Prediction of High Risk of Deviations in Home Care Deliveries

Anders L Madsen ALM@HUGIN.COM
HUGIN EXPERT A/S and Department of Computer Science, Aalborg University, Denmark

Kristian G Olesen KGO@CS.AAU.DK
Department of Computer Science, Aalborg University, Denmark

Heidi Lynge Løvschall HEIDI.LYNGE.LOEVSCHALL@HJORRING.DK
Digitalization and Welfare Technology, Municipality of Hjørring, Denmark

Nicolaj Søndberg-Jeppesen NSJ@HUGIN.COM
HUGIN EXPERT A/S, Aalborg, Denmark

Frank Jensen FJ@HUGIN.COM
HUGIN EXPERT A/S, Aalborg, Denmark

Morten Lindblad M@DATAPROCES.DK
Dataproces Group A/S, Nibe, Denmark

Mads Lause Mogensen MM@TREATSYSTEMS.COM
TREAT Systems ApS, Aalborg, Denmark

Trine Søby Christensen TSC@LSI-ND.DK
Life Science Innovation North Denmark, Aalborg, Denmark

Abstract

This paper presents a real-world application of Bayesian networks to support existing home care quality supervision. In Denmark home care is delivered by municipalities, where the individual citizen is free to select the service provider, private or public. The aim of our work is to support the home care control process by identifying significant deviations automatically, pointing to reasons for a significant deviation and identifying future home care deliveries where there is a high probability of deviation between granted and delivered care to the individual citizen. Home care is granted as packages of time measured in minutes and we define a too high delivery rate as larger than 150%. In the municipality under study in this work (municipality of Hjørring), the supervision of home care delivery is a manual and time consuming process prone to human error. This paper presents the results of efforts to automate parts of the supervision using Bayesian network modelling and data analysis. The results of the pilot study show significant potential in applying Bayesian network modelling and data analysis to this challenge for the benefit of the municipality, the employees and the citizens.

Keywords: Bayesian network; learning from data; real-world application; software.

1. Introduction

Methods and techniques originating from Artificial Intelligence have matured to a point where they are applied in an increasing number of domains. Improved methods combined with a continuously growing body of data, have resulted in new applications, and have entered areas such as, e.g., accounting and administration. There is also a growing awareness in the public sector — an area that has particular requirements in the sense that decisions and rulings have to be backed up by arguments and reference to relevant legislation. This need for transparency points in the direction of explicit models of the domain under consideration — one of the virtues of Bayesian networks (Pearl, 1988; Jensen and Nielsen, 2007; Kjærulff and Madsen, 2013).

Bayesian networks are well-suited for handling uncertainty and noisy data in highly structured domains. They are efficient tools for combining information from literature, empirical data and expert knowledge into a single model representation taking uncertainty into account. Bayesian networks have a long history in health care. Recent work has focused on automatic construction of models for specific diseases such as kidney stones (Kazemi and Mirroshandel, 2018) or ulcers (Kaewprag et al., 2017), methodologies, such as image classification (Arias et al., 2016), general models for diagnosis based on electronic patient journals (Shen et al., 2018), and for models for risk prediction (Arora et al., 2019). A recent study by Kyrimi et al. (2020) provide a comprehensive analysis of Bayesian networks in health care and conclude that there are still unexploited potentials to be gained. Goodson et al. (2008) reports on the use of Bayesian networks for nursing home care quality, which is the application reported in the literature that appears to be closest to our work.

We report on a case study from the public administration in Denmark within the domain of home care. This is a field where substantial data is available, and the study combines analysis of this data with experience from practitioners to assist in various tasks, such as control, planning and explanation of deviations between planned and actual activities. A home care (also known as domiciliary or home help) worker or assistant provides help or care in people's own homes to elderly people or those with disabilities or special needs. They perform tasks related to care functions such as, for instance, cleaning, helping with shopping, helping with taking medicine. The aim of the present study is to apply Bayesian networks to predict a high risk of deviations in home care deliveries in a municipality. Granted home care is organised into packages that are measured in minutes. As some deviation between planned and actual activities are expected, we focus on a too high delivery rate, which is later defined more formally as a delivery that is more than 150% larger than the granted amount. The Bayesian network should support the municipality in reaching its objective to deliver the right amount of help at the right time to ensure that citizens can live in their own homes for as long as possible.

2. Preliminaries and Notation

A Bayesian network represents an efficient factorization of a joint probability distribution into a set of conditional distributions when the graph of the model is sparse. We consider models with discrete variables only and the use of function nodes. A Bayesian network with function nodes (Madsen et al., 2014) is an extension of Bayesian networks to include function nodes. A function node represents a numerical value computed from a mathematical expression. An expression may include probabilities computed through inference in a Bayesian network and may be used to define constants in mathematical expressions defining the content of a CPD. In this paper, we use function nodes to compute values based on the results of inference in a Bayesian network.

To guide the user in the process of identifying a possible explanation of a too high delivery rate, we will use Bayes Factor K (Good, 1960) defined as $K = \frac{P(\epsilon \mid h_1)}{P(\epsilon \mid h_2)}$, where ϵ is the set of evidence, h_1 is the hypothesis of too high delivery rate, and h_2 is the hypothesis of not too high delivery rate. This value is computed for each individual observation ϵ_i to obtain an indication of the weight of ϵ_i as an explanation of too high delivery rate. The value of K is interpreted using the scale suggested by Jeffreys (1961) and slightly modified by Lee and Wagenmakers (2013) shown in Table 1. The

K	Strength of evidence	Color
$K < 1$	Negative	RGB(255,255,255)
$1 \leq K < 3$	Barely worth mentioning	RGB(255,238,238)
$3 \leq K < 10$	Substantial	RGB(255,204,204)
$10 \leq K < 30$	Strong	RGB(255,170,170)
$30 \leq K < 100$	Very Strong	RGB(255,85,85)
> 100	Decisive	RGB(255,0,0)

Table 1: Interpretation of Bayes Factor K and colour coding.

last column of the table illustrates how the strength of evidence is illustrated to the user using RGB color model to set the intensity of the text background color.

The Bayesian network models are either constructed by hand using expert knowledge or from data as a Tree-Augmented Naive Bayes model (TAN) (Friedman et al., 1997), where parameter estimation is performed using the Expectation-Maximization (EM) algorithm (Lauritzen, 1995).

3. Domain of Application and Methodology

Denmark is organized into 98 municipalities and the municipality of Hjørring is a typical province centered around the city Hjørring in the northern part of Jutland. The municipalities have some degrees of freedom to form their organisation. Typical organisations include an authority to control that rules and regulations are followed, an economy section to deal with budgets and finances, a section for social welfare that grants help to citizens and sections that provide actual services. Municipalities in Denmark have a high level of digitalization, which enables automation and the use of artificial intelligence technology for data analysis. Hjørring is considered to be an innovative and forward thinking municipality with a keen interest in applying new technologies to solve problems more efficiently to the benefits of its citizens.

Package	Minutes per week		
(p)	Lower limit ($l(p)$)	Upper limit ($u(p)$)	Median ($m(p)$)
Mini Evening	1	100	63
Small Evening	101	150	126
Medium Evening	151	260	203
Large Evening	261	420	336
Maxi Evening	421	559	490
Mini Day	1	105	56
Small Day	106	230	168
Medium Day	231	440	329
Large Day	441	720	588
Maxi Day	721	1,365	1,050

Table 2: Packages for personal assistance and their sizes (minutes).

Area	Citizens	Recipients of home care
Hjørring	65,257	2,105
Denmark	5,781,190	145,408

Table 3: Number of citizens and recipients of home care in 2018 (Statistics Denmark, 2020).

Allocated help	Males	Females
All	722	1,383
Less than 2 hours	428.5	913.5
2 - 3.9 hours	117	176
4 - 7.9 hours	92.5	155.5
8 - 11.9 hours	44.5	76.5
12 - 19.9 hours	23	34.5
More than 20 hours	16.5	27

Table 4: Recipients of help according to allocated time (hours per week) and sex in 2018 (Statistics Denmark, 2020).

According to Danish law the municipalities have an obligation to provide personal assistance and care to citizens with reduced physical or psychological abilities. Home care is a complex enterprise involving several parts of the organisation. Help is granted by social welfare. An official visits the citizen and allocate assistance according to identified needs. Help is allocated in terms of packages where a package covers a certain range of assistance measured in minutes per week. Table 2 show examples of standard packages for personal care in the municipality of Hjørring.

The number of citizens and recipients of home care are shown in Table 3 together with the same numbers for the entire country. As can be seen Hjørring is approximately an average municipality according to population with a slight over-representation of recipients of home care. A more detailed overview of recipients is shown in Table 4. The majority of citizens receives less than 2 hours of help per week and quite few receives more than 12 hours. As a general rule of thumb, the cost (annually) of five minutes of home care weekly is approximately EUR 2,000. Citizens with a permanent need for substantial assistance are typical offered accommodation in a nursing home. It is striking that there is almost twice as many women as men among recipients of help, this is a reflection of differences in life expectancy.

In order to simplify the presentation and to make it as precise as possible, we will define a number of concepts related to the delivery of home care. Home care is *granted* to a citizen in terms of a *package* of help defined as a weekly service. The amount of help granted is measured in minutes where each package size has weekly lower and upper limits on time (Table 2). The content of the package is determined subsequently. Once home care is granted by the authorities, it is delivered through home visits by one or more home care workers. The actual delivery is planned by a planner and the planned amount of help for a visit may differ from the granted amount of help as long as the total delivered help stays within the limits of the package on average over a time period of, for instance, three weeks. A *delivered* amount of help is measured in minutes.

Definition 1 (Grant) *A grant is an allocation of services in the form of a package p. The size of a grant g is measured in minutes and denoted $g_t(p)$.*

Given a package p with lower limit $l(p)$ and upper limit $u(p)$ an amount $l(p) \leq g_t(p) \leq u(p)$ is assigned to the citizen on a weekly basis, see Table 2. The execution or planning of the home care may be different from the granted, e.g., a citizen and a home care provider may agree that the service is delivered only every second week. This is important to know when automating the control. The results become less stable as the time period for the analysis is decreases. A three weeks period is considered absolute minimum.

Definition 2 (Delivery) *A delivery is the visit v of a home care worker performing a care function. The duration of the visit v is measured in minutes and denoted v_t.*

Definition 3 (Delivery rate) *The delivery rate $\delta(w, V, p)$ for a set of deliveries $V = \{v^{(1)}, \ldots, v^{(n)}\}$ for a specific package p of size $g_t(p)$ is computed as*

$$\delta(w, V, p) = \frac{\sum_i v_t^{(i)}}{w \cdot g_t(p)},$$

where w is the number of weeks.

Definition 4 (Too high delivery rate) *A too high delivery rate for package p, deliveries V and weeks w is defined as $\delta(w, V, p) >= 150\%$.*

A citizen may be assigned multiple packages, e.g., a day and an evening package, but may not receive two day and two evening packages (a citizen may change package size within one week, which complicates the analysis).

The authorities have an obligation to control the delivery of home care services in order to reduce and correct deviations between granted and delivered service, as well as understand the reasons for deviations. Today the quality control is a manual and time-consuming process. An Bayesian network model linking different observations on the home care deliveries to a variable indicating a too high delivery rate may be a valuable decision support tool in the control process, as it could help pointing to possible reasons for a deviation. The controller can use this information as a tool to identify cases that should be investigated further and cases where there is an acceptable explanation of the deviation. Both options will help the controller to focus on the cases where action should be taken to reduce future deviations.

The aim of this study is to investigate the use of Bayesian networks to automate parts of the existing control process. The models are delivered as part of a web-based software system to support the employees at the municipality of Hjørring in performing their work as efficiently as possible.

3.1 Methodology

The study follows the CRISP-DM methodology (cross-industry process for data mining) as well as the guidelines provided by (Madsen et al., 2015). CRISP-DM is aimed at structuring projects involving substantial elements of data mining and machine learning. The resulting approach has six phases. The phases are performed iteratively covering the following issues: (1) Business Understanding, (2) Data Understanding, (3) Data Preparation, (4) Modelling, (5) Evaluation, and (6)

Deployment. The process has been beneficial for both knowledge engineers, domain experts and other staff from the municipality as a better understanding of the data held by the municipality has been obtained. The work has been highly explorative and one of the important results is an understanding of the opportunities and limitations of the data and the quality of the data. The CRISP-DM technique provides a systematic approach for conducting projects. A single iteration may suffice to establish a proof-of-concept, but typically repeated passes lead to a better understanding of the business and eventually to automated systems that support the mission of the organization.

4. Data Analysis

Data is extracted from multiple systems and aggregated on a weekly basis. Data is, in principle, delivered as two separate files for each week. One file includes master data on each citizen (e.g., social security number, civil status, gender, age, address) who has received a delivery and another file with information on the deliveries to each citizen in the specific week. This information is collected from the Business Intelligence systems of the municipality, but sensitive data, such as detailed health, registrations are not available.

The two data files are combined into a single file through a join operation using the social security number as unique key. The left part of the join operation is the master data on the citizen, while the right part is the list of deliveries to the citizen with information on date of delivery, package, start and end dates (on the package grant), amount of time granted, planned time and used time for this delivery. There a usually multiple deliveries for each citizen. The data is folded to a single vector with features computed from the raw data, possibly across multiple weeks. The folding of the data to a single row takes care of the time dimensionality of the deliveries through feature engineering. There are features for number of packages, number of visits, amount of time used, number of extra-visits, time used on emergency visits, and similar.

Data is noisy and quality is not perfect with, for instance, inconsistencies between files, reporting errors, missing information and procedures for reporting not being followed. Data is collected on a weekly basis, and the granularity could be increased to a daily basis, but this would introduce too much jitter due to fluctuations in the length of daily visits and precision in the registrations. By aggregating data on a weekly basis such variations are smoothed out. The user selects the number of weeks to include in the data. This is usually a multiple of three due to how some packages are granted. The objective is to increase robustness of the results ensuring that short-term deviations are not a major issue, e.g., the citizen may be on vacation or the citizen and the service provider may agree to change a granted weekly visit of 15 minutes to a biweekly visit of 30 minutes. This complicates the analysis significantly and means that the analysis should be performed over some period of time (e.g., three or six weeks) to cope with such fluctuations in the data.

5. Models

The objective of the control obligation of the authorities is to identify (*long-term*) deviations between granted and delivered home care. The analysis is always performed for a period selected by the user and data is aggregated for the selected period. To support the control process and to automate the identification and prediction of deliveries with a high risk of deviations, a set of four different types of models have been developed:

Action lists This is a set of eleven models that implements simple rules (e.g., one model produces the list of citizens receiving two or more Day packages, which is not allowed while another model produces the list of citizens receiving more than five deliveries without a granted package).

Package fitness This model computes the delivery rate $\delta(w, V, p)$ under package p for weeks w using $l(p)$ or $u(p)$ in Definition 3 (instead of $g_t(p)$) depending on $\sum_i v_t^{(i)} < l(p)$ or $\sum_i v_t^{(i)} > u(p)$. Figure 1 shows the structure of the model where hexagon shaped vertices represent function nodes. Recall that a function node simply defines a numeric value that is computed from a mathematical expression. Function nodes can also represent an *input* value. For instance, the value of the function node *Delivered* represents time delivered to the citizen under a specific package for a selected period.

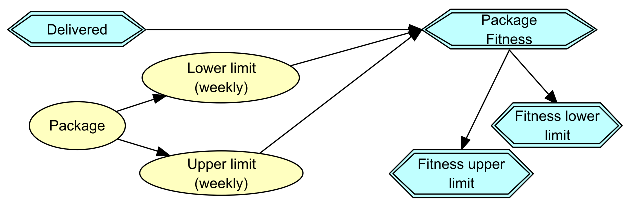

Figure 1: Model for assessment of package fit.

Delivery rate This model is estimated as a classification model from the aggregated data of the selected period. The purpose is not to classify deliveries (as the delivery rate is known), but to build a model to point to potential explanations of a too high delivery rate using a colour coding for each observation based on Bayes Factor K. The colour intensity of an observation ϵ_i in the user interface is determined by the value of K for ϵ_i.

A Boolean variable representing a too high delivery rate is the root variable of a TAN (Friedman et al., 1997) constructed with the HUGIN software (Madsen et al., 2005). The explanatory variables are the features. With the model constructed, the virtues of Bayesian Networks can be exploited, in particular the ability to point out the most influential indicator variables. This may help in identification of possible reasons for the deviation in a specific case and may lead to further insight into the current practice. There could, e.g., be problems in cases with a specific service provider or deviation could be due to a recent change in number or nature of allocated services.

Delivery rate prediction This model uses the same indicators as the model for explaining *delivery rate*, but with the purpose of predicting $\delta(w, V, p) \geq 150\%$ under package p for the next period. In this model, the delivery rate for the selected period becomes an indicator and the prediction is made for the following period of the same length as the selected period. The variables of the models include the target variable specifying if $\delta(w, V, p) \geq 150\%$ combined with indicators for the previous period.

Figure 2 shows an example of the structure of a prediction model where indicators are coloured to illustrate groups of similar variables. For instance, the lightblue variables represent master data, the lightgreen variables represent information about packages (including delivered

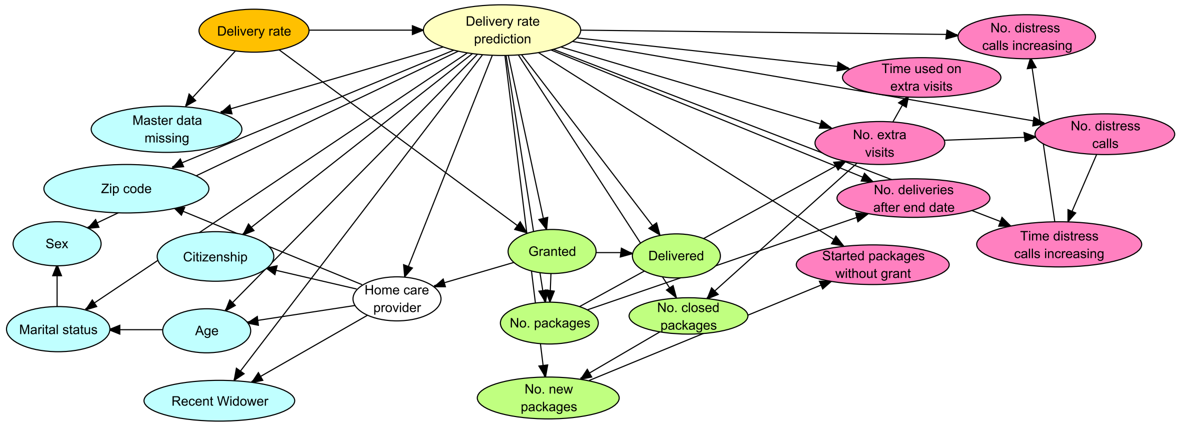

Figure 2: Model for prediction of future deviations in the delivery rate.

and granted time), and the purple variables represent signs of problems (e.g., extra visits and distress calls). Based on this model, cases with a high risk of deviation may be detected in advance. A possible explanation of such a forecast could be an increasing number of emergency calls. If disclosed in advance preventive actions for the citizen may be initiated. The results of an analysis are presented to the user as a table, where each line is a delivery (over the selected period) and the columns show the individual findings on the indicators and additional information (e.g., social security number and master data).

Home care includes a number of numeric entities (e.g., age and time) that must be considered. The organization of care into packages (as shown in Table 2) makes it natural to use discretization in the models for action lists and package fitness. For the models on delivery rate and delivery rate prediction, we use different types of discretization including expert assessment and a supervised approach (Fayyad and Irani, 1993). For instance, zip code is discretized based on official Danish zip code system (we group zip-codes according to the five regions of Denmark) while age is discretized using the minimum entropy principle of Fayyad and Irani (1993).

The models consider data from different angles and provide a suite of tools to support the control process. One of the main advantages of encoding the analysis as a Bayesian network model (with function nodes) is that it is easy for BI staff to maintain the models. The updated models can be deployed without help from IT staff. The use of function nodes makes it possible to encapsulate all calculations in a single model reducing the risk of errors as there is a clear separation between IT-integration and the business logic.

It has been an important requirement from the municipality that the analysis should use the data that was readily available and support a continuous refinement as additional data sources are identified and harvested.

6. Experimental Analysis

6.1 Data

Table 5 shows information on the data used in the experiment. The data covers a 14 weeks period from week 37, 2019 to week 50, 2019. In the table, we let D_D denote the set of deliveries (after

Week	Citizens in D_D	Citizens NOT in D_M	Granted (m)	Delivered (m)	Delivery rate
37	2,215	6	385,365	427,517	1.11
38	2,220	6	385,713	425,499	1.1
39	2,221	6	384,503	427,275	1.11
40	2,226	6	381,153	420,036	1.1
41	2,218	6	379,542	414,089	1.09
42	2,220	6	380,134	415,955	1.09
43	2,233	6	382,186	419,192	1.1
44	2,225	9	381,758	412,870	1.08
45	2,209	8	375,173	397,424	1.06
46	2,202	8	372,904	393,536	1.06
47	2,206	8	369,340	387,812	1.05
48	2,201	9	373,794	385,110	1.03
49	2,196	8	372,157	383,981	1.03
50	2,194	8	370,910	392,263	1.06

Table 5: Information on the data by week.

preprocessing) and D_M denote the set of master data. The table shows, for instance, that in week 42 there are 2,220 unique citizens contained in the deliveries data D_D and only 2,214 unique citizens in the master data D_M, i.e., six citizens have had one or more deliveries without an associated granted package. This could, for instance, be a distress call. The total amount of granted time is 380,134 minutes, and the total delivered time is 415,955 producing a delivery rate of 1.09.

6.2 Setup

The policy of the municipality, is that the control or analysis of deliveries is performed for periods of three weeks. This gives a certain level of stability in the results. For shorter periods, the delivery, for instance, may be fluctuating due to, say, vacation.

The top three rows of Table 6 show the periods considered in the experimental analysis. The lines are denoted Past, Present and Future that each covers three weeks. At each analysis, we assume that the user is considering the period Present, i.e., this period is selected by the user.

If the user is *analysing* the delivery rates, then a model is constructed based on the Present weeks. If the user is *predicting* the delivery rates, then a model is constructed from data on the indicators for the Past and the delivery rate for Present. This model is used to predict the delivery rate for the future using the data on the indicators for the Present. This means that a new model is created each time a different period is selected by the user.

6.3 Results

Table 6 shows the results of the experimental analysis considering two different models, i.e., the Naive Bayes Model (NBM) and the TAN) across the test period. Notice that we are evaluating models constructed for analysing deliveries and predicting too high delivery rate.

The row *Accuracy — explain* specifies the classification accuracy of the model on the training data. This relates to analysing deliveries. The analysis covers three weeks (e.g., in the column

Past	W37-39	W38-40	W39-41	W40-42	W41-43	W42-44
Present	W40-42	W41-43	W42-44	W43-45	W44-46	W45-47
Future	W43-45	W44-46	W45-47	W46-48	W47-49	W48-50
Naive Bayes Model						
Accuracy — explain	0.89	0.89	0.87	0.85	0.87	0.86
Accuracy — predict	0.89	0.88	0.89	0.88	0.88	0.89
$\delta(w, V, p)$ — predicted	2.24	1.69	1.52	1.37	1.1	1.64
Tree-Augmented Naive Bayes Model						
Accuracy — explain	0.94	0.95	0.95	0.94	0.94	0.94
Accuracy — predict	0.91	0.91	0.91	0.9	0.91	0.91
$\delta(w, V, p)$ — predicted	1.98	2.26	2.32	2.05	2.08	1.8

Table 6: Performance by week for two different models.

Present W41-43, the analysis covers weeks 41 to 43 aggregated). Please note that the purpose of the model in this case is not to classify each delivery, but to identify potential explanations of a too high delivery rate in order to make the control more efficient. In the user interface of the system, the deliveries are shown as a table sorted by delivery rate, where a colour coding scheme based on Bayes Factor (see Table 1) is used. The intensity of the background colour of an observation in the table is determined by the value of K. If $K < 3$ for an observation ϵ_i, the background is white. The row *Accuracy — predict* specifies the classification accuracy of the model when predicting delivery rates (the accuracy is measured on a hold-out dataset). This relates to predicting too high delivery rate. The analysis covers a period of three weeks (e.g., in column *Present W41-43*, the model is trained on data from deliveries for weeks 38 to 40 aggregated to predict delivery rates in weeks 41 to 43. The accuracy is evaluated on weeks 41 to 43 to predict weeks 44 to 46). The row $\delta(w, V, p)$ *— predicted* specifies the total delivery rate of the deliveries predicted as too high.

It is clear from the results that a relatively high accuracy is obtained in both models (using a threshold of 0.5). The TAN performs better than the NBM in almost all cases. Also, the deliveries predicted to be too high have a high delivery rate. In many cases the delivery rate is larger than 200% for the TAN.

The delivery rate experiments show that for TAN the number of extra visits is consistently among the top three indicators having highest mutual information relative to the target across the test cases. The other indicators in top three (across the test cases) are time used on extra visits, number of granted packages and number of new packages. The delivery rate prediction experiments show that using the supervised discretization on all variables eliminates a large share of indicators (due to a single state) and reduces the number of true positives but increases precision only slightly. For instance, for *Present W41-43* the accuracy is 0.92 with 24 true positives compared to accuracy 0.91 and 38 true positives for the case where expert assessments are also used.

7. Conclusion and Future Work

The current control process (in Hjørring and many other municipalities like it) is manual and hindsight focused. There is no system to support prediction of expected home care needs. The people with access to the data is focused on identifying *too high deviations* and contacting the service provider for explanations. This system automates parts of the process and enables other staff members (outside the economic control unit) to analyse the data in other to take preventive actions and identifying systematic errors in the delivery process. We have found that the involvement of users and other stakeholders in the process from the beginning and the intuitive nature of Bayesian networks have been vital to the success of the project.

The prototype was tested on live data in the municipality over the period from November 2019 until March 2020. The main feedback from the municipality was that especially the analyses of package fit added value. The explicit reports of gaps between granted and delivered services made the control task considerably easier. In particular, it turned out that there were many cases where citizens had been allocated services, that was not provided. That control reduced the number of allocated packages with approximately 5%. On the other hand, there was also a number of cases where the actual deliveries exceeded what was granted. In some cases, this lead to an adjustment of the allocated services, but in more cases the deliveries were reduced to what was actually granted. This happened in a constructive collaboration between the requester and the provider.

During the test period, the system lead to a $1 - 2\%$ reduction of delivered services. With an annual budget of around 130 million Danish kroner this is a quite satisfactory outcome of a pilot study, even if it includes some one-off savings.

The potential of the system is substantial. In Denmark there is 98 municipalities with the same duties in terms of control and all of them under a constant pressure to make the business more efficient and to harvest all possible savings while not compromising quality of services. Future and ongoing work includes further developing the explanation and prediction models by identifying potential new data sources.

Acknowledgments

The project is partly financed by regional funds of the European Union through Life Science Innovation North Denmark's Sundhedsteknologiske Service Program.

References

J. Arias, J. Martínez-Gómez, J. A. Gámez, A. G. S. de Herrera, and H. Müller. Medical image modality classification using discrete bayesian networks. *Computer Vision and Image Understanding*, 151:61–71, 2016.

P. Arora, D. Boyne, J. J. Slater, A. Gupta, D. R. Brenner, and M. J. Druzdzel. Bayesian networks for risk prediction using real-world data: A tool for precision medicine. *Value in Health*, 22(4): 439–445, 2019.

U. M. Fayyad and K. Irani. Multi-interval discretization of continuous-valued attributes for classification learning. In *Proc. of the 12th International Joint Conference on Artificial Intelligence*, page 1022–27, 1993.

N. Friedman, D. Geiger, and M. Goldszmidt. Bayesian network classifiers. *Machine Learning*, pages 1–37, 1997.

I. J. Good. Weight of evidence, corroboration, explanatory power, information and the utility of experiment. *J. R. Stat. Soc.*, B22:319–331, 1960.

J. Goodson, W. Jang, and M. Rantz. Nursing home care quality: Insights from a bayesian network approach. *J. R. The Gerontologist*, 48(3):338–348, 2008.

H. Jeffreys. *Theory of probability*. New York: Oxford University Press, 3rd edition, 1961.

F. V. Jensen and T. D. Nielsen. *Bayesian Networks and Decision Graphs*. Springer, 2nd edition, 2007.

P. Kaewprag, C. Newton, B. Vermillion, S. Hyun, K. Huang, and R. Machiraju. Predictive models for pressure ulcers from intensive care unit electronic health records using bayesian networks. *BMC Medical Informatics and Decision Making*, 17:81–91, 2017.

Y. Kazemi and S. A. Mirroshandel. A novel method for predicting kidney stone type using ensemble learning. *Artificial Intelligence in Medicine*, 84:117–126, 2018.

U. B. Kjærulff and A. L. Madsen. *Bayesian Networks and Influence Diagrams: A Guide to Construction and Analysis*. Springer, 2nd edition, 2013.

E. Kyrimi, S. McLachlan, K. Dube, M. R. Neves, A. Ali Fahmi, and N. Fenton. *A Comprehensive Scoping Review of Bayesian Networks in Healthcare: Past, Present and Future*, 2020. Preprint, https://www.researchgate.net/, visited May 2020.

S. L. Lauritzen. The EM algorithm for graphical association models with missing data. *Computational Statistics & Analysis*, 19:191–201, 1995.

M. D. Lee and E.-J. Wagenmakers. *Bayesian modeling for cognitive science: A practical course*. Cambridge, England: Cambridge University Press., 2013.

A. L. Madsen, F. Jensen, U. B. Kjærulff, and M. Lang. HUGIN - The Tool for Bayesian Networks and Influence Diagrams. *International Journal on Artificial Intelligence Tools 14*, 3:507–543, 2005.

A. L. Madsen, F. Jensen, M. Karlsen, and N. Søndberg-Jeppesen. Bayesian Networks with Function Nodes. In *Proc. of PGM*, pages 286–301, 2014.

A. L. Madsen, N. Søndberg-Jeppesen, N. Lohse, and M. S. Sayed. A Methodology for Developing Local Smart Diagnostic Models Using Expert Knowledge. In *Proc. of the 2015 IEEE International Conference on Industrial Informatics*, page 6, 2015. ISBN 978-1-4799-6648-6.

J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Series in Representation and Reasoning. Morgan Kaufmann, 1988.

Y. Shen, L. Zhang, J. Zhang, M. Yang, B. Tang, Y. Li, and K. Lei. Cbn: Constructing a clinical bayesian network based on data from the electronic medical record. *Journal of Biomedical Informatics*, 88:1–10, 2018.

Statistics Denmark. 2020. Statistics Denmark's data repository at https://www.statistikbanken.dk, visited May 2020.