integrates the inputs from incoming synapses and, if those inputs are above some threshold, 'spikes', sending an impulse along outgoing synapses. Yet nothing more than a conglomeration of such interacting neurons produces human intelligence.

Again, vastly complex human societies, for example economies, arise from the interactions of individual people going about their everyday lives, buying and selling goods in order to satisfy their needs. As can be seen in this second example, the agents themselves can adapt and change to their environment, learning from past experience and tailoring their responses based on their knowledge. Rather than passively responding to events, they take advantage of any knowledge that they acquire.

The key feature of complex adaptive systems, then, is the emergence of global patterns from the local interactions between agents, a property often described as self-organisation; local interactions between agents that are close together leads to the emergence of global orderings. In interesting systems the global patterns cannot be predicted in advance; even relatively simple rules provide very varied, or even chaotic, behaviour; see for example cellular automata [16,17].

Complex Adaptive Systems for Health Informatics

So what can health informatics learn from complexity theory? The first and most obvious point is that any healthcare environment is a complex adaptive system, comprised of many multi-layered and varied interactions between agents such as patients, doctors and nurses. Models of these interactions can aid in the understanding of the dynamics between healthcare professionals, and can look at how these interactions vary when systems are changed. This can complement other research, for example looking at the introduction of computer-based clinical information systems [18].

However, complex adaptive systems can lead to insights in many other places. For example, in the introduction it was pointed out that the cycle of clinical quality review and improvement meant that the environment in which clinical care is given changes. This is hard to model using more traditional methods, but complex adaptive systems evolve continuously and are hence inherently dynamic. Furthermore, the feedback cycles in the system can be identified and studied.

Other places where complex adaptive systems can be usefully employed in health informatics include:

- Disease spread modelling [19]
- Choice modelling [20]
- Population level action against disease [19]
- Management [21]

Proposed Research

As this very brief overview has shown, there are many areas in healthcare where understanding the interactions between agents is crucial to understanding the underlying processes, and where the environment under study is varying over time. This means that traditional methods of analysis, such as statistical modelling, are not suitable.

Theory from complex adaptive systems could be used to model and investigate these interactions. As yet, there has been little or no work in this area, yet the potential benefits are huge.

Machine Learning

Machine learning and artificial intelligence have already been used in areas such as disease diagnosis and the analysis of medical images and signals. However, we believe that the application of machine learning techniques is significantly underdeveloped in general health informatics, and especially in the area of clinical quality. The successes that have been had applying techniques from this area to bioinformatics suggests that that there is much scope for investigation in the area.

Typical machine learning methods adapt to more closely model any data that they are presented with. In this way they can actually be used to model the individual adaptive agents within a complex adaptive system, but in addition, they are useful in their own right. Other applications of machine learning include identifying the dominant features, or combinations of features, that affect the outcomes of a process and performing adaptive regression and classification. For a review of methods, see for example [22,23].

Machine Learning in Health Informatics

An example of the benefits of machine learning techniques for clinical quality analysis can be seen in [4]. In that paper they consider the question of detecting abnormal death rates among general practitioners based on features such as age, gender and quarterly capitation figures for age group and locality by using Statistical Process Charts (SPC).

When data about a particular practitioner is found to be abnormal by the SPC it is analysed by hand and additional factors are taken into consideration. Clearly, such a process can only be considered for small-scale studies. Machine learning techniques such as novelty detection [24,25] could be usefully employed in such problems [5]. Novelty detectors have been employed in diagnostic systems, most notably the detection of masses in mammograms. The novelty detector learns a model of 'normality' for example, mammograms that have been inspected and do not show any examples of masses. In operation the filter then highlights any input that differs in some way from those that have been seen previously; for an overview see [25].

Proposed Research

There are a huge number of machine learning methods available to the health informatician. In addition to the methods mentioned above, of particular interest to health informatics are methods that identify connections between input features and suggest networks of interaction, such as Bayesian networks [26], and principled methods of dimensionality reduction such as Principal Component Analysis [27], Independent Component Analysis [28] and Local Linear Embedding [29].

A suitable program for research will critically examine those places where standard statistical techniques are used now and will ask whether or not they are appropriate, or whether non-linear methods of learning should be used, in conjunction with methods of choosing suitable input features using correlation analysis and dimensionality reduction.

Discussion and Conclusions

The amount of clinically-relevant data that is produced and stored is growing day by day, and finding methods of reliably analysing this data automatically is one of the major tasks of health informatics. In this paper we have identified a number of problems that make traditional methods of statistical modelling unsuitable:

- large, noisy datasets with missing entries
- · mixture of quantitative and qualitative data
- data with low time granularity, but high space granularity
- · dynamic cycles with feedback loops
- large numbers of interacting agents with many multilevel and varied interactions
- complex (non-linear) interactions between features that are not known a priori.

We have then gone on to propose that the solutions to these methods can be found in the fields of complex adaptive systems and machine learning. Complex adaptive systems explicitly model the interactions between groups of learning agents, with the overall behaviours of the system emerging from these interactions. In this way we can identify feedback cycles, and examine whether or not stable solutions emerge from the interactions of the agents from the system is perturbed from equilibrium.

While machine learning has been used for medical tasks, such as the detection of masses in mammograms and classification, the extension to clinical quality data has been lacking. We have shown that it has benefits over statistical methods such as Statistical Process Charts in that the features that will be used as inputs do not need to be selected as carefully, and have highlighted novelty detection as one method that would be particularly useful for highlighting deviations from good clinical practice. Furthermore, we have pointed out the importance of dimensionality reduction, in order to reduce the computational costs and difficulty in dealing with the data, and Bayesian networks for learning the links between sets of correlated inputs and outputs from a model.

In addition, we have drawn parallels between health informatics – which is now having to deal with the problem of turning data into useful information that can be used to improve clinical quality – with bioinformatics, which has built up a symbiotic relationship with machine learning over the past few years, and which suffers from many of the same problems – large amounts of noisy data, feedback cycles between genes, and a shortage of temporal resolution.

The intelligent application of complex adaptive systems and machine learning to clinical data has the potential to revolutionise health informatics. In this paper we have suggested a few areas that are particularly suited to these methods and identified some potential techniques. Bringing this about is the task of the wider community.

Acknowledgements

SM is funded by the MIAS IRC project, EPSRC grant number GR/ N14248/01.

References

- [1] Department of Health. Building the Information Core -Implementing the NHS Plan. London, The Stationery Office
- [2] Balas A and Ingram D, Health Information Infrastructure for Managing Knowledge: Public Health Response to Citizen Expectations in the U.K. and the U.S<u>Maybe 'An Evolving Response to Citizen...</u> Commonwealth Fund-Nuffield Trust Conference, 2003.
- [3] Bird SM, Cox D, Farewell DT, Goldstein H, Holt T, Smith PC. Report of the RSS Working Party on Performance Monitoring in the Public Services. Royal Statistical Society, London 2003.
- [4] Aylin P, Best N, Bottle A, Marshall C. Following Shipman: a pilot system for monitoring mortality rates in primary care. Lancet 2003;362:485-491.
- [5] Buchan I, Marsland S, Beatty P. Re. Following Shipman: a pilot system for monitoring mortality rates in primary care. Lancet (in press 2003).
- [6] Bloor K, Freemantle N, Khadjesari Z, Maynard A. Impact of NICE guidance on laparoscopic surgery for inguinal hernias: analysis of interrupted time series. BMJ 2003; 326: 578.
- [7] Baldi P and Brunak S. Bioinformatics: The Machine Learning Approach, MIT Press, 1998.
- [8] Campbell, C and Mukherjee S. Neural Information Processing Systems Workshop on "Machine Learning Techniques for Bioinformatics", 2001.
- [9] Pavlovic V, Garg S, Kasif, S. International Conference on Machine Learning Workshop on "Machine Learning in Bioinformatics", 2003.
- [10]Prigogine I and Nicolis G, Self-Organization in Non-Equilibrium Systems: From Dissipative Structures to Order Through Fluctuations, J. Wiley & Sons, 1977.
- [11]Prigogine I and Nicolis, G, Exploring Complexity, W. H. Freeman & Co., 1989.
- [12] Veropoulos K, Campbell C, and Learmonth G, Image processing and neural computing used in the diagnosis of tuberculosis. Colloquium on Intelligent Methods in Healthcare and Medical Applications, 1998.
- [13] Neves J, Alves V, Nelas L, Romeu, A and Basto S, An information system that supports knowledge discovery and data mining in medical imaging, Proceedings of Workshop on Machine Learning in Medical Applications, 37-42, 1999.
- [14] Shwe M, Middleton B, Heckerman D, Henrion M, Horvitz E, Lehmann H and Cooper G. Probabilistic diagnosis using a reformulation of the INTERNIST-1/QMR knowledge base I: The probabilistic model and inference algorithms. Methods of Information in Medicine 30: 241-255, 1991.

- [15] Jaakkola T and Jordan MI, Variational Probabilistic Inference and the QMR-DT Network, Journal of Artificial Intelligence Research 10:291-322, 1999.
- [16] Gutowitz H, Cellular Automata and the Sciences of Complexity, Part I. Complexity 1(5), 1996.
- [17] Crutchfield JP, The Calculi of Emergence: Computation, Dynamicsm and Induction, *Physica D*, 75:11-54, 1994.
- [18]Munir SK and Kay S, Organisational culture matters for systems integration in health care. In press, proceedings AMIA, 2003.
- [19]Terrison B, http://www.complexityprimarycare.org/ Ex02SpeakerNotes.htm, 2001. Accessed at 14/09/2003.
- [20] The Economy as an Evolving Complex System II. Proceedings Volume XXVII, Santa Fe Institute Studies in the Science of Complexity, Arthur WB, Durlauf SN, Lane WA, Addison-Wesly, 1997.
- [21]Mitleton-Kelly E, Complexity: Partial Support for BPR. In Systems Engineering for Business Process Change. Edited by Henderson P, Springer-Verlag, 2000, pp. 24-37.
- [22] Hastie T, Tibshirani R, Friedman T, The Elements of Statistical Learning, Springer Verlag 2001.
- [23]Bishop C, Neural Networks for Pattern Recognition, OUP 1995.
- [24] Tarassenko L, Hayton P, Cerneaz N, and Brady M, Novelty detection for the identification of masses in mammograms. Proceedings of the 4th IEE International Conference on Artificial Neural Networks, pages 442 – 447, 1995.
- [25]Marsland S, Novelty Detection in Learning Systems, Neural Computing Surveys, 3:157--195, 2003.
- [26]Heckerman D, A Tutorial on Learning with Bayesian Networks. In Learning in Graphical Models, M. Jordan, ed.. MIT Press, pages 301 – 354, 1999.
- [27] Jolliffe IT, Principal Component Analysis, Springer Verlag,
- [28] Hyvärinen A, Karhunen J and Oja E, Independent Component Analysis, John Wiley & Sons, 2001.
- [29] Roweis S and Saul L, Nonlinear dimensionality reduction by locally linear embedding. Science 290(5500):2323— 2326, 2000.

Address for correspondence

Dr Stephen Marsland
Imaging Science and Biomedical Engineering
Stopford Building
University of Manchester
Manchester M13 9PT, UK.
stephen.marsland@man.ac.uk

Phone: +44 (0)161 275 5140 Fax: +44 (0)161 275 5145