International Journal of Medical Informatics 62 (2001) 79–99

International Journal of
Medical
Informatics

www.elsevier.com/locate/ijmedinf

Modeling a medical environment: an ontology for integrated medical informatics design

John R. Hajdukiewicz [a,b], Kim J. Vicente [a,b,*], D.J. Doyle [a], Paul Milgram [a,b], Catherine M. Burns [c]

[a] *Institute of Biomaterials & Biomedical Engineering, University of Toronto, Toronto, Canada*
[b] *Department of Mechanical & Industrial Engineering, University of Toronto, 5 King's College Road, Toronto, Ontario M5S 3G8, Canada*
[c] *Department of Systems Design Engineering, University of Waterloo, Waterloo, Canada*

Received 15 September 2000; received in revised form 7 January 2001; accepted 9 January 2001

Abstract

Modern medical environments have seen an increase in technological complexity and pressures of handling more patients with fewer resources, resulting in higher demands on medical practitioners. Medical informatics designers will have to focus on the problem of organizing medical information more effectively to enable practitioners to cope with these challenges. This article addresses this research problem for the particular area of medical problem solving in patient care. First, we describe a traditional modeling approach for medical reasoning used as a basis for developing some decision support systems. We argue these models may be faithful to what is known about biomedical knowledge, but they have limitations for human problem solving, especially in unanticipated situations. Second, we present an ontological framework, known as the abstraction hierarchy (Rasmussen, IEEE Trans. Man. Cybernetics 15 (1985) 234–243), for integrating patient representations that are faithful to existing biomedical knowledge and that are consistent with what is known about human problem solving. Through an example of a critical event in the operating room, we reveal how this framework can support medical problem solving in unanticipated situations. Third, we show how to use these representations as a frame of reference for mapping medical roles, responsibilities, sensors, and controls in an operating room context. Finally, we provide some insight for medical informatics designers in using this framework to design novel training programs and human–computer displays. © 2001 Elsevier Science Ireland Ltd. All rights reserved.

Keywords: Medical informatics; Design framework; Decision support; Abstraction hierarchy

---

1. Introduction

Like many other complex sociotechnical environments, contemporary medical domains are undergoing vast and rapid changes.

---

* Corresponding author.
E-mail address: benfica@mie.utoronto.ca (K.J. Vicente).

PII: S1386-5056(01)00128-9

These changes are part of a larger societal trend requiring organizations to 'do more with less' [2]. In the case of medicine, there is an increased demand to provide higher quality care to sicker patients in the face of increasing economic and work pressures. As a result, devices and technologies have increased, both in number and in complexity. An unintended, but nevertheless very real, side effect of these changes is an increase in the cognitive demands experienced by medical personnel [3,4]. Medical informatics designers will have to take specific measures if medical practitioners are going to be able to cope effectively with this challenging increase in demands.

Various directions have been pursued in the medical informatics community to tackle this problem [5–15]. Some researchers have focused on developing expert and decision support systems that code, filter, and interpret sensed data using algorithms based on heuristics and traditional biomedical models [5–10]. These systems are envisioned to detect and diagnose patient events; however, they are currently constrained by the limited capabilities of the programmed software and sensor technology [5,7]. As a result, these approaches may not support critical aspects of medical problem solving, especially in uncertain and unanticipated cases. Other researchers have focused on studying the interaction between people and medical technology (e.g. human–computer interaction), looking at ways of making the technology easier to use and understand in particular medical environments [11–15]. These approaches help practitioners perform known tasks and perceive information more effectively and efficiently. However, they may not support medical problem solving in unforeseen situations.

One way to think about this design challenge is to view medical informatics as providing a bridge between the domain of medicine and practitioner psychology [15]. On the one hand, computer support systems must be faithful to the best existing medical knowledge. On the other hand, such systems must also be faithful to what we know about human cognition. Traditionally, these two types of knowledge have tended to reside in different researchers, creating a symmetric set of problems. Medical researchers are familiar with medical constraints, but do not generally know as much about human cognition, resulting in the design of medical informatics systems that are not as easy as they could be for people and organizations to use. Conversely, psychological researchers are familiar with cognitive constraints but do not generally know as much about medicine, resulting in the design of medical informatics systems that are not as tailored to medical environments as they could be. To help solve these problems, it would be useful to have an integrated framework for medical informatics design that systematically bridges these two worlds in a principled and deliberate fashion.

This article addresses this research problem for the particular area of medical problem solving in patient care. First, a traditional approach to modeling biomedical knowledge for medical reasoning, found in numerous medical textbooks and decision support systems, is described. We argue that these models may be faithful to what is known about biomedical knowledge, but they have limitations for medical problem solving especially in unanticipated situations. Second, an ontological framework that can be used to bridge the medicine–psychology gap is described. We show that this framework, known as the abstraction hierarchy [1], can be used to develop representations of the patient that are faithful to existing biomedical knowledge. Also, through an example of a critical event in the operating room, we demonstrate that this framework can support medical problem solving. Third, we show how these representations can be used as a frame of reference in mapping operating room roles, responsibilities, sensors, and controls. Finally, by developing such representations, we suggest how medical informatics designers can identify information requirements that could be used to develop innovative training programs and human–computer displays.

2. Representing biomedical models

The need for an integrated design framework can be illustrated by examining the ways in which biomedical models have typically been represented (e.g. in medical textbooks and clinical manuals). While these models are faithful to what is known about how the body works, the structure is not very compatible with how people solve problems. This section describes a traditional approach to representing biomedical knowledge and discusses an alternative approach for structuring the same information in a way that is compatible with medical problem solving.

2.1. Motivation: a traditional approach

Various biomedical models have been created to support medical reasoning and training, and inform the design of medical decision support systems (i.e. information requirements) [5]. A large majority of these models depict a causal and relational web of interacting physiological and process variables. One example of a traditional model is shown in Fig. 1, created from discussions in Sherwood [16], Mohrman and Heller [17], and Gomez [18]. This example shows the physiological factors and relations affecting mean arterial blood pressure.

Consider how this network may be used for medical problem solving in a loss of blood pressure scenario during an operation.

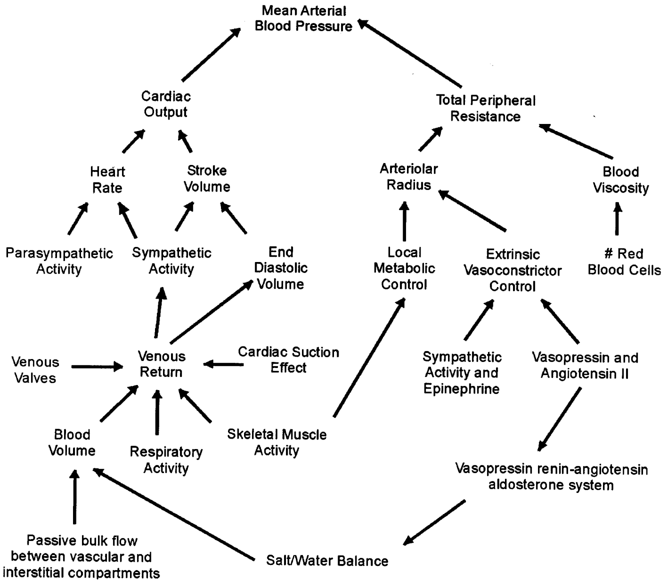

Fig. 1. A traditional approach to modeling biomedical knowledge (adapted from Sherwood [16], Mohrman and Heller [17], and Gomez [18]).

The anesthesiologist may look at the electrocardiograph (ECG) waveform to infer heart function (e.g. heart rate). With a slightly increasing heart rate, the physician may initiate an indirect action to stop blood pressure from falling by administering epinephrine. From the diagram, epinephrine acts to constrict the arterioles, resulting in an increase in total peripheral resistance and mean arterial blood pressure. In this case, blood pressure should return to a normal condition.

There are two advantages to such a representation. First, by mapping the factors that affect biomedical processes, one may reason about biological function in a clinical and training setting. Second, these models focus on key determinants or drivers of biological function under typical conditions, thereby providing an indication of the possible impact of an intervention.

There are, however, a number of limitations to this representation as well. First, the flat network representation can become quite complex. As more variables and detail are integrated into this representation, it becomes more difficult to determine the coupled relations and the effect of an intervention. For example, by using this structure to model cardiovascular, respiratory, and nervous system interactions, it becomes harder to diagnose problems. Second, within this structure, the variables and associations can be qualitatively different. For example, arteriolar radius is a physical characteristic whereas circulation is a biological process. To combine these two is to lump together apples and oranges. Third, the relations between elements in the network are not the same. For example, cardiac output is the product of heart rate and stroke volume; however, venous return is affected by, but not the product of, venous valves and respiratory activity. Fourth, the structure does not represent the entire picture in terms of biomedical knowledge, potentially leaving out important factors depending on the context. While the representation focuses on key factors in a normal case, in atypical cases, many other important factors are not mapped. For example, in Fig. 1, the effect of blood loss due to hemorrhaging is not factored in.

In summary, many traditional models of biomedical knowledge can be conceptually muddled, showing different relationships with the same undifferentiated representation. In addition, it becomes harder to reason using these models as more variables are introduced. These models are faithful to biomedical knowledge but can be very complex. As a result, these types of models make it difficult for practitioners to solve problems. These limitations are not generally acknowledged in the medical literature, as the modeling approach is still a standard way of representing biomedical knowledge.

2.2. The abstraction hierarchy: an alternative approach

In order to overcome a few of the limitations of this traditional approach to modeling biomedical knowledge, an alternative approach using the abstraction hierarchy (AH) framework [1] to create a work domain model (WDM) of the human body provides an innovative direction. The AH framework provides a useful approach for systematically structuring domain knowledge in a way that is compatible with how people can solve problems. Empirical evidence that supports these benefits of the AH framework is discussed in Rasmussen et al. [19] and Vicente [20] for other complex work environments (e.g. process control, aviation) and for various design purposes (e.g. display design, training, sensor selection). The AH has been shown to support human problem solving by structuring the domain in a goal-relevant

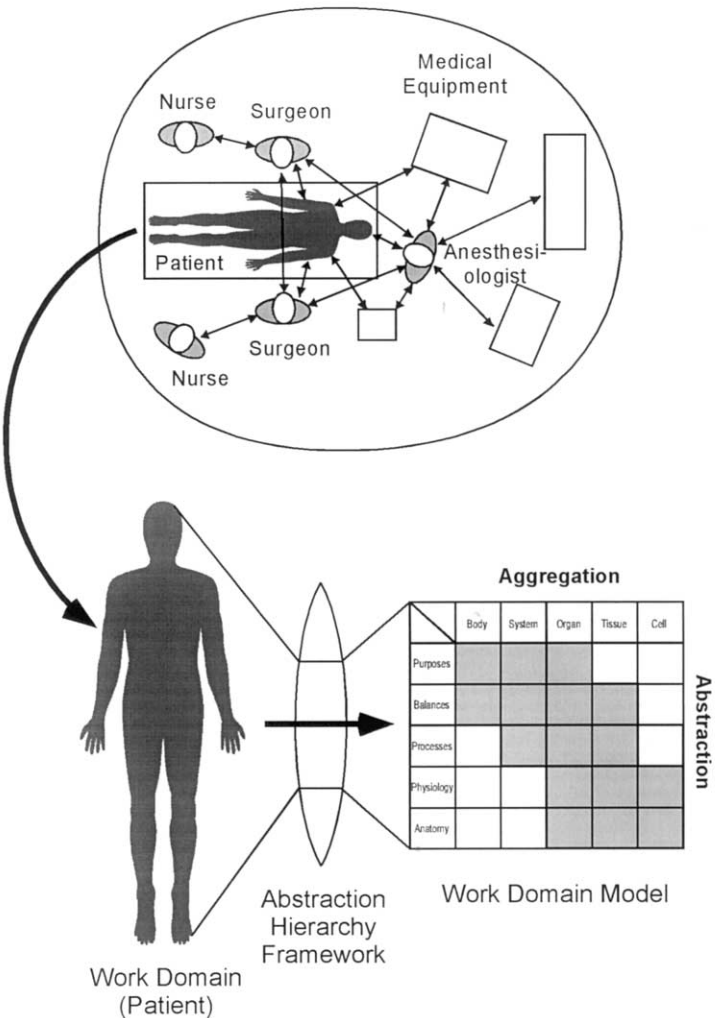

Fig. 2. Creating a work domain model of the human body (i.e. patient) in the operating room.

manner at various levels of function and detail.

Fig. 2 illustrates the process of generating a WDM of the patient (i.e. human body) in an operating room. In this case, the operating room consists of a team of medical personnel (nurses, surgeons, and an anesthesiologist) who interact with each other, the patient, and medical equipment to perform a surgical procedure. A work domain is an aspect or object in this environment that is controlled and, due to its complexity and purpose, can require problem solving by the medical personnel. For example, a work domain could be the patient or a complex medical device (e.g. anesthesia workstation). Here we chose the patient to be the work domain. The AH framework acts as a lens for generating the patient WDM, a functional representation of the problem space. This representation maps higher-level functions of the patient work domain (i.e. purposes) with lower-level action opportunities. The patient WDM is divided into different levels of abstraction (AH) and aggregation (part–whole hierarchy–PWH). Each cell in this patient WDM matrix defines a complete and different causal representation of the same patient work domain, uniquely defined by the particular levels of abstraction and aggregation. Next, each dimension is discussed in more detail.

2.2.1. *Abstraction hierarchy (AH)*

The AH dimension outlines the various goal-relevant, structural means–ends relations of the human body, with familiar references to medical practice. Each level provides a different 'language', or set of concepts, for modeling the same patient work domain. The lower levels include the anatomical and physiological structures, which represent the available resources and component functions in the body. The higher levels include the purposes of the human body. Changes in configurations of the anatomical structures propagate bottom-up to affect the higher-level functions. The purposes propagate top–down, providing reasons for lower-level functions. In the following, each level of abstraction is discussed in turn [21,22].

Purposes: This level contains the physiological purposes governing the interaction between the patient and the medical environment. Examples include homeostasis (maintenance of internal environment), adequate perfusion, circulation, oxygenation, ventilation, and circulatory volume.

Balances: This level contains the concepts necessary for setting priorities and allocating resources to the generic physiological processes. Examples include physiological balances in salt/water, oxygen supply/demand, electrolytes, and conservation relationships (e.g. mass, energy and momentum).

Processes: This level contains generic physiological processes that are to be coordinated irrespective of the underlying physiology and component configuration. Examples include circulation, perfusion, oxygenation, ventilation, metabolism, storage, diffusion, osmosis, binding, chemical release, and heat transfer.

Physiology: This level contains the physiological functions available to establish and maintain the processes. Examples include the functioning of organs and other body components. This is also the level where actions can be performed to reconfigure physiological state (e.g. administer a vasodilator to increase blood flow through the arterioles).

Anatomy: This level includes specific anatomical structures. Examples include the location, appearance, form, and material structure of the human heart.

In Fig. 3, the structural means–ends links between the different levels of abstraction are shown for parts of the patient cardiovascular system. The lower levels include the cardiac and circulatory functions necessary to support the higher-level purposes of adequate circulation and blood volume; the higher levels provide reasons for lower level functions.

Problem solving can occur by shifting the mental focus across these levels of abstraction. Information will be required from the AH level currently in the practitioner's mental focus, including the functional structure, state, and what needs to be controlled (i.e. the WHAT). For example, in Fig. 3, the current task may be to control systemic circulation. Information is also required from the AH level above, which indicates the reason of the control decision (i.e. the WHY). For example, in Fig. 3, the reasons for controlling systemic circulation are to support the functions of mass transfer and balance to the organs. Finally, information is required from the AH level below, which indicates the physiological resources available for implementing the decision (i.e. the HOW). For example, in Fig. 3, heart rate, rhythm, and arteriolar radius can be resources to control systemic circulation. As the practitioner's mental focus shifts across levels of abstraction during the problem solving episode, so does the WHY, WHAT, and HOW formulation, like a 'sliding window' that moves up and down [1].

2.2.2. *Part–whole hierarchy (PWH)*

The PWH dimension outlines the structural organization of the human body, aggregating body components into less detailed objects. This is one method of coping with and managing complexity in a real environment, by changing the span of attention from more detailed to less detailed representations [1]. Components are grouped or 'chunked' together to form a pyramid structure, with the top representing the whole patient work domain and the bottom representing the component parts.

In aggregating the organization of the body into a part–whole structure, a common method found in the medical informatics literature [23] and medical practice is utilized. The body has been generally organized according to the taxonomy noted below [16,24].

Body: A collection of various organ systems structurally and functionally linked as an entity separate from the physical environment.

System: A collection of organs that perform related functions that interact to accom-

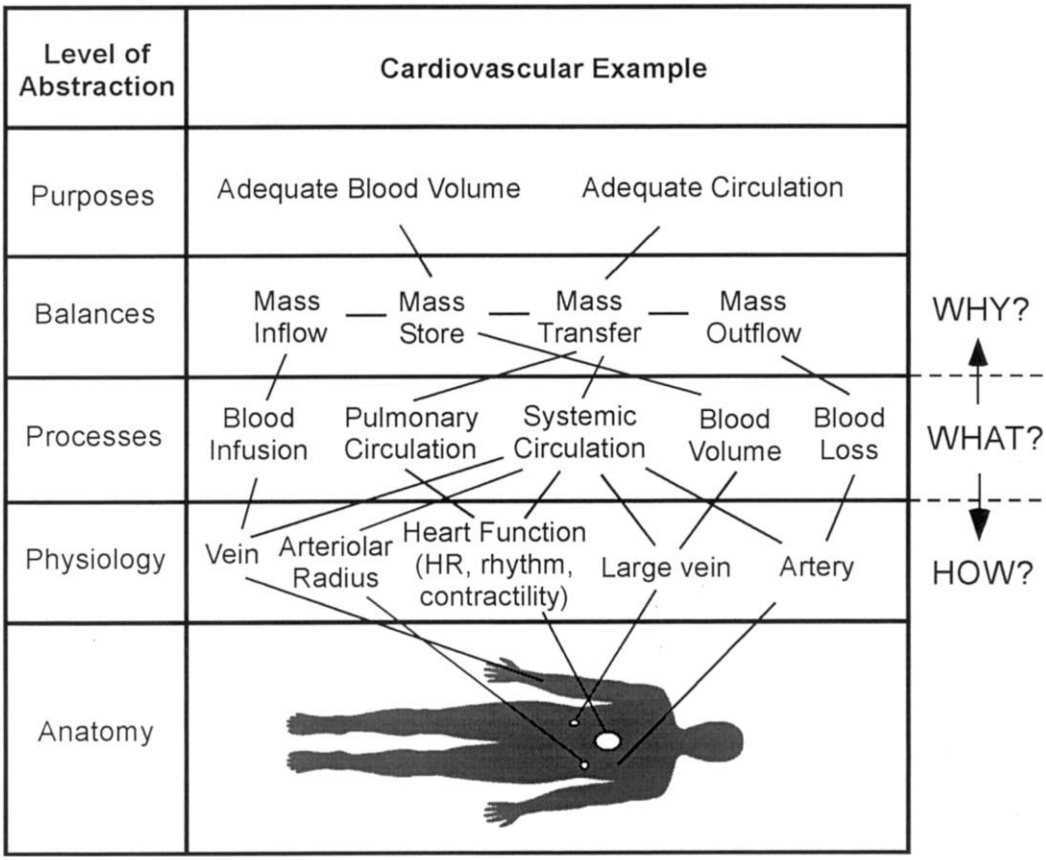

Fig. 3. Partial abstraction hierarchy of the patient cardiovascular system.

plish a common activity that is essential for the survival of the human body.

Organ: Two or more types of primary tissue organized to perform particular functions.

Tissue: A collection of cells of similar structure and function with varying amounts of extracellular material. There are generally four types of primary tissue: muscle, nervous, epithelial, and connective.

Cell: This is the smallest unit capable of carrying out processes associated with life.

Cells are at the bottom of the PWH. These are aggregated up the part–whole structure to form tissues, organs, systems, and finally the whole body.

2.2.3. WDM of the human body — a functional representation of the problem space

The AH and PWH may be combined into a matrix that defines the WDM (i.e. structure and content). Movement across the map can depict a problem solving path taken by the problem solver through the work domain, on the basis of detail (level of aggregation) and goal-relevant function (level of abstraction) [19]. This path also represents shifts in focus and attention span, as the problem solving evolves dynamically.

Applying these principles to the human body, a WDM combines the AH and PWH dimensions (Fig. 4(a)). To illustrate more detail, a vital system that is monitored and controlled during surgery is presented in Fig. 4(b), the patient cardiovascular system, with its purposes of providing adequate transport and distribution of blood throughout the body (i.e. circulation) and the maintenance of blood volume. Only the system and organ levels of aggregation, with their associated sub-levels, are presented because these levels were determined to be most relevant for the purposes specified. In general, the structure of the human body, as represented by the WDM, is similar from person to person, regardless of the context where work is done (e.g. operating rooms, intensive care units, and outpatient clinics). The difference occurs when specific contexts map onto the WDM and demand different information and control requirements of the structure. Examples from the operating room are discussed later in this article.

As mentioned earlier, each cell in Fig. 4 is a different model of the same patient work domain. Accordingly, within each cell, we find a model consisting of different objects or functions connected by causal relations. Fig. 5 illustrates the causal arrangements for selected parts of the human body that are reasonable to illustrate, given the complexity of the cardiovascular system (i.e. levels of abstraction, balances and processes; levels of aggregation, system, sub-system, organ).

2.2.4. Patient WDM as a frame of reference for medical problem solving

The patient WDM can be used as a frame of reference to model and map problem solving behavior in a clinical setting. Fig. 6 illustrates a problem solving trajectory for the onset of a critical event in the operating room. An anesthesiologist's behavior was observed in a high-fidelity simulator, and his actions and verbalizations were first transcribed (Table 1) and then mapped onto the patient WDM [25]. Each numbered node in the diagram refers to a particular patient monitoring or control activity, described in Table 1. For the purposes of this example, only the trajectories related to cardiovascular monitoring and control were considered. In this particular scenario, the patient was undergoing a total knee transplant under spinal anesthesia. The tourniquet was released, resulting in a series of events: hypotension, blood loss, myocardial ischemia, ventricular

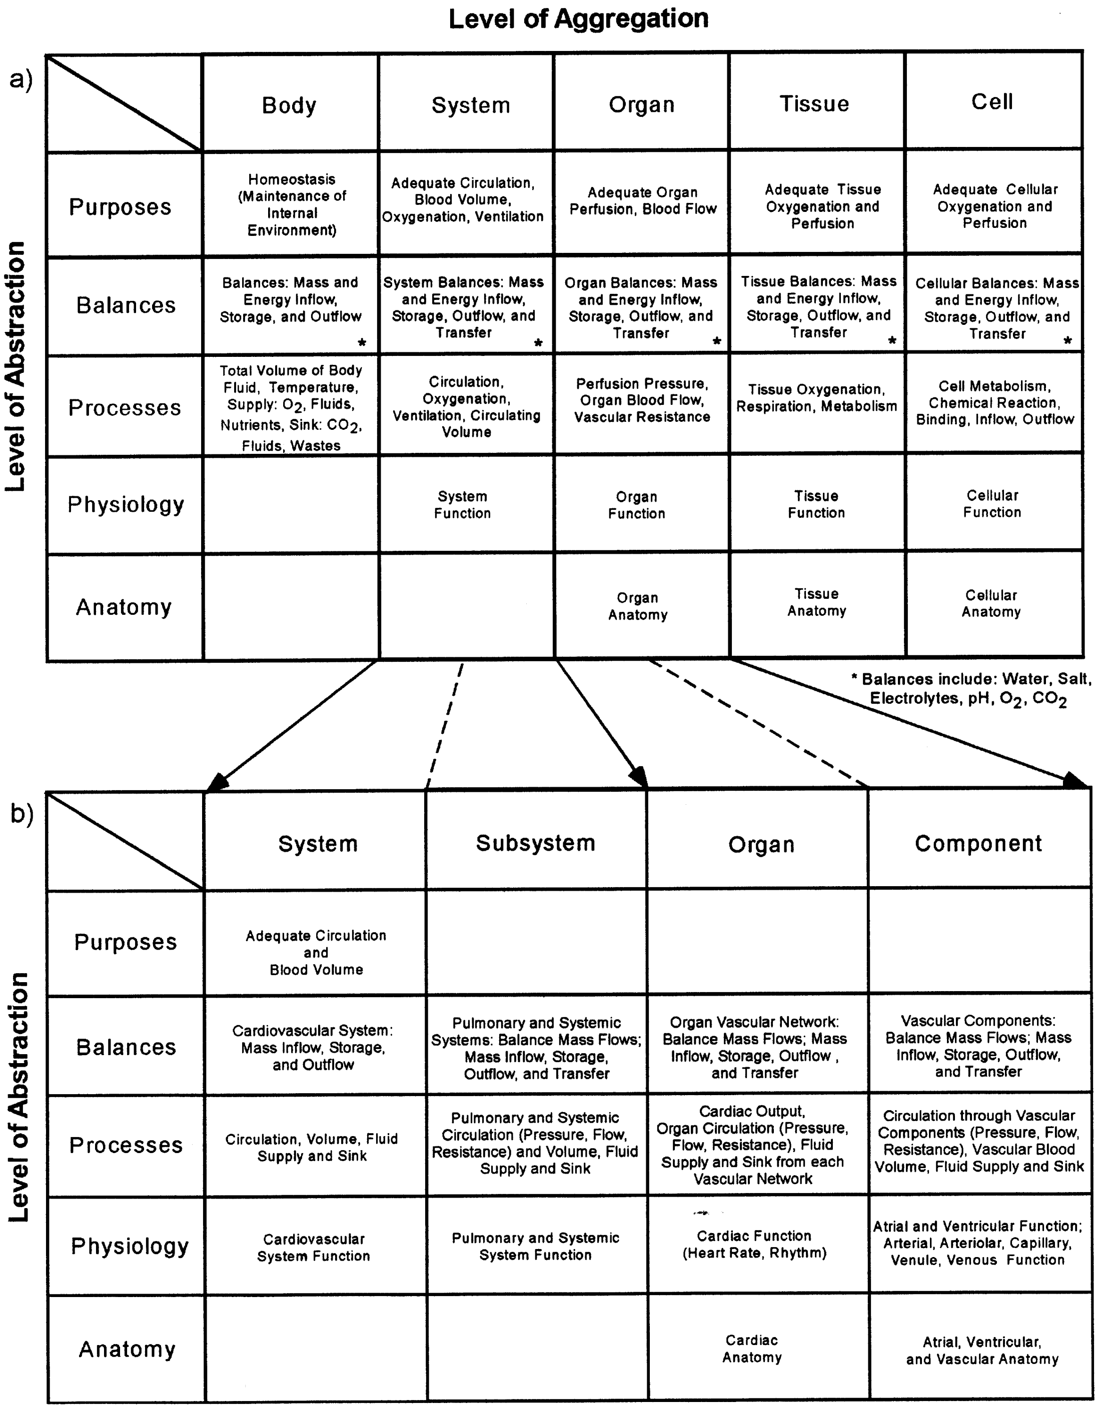

a)

	Body	System	Organ	Tissue	Cell
Purposes	Homeostasis (Maintenance of Internal Environment)	Adequate Circulation, Blood Volume, Oxygenation, Ventilation	Adequate Organ Perfusion, Blood Flow	Adequate Tissue Oxygenation and Perfusion	Adequate Cellular Oxygenation and Perfusion
Balances	Balances: Mass and Energy Inflow, Storage, and Outflow *	System Balances: Mass and Energy Inflow, Storage, Outflow, and Transfer *	Organ Balances: Mass and Energy Inflow, Storage, Outflow, and Transfer *	Tissue Balances: Mass and Energy Inflow, Storage, Outflow, and Transfer *	Cellular Balances: Mass and Energy Inflow, Storage, Outflow, and Transfer *
Processes	Total Volume of Body Fluid, Temperature, Supply: O_2, Fluids, Nutrients, Sink: CO_2, Fluids, Wastes	Circulation, Oxygenation, Ventilation, Circulating Volume	Perfusion Pressure, Organ Blood Flow, Vascular Resistance	Tissue Oxygenation, Respiration, Metabolism	Cell Metabolism, Chemical Reaction, Binding, Inflow, Outflow
Physiology		System Function	Organ Function	Tissue Function	Cellular Function
Anatomy			Organ Anatomy	Tissue Anatomy	Cellular Anatomy

* Balances include: Water, Salt, Electrolytes, pH, O_2, CO_2

b)

	System	Subsystem	Organ	Component
Purposes	Adequate Circulation and Blood Volume			
Balances	Cardiovascular System: Mass Inflow, Storage, and Outflow	Pulmonary and Systemic Systems: Balance Mass Flows; Mass Inflow, Storage, Outflow, and Transfer	Organ Vascular Network: Balance Mass Flows; Mass Inflow, Storage, Outflow , and Transfer	Vascular Components: Balance Mass Flows; Mass Inflow, Storage, Outflow, and Transfer
Processes	Circulation, Volume, Fluid Supply and Sink	Pulmonary and Systemic Circulation (Pressure, Flow, Resistance) and Volume, Fluid Supply and Sink	Cardiac Output, Organ Circulation (Pressure, Flow, Resistance), Fluid Supply and Sink from each Vascular Network	Circulation through Vascular Components (Pressure, Flow, Resistance), Vascular Blood Volume, Fluid Supply and Sink
Physiology	Cardiovascular System Function	Pulmonary and Systemic System Function	Cardiac Function (Heart Rate, Rhythm)	Atrial and Ventricular Function; Arterial, Arteriolar, Capillary, Venule, Venous Function
Anatomy			Cardiac Anatomy	Atrial, Ventricular, and Vascular Anatomy

Fig. 4. A work domain model of: (a) the human body; and (b) the cardiovascular system [21]. Reprinted with permission from *Proceedings of the Human Factors and Ergonomics Society 42nd Annual Meeting**, 1998. Copyright 1998 by the Human Factors and Ergonomics Society. All rights reserved.

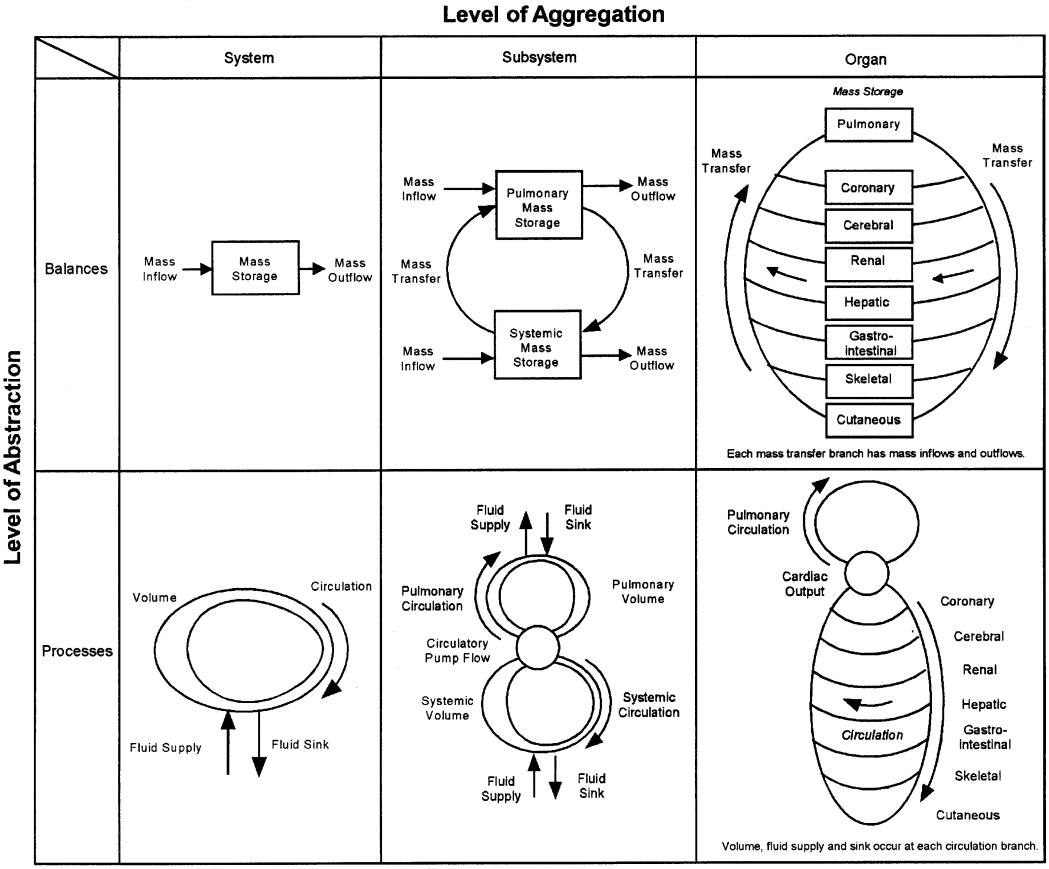

Fig. 5. A causal representation of the work domain of the human body at a particular level of abstraction and aggregation [25]. Reprinted with permission of the author.

tachycardia, fibrillation, and finally cardiac arrest. During the critical event, the anesthesiologist sampled the monitors for patient variables, interacted with the surgeon during blood loss, and performed interventions as the situation progressed.

Using the patient WDM, three observations were noticed about the anesthesiologist's trajectory during the simulator session. First, the problem solving route taken during the simulation was cyclical, generally moving between the higher levels and lower levels of abstraction and aggregation. In the scenario, the cycles corresponded to the anesthesiologist verifying information and monitoring the progress of the interventions initiated. Second, the problem solving trajectory expanded to include more levels of abstraction and aggregation as the critical event was occurring, and contracted as the patient condition became more evident. Third, many trajectory nodes were at the lower levels of abstraction. This behavior was primarily the result of limited access to the patient work domain with the majority of available monitored variables at the physiology and processes levels of abstraction, and organ level of aggregation. It is important to note that this trajectory is specific to the medical scenario utilized. In order to determine any general conclusions regarding the behavior of anesthesiologists in the operating room, additional studies are required.

However, this example demonstrates the potential usefulness of the patient WDM for supporting medical problem solving. First, the framework supports reasoning when dealing with unanticipated events by providing the necessary structural constraints for proper biological function. Because specific scenarios, events, actions, or procedures are not modeled, this approach is inherently tailored to systems that can change (expectedly or unexpectedly) and are adaptive, as demonstrated in Fig. 6. Second, the framework provides a psychologically relevant basis for problem solving by providing a goal-relevant model of the patient work domain [1]. The AH is a complete functional map of the human body at different levels of abstraction and aggregation. One may think about the human body in terms of its purpose and low level of detail, or in terms of the physical characteristics of physiological components at a high level of detail. As shown in Fig. 6, problem solving can be seen as reasoning through this map (between the top-left and bottom-right cells) as one monitors or controls the patient's state. Thus, the AH framework provides a useful approach for systematically structuring biomedical knowledge in a way that is compatible with practitioners' problem solving processes.

2.3. Comparing the traditional approach with the abstraction hierarchy approach

The WDM of the human body developed using the AH framework (Figs. 2–5) may be compared with causal and relational diagrams of the human body using a traditional modeling approach (Fig. 1). First, the causal

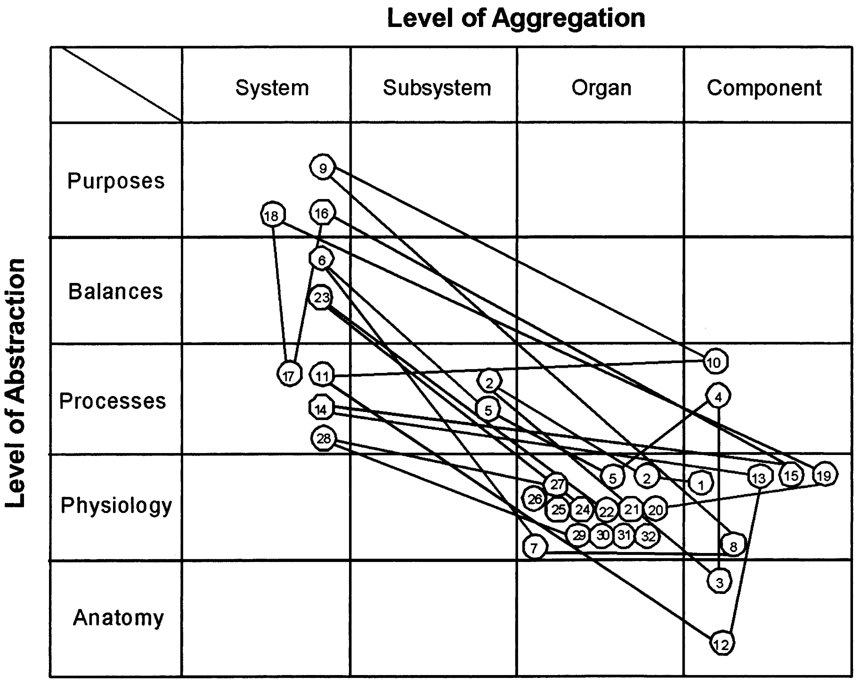

Fig. 6. Mapping problem solving behavior onto the patient work domain model during a surgical procedure.

Table 1
Transcription of reasoning trajectory in Fig. 6 (adapted from [25])[a]

No.	*Patient Variables* [*HR*, *BP*, *SaO$_2$*] *& Status*, Environment Description	Verbal comments
1.	*[86, 139/89, 97] Normal* **A** asks **S** about procedure status. **S** starts to release tourniquet.	**A** → **S:** ‘So are you about to release the tourniquet?’ **S** → **A:** ‘Yes.’
2.	*[86, 139/89, 97] Normal* **A** records patient variables on chart.	
3.	*[87, 124/75, 97] BP dropping* **A** queries **S** about bleeding. **S** looks at surgical field.	**A** → **S:** ‘How’s it looking down there? Any bleeding?’
4.	*[87, 124/75, 97] BP dropping* **S** replies that there is not much bleeding.	**S** → **A:** ‘No real bleeding, just a little bit, not much. It’s under control.’
5.	*[92, 119/74, 97] Minor Hypotension* **A** scans patient monitors.	
6.	*[95, 116/73, 97] Minor Hypotension* **A** asks **S** about blood loss.	**A** → **S:** ‘How much blood do you think you’ve lost?’ **S** → **A:** ‘About 100 so far.’
7.	*[100, 113/72, 97] Minor Hypotension* **A** scans patient monitors.	**A:** ‘Starting to get a little tachycardic here.’
8.	*[105, 108/71, 97] Hypotension* **A** increases IV flow.	**A:** ‘I’m going to increase the IV for the moment, it’s wide open.’
9.	*[105, 108/71, 97] Hypotension* **A** checks with **P**.	**A** → **P:** ‘How are you doing **P**?’ **P** → **A:** ‘I feel nauseated.’
10.	*[130, 78/69, 97] Major Hypotension* **A** asks **S** about blood loss again. **A** twists **P**’s head, as **P** becomes sick.	**A** → **S:** ‘Are you sure you’re … what’s your bleeding like down there?’ **S** → **A:** ‘It’s almost all done, no problem at all.’ **P** → **A:** ‘Oh, I think I’m going to throw up.’ **A** → **P:** ‘OK.’
11.	*[162, 77/67, 97] Major Hypotension* **A** asks **S** about blood loss again.	**A** → **S:** ‘We’re into some major hypotension here. Are you sure you’re not bleeding?’
12.	*[162, 77/67, 97] Major Hypotension* **S** verifies there is minimal blood loss.	**S** → **A:** ‘Trust me, I’m looking at the wound, I don’t see a lot of blood.’
13.	*[162, 80/60, 97] Major Hypotension* **A** is squeezing fluid through IV.	**A** → **S:** ‘I’m squeezing fluid here.’ **S** → **A:** ‘What’s going on?’
14.	*[162, 80/60, 97] Major Hypotension* **A** informs **S** of the patient’s status.	**A** → **S:** ‘He’s quite hypotensive. HR is way up. I haven’t done anything at the moment.’

Table 1 (*Continued*)

No.	*Patient Variables* [*HR*, *BP*, *SaO_2*] *& Status*, Environment Description	Verbal comments
15.	*[162, 81/60, 97] Major Hypotension* **A** puts in drug connection to IV.**A** squeezes fluid in IV and instructs **N** to turn up tourniquet.	**S→A:** 'We took the tourniquet down.' **A→S:** 'OK, can you put it back up?' **A:** 'I'm going to give him 5mg of ephedrine.' **S→A:** 'I have no control over tourniquet.' **A→N:** '**N** could you turn up the tourniquet?' **A:** 'I'm squeezing in…'
16.	*[158, 86/65, 97] Major Hypotension* **A** takes stethoscope to **P'**s chest, while asking **N** to check for blood supplies.	**P→A:** 'I'm sick to my stomach.' **A→P:** 'OK.' **A→N:** 'N, do have any blood in the fridge? We probably will need some.' **N→A:** 'I'll check.'
17.	*[150, 101/77, 97] Major Hypotension* **S** inquires on status.	**S→A:** 'What seems to be the problem there?' **A→S:** 'He's looking to me to be a bit hypotensive. He's getting a little better. Certainly his coronary artery disease does not like what we've done.'
18.	*[136, 107/80, 97] Hypotension* **A** squeezes IV fluid and asks **N** for blood.	**A→N:** 'Can you order 4 units of packed cells? I'm squeezing fluid.'
19.	*[130, 104/83, 97] Ventricular Tachycardia* **A** scans patient monitor, recognizes heart state shift to ventricular tachycardia, and recognizes he needs help.	**A:** 'OK now we have ventricular tachycardia. I'm going to need some help at this point.'
20.	*[176, 98/80, 97] Ventricular Tachycardia* **A** adjusts anesthesia machine and instructs **N** to get **Ah**. **A** places oxygen mask on **P** to oxygenate organs.	**A→N:** 'Can you get me some help. Dr. **Ah** is next door.' **A→P:** 'Here's some oxygen to breath.' **A→N:** 'Can you get Dr. **Ah** quickly?
21.	*[176, 10/10, -] Fibrillation, Cardiac Arrest* **A** notices **P** is in cardiac arrest and asks **S** to perform CPR.	**A:** 'Oh, this is interesting. OK, we have cardiac arrest.' **A→S:** 'Can you start CPR?'
22.	*[43, 88/68, 97] Cardiac Arrest* **S** starts compressions.	
23.	*[43, 88/68, 97] Cardiac Arrest* **S** inquires on problems.	**S→A:** 'What happened? Everything looked fine. Blood loss was 150 tops. Did you give the antibiotics?'
24.	*[43, 88/68, 97] Cardiac Arrest* **A** informs surgeon of anesthesia status.	**A→S:** 'No, I didn't even give the antibiotics. Chest was clear when I listened to it.'
25.	*[98, 99/63, 97] Cardiac Arrest* **Ah** comes in. **A** keeps oxygen mask on **P** and explains situation to **Ah**.	**A→Ah:** 'Cardiac arrest here. I think it was precipitated by blood loss, but he's got underlying coronary artery disease.'

Table 1 (*Continued*)

No.	*Patient Variables* [*HR*, *BP*, SaO_2] *& Status*, Environment Description	Verbal comments
26.	*[92, 93/69, 97] Cardiac Arrest* **A** instructs **Ah** to administer drug and **N** to get crash cart.	**A → Ah:** 'I was wondering if you can give him some lidocane.' **A → N:** '**N**, can you get the crash cart?'
27.	*[92, 104/69, 97] Cardiac Arrest* **Ah** verified dose with **A** and puts in lidocane through IV.	**Ah → A:** 'What do you want — 100 mg?' **A → Ah:** '100 mg.'
28.	*[0, 120/75, 97] Cardiac Arrest* **A** notices IV running out and **Ah** puts another bag on.	**A → Ah:** 'I will need some more IV fluid.'
29.	*[0, 122/71, 97] Cardiac Arrest* **Ah** notes the lidocane is in **P**. **A** asks to stop CPR and check **P**'s rhythm.	**Ah → A:** 'Lidocane is in.' **A → S:** 'OK, could you stop CPR and see what the rhythm looks like.'
30.	*[102, 17/14, -] Cardiac Arrest* **A** checks monitor to see if there is sinus rhythm, and none found. **A** asks someone to defibrillate **P**.	**A:** 'OK, could we defibrillate him?' **A → N:** '**N** are you certified to defibrillate?'
31.	*[98, 192/21, 97] Sinus Rhythm* **Ah** defibrillates **P** – sinus rhythm.	**A:** 'We have sinus rhythm back and we have a pulse.'
32.	*[158, 234/158, 97] Sinus Rhythm* **A** instructs **Ah** to start lidocane drip. **A** notices **P** is starting to breath.	**A → Ah:** 'OK **Ah**, can you set up a lidocane drip.' **A:** 'Oh he's starting to breath spontaneously now.'

[a] No., identification number used in Fig. 2; HR, heart rate; BP, arterial blood pressure; SaO_2, blood oxygen saturation; CPR, cardiopulmonary resuscitation; IV, intravenous; **A,** Anesthesiologist; **S,** Surgeon; **N,** Nurse; **Ah,** Assistant Anesthesiologist.

and relational network diagrams generally mix different levels of abstraction and aggregation into one representation. For example, in Fig. 1 lower levels of abstraction and aggregation (e.g. arteriolar radius and venous valves) are intertwined with higher levels of representation (e.g. salt/water balance). In contrast, the WDM explicitly structures the human body with different stratified representations using three types of links: means–ends, part–whole, and causal. Causal relations using the AH framework are specific to a particular level of abstraction and aggregation. Second, the WDM is organized based on goal-relevant function. In contrast, the causal and relational diagrams are generally not segregated by function, and incorporate control activity affecting behavior. Third, the causal and relational diagrams mix behavior with structure (e.g. sympathetic activity affecting heart rate). The WDM does not include behavior, but provides possibilities and constraints for action.

These differences have implications for medical problem solving. First, by modeling the patient in a goal-relevant manner, one may reason about the patient's condition and perceive intervention opportunities more easily, especially in unanticipated situations. Second, by explicitly structuring the patient using structural means–ends, part–whole, and causal links, one may navigate in this problem space more easily by choosing the appropriate level of focus and span of attention as required [19].

3. Using the WDM for integrated medical informatics design

In this section, we discuss how the patient WDM could be used for systems design in real medical settings. These discussions are based on field and high-fidelity simulator studies conducted in the operating room [25].

3.1. Practitioner roles and responsibilities

The WDM can be used as a shared frame of reference illustrating the relative areas of responsibility, and thus information needs, for various medical personnel. Practitioners work individually or in teams to treat patients in different medical settings. For example, during an operation, the patient is monitored and controlled by medical teams, who have varying goals, objectives, and obligations. The personnel have responsibilities that focus their attention to various parts of the patient work domain. To illustrate this concept, Fig. 7 depicts two main areas of responsibility in the operating room: the surgical aspect and the supporting anesthesia.

The surgical team concentrates primarily on the lower levels of abstraction and at the

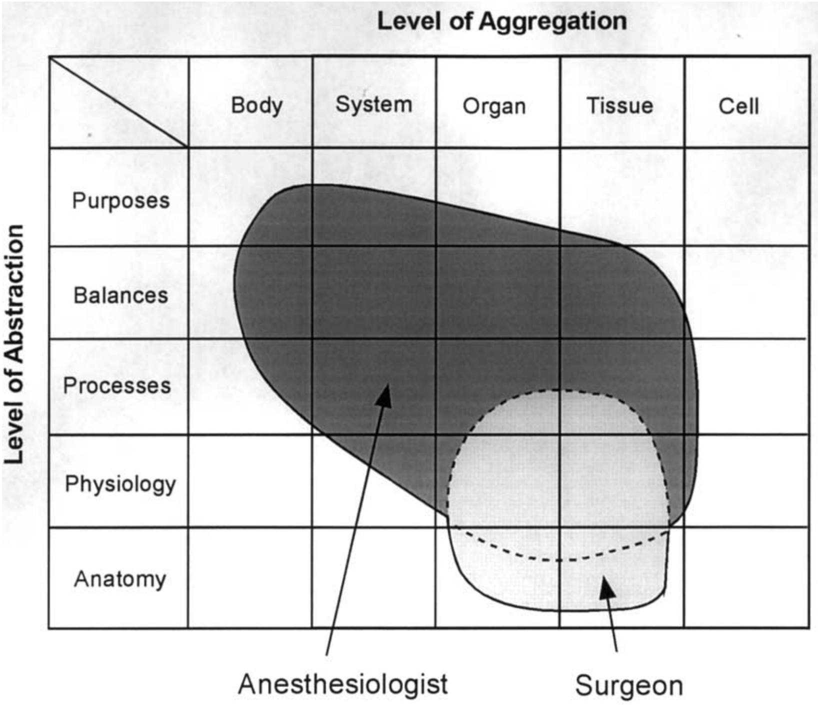

Fig. 7. Allocating team responsibilities onto the patient work domain model [25]. Reprinted with permission of the author.

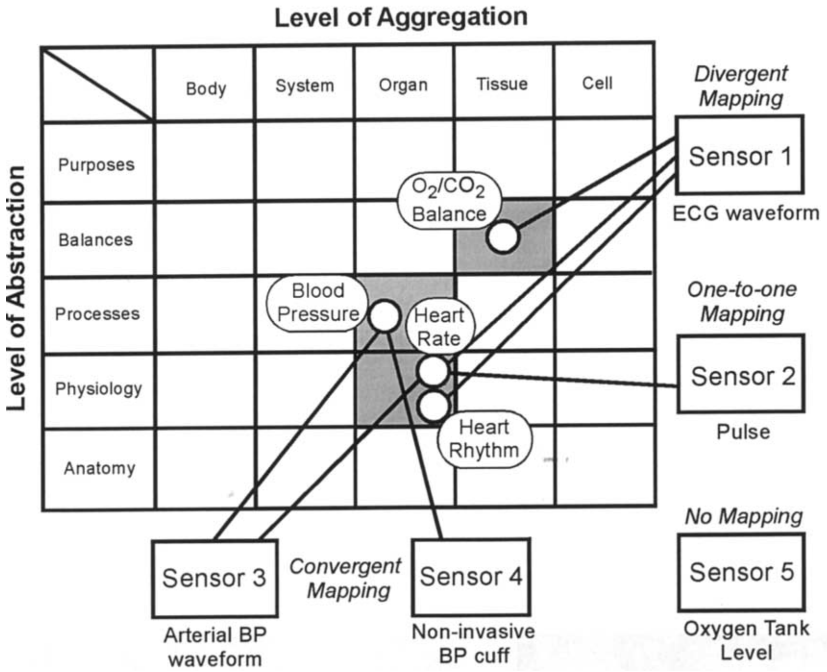

Fig. 8. Mapping operating room sensors onto the patient work domain model.

organ and tissue aggregation levels during the operation (e.g. modifying the patient's current anatomical structure to meet particular surgical goals). The anesthesiologist and anesthesia assistants generally concentrate at the higher levels of abstraction and a broader range of aggregation levels, ensuring patient safety during the operation. The allocation regions on the patient work domain can overlap, requiring the sharing of information regarding planned events, complications, and outcomes, important in formulating goals and coordinating tasks.

These allocation regions will vary for different medical settings. For example, in planning the surgery, both the anesthesiologist and surgeon determine how fit the patient is. This can involve problem solving at all levels in the patient WDM by both practitioners. However, once roles and allocation regions have been determined, information and control requirements (as defined by the patient work domain variables) for each team member emerge. These differing requirements have implications for the design of medical informatics systems that are integrated, yet tailored to different roles and responsibilities.

3.2. Sensors and controls

The patient WDM can also be used as a frame of reference for defining sensors and controls in a medical setting.

3.2.1. Monitoring patient state with sensors

In monitoring, information regarding the patient's status is incomplete and uncertain with respect to different levels in the patient WDM. Only a few variables can be sensed (e.g. ECG waveform, heart rate, blood pressure, blood oxygen saturation, patient's pulse, and skin colour). In addition, practitioners generally cannot see the underlying structures of the patient. Indirect measures, including analytical derivations, are used as surrogate information sources [22,25]. Nevertheless, the patient WDM can still serve as a unifying frame of reference for organizing monitored information from this dynamic work domain.

Fig. 8 shows four types of mapping between the patient WDM and operating room sensors: one-to-one, convergent, divergent, and no mapping. With a one-to-one mapping, one sensor maps onto one patient variable. For example, checking a patient's pulse provides information about heart rate. With convergent mapping (or redundancy), many sensors map onto one patient variable. Practitioners use this method to reduce the high level of uncertainty in measurements from the environment (e.g. artifact, noise, and calibration errors). For example, heart rate can be determined directly from the ECG signal as well as indirectly from other monitor signals (e.g. arterial blood pressure waveforms). With divergent mapping, some sensors provide evidence for many patient variables. For example, the ECG waveform provides evidence for heart rate, heart rhythm, and adequate myocardial oxygenation, among others. Finally, with no mapping, some sensors do not map onto any patient variables. The pressure in an unused oxygen tank at the anesthesia workstation is an example.

3.2.2. Action opportunities to manage patient state with interventions

In controlling the patient's clinical status, medical practitioners can use various interventions (e.g. drug, manual, or equipment) to change the configuration, augment the function, or replace the function of body components as outlined in the WDM. As with monitoring equipment, many of the controls are not direct and may affect many parts of the patient (e.g. the antihypertensive labetolol affects heart rate and vascular resistance). First, the configuration of physiological components in the body may be controlled. For example, the anesthesiologist may use a vasoconstrictor to increase mean arterial blood pressure (processes level). The drug decreases the arteriolar radii (i.e. a change at the physiology level) to achieve the desired response. Second, a physiological component function may be enhanced. For example, the inflow of fluids from capillaries may not be adequate to maintain the immediate requirement of circulating volume. In some cases, the anesthesiologist may squeeze saline through the intravenous line(s) to augment this function. Finally, the function of physiological components may be substituted by using medical equipment. For example, in cardiac surgery, part of the purpose of the cardiopulmonary bypass machine is to function as a blood pump (i.e. heart at the physiology level) while the patient's heart is operated on.

3.3. Training programs

Another application of the patient WDM in medical informatics is to support medical training. The usefulness of the AH framework for training have been previously studied for process control [26].

The patient WDM may be a useful training tool for assessing the patient before, during, and after the operation. Most surgical patients have some form of depressed function in parts of their body systems, as a result of previously developing diseases or the selection of anesthesia for the operation. The patient WDM can be used to assess these situations by mapping the patient's case and history onto the WDM. For example, in the case of coronary artery disease, the condition propagates bottom up through the patient WDM, constraining the functions at higher levels of abstraction. At the anatomy level, a patient with coronary artery disease has plaque build up on specific cardiovascular regions to the heart. This condition has the effect of impeding arterial function (physiology level), blood flow (processes level), mass flow to the heart (balances level), and adequacy of circulation (purposes level). Using the WDM model, practitioners may be able to better manage these situations and choose particular interventions as required. For example, Muravchick [27] discusses possible strategies for compensating for this constrained function. Anesthesiologists may decide to use drugs that minimize the cardiac demand for oxygen during surgery, or use a therapeutic agent (e.g. vasodilator) to increase blood flow to the heart. These interventions are selected to maintain the balance between oxygen supply and demand (balances level). This application of the patient WDM may be useful as a frame of reference for both training for problem solving in real medical situations and in case-based learning with hypothetical medical situations.

3.4. Integrated human–computer displays

The WDM can also be used as a basis for display design [28]. The AH framework can be used as a modeling tool to represent the work domain in terms of information content and structure. Fig. 9 shows conceptually how the AH can be implemented for medical work environments, with extensions to medical teams. The patient WDM can be used to determine, integrate, and organize the information requirements for human–computer displays, in a way that is compatible with how medical practitioners can do problem solving in the context of a medical environment (i.e. medical team in an operating room). For example, the information requirements for surgeons are different compared with anesthesiologists, although both need the same information from overlapping regions of the patient WDM. Once these requirements are determined for each person or team, an appropriate form may be designed

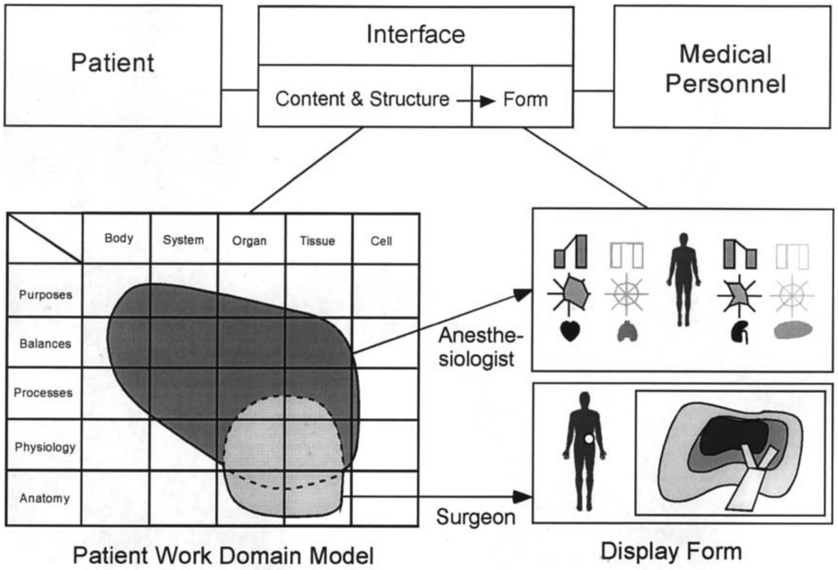

Fig. 9. Directions for integrated medical informatics design for human–computer displays.

that is compatible with the capabilities of the medical personnel (e.g. using perceptual features instead of numbers, requiring less analytical reasoning). For example, as shown in Fig. 9, surgeons may primarily need a view of the surgical field and an anatomical map. Anesthesiologists, on the other hand, may need integrated higher-level representations of how well the patient is coping throughout the surgical procedure. This is indicated in Fig. 9 as a series of generic integrated displays.

There are several examples in the literature where the AH was used for display design in various work domains (e.g. aviation, medicine, process control) [20]. In medical informatics, Sharp and Helmicki [22] designed a human–computer display for a neonatal intensive care unit. They developed a WDM for tissue oxygenation, and created a display based on this model. An experiment was conducted in a simulated clinical environment. Despite the constraints of limited sensor technology and information availability, the new display compared favorably to the existing display.

4. Conclusions

This research demonstrates that the AH framework can be used as an ontological framework to bridge the medicine–psychology gap by structuring biomedical knowledge in a way that is compatible with practitioner problem solving. We also showed that the patient WDM can be applied to real medical settings and provide insight into integrated medical informatics design. By developing representations using the AH framework, medical informatics designers can identify information requirements that can then be used to design novel training programs and human–computer displays.

There are two key advantages of this approach compared with a traditional modeling approach. First, the AH framework provides the necessary medical informatics for dealing with unanticipated events by providing the structural constraints for proper biological function. In contrast, causal and relational diagrams from a traditional approach, incorporated in some medical decision support systems, emphasize the anticipated and may leave out factors necessary for problem solving in unanticipated situations. Second, the AH approach explicitly structures the patient WDM using specific types of links (i.e. means–ends, part–whole, and causal) that make it easier to navigate in the problem space of patient care. Causal and relational diagrams from a traditional approach can be conceptually muddled, showing different types of links within the same flat representation, making it relatively harder to solve problems.

Despite these contributions, this research has several limitations that motivate several research topics. First, the models developed in this article were limited in scope and did not include the integration of all vital body functions. Further research is required to broaden the scope and increase the detail of analysis. Second, no specific human–computer display designs were introduced for the operating room using the results of this research. However, some researchers have designed displays for other medical environments (e.g. neonatal ICU) using this approach [22]. Empirical evidence has shown there are benefits over existing display design approaches. The next step is to continue with display design for medical team environments, incorporating the results of this research. Third, the AH framework does not explicitly model aspects of medical decision making beyond problem solving (e.g. rule-based knowledge in clinical reasoning and strategy formation). Other modeling techniques need to be incorporated, taking into account these considerations [19,20].

Acknowledgements

We would like to thank the medical staff at The Toronto Hospital, and Dr Matt Kurrek at the Sunnybrook Health Sciences Centre in Toronto for their assistance during the data collection process. In addition, we would like to acknowledge support from the Natural Sciences and Engineering Research Council of Canada.

References

[1] J. Rasmussen, The role of hierarchical knowledge representation in decision making and system management, IEEE Trans. Sys. Man Cybernetics 15 (1985) 234–243.

[2] R. Pool, Beyond Engineering: How Society Shapes Technology, Oxford University Press, New York, 1997.

[3] F.K. Orkin, Patient monitoring in anesthesia as an exercise in technology assessment, in L.J. Saidman, N.T. Smith (Eds.), Monitoring in Anesthesia, Butterworth-Heinemann, Boston, MA, 1993, pp. 439–455.

[4] D.R. Westenskow, Alarms, in: L.J. Saidman, N.T. Smith (Eds.), Monitoring in Anesthesia, Butterworth-Heinemann, Boston, MA, 1993, pp. 475–490.

[5] F.A. Mora, G. Passariello, G. Carrault, J.-P.L. Pichon, Intelligent patient monitoring and management systems: A review, IEEE Eng. Med. Bio. Mag. 1993, 23–33.

[6] J. Stausberg, M. Person, A process model of diagnostic reasoning in medicine, Int. J. Med. Informatics 54 (1999) 9–23.

[7] E.S. Berner, R.S. Maisiak, C.G. Cobbs, O.D. Taunton, Effects of a decision support system on physicians' diagnostic performance, J. Am. Med. Informatics Assoc. 6 (1999) 420–427.

[8] C. Cao, T.-Y. Leong, A.P.K. Leong, F.C. Seow, Dynamic decision analysis in medicine: A data-driven approach, Int. J. Med. Informatics 51 (1998) 13–28.

[9] N. Maglaveras, T. Stamkopoulos, K. Diamantaras, C. Pappas, M. Strintzis, ECG pattern recognition and classification using non-linear transformations and neural networks: A review, Int. J. Med. Informatics 52 (1998) 191–208.

[10] J. Vila, J. Presedo, M. Delgado, S. Barro, R. Ruiz, F. Palacios, SUTIL: Intelligent ischemia monitoring system, Int. J. Med. Informatics 47 (1997) 193–214.

[11] B. Thull, H.-J. Popp, G. Rau, Man-machine interaction in critical care settings, IEEE Eng. Med. Bio. Mag. 1993, 42–49.

[12] J.H. Siegel, Integrated approaches to physiologic monitoring of the critically ill, in: J.S. Gravenstein, R.S. Newbower, A.K. Ream, N.T. Smith (Eds.), An Integrated Approach to Monitoring, Butterworths, Boston, MA, 1983, pp. 41–57.

[13] J. Starren, S.B. Johnson, An object-oriented taxonomy of medical data presentations, J. Am. Med. Informatics Assoc. 7 (2000) 1–20.

[14] T.G. Peterson, Design of the four-component anesthesia workstation, J. Clin. Monitoring 10 (1994) 335–338.

[15] V.L. Patel, D.R. Kaufman, Medical informatics and the science of cognition, J. Am. Med. Informatics Assoc. 5 (1998) 493–502.

[16] L. Sherwood, Human Physiology, second ed., West Publishing Company, New York, 1993.

[17] D.E. Mohrman, L.J. Heller, Cardiovascular Physiology, fourth ed., McGraw-Hill, New York, 1997.

[18] M.N. Gomez, Planning for monitoring in healthy patients, in: M.C. Rogers, J.H. Tinker, B.G. Govino, D.E. Longnecker (Eds.), Principles and Practice of Anesthesiology, vol. 1, Mosby Year Book, St. Louis, MO, 1993, pp. 74–85.

[19] J. Rasmussen, A.M. Pejtersen, L.P. Goodstein, Cognitive Systems Engineering, Wiley, New York, 1994.

[20] K.J. Vicente, Cognitive Work Analysis: Toward Safe, Productive, and Healthy Computer-based Work, Erlbaum, Mahwah, 1999.

[21] J.R. Hajdukiewicz, D.J. Doyle, P. Milgram, K.J. Vicente, C.M. Burns, A work domain analysis for patient monitoring in the operating room, Proceedings of the Human Factors and Ergonomics Society 42nd Annual Meeting, 1998, pp. 1038–1042.

[22] T.D. Sharp, A.J. Helmicki, The application of the ecological interface design approach to neonatal intensive care medicine, in Proceedings of the Hu-

man Factors and Ergonomics Society 42nd Annual Meeting, 1998, pp. 350–354.

[23] C. Rosse, J.L. Mejino, B.R. Modayur, R. Jakobovits, K.P. Hinshaw, J.F. Brinkley, Motivation and organizational principles for anatomical knowledge representation: The digital anatomist symbolic knowledge base, J. Am. Med. Informatics Assoc. 5 (1998) 17–40.

[24] A.C. Guyton, Textbook of Medical Physiology, seventh ed., W.B. Saunders, PA, 1986.

[25] J.R. Hajdukiewicz, Development of a Structured Approach for Patient Monitoring in the Operating Room, Unpublished M.A.Sc. Thesis, University of Toronto, 1998.

[26] C.N. Hunter, K.J. Vicente, F. Tanabe, Can 'theoretical' training improve fault management performance?, Proceedings of the 1996 American Nuclear Society International Topical Meeting on Nuclear Plant Instrumentation, Control and Human–Machine Interface Technologies, 1996, pp. 683–690.

[27] S. Muravchick, The Anesthetic Plan: From Physiological Principles to Clinical Strategies, Mosby Year Book, St. Louis, MO, 1991.

[28] K.J. Vicente, J. Rasmussen, Ecological interface design: Theoretical foundations, IEEE Trans. Sys. Man Cybernetics 22 (1992) 589–606.