Proceedings of the 2009 IEEE International Conference on Systems, Man, and Cybernetics
San Antonio, TX, USA - October 2009

A Conceptual Model for a Value-Driven Learning Healthcare System

Ying Tang
Dept. of Electrical and Computer Engineering
Rowan University
Glassboro, NJ 08028
Email: tang@rowan.edu

Lenard Pogach
Center for Healthcare Knowledge Management
DVA New Jersey Veterans Health Care System
385 Tremont Ave., East Orange, NJ 07018
Email: leonard.pogach@va.gov

***Abstract*—Making context based and pertinent clinical knowledge to the point of care that is most appropriate for individuals becomes an unprecedented challenge to bring American closer than ever to the promise of personalized health care. This paper, from the engineering perspective, presents a new conceptual framework that keeps patients in focus and continuously incorporates new knowledge to improve quality and value in healthcare. A colored Petri net model linked with a hybrid Bayesian network is then proposed to address two imperative issues in the development of a learning healthcare system: (1) mathematical representation of the complex health care processes to ensure shared decision-making; and (2) exploration and exploitation of knowledge about patients and health outcome measures for continuous improvement.**

***Keywords*—Value-driven healthcare system, learning healthcare system, colored Petri net, hybrid Bayeisan netwok.**

I. Introduction

"Care that is important is often not delivered, and care that is delivered is often not important [8]". Due to the overwhelming amount of healthcare information that can be relevant in clinical reasoning processes, making context based and pertinent clinical knowledge to the point of care, that is most appropriate for individual needs and conditions, becomes an unprecedented challenge to bring American closer than ever to the promise of personalized health care. Current practice in healthcare relies mostly on two approaches of knowledge transfer: clinic guidelines and case based reasoning [5]. The former, although considered as "best practice" of solving clinical problems, is usually developed in ideal clinical settings with very high generality. The latter uses the similar patient practice as a guideline. However, the human diversity combined with ethical and empathy considerations can make a clinical guideline, either developed in ideal settings or derived from similar patient cases, more or less unsuitable for use in the treatment of a patient [5]. Therefore, how to draw on the best evidence for clinical decision making as elaborated in Figure 1 calls for a new conceptual framework to integrate the aforementioned two sources of information with patient factors. Meeting this challenge is the ultimate goal of our project. In particular, we, adapting from industry approaches, propose a value-driven learning healthcare system that has the ability to evaluate multiple, but constrained inputs of care to incrementally improve intermediate outcomes at the population level. The rest of the paper is organized as follows. Section II presents the background and motivations for value-driven healthcare. Section III overviews the generic model for a value-driven learning healthcare system. Section IV focuses on the decision-making models using Petri nets and Bayesian network, followed by the conclusion in Section V.

Figure1: Complex clinical decision-making with multiple inputs of care

II. Motivation for Value-driven Healthcare

The risks and benefits of healthcare control differ substantially based on patients' disease duration, life expectancy and comorbid conditions. Additionally, the specific medications used to achieve control differ in efficacy, side effects and costs. Therefore, the transparency of how these factors are presented, when healthcare is delivered or healthcare decision is made, is of extreme importance.

978-1-4244-2794-9/09/$25.00 ©2009 IEEE

A. *Shared Decision Making*

As pointed by the U.S. Department of Health & Human Service, the healthcare system in America is a disconnected collection of large and small medical business, health care professionals, treatment centers, hospitals, and all who provide support for them [9]. Such disconnection makes the needed information hard to share and compare at the time of decision making by clinicians and patients, resulting in significant amount of inappropriate care. For example, a health plan that seeks to "maximize" quality adjusted life years may have greater pharmaceutical utilization, causing higher premiums in contrast to a plan that has a more restricted formulary. Other plans could compete on competitive strategies such as leveraging advanced technology or customer service to achieve best value for the money or customer satisfaction. Definitely, a better care comes from shared decision-making processes with systematic reviews and synthesis of the evidence regarding benefits, risks, costs, alternatives and uncertainties of health care interventions as well as an assessment of the trade-offs between effectiveness, cost, the patient's perspective [7].

B. *Efficiency Measures*

The current industry standard for the "quality" of care has been narrowly focused on laboratory tests and intermediate outcome measures. Using diabetes care as an example, the thresholds that have been long recommended for internal quality improvement and performance measures are hemoglobin A1c (HbA1c) < 7%, blood pressure less than 130 mm/Hg, and total cholesterol less than 200 mg/dl [4]. In fact, only one in fourteen U.S adults with diabetes meets these thresholds. Apparently, this "one size fits all quality approach" represents both a flawed understanding of the epidemiology of diseases and a misapplication of the six sigma approach to chronic care treatment [1, 6]. A new measurement paradigm should provide suitable clinic guidelines that respect patient preferences and physician judgment and carefully consider the individual patient's absolute risk of benefits and harms. In addition, the model should be able to exploit any new knowledge that might become available during medical treatment processes to refine the guidelines for future decision makings.

III. GENERIC MODEL FOR A VALUE-DRIVEN LEARNING HEALTHCARE SYSTEM

To offer health care transparency and provide the care most appropriate to individual patients, this paper presents a value-driven learning healthcare system. As depicted in Figure 2, it consists of the following interactive and cooperating constituent units: (1) Database; (2) Inspection Unit; (3) Learning Health Controller; and (4) Temporal Observation Unit.

A. *Database*

Database is the information center of a value-driven learning healthcare system. It consists of three parts, Clinic Guideline Database that archives general solutions/suggestions for different types of clinic problems; Patient Database that has the detailed patient health records; and Resource Database that keeps the knowledge of local organizational settings (e.g., hospital and clinics resource information) and insurance information.

B. *Inspection Unit*

Before any medical decision to be made towards a particular disease treatment, a patient will take certain laboratory tests that are suggested by the clinical guidelines during the inspection process. The test results are then fed into the database to keep the patient record up-to-date.

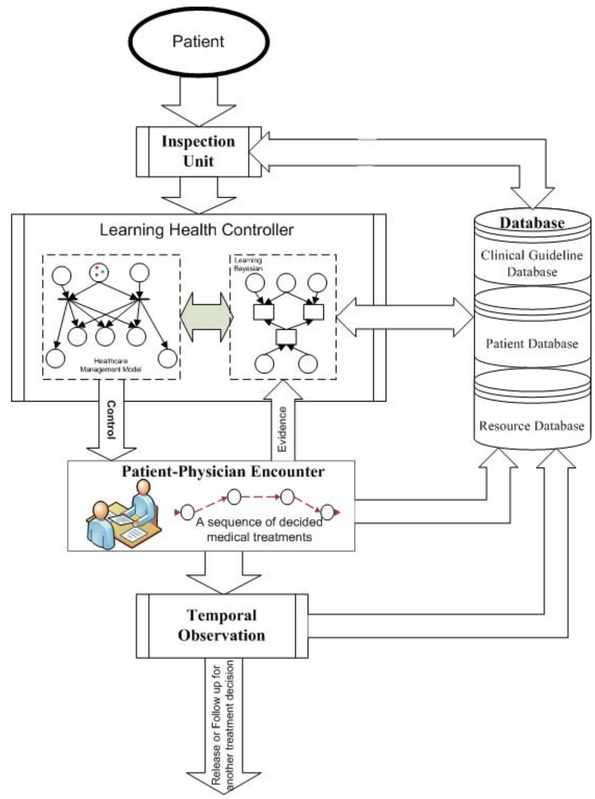

Figure 2: A generic model for a value-driven learning healthcare system

C. *Learning Health Controller*

In the context of uncertainty, our goal is to offer a framework that allows patients to search for a reasonable and acceptable (to the individual) clinical solution, not necessarily an "optimal" one. For this to be successful, two main problems are considered in our design. The first is an effective modeling of shared decision makings that would provide the full knowledge of feasible actions to be performed at any instant of care processes. The second is to acquire (1) information that characterizes benefits, risks and costs of treatments, as well as patients' condition; and (2) assessment rules that are imperative to determine what is considered as "quality" care to individual patients. A colored Petri net model in conjunction with a hybrid Bayesian network is introduced for the two problems, which will be discussed in detail in the next section.

D. Temporal Observation Unit

According to the evaluation made in Learning Health Controller, a series of clinical treatments are decided and conducted at the patient-physician encounter. Meanwhile, selections of target values for each intermediate outcome are set, which may be adjusted based upon changes in patients' condition, experience with side effects, or concerns over cost. During the temporal observation, any deviations will be fed into the database to refine the guidelines.

The rest of the paper focuses on the health decision-making in the value-driven learning healthcare system.

IV. Learning-based Healthcare Decision-Makings

From the modeling standpoint, this paper invokes a special class of Petri net models, i.e., Color Petri net (CPN), where a place can contain multiple tokens distinguished by a color associated with each token. This model allows for a natural representation of the critical issues for a health decision-making: (1) patients with different comorbid condition, insurance plans, and healthcare preferences as colors of tokens; (2) value-driven decisions as transition firings. In addition, the paper links the CPN with a hybrid Bayesian network to catch the uncertainty in health decision making. Moreover, the framework is designed in a way to incorporate any newly available information into the system for continuous improvement. The foundation of the proposed CPN model is first given in Subsection IV-A, followed by the learning Bayesian network and adaptive health decision heuristic in Subsections IV-B and IV-C, respectively.

A. Colored Petri Net

Petri nets, as a graphical and mathematical tool, provide a uniform environment for modeling, formal analysis, and design of discrete event systems [10]. A Petri net may be identified as a particular kind of bipartite directed graph populated by three types of objects. They are places, transitions, and directed arcs connecting places to transitions and transitions to places. Pictorially, places are depicted by circles and transitions by bars. A place is an input (output) place to a transition if there exists a directed arc connecting the transition (place) to the place (transition). In order to study dynamic behavior of the modeled system, each place may potentially hold either none or a positive number of tokens, pictured by small solid dots. Formally, a Petri net can be defined as follows:

Definition 4.1 A Petri net is defined as a five-tuple: PN = (P, T, I, O, M)

1. $P = \{p_1, p_2, \ldots, p_n\}$ is a finite set of places.
2. $T = \{t_1, t_2, \ldots, t_s\}$ is a finite set of transitions, $P \cup T \neq \varnothing$, $P \cap T = \varnothing$.
3. $I: P \times T \rightarrow \{0, 1\}$ is an input function that defines the set of directed arcs from P to T, where $I_{ij} = 1$, if p_i is an input place for t_j, otherwise 0.
4. $O: P \times T \rightarrow \{0, 1\}$ is an output function that defines the set of directed arcs from T to P, where $O_{ij} = 1$, if p_i is an output place for t_j, otherwise 0.
5. $M: P \rightarrow \{0, 1, 2, \ldots\}$ is a marking vector whose i^{th} component represents the number of tokens in the i^{th} place. An initial marking is denoted by m_0.

Considering the heterogeneity of patients' physical condition and health plan options, colors are introduced into the aforedefined Petri net model and are associated with tokens. The color of a token a that represents a patient waiting for a clinical treatment decision in the system is defined as $c(a)$:

$$c(a) = \langle \mathrm{ID}_a, \mathrm{PrimaryDisease}, (\mathrm{Comorbid\ Condition})_a, \mathrm{InsuranceType}_a, (\mathrm{Quality\ Measures})_a \rangle$$

where ID_a is the patient identification number. The second, third and fourth components identify the patient's primary disease, a set of other illnesses that the patient might have, and the insurance plan that the patient is currently enrolled. The last component, directly coming from the inspection process, is a list of intermediate measurement outcomes pertinent to the patient's primary disease. For example, the token color of a patient for diabetes treatment is ⟨1002, *Diabetes*, (hypertension, hyperlipidemia, stage 1 congestive heart failure, osteoarthritis), *Medical Part D Prescription Drug Plan,* (BP: 154/92 mmHg, HbA1c: 8.9%, cholesterol: 145 mg/dl)⟩, indicating that the patient has multiple cardiovascular risk factors in addition to his diabetes, and all of the important measures of his diabetes situation are way above normal according to the thresholds set by the National Committee for Quality Assurance [4]. The detailed definition of a Colored Petri Net (CPN) in modeling health decisions is given below.

Definition 4.2 A Colored Petri Net (CPN) is an extension of PN with twelve-tuple:

$CPN = (P, T, I, O, C, \Gamma^+, \Gamma^-, m_0, \lambda, \tau, w, \rho)$

- P is partitioned into P_{PWP}, representing patients waiting for process (e.g., waiting for a treatment decision); P_{PIP}, patients in process, P_{TO}, available treatment options, and P_{SE}, potential health risks of treatments.
- T is partitioned into three subsets: T_I, representing patients being admit/re-admit into the health system; T_O, patients being released from the system; and T_F, the flow of patients from waiting-for-process to in-process.
  - [1] $\exists\ t \in T_I$, ${}^\bullet t = \varnothing$. This type of transitions is usually called source transitions, which has no input places. In our model, it represents the patient admission process.
  - [2] $\forall t \in T_O$, $t^\bullet = \varnothing$, ${}^\bullet t \cap P_{PWP} = \varnothing$, ${}^\bullet t \cap P_{TO} = \varnothing$, ${}^\bullet t \cap P_{SE} = \varnothing$, and $|{}^\bullet t \cap P_{PIP}| = 1$. This type of transitions is usually called sink transitions, which has no successor but only one predecessor place p in P_{PIP}. In our model, it represents the patient being released from the health system.

- C: $P \cup T \rightarrow \Sigma$ is a color function that defines color domains of places and transitions, where Σ is nonempty finite set of colors.
- The I^+ and I^- are the post-incidence and pre-incidence matrices, respectively. In particular, $I^+[p, t]$ ($I^-[p, t]$) associates a set of colors of $C(t)$ (the set of colors of $C(p)$) to a set of colors of $C(p)$ (a set of colors of $C(t)$), where $t \in {}^{\bullet}p$ ($p \in {}^{\bullet}t$). The pre- and post-incidence matrices are defined as follows:
  [1] $\forall(p, t) \in I$: if $p \in P_{PWP} \cup P_{PIP}$, $I^-[p, t] = f_{id}$; if $p \in P_{TO}$, $I^-[p, t] = \langle\bullet\rangle$; Otherwise, $I^-[p, t] = 0$. f_{id} stands for "the function makes no transformation in the elements", and $\bullet$ represents the neutral color. In our model, the colors of tokens that represent the availability of treatment options or side effects are neutral.
  [2] $\forall(t, p) \in O$:
  - if $p \in P_{PWP} \wedge {}^{\bullet}t = \varnothing$, $I^+[p, t]$ results from the direct measures of the patient in the inspection process; otherwise, $I^+[p, t] = f_{id}$;
  - if $p \in P_{PIP}$, $I^+[p, t] = f_c$, where f_c is a function that is decided through the learning HBN to update the color;
  - if $p \in P_{PIP}$, $I^+[p, t] = f_{id}$; and
  - if $p \in P_{TO} \cup P_{SE}$, $I^+[p, t] = \langle\bullet\rangle$.
- m_0 satisfies:
  [1] $\forall p \in P_{\text{PIP}} \cup P_{PWP} \cup P_{SE}$, $m_0(p) = 0$
  [2] $\forall p \in P_{\text{TO}}$, $m_0(p) = \langle\bullet\rangle$.
- λ: $T_F \rightarrow \{0\} \cup \Re^+$ is an out-of-pocket cost function associated with a flow transition t with a particular color. Given the health plan information in the color and the treatment being considered, the out-of-pocket cost function is assumed known *a priori*.
- τ: $P_{PIP} \rightarrow \{0\} \cup \Re^+$ is the expected benefit of a patient undergoing certain treatment. The value is derived through the learning HBN
- w: $P_{SE} \rightarrow \{0\} \cup \Re^+$ is the expected health risk of a treatment $p \in P_{SE}$. The value, determined through the learning HBN, is heavily dependent on the treatment and the patient physical condition.
- ρ: $T_O \rightarrow [0, 1]$ is a probability value associated with a transition $t \in T_O$. It is updated through the learning HBN.

An example of CPN for a simplified decision-making process in diabetes care is given in Figure 3. Patients are admitted into the system through "admission/re-admission" transitions t_1 and t_6 or t_7, respectively. Tokens with the corresponding colors that reflect the patients' physical condition and health plan options are added to the patient place p_1, accordingly. A rather simple scenario is considered in the model, where only two composite treatment options (i.e., antihypertensive medication plus a generic statin versus oral hypoglycemic agent, represented by p_2 and p_3 respectively) are available to the patients, each with different side effects represented by p_4, p_5, and p_6. Taking the action of choosing one treatment (e.g., t_2 corresponds to p_2) over the other leads the patient to the in-therapy stage (e.g., p_7 representing patient undergoing the treatment p_2). Once a patient is in the stage of "in-therapy" (i.e., p_7 and p_8), the patient is either released from (e.g., t_4 and t_5) or re-admitted into (e.g., t_6 and t_7) the system with certain probabilities.

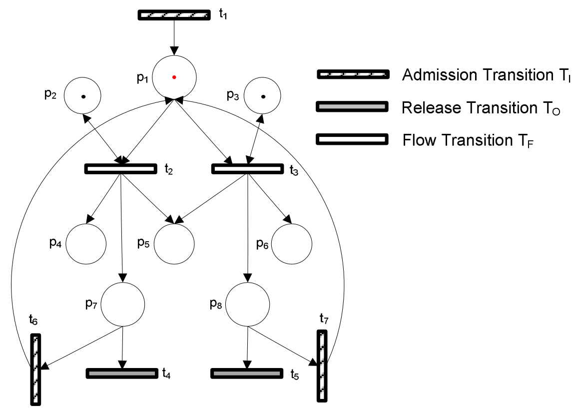

Figure 3: An example of a CPN

B. Learning Hybrid Bayesian Network (HBN)

While the CPN has full knowledge of the clinical actions that can be performed, it does not necessarily know which action is the most appropriate to take without acquiring a set of pertinent data. Such data, imperative to healthcare reasoning, is contingent on the prevailing medical condition and health plan option of the patient being considered. To alleviate this concern, an HBN is used to ascertain the uncertain parameters that describe the important aspects of care. The HBN is a directed acyclic graph which shows how the quantitative and qualitative aspects of the healthcare decision-making process are related. The quantitative nodes are continuous valued nodes, whereas the qualitative nodes are the discrete valued nodes. The probabilistic relationships in the graph are represented as directed arcs, where the parent nodes influence the value of the child nodes. Using this construct, the relationships between all of the parameters of the health decisions can be modeled. When some of the parameters are known, inference can be performed on the HBN, allowing the unknown parameters in the HBN to be ascertained via their probabilistic relationships with the known values in the HBN. The detailed description of the HBN is given below.

Definition 4.3 A Hybrid Bayesian network (HBN), $G = (V, E)$, is a directed acyclic graph [2]

- $V = \{v_1, v_2, \ldots, v_n\}$ is a set of vertices corresponding to random variables in the domain $V = \Delta \cup \Gamma$. The vertices in set Δ represent qualitative variables, and those in set Γ, quantitative variables.
- $E = \{(v_i, v_j) | \ v_i, v_j \in V, \ i \neq j\}$ is a set of ordered pairs of vertices indicating a probability dependency between the parent node v_i and the child node v_j. The network encodes a joint probability distribution of the domain variables $\mathbb{P}(v_1, v_2, \ldots, v_n)$, where ψ_i denotes the parents of node v_i.

$$\mathbb{P}(v_1, v_2, \ldots, v_n) = \prod_i \mathbb{P}(v_i|\psi_i)$$

In our approach, a HBN can be directly derived from its corresponding CPN. First, we need identify variables of interest in the CPN for the healthcare decision-making, which include the benefits, risks, and costs of each treatment as well as patients' physical condition and health plan options. As matter of fact, the former set of variables regarding the benefits and risks of a treatment for a patient is contingent on the patient's physical condition. Their causal relationships are then explicitly modeled in the HBN. However, the cost of a treatment for a patient is not modeled in the HBN due to its deterministic nature given the patient health plan option. Figure 4 is an example of the HBN in relation to the CPN in Figure 3. $\forall p \in P_{PWP}$ in the CPN, there is a correlative continuous vertex in the corresponding HBN that represents the prevailing condition of the patient's primary disease before any treatment. Depending on the color of a patient token in $p \in P_{PWP}$, several other nodes can be added in the HBN. For instance, the quality measures for diabetes (i.e., BP, HbA1c, and Cholesterol) are introduced into the HBN as continuous nodes and the patient's comorbid condition (e.g., heart failure and hypertension, etc.) as discrete nodes. In a similar fashion, each treatment option represented in the CPN has a discrete node in the HBN corresponding to its probability of success. Each health risk of a treatment in the CPN has a corresponding continuous node in the HBN, indicating the patient's health risk. Apparently, the health risk of a patient undergoing a treatment is influenced by the success of the treatment and the prevailing patient's physical condition. All together further impact the patient's physical condition after the treatment. Finally, $\forall t \in T_O$ in the CPN, there is a correlative discrete vertex in the corresponding HBN that represents the probability of a patient being released from the system. Finally, the health benefit of a particular treatment to a patient is modeled as a differential equation of the variation in the patient's physical condition before and after the treatment.

Once the structure of the HBN has been determined, the probabilistic relationships between each of the nodes must be ascertained. Any of the nodes which do not have a parent node, such as the quality measures of a patient's diabetes or a patient's comorbid condition, is either simply a discrete Bernoulli distribution or a univariate Gaussian distribution. Any of the continuous valued nodes with discrete and/or continuous parents, such as the nodes pertaining to the prevailing condition of a patient's primary disease, is represented as conditional Gaussian. For the node which represents the probability of a patient being released from the system, a logistic function [3] is used, since it is a discrete valued node with a continuous parent.

The use of inference on a HBN then allows us to ascertain the posterior probabilities of nodes in the HBN, given some evidence. For example, when a patient is diagnosed with Type 2 diabetes and hypertension, he is at greater risk of cardiovascular complications and renal disease. Since the relationship between a patient's physical condition (including the comorbid condition) and his health risks of undertaking a treatment is well defined in the HBN, inference can be used to update probabilities that bring the estimated patient's health risks closer to the real practice. A cursory overview of a patient in the initial stage or any clinical treatment action can yield valuable evidence for use with inference. The results from inference can then be feedback to the CPN and used to refine health guidelines and to improve future health decisions.

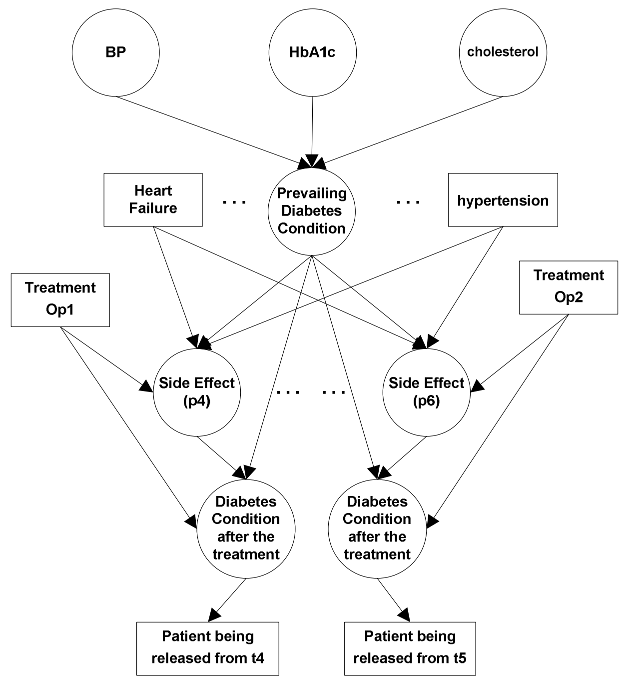

Figure 4: A HBN derived directly from the CPN in Figure 3. A selection of comorbid condition nodes and health risk nodes are not represented (but replaced by "$\cdots$") for the sake of clarity.

C. Adaptive Health Decision Heuristic

With the natural setting of integrating the CPN model with a HBN, the health reasoning problem can be solved using a closed-loop control as depicted in Figure 2. At each patient-physician encounter, any knowledge that is known about the patient, including his physical condition and health plan options, is sent to the HBN for use in inference. Once the inference is performed, the values of the uncertain parameters are updated (e.g., the risks of a patient under certain treatment) and fed back to the CPN. With the newly updated knowledge of health treatment options, an informed decision can then be made by evaluating the expected health value of a treatment transition $t \in T_F$ in the CPN as defined in Eq. 4.1, where $q \in t^{\cdot} \wedge q \in P_{SE}$ and $p \in t^{\cdot} \wedge p \in P_{PIP}$.

$$HealthValue(t) = \tau(p) - \lambda(t) - \sum w(q) \tag{4.1}$$

This process continues until the patient is released from the system with certain probability ρ. The readmission of the patient with the probability 1-ρ resumes the reasoning process again.

V. Conclusion

This paper presents a learning-based strategy to explicitly model and effectively manage involved uncertainty in healthcare decision-making. More specifically, a colored Petri net model linked with a hybrid Bayesian network is proposed to address two imperative issues in the development of a value-driven learning healthcare system: (1) mathematical representation of the complex health care processes to ensure shared decision-making; and (2) exploration and exploitation of knowledge about patients and health outcome measures for continuous improvement. To benchmark the proposed conceptual model in large-scale health practices is the focus of our future research.

References

[1] Aron DC. and Pogach L., "One size dose not fit all: a continuous measure for glcemic control in diabetes: the need for a new approach to assessing glycemic control," *The Joint Commission Journal on Quality and Patient Safety*, 2007; 33: pp. 636-643.

[2] Grochowski, D. and Tang, Y., "A Machine Learning Approach for Optimal Disassembly Planning," *International Journal of Computer Integrated Manufacturing,* Vol. 22, No. 4, 2009, pp. 374-383.

[3] Murphy, K. P. "A Variational Approximation for Bayesian Networks with Discrete and Continuous Latent Variables", *Proc. of the Conf. on Uncertainty in AI*, pp. 457-466, 1998.

[4] National Committee for Quality Assurance. HEDIS 2007 Vol.2, Technical manual, Washington D. C., 2006

[5] Pedersen, K. V., "Using data mining techniques for optimizing medical treatment processes," ARC Research Network in DM&KD Health Data Mining Workshop, University of South Australia, Mawson Lakes, April 11, 2005.

[6] Pogach, L., Englegau, M., and Aron D., "Measuring progress towards achieving A1c goals: pass/fail or partial credit," *JAMA*, 2007; 297: pp. 520-523

[7] Strite, S. and Stuart, M. E., "What is an evidence-based, value-based health care system? (Part I)," *The Physician Executive*, Jan-Feb., 2005, pp. 50-54.

[8] The learning healthcare system: workshop summary (IOM Roundtable on Evidence-Based Medicine). http://books.nap.edu/catalog.php?record_id=11903#toc accessed March 31st , 2009

[9] US Department of Health & Human Services. http://www.hhs.gov/valuedriven/ accessed March 31st, 2009.

[10] Zurawski, R. and Zhou, M. C., "Petri nets and indutrial applications: a tutorial," *IEEE Transactions on Industrial Electronics,* Vol. 41, No. 6, 1994, pp. 567-583.