Chronic Disease Management: a Business Intelligence Perspective

Leelani Kumari Wickramasinghe[1]   **Damminda Alahakoon**[2]   **Michael Georgeff**[1]

Peter Schattner[1]   **Daswin De Silva**[2]   **Oshadi Alahakoon**[2]   **Akuh Adaji**[1]

Kay Jones[1]   **Leon Piterman AM**[1]

[1] Department of General Practice
Monash University,
Melbourne, Australia.
Email: `Leelani.Wickramasinghe@monash.edu`

[2] Clayton School of Information Technology
Monash University,
Melbourne, Australia.
Email: `Damminda.Alahakoon@monash.edu`

Abstract

Chronic disease management is one of the main areas in healthcare that Health Knowledge Management (HKM) can provide beneficial outcomes. Information Communication Technology (ICT) enabled Chronic disease management network (cdmNet) delivers comprehensive chronic disease management solutions by incorporating all key processes of the Chronic Care Model (CCM) developed by Wagner and his group. Through its roll-outs, cdmNet has accumulated detailed data about the chronic disease management process. This paper presents a new Business Intelligence (BI) module developed to analyse, visualise and extract knowledge from the cdmNet data. The aim of the BI module is to facilitate the short-term and long-term decision making and improve understanding of collaborative care models, policy and economic models underlying chronic disease management. The paper contains preliminary results obtained from applying the BI module to the cdmNet data.

Keywords: Chronic disease management, data mining, business intelligence

1 Introduction

Healthcare Knowledge Management (HKM) is defined as the systematic creation, modeling, sharing, operationalization and translation of healthcare knowledge with the emphasis to improve the quality of patient care [Sibte and Abidi, 2008]. The aim of HKM is to provide high quality, well-informed and cost-effective patient care decisions to healthcare stakeholders such as government, healthcare professionals and even patients themselves. Due to the increasing number of patients with chronic disease and the associated medical care costs [Anderson and Wilson, 2006, Anderson and Johnson, 2004], chronic disease management is one among many other areas of healthcare, that HKM can provide beneficial outcomes.

The Chronic Care Model (CCM) developed by Wagner and his group [Wagner et al., 2001a] em-phasises collaboration among care providers and the patient in creating and maintaining a care plan for patients with chronic disease. A recent Australian initiative, chronic disease management network (cdmNet) [Georgeff et al., 2010, Georgeff and Hilton, 2010] is an ICT implementation to put the CCM into practice. cdmNet focuses on the processes which are key to the CCM by changing chronic disease management domain from point-to-point, episodic, referral to a continuous, collaborative, networked model. Through its already rolled-out phases and future roll-outs, cdmNet, for the first time has accumulated detailed data about the process of chronic care management. The knowledge concealed within the raw data have the potential to provide an evidence base for the CCM and could be the key to informed health policy for patients with chronic disease.

Business Intelligence (BI) refers to computer-based techniques to spot, drill down (for detailed information), roll up (for abstract information) and analyse business data [Larson, 2008]. BI technologies are incorporated with functions such as reporting, online analytical processing, analytics and data mining to provide historical, current and predictive views of business operations.

The aim of this paper is to demonstrate the application of BI techniques to systematically create, model, share and translate data generated from chronic disease management process which enables data driven decision making. Application of BI techniques to the cdmNet data is threefold: (1) to interpret prevailing chronic disease management process from stakeholders' perspectives; (2) to extract interesting patterns hidden within data; and (3) to provide a basis for validating the Chronic Care Model. All these are key for: (1) short-term and long-term decision making; and (2) improving understanding of collaborative-care models, policy, and economic models underlying chronic disease management. The cdmNet-BI module presented in the paper aims to fulfil the aforementioned three criteria, thus providing guidance to improve the quality of chronic disease patient care.

The paper is organised as follows: Section 2 provides background material on chronic disease management and the cdmNet system. The cdmNet Business Intelligence (cdmNet-BI) module is described in Section 3. Section 4 presents the preliminary results obtained by applying this BI module to the cdmNet data. Concluding remarks and future work are discussed in Section 5.

2 Background

The basis of this paper is chronic disease management and the cdmNet system which accumulates chronic disease process related data. Brief introductions to chronic disease management and cdmNet system are provided in Sections 2.2 and 2.1 respectively.

2.1 Chronic Disease Management

Chronic diseases such as type 2 diabetes, asthma, and arthritis generally cannot be cured completely. Hence they persist through patients' lifetime and require adequate management [2007]. In Australia, the management of a patient with a chronic disease, tends to emphasis on following the best-practice clinical guidelines. Once an individual is identified with a chronic condition that requires ongoing management, a General Practitioner (GP) may tend to create a General Practice Management plan (GPMP) (a care plan in the context of this paper) from their assessment of the patient's underlying conditions and care objectives. Care plans include goals, strategies and specific tasks. These will include services by other care providers (e.g., diabetic education, podiatry service), clinical interventions by the GP (e.g., tests such as lipid tests ordered) and any medication required for managing the chronic condition. Once the GPMP is created, the GP identifies possible healthcare providers who can provide the services listed in GPMP and creates a formal document called a team care arrangement (TCA). TCA includes the details of healthcare providers who provide each service identified in GPMP. GPMPs and TCAs are known as Medicare Benefits Schedule (MBS) items [2009]. Over the course of the GPMP/TCA lifetime the GP reviews the execution of the plan, follows up whether the patient is attending appointments (either directly or through a carer), whether key tests or treatments have been performed, and monitors patient's condition for changes.

Prior to cdmNet and in the regions in which cdmNet is not yet rolled-out, GPs use chronic disease management care plan templates available in GPs' desktop application for care plan creation. Template based chronic disease management is not efficient and does not comply with the CCM due to following limitations:

- Templates provide a guideline, requiring some manual personalisation that demands GP's time;

- Lack of care coordination; and

- Either limited or no support for patients in adhering to care plans ensuring appointments are made, visits are attended and medications are renewed.

cdmNet provides solution to these issues by [Georgeff and Hilton, 2010]:

- Creating best practice, personalised care plans;

- Distributing care plans to the patients' care team and to the patients;

- Monitoring continuously the care plan, medication renewals, and appointments;

- Ensuring timely follow up;

- Facilitating collaboration by sharing the health record, care plan, and progress against the care plan among the care team and with the patient; and

- Supporting patient self-management by sending alerts, reminders, and notifications.

cdmNet was trialled in the Barwon South Western Region (BSWR) of Victoria and the Eastern Goldfields Region (EGR) of Western Australia. These trials involved 97 GPs, 268 other healthcare providers (including practice nurses), and 733 patients with diabetes. cdmNet is currently being rolling out, covering a population of over 1.2 million, in Melbourne metro region and regional Victoria, Queensland, Tasmania and Western Australia.

2.2 Chronic Disease Management Network (cdmNet)

The aim of cdmNet is to reduce the problems associated with template based chronic disease management by implementing most of the elements of the CCM. cdmNet identifies processes as the key to the CCM and breaksdown the CCM processes to stages as follows [Georgeff and Hilton, 2010]:

- Planning: identifies chronic disease sub-population and creates best-practice, personalised care plans;

- Collaboration: identify healthcare providers to be part of the care plan, develop agreements and distribute care plans among the care providers; and

- Monitoring: track compliance with care plan by care team and patient in real-time, monitor key patient health parameters (e.g., blood glucose) and support adherence by sending alerts, reminders and notifications.

Through these processes, cdmNet accumulates data on:

- Patients: demographics (e.g., gender, age, marital status), lifestyle (e.g., drinking status, smoking status), measurements (e.g., blood pressure, body weight), and medications (e.g., name, strength, dose);

- Items on care plans: service, goal, target, care provider, frequency;

- Appointments with providers: With whom, when, service, status;

- Information about MBS items: type, status;

- Use of cdmNet adherence support services: type (e.g., email, sms), recipient, when; and

- cdmNet web page accesses: accessor, when, which page

The aim of cdmNet-BI module is to harvest tacit knowledge concealed within this detailed data about the process of chronic care management for decision support. Several key performance indicators, such as the use of cdmNet for care planning by GPs, have been established to ascertain the performance of cdmNet.

3 cdmNet Business Intelligence Module (cdmNet-BI Module)

The purpose of the cdmNet-BI module is to convert above mentioned transactional data gathered in and/or generated from cdmNet operational processes to knowledge. It consists of three sub-modules: (1)

pre-processing, (2) dashboard, and (3) data mining as shown in Figure 1. The pre-processing sub module, as described in Section 3.1 converts transactional data to a format that can be used in dashboard and data mining sub modules. The dashboard and the data mining sub modules aim to provide solutions for closed and open questions related to chronic disease management process. Closed questions are answered by distinct values without any uncertainty from the cdmNet data. Solutions to open questions contain certain degree of uncertainty and may have many possible answers.
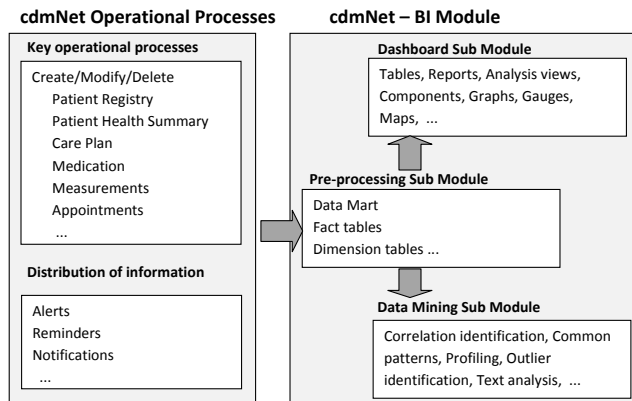

Figure 1: cdmNet-BI module.

3.1 Pre-processing Sub Module

The cdmNet data generated by day-to-day transactions, as any other business data are stored in databases which adopt database normalization rules. When analysing the data, speed of access is the main criteria to be considered. In addition, depending on the analyses to be performed, it may not require all transactional data. Incorporating these requirements, the pre-processing sub module performs two types of functions:

1. Identifies data required for analysis; and

2. Converts data identified in Item (1) above to gain speed of access to suit analysis purposes.

After identifying the data required for analysis, the pre-processing sub module creates a data mart to contain only the required data and populates it using extract, transform, and load (ETL) process [Larson, 2008]. Data mart is a denormalised (repeated) relational database designed for speed of access [Larson, 2008]. When data are denormalised, analysis and reporting can be performed using only a few table joins thus increasing the speed.

A data mart structure is based on four elements: measures, dimensions, attributes and hierarchies [Larson, 2008]. A measure is a numeric quantity that represents some aspect of the operational system. Number of care plans generated, number of over weight diabetes patients, and number of users who access cdmNet web pages are few examples of measures in this application. The measures are stored in tables called fact tables; for example CareplansFact, PatientWeightFact, and PageAccessesFact.

Dimensions are the perspectives in analysing data. For example, often stakeholders are not interested in the total number of care plans generated by the system. Instead they prefer to slice and dice this total into its constituent parts. For example, commonwealth government may be interested in comparing

the performances of each state while a particular state government is interested in comparing different general practice divisions within the state. These contributing components such as state and general practice division are known as dimensions.

Attributes refer to any additional information stored about dimension members. For example, general practice division name is an attribute of the general practice dimension. Usually, a dimension is part of a large structure known as a hierarchy. For example, state and division may have their own hierarchies. In summary, all data mart components are created and populated by the pre-processing sub module.

3.2 Dashboard Sub Module

The Dashboard sub module is designed to provide two functionalities:

- to provide solutions to closed questions; and

- to present solutions in line with decision makers' objectives.

Few examples of chronic disease closed questions are:

1. For each month how many care plans are created in a particular state (say Victoria) by a particular GP (say Dr. John Smith) for a given time period (say from 2010-01-01 to 2010-06-30)?

2. Which state has generated the most number of care plans?

3. How many chronic disease patients are listed within a given GP and what are their demographics?

4. Compare the number of care plans created and reviewed by a given GP with state, division and organisation averages; and

5. Compare care planning performances between and among states, divisions and GPs.

As an end result, the dashboard sub module represent all key performance indicators in respect to stakeholders. Easy access and visualization are the main considerations of this sub module. Fulfilling these requirements, the sub module creates portals for different stakeholders. Portals include tables, reports, analysis views, components (radio, select, check), graphics (bar charts, pie charts), gauges (dial charts, traffic light) and maps. All these user interface components visually illustrate key performance indicators.

3.3 Data Mining Sub Module

As any other real world system, chronic disease management process is concerned with open questions, which may have multiple possible solutions with varying degrees of certainty. Few examples of such open questions are:

1. What are the measures that determine health outcomes in chronic disease management? Possible measures can be:

   - changes in patient measurements such as body weight, blood pressure; and/or

   - changes in number of hospitalisations or work force participation;

2. Which variables determine the health outcomes identified in Item (1) above? Is it

- the chronic disease?
- number and type of service in care plan?
- patient's adherence to the care plan?
- regular reviews and follow-ups by GPs?
- patient's interaction with cdmNet web site?

Hence, the aim of the data mining sub module is twofold: (1) to provide possible solutions for open questions; and (2) to extract interesting patterns from data. The outcomes of the latter may in turn provide guidance to the former aim.

Data mining is defined as non-trivial discovery of novel, valid, comprehensible and potentially useful patterns from data [Fayyad et al., 1996]. This sub module aims to use both directed and undirected knowledge discovery techniques [Berry and Linoff, 2004] to provide functionalities such as clustering, classification, prediction, association mining, and text mining [Smith, 2000].

In directed knowledge discovery process, there is a goal which directs the knowledge discovery process. For example, we can select certain variables, such as patients' demographics (gender, age, weight) and lifestyle (smoking status, drinking status) as inputs and blood glucose level as output. A model can be trained to learn the relationship between inputs and output using a subset of the cdmNet data. The rest of the data can be used to validate the knowledge discovered. Once the model is trained to a certain confidence level, it can be used to predict blood glucose level of a new patient when demographics and lifestyle are known.

Undirected knowledge discovery techniques are used to investigate the underlying structure of data or for pattern recognition. These techniques neither distinguish between variables and outcomes nor contain a specific goal. All data are considered as parameters and the technique determines relationships that may exist among the parameters. Currently, the data mining sub module uses clustering as the undirected knowledge discovery technique. Clustering algorithms classify data into groups based on the natural structure of data without obtaining additional information about data from a data analyst [Smith, 2000]. Specifically, self-organizing map (SOM) [Kohonen, 1982a,b] is used as a clustering technique as SOM has the ability to group data patterns based upon the natural associations found in data [Smith, 2000]. In addition, SOM provides cluster visualisation by mapping high dimensional input space (all parameters) to low dimensional feature space. Feature space only contains variables that significantly contributed towards cluster separation.

4 Application of the cdmNet-BI Module to cdmNet Data

After evaluating a number of BI tools, ETL functionalities of the pre-processing sub module and the dashboard sub module are built using Pentaho BI suite community edition [2008]. Pentaho is a free, open source and Java based BI tool. The cdmNet system is built using Java technologies. Thus Pentaho makes it convenient to integrate the BI module with the cdmNet system. Data mart is built using PostgreSQL [1996] database management system as of the cdmNet system. Currently, for clustering purposes the data mining sub module uses Viscovery SOMine [2010] which is reviewed as a commercial standard tool for the SOM algorithm.

4.1 Results from the Pre-Processing Sub Module

This section describes the dimensions that can be used to represent the measure, number of care plans. We identified 7 dimensions: state, postcode, division, practice, GP, time and item to represent the measure (fact), 'number of care plans'. These dimensions provide single and combinations of perspectives to the given measure. An architecture called star schema is generally used for the design of a data mart [Kimball et al., 2002]. A star schema that represents this fact table and its dimensions is shown in Figure 2.
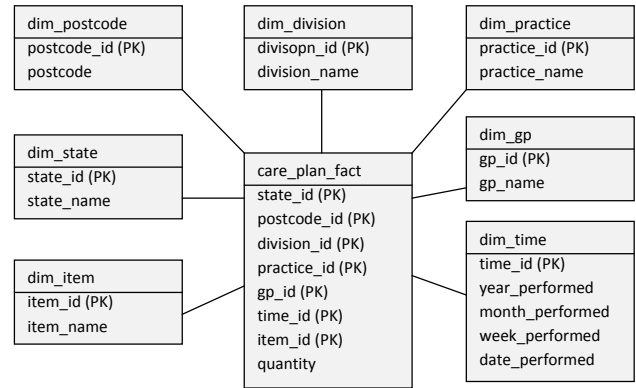

Figure 2: Star schema with dimension and fact tables.

A data mart using PostgreSQL was created to reflect the fact and dimension tables and populated with the cdmNet data.

4.2 Results from the Dashboard Sub Module

This section illustrates several screens developed using the dashboard sub module. An analysis view with drill down (for detailed information) and roll up (for abstract information) capabilities is shown in Figure 3. It is a visual representation of all dimensions and the fact described in Section 3.1. Its navigational capabilities provide stakeholders with the ability to investigate along one or many dimensions depending on their requirements.

A portal developed for the GPs is shown in Figure 4. The table at the top left hand corner shows cumulative monthly revenue and the corresponding line chart is shown at the top right hand corner. The pie chart at the bottom right hand corner shows the percentage of care plan items created by the GP during the current month. The bar chart at the bottom left hand corner compares the GP's performance against the average of all GPs, average of top GPs (who have created more than a specific number of care plans) and the national average.

### 4.3 Results from the Data Mining Sub Module

Preliminary results obtained from the data mining sub module to identify the factors that contribute to diabetes chronic condition is described in this section. HbA1c or haemoglobin A1C is a metabolic measure used to diagnose diabetes. HbA1c is a molecule created in red blood cells when glucose sticks to them. A normal non-diabetic HbA1c is 3.5-5.5%. With diabetes, about 6.5% is considered normal. Achieving HbA1c $< 7\%$ is considered as a diabetes management goal [Wagner et al., 2001b].

| State	Postcode	Division	Practice	GP	Time	Item	Quantity
⊟All States	⊞All Postcodes	⊞All Divisions	⊞All Practices	⊞All GPs	⊞All Times	⊞All Items	4,031
SA	⊞All Postcodes	⊞All Divisions	⊞All Practices	⊞All GPs	⊞All Times	⊞All Items	149
VIC	⊞All Postcodes	⊟All Divisions	⊞All Practices	⊞All GPs	⊞All Times	⊞All Items	2,539
		Dandenong Casey	⊞All Practices	⊞All GPs	⊞All Times	⊞All Items	4
		General Practitioners Association of Geelong	⊞All Practices	⊞All GPs	⊞All Times	⊞All Items	2,400
		Monash Division of General Practice	⊟All Practices	⊞All GPs	⊞All Times	⊞All Items	103
			East Bentleigh Medical Group	⊞All GPs	⊟All Times	⊞All Items	103
					⊞2008	⊞All Items	4
					⊞2009	⊞All Items	45
					⊟2010	⊞All Items	54
					⊞2010 January	⊟All Items	11
						Allied Health Referral Form	2
						GPMP	3
						GPMP Review	2
						HMR	2
						TCA	2
					⊞2010 February	⊞All Items	10
					⊞2010 March	⊞All Items	3
					⊞2010 April	⊞All Items	3
					⊞2010 May	⊞All Items	1
					⊞2010 June	⊞All Items	14

Figure 3: Analysis view with drill down and roll up capabilities.

Cumulative Revenue From Items and Reviews

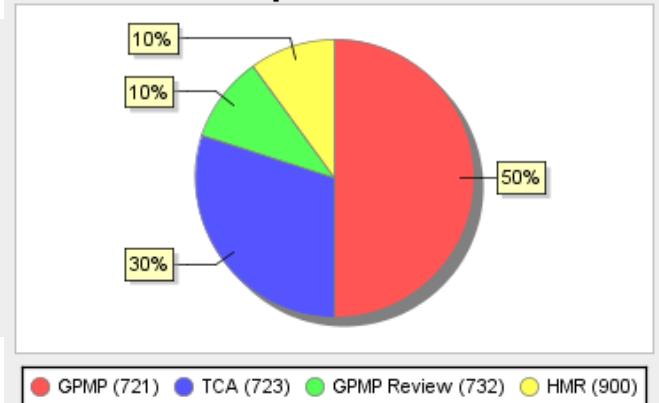| Year-Month	GPMP (721)	TCA (723)	GPMP Review (732)	TCA Review (732)	ACoC	HMR (900)
2010 January	$ 133.65	$ 105.90	$ .00	$ .00	$ .00	$ 143.40
2010 February	$ 133.65	$ 105.90	$ .00	$ .00	$ .00	$ 143.40
2010 April	$ 133.65	$ 211.80	$ 133.60	$ .00	$ .00	$ 143.40
2010 June	$ 267.30	$ 211.80	$ 133.60	$ .00	$ .00	$ 143.40
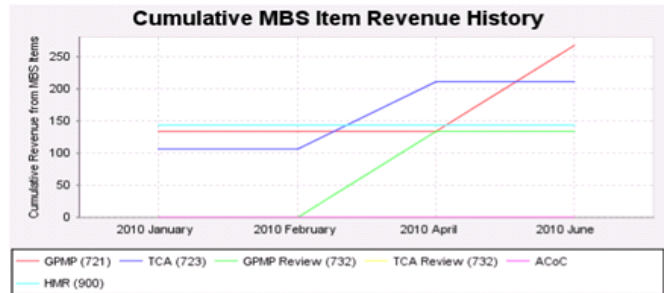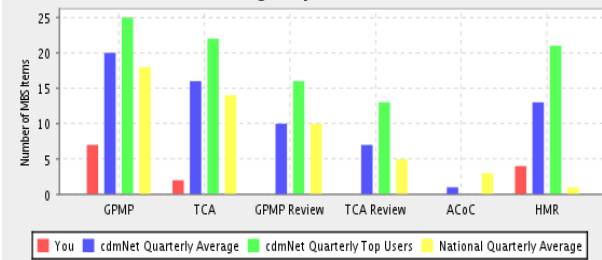
Figure 4: A screen from the GP portal. The table and the line chart show the cumulative monthly revenue of the GP from care plan items for a given time period. The pie chart shows the percentage of care plan items created by the GP during the current month. The bar chart compares the GP's performance against the average of all GPs, average of top GPs (who have created more than a specific number of care plans) and the national average.

Chronic diseases such as diabetes can be directly related to patients' lifestyle and demographics. This experiment investigates the impact of patients' lifestyle and demographics on HbA1c. It was carried out based on three categories of parameters: (1) demographics; (2) lifestyle; and (3) metabolic measures. The parameters considered for each category include:

1. Demographics: gender, age, marital status;

2. Lifestyle: drinking status, smoking status; and

3. Metabolic measure: HbA1c.

To determine the impact, for each patient, HbA1c measured over a time period is considered. That is, the experiment includes patients whose HbA1c level is measured more than once. In the original data set, there are 657 patients with HbA1c readings, but only 338 have measured it more than once. The total number of readings for a given patient varies between 2 and 42. Therefore, data corresponding to these 338 patients were used. In addition to the parameters discussed in Section 4.3, this experiment includes a number of parameters which describe statistical measures of HbA1c and trend in consecutive HbA1c measurements as follows:

- Parameters relate to statistical measures of HbA1c: average HbA1c, minimum HbA1c, maximum HbA1c, and standard deviation of HbA1c; and

- Parameters relate to trend in consecutive HbA1c measurements: percentages of increased consecutive readings, decreased consecutive readings, and stable consecutive readings

For each patient, trend in consecutive HbA1c readings is calculated as:

$$percentage\ of\ increased\ consecutive\ readings = count(V_{t+1} > V_t)/N * 100$$

where V denote HbA1c reading, t denote time and N denote total number of HbA1c readings of that particular patient. Similar calculations are carried out for percentage decrease and percentage stable.

The experiments were carried out using clustering as the knowledge discovery technique, self-organizing map as the clustering technique and Viscovery SOMine application as the tool to obtain SOM clusters as mentioned in Sections 3.3 and 4. A few preprocessing steps were carried out to encode non-numeric data before applying the Viscovery SOMine application as follows:

- gender: binary data (0 - male, 1 - female);

- drinking status: binary data (0 - non drinker, 1 - drinker);

- smoking status: 0 - non smoker, 0.5 - ex-smoker, 1 - smoker; and

- marital status: 1 of N for single, defacto, married, widowed, divorced and separated;

Parameters that take continuous values are transformed using Sigmoid or Logarithmic transformations to obtain normal distributions.

- Sigmoid transformation: age, average HbA1c, minimum HbA1c, maximum HbA1c, percentage increase, percentage decrease, percentage stable; and

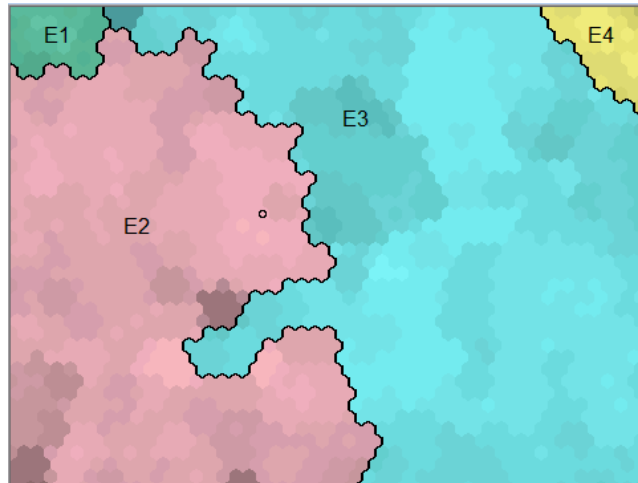- Logarithmic transformation: standard deviation of HbA1c

Figure 5: Visual representation of formed clusters.

A visual representation of clusters generated by the Viscovery SOMine application is shown in Figure 5. It consists of four clusters: E1, E2, E3 and E4. Table 1 contains the average values of each parameter contributing towards the formation of each cluster. Using the values in Table 1, an analysis of each cluster is carried out in Table 2 to determine the reason for the formation of clusters. Such reasoning has the potential of discovering knowledge.

According to Table 2, clusters E1 and E2 include patients with HbA1c > 7%. Patients in E1 has a tendency to increase HbA1c, while patients in E2 have fluctuating readings. The data highlight that while heavy drinking and previous smoking contribute to high HbA1c, whether HbA1c continues to increase or fluctuate is effected by gender and marital status.

Clusters E3 and E4 have patients with HbA1c < 7% and they continue to have stable readings. The data indicate that both moderate or no alcohol and non smoking determines low HbA1c. Age and marital status are the determining factors for seperating the two cluters.

Even though these patterns are not unknown or unforeseen, this technique:

- provides a guideline for stakeholders to identify patients who need more support (education, tests, other health professional visits) to manage their diabetes from patients from patients demographics and lifestyle; and

- identifies hypothesis for further testing, for example married males around 60 years of age who are heavy drinkers and ex-smokers find it hard to control their diabetes.

5   Conclusions and Future Work

This paper presented a Business Intelligence (BI) module developed to analyse, visualise and extract knowledge from the data accrued in the cdmNet system. BI module consists of three sub modules: (1) pre-processing; (2) dashboard; and (3) data mining. The purpose of the pre-processing sub module is to convert the cdmNet data to a form suitable for fast access. The dashboard sub module provides interfaces with drill down, roll up and graphic components. It enables stakeholders to navigate through information to understand the prevailing chronic disease management from individual stakeholder perspectives. The data mining sub module aims to extract patterns from

Table 1: Cluster Summary.

	E1	E2	E3	E4
gender	0.522	0.236	0.518	0.685
single	6.94E-019	0.0467	0.112	0
defacto	0.000429	0.0208	0.00577	7.38E-020
married	0.0259	0.384	0.324	7.32E-013
widowed	5.9E-018	5.52E-018	0.00121	0.977
divorced	0.974	0.000275	1.56E-005	0
separated	2.17E-019	0.0152	1.02E-019	7.38E-020
drinking status	0.998	0.816	0.738	0.29
smoking status	0.425	0.403	0.287	0.129
age	63.6	61	68.2	79.6
average HbA1c	**7.41**	**8.17**	**6.6**	**6.78**
minimum HbA1c	6.25	7.15	6.14	6.17
maximum HbA1c	10	9.55	7.19	7.48
standard deviation	1.58	1.04	0.414	0.445
% increase	**59.6**	**42.5**	**39.9**	**33.2**
% decrease	**23.9**	**46.5**	**39.2**	**31.7**
% stable	**15.7**	**10.8**	**20.6**	**33.5**
cluster quantization error	0.000349	1.3E-005	0.000116	2.39E-005

Table 2: Cluster Analysis.

Cluster	Patients's demographics and lifestyle	Average HbA1c	HbA1c variation
E1	divorced, elderly (average age of 64) males and females who are heavy drinkers and ex-smokers	> 7%	high tendency to increase
E2	males, considerable number of married and mostly with unknown marital status who are heavy drinkers and ex-smokers with an average age of 61	> 7%	Fluctuates heavily above normal
E3	elderly (average age of 68) males and females, considerable number of married and mostly with unknown marital status who are moderate drinkers and non smokers	< 7%	Stable
E4	very elderly (average age of 80) widowed females who neither drink nor smoke	< 7%	Stable

the cdmNet data which may potentially provide solutions to many open questions underlying chronic disease management.

The paper contains only the preliminary results obtained using the BI module. Our future work include enhancing the data mining module: (1) to include other data mining functionalities such as classification, prediction, association mining, and text mining; (2) to identify measures that determine healthcare outcomes in chronic disease management; and (3) to identify correlation among all variables found in the cdmNet data and the potential health outcomes.

Acknowledgements

References

Postgresql, 1996. URL http://www.postgresql.org/. Accessed: 01. March 2010.

Chronic disease management, 2007. URL http://www.adgp.com.au/site/index.cfm?display=333. Accessed: 12. November 2007.

Pentaho community wiki home, 2008. URL http://wiki.pentaho.com/display/COM/Community+Wiki+Home. Accessed: 01. March 2010.

Chronic disease management (cdm) medicare items, 2009. URL http://www.megpn.com.au/Docs/MBS%20items/Overview%2520Fact%2520Sheet.pdf.

Viscovery, 2010. URL http://www.viscovery.net/publications/books. Accessed: 01. May 2010.

G. Anderson and K. Wilson. Chronic disease in california: Facts and figures, 2006.

H. R. Z. T. Anderson, G. and N. Johnson. Chronic conditions: Making the case for ongoing care, 2004.

M. Berry and G. Linoff. *Data mining techniques: for marketing, sales, and customer relationship management.* Wiley New York, 2004.

U. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy. *Advances in knowledge discovery and data mining.* MIT press, 1996.

M. Georgeff and J. Hilton. Why is telemedicine not more widely used? *Feature*, pages 32 – 41, 2010. URL www.pulseitmagazine.com.au.

M. Georgeff, J. Hilton, L. P. AM, C. Chapman, L. Segal, M. Paterson, P. Schattner, K. Jones, A. Adaji, D. Mortimer, T. D. AM, B. Costa, K. Fitzgerald, B. Rasmussen, and K. Sweeny. Cdm-net:a broadband health network for transforming chronic disease management. Final report, Precedence Health Care, 2010.

R. Kimball, M. Ross, and R. Merz. *The data warehouse toolkit: the complete guide to dimensional modeling.* Wiley, 2002. ISBN 0471200247.

T. Kohonen. Analysis of a simple self-organizing process. *Biological cybernetics*, 44(2):135–140, 1982a.

T. Kohonen. Self-organized formation of topologically correct feature maps. *Biological cybernetics*, 43(1):59–69, 1982b.

B. Larson. *Delivering Business Intelligence with Microsoft SQL Server 2008.* McGraw-Hill Osborne Media, 2008.

S. Sibte and R. Abidi. Healthcare Knowledge Management: The Art of the Possible. 4924:1–20, 2008.

K. Smith. Introduction to Neural Networks and Data Mining: For Business Applications. *Journal-Operational Research Society*, 51(6):771–772, 2000.

E. Wagner, B. Austin, C. Davis, M. Hindmarsh, J. Schaefer, and A. Bonomi. Improving chronic illness care: translating evidence into action. *Health Affairs*, 20(6):64, 2001a.

E. Wagner, N. Sandhu, K. Newton, D. McCulloch, S. Ramsey, and L. Grothaus. Effect of improved glycemic control on health care costs and utilization. *Jama*, 285(2):182, 2001b.