Application and Testing of Modifications to the TREAT Sepsis Network

Logan Ward

Group Number 1089

June 1, 2012

Abstract

Historically, sepsis has been a condition which is difficult to both diagnose and treat. It is a leading cause of death in part because many patients receive inappropriate antibiotic treatment. Although understanding of the pathophysiology of sepsis has improved, the search continues for a magic-bullet biomarker which can confirm diagnosis in a timely and cost-effective manner. In the area of sepsis, TREAT has been successful so far in classifying septic patients in terms of severity, allowing higher risk patients to undergo more expensive tests which may help to identify the cause of sepsis and thus recommend appropriate treatment. However, the sepsis network in TREAT is old, and potentially out of date, and it is thought that further improvements could be made to this part of the network.

The project builds on previous work completed by the author on sepsis diagnosis, with the implementation of modifications to the sepsis network. Several sets of modifications are presented, from structural changes such as the addition of new nodes, redesign of sections of the network and change to the states within nodes along with the replacement of discrete chance with continuous chance nodes. Automatic learning within causal probabilistic networks (CPNs) is also investigated, with descriptions of the learning process within the Hugin development environment as included. Methodologies for learning within the sepsis network are presented along with a description of the database requirements for such a process to succeed.

Contents

1	Intr	roduction	2			
2	Background					
	2.1	Sepsis: History, Diagnosis and Treatment	3			
	2.2	Use of Clinical Decision Support Systems	5			
	2.3	Causal Probabilistic Networks and Their Uses in Medicine	6			
	2.4	The TREAT CPN	8			
	2.5	Previous Work	12			
3	Literature Search					
	3.1	Differentiation between bacterial and viral infections	15			
	3.2	Differentiation between sepsis and non-infectious SIRS, and the states of each	17			
4	Sepsis Network Modifications					
	4.1	Strategy	19			
	4.2	Modifications Made - Structural Changes	20			
	4.3	Modifications Made - Continuous Distributions	27			
5	EM	Learning for Optimisation	37			
	5.1	Hugin EM Learning Tutorial	37			
	5.2	Patient Database used for Learning	40			
	5.3	Data Preparation and Mapping	40			
6	Dis	cussion	43			
7	Cor	nclusion	44			

Intended Audience

This report is intended for those interested in medical decision support, and in particular, causal probabilistic networks. It is suitable to be read by students in the field of biomedical engineering, as well as those with a particular interest in the subject area. It is expected that the reader will have at least some medical knowledge and a basic understanding of probability calculus.

1 Introduction

The Treat CPN is a causal probabilistic network designed to identify infectious diseases and recommend optimal antibiotic treatment based on coverage, economical cost and ecological cost. By identifying the causative agent, broad spectrum antibiotics, and the associated risk of creating resistance, can be avoided. Treat is currently in use in Hvidovre Hospital and to be installed at Sygehus Lilleblt, with plans for implementation in several more hospitals in Denmark and internationally. Since its inception, Treat has improved both the percentage of appropriate antibiotic treatments and the coverage of such treatments (compared with a physician-only regime) while reducing the ecological cost. The Treat network is continually being modified and optimized to ensure that the best product is available. One area that has been identified for possible improvements is the sepsis network.

Historically, sepsis is a very difficult disease to diagnose and also to treat. Despite extensive research and clinical trials, sepsis remains a leading cause of death, with mortality rates of 18-29% reported in US epidemiological studies [1, 2]. Early appropriate antibiotic treatment is essential, since it approximately halves the mortality. Relative mortality increases by 8% for each hours delay after onset of infection. Early appropriate antibiotic treatment requires that the infection is detected and subsequently that appropriate treatment is administered. A CPN such as Treat has great potential for the early detection of sepsis, and is able to then stratify patients according to risk. High-risk patients can then be given further, more expensive diagnostic tests, such as PCR tests to identify the pathogens, while lower risk patients can be managed accordingly.

This project describes steps taken to improve Treat's sepsis network, and the results of the improvements. The work builds on knowledge obtained through a previous Aalborg University project which focused on the diagnosis on sepsis, and identified ways in which the sepsis network could be improved.

2 Background

This section aims to provide background information on sepsis, medical decision support, and the TREAT system. In particular, the history of sepsis, information about diagnosis and treatment, the history of clinical decision support systems and their potential for use in modern clinical settings, information about causal probabilistic networks, how they work and their uses, a description of the TREAT CPN, its background in terms of development and clinical use and effectiveness, and a summary of the previous semesters project on sepsis diagnosis.

2.1 Sepsis: History, Diagnosis and Treatment

Despite the best efforts of clinicians and researchers, sepsis remains a leading cause of death in critically ill patients [3, 4, 5]. It is clear from the literature that although sepsis is a very well researched and document subject within medical circles, it is often a very difficult condition to diagnose and to treat [3, 4, 5, 6, 7]. Sepsis is defined as a systemic inflammatory response due to the presence of an infection [8]. In 1992, a consensus conference held by the American College of Chest Physicians decided that a patient will be diagnosed with sepsis if they present with SIRS (Systemic Inflammatory Response Syndrome), which is readily diagnosed by the presence of at least two of:

- Fever or hypothermia
- Tachycardia
- Tachypnea or hyperventilation
- Leukocytosis or leukopenia

in addition to a confirmed infection [9]. However, due to the lack of specificity given by this list, the criteria (again in addition to a confirmed infection) were extended following an international consensus conference in 2001 [10]:

- General variables
 - Fever or hypothermia

- Tachypnea
- Altered mental state
- Unexplained hyperglycemia
- Inflammatory variables
 - Leukocytosis or leukopenia
 - Increased CRP (>2 SD above normal)
 - Increased PCT (>2 SD above normal)
- Tissue perfusion variables
 - Unexplained hyperlactatemia
 - Decreases capillary refill or skin mottling
- Organ dysfunction variables
 - Unexplained hypoxemia
 - Acute oliguria
 - Coagulation abnormalities
 - Ileus
 - Hyperbilirubinemia
 - Thrombocytopenia

More specifically, severe sepsis is defined as sepsis with an associated acute organ failure, and septic shock is severe sepsis where the patient remains hypotensive despite adequate fluid resuscitation. These diagnoses do not define different illnesses, but instead define more serious conditions of the same illness.

The high mortality rates associated with sepsis, in addition to its prevalence in intensive care patients, has lead to extensive research into potential biomarkers and methods of treatment. A review published in 2010 [5] describes 178 different biomarkers that have been evaluated for use in sepsis. However, out of these 178 biomarkers, few were tested for diagnosis, with most evaluated as prognostic markers. None of the biomarkers listed were able to achieve the sensitivity or specificity that would be required for routine clinical use. The main problem with the use of biomarkers in sepsis is that they are not able to accurately distinguish sepsis from other diseases or conditions. Although no single biomarker has been effective, the author notes that there is potential that a combination of markers would be more successful.

Diagnosis is not the only problematic area in sepsis studies, with effective treatment also proving elusive. Therapies targeting some of the physiological response to sepsis, for example, modulating the effects of inflammatory mediators such as tumor necrosis factor alpha ($\text{TNF}\alpha$) or interleukin-1 (IL-1), have been ineffective, and in some cases damaging as the host's immune response to the microbial invasion is mitigated [6]. In 2008, a set of international guidelines for the treatment and management of severe sepsis was published [11] as an update to the guidelines formed by a group of 11 organizations - Surviving Sepsis Campaign guidelines for management of severe sepsis and septic shock, published in 2004. Recommendations and suggestions are presented, along with the rationale for such decisions given the case presented with. However, this is not a definitive guide on sepsis treatment, as progress is still being made on new treatments and diagnostic techniques.

2.2 Use of Clinical Decision Support Systems

Clinical decision support systems have at times met with controversy, with few published details of clinical success. Garg [12] identified 100 decision support systems in a 2005 review, including systems for diagnostics, reminders, disease management and drug-dosing or prescription. Although many of the systems (64%) reported improvements in practitioner performance, the effect on patient outcome is still largely unclear. Although there is seemingly a lack of success in the field of medical decision support, the absence of widespread success does not imply that the idea of implementing such a system is flawed, only that in many cases the design and/or implementation of such systems have been.

There are three main variants of clinical support systems; knowledge-based systems (includes rulebased and fuzzy logic systems), model-based systems (including differential equations and causal probabilistic networks, CPNs) and data-based systems (including neural networks and regression models). Model-based systems have the advantage that they form a hybrid approach where the structure/model is based on knowledge and the system can be tuned by data. The hypothesis is that model-based systems are able to utilize the benefits of both knowledge- and data-based systems while eliminating the problems with each.

The Center for Model-based Medical Decision Support (MMDS) at Aalborg University has developed several MMDS systems, including models based on differential equations and CPNs. The differential equation, or compartment models are Glucosafe for blood glucose control in the intensive care unit (ICU) and INVENT, a system used to determine optimal ventilator settings. MUNIN, a neuromuscular network, and TREAT, a network for optimal therapy for infectious diseases are examples of CPN-based systems.

2.3 Causal Probabilistic Networks and Their Uses in Medicine

Causal probabilistic networks or CPNs, also known as Bayesian networks, are sets of causally-linked stochastic variables which form a network. Each stochastic variable is represented by a node, each of which has a set of possible values or states. Relationships between nodes can be represented by directional links showing the direction of causality. Each node has a table of conditional or a priori probabilities, based on the possible states of, and whether it has parents.

The simplest form of CPN is the idiot-Bayes (or naive-Bayes) network where there is a single parent node, simply connected to a set of child nodes. The children are assumed to be independent of each other and conditional on the parent. In more complex cases, a naive-Bayes approach may not be appropriate and nodes can be multiply connected, where nodes may have more than one parent on which they are conditional.

TREAT and MUNIN, both products of the Aalborg University MMDS Department. The MUNIN CPN contained approximately 1000 nodes which were organised into four layers. The top layer consists of the various neuromuscular diseases that MUNIN can diagnose, two layers of intermediate nodes, and a bottom layer of observable findings, as shown in Figure 1. The intermediate nodes are unobservable, and represent pathophysiological concepts.

The TREAT CPN, which will be discussed further in Section 2.4, is a little more complicated than MUNIN. TREAT is designed for the identification and optimal treatment of infectious diseases, and contains around 8000 nodes. There are several substructures within the TREAT network which is



Figure 1: Structure of the MUNIN CPN

split into infection sites, and then each infection site into sections dealing with treatment, pathogen prevalence, symptoms and risk factors.

The structuring of CPNs and formation of conditional (or a priori) probability tables is only one part of the solution. There must exist an engine to perform inferences within the network based on evidence entered, for this is where CPNs can become powerful. Given a accurately defined causal structure and appropriate probability tables, evidence entered as findings can be used to infer the cause. A major step towards making this possible for complex medical models was the development of efficient algorithms for inference in both singly connected CPNs (Pearl [13]) and multiply connected CPNs (Lauritzen and Spiegelhalter [14]).

The development of CPNs is made easier by software designed for the purpose. Hugin, developed by the company HUGIN EXPERT A/S, has existed since 1989 and provides a graphical user interface (GUI) for CPN development as well as the inference engine for efficient updating of probabilities [15]. Hugin provides tools for developing networks with both discrete and conditional chance nodes. Conditional chance nodes allow normal distributions to be specified as a mean and variance rather than specifying conditional probabilities for each state. Methods are also provided for automatic learning within CPNs, where probability tables and distributions can be learned from databases. The learning process is discussed further in Section 5.

2.4 The TREAT CPN

TREAT is a DSS that was developed after its creators identified a need to provide improved empirical antibiotic treatment, and the fact that a computerized system could better utilize information available to physicians on presentation at a hospital. A causal probabilistic network was chosen as the platform for the model to best reflect the knowledge obtained through certain evidence. TREAT aims to increase appropriate empirical antibiotic usage, which has been shown to improve patient outcomes [16, 17], while reducing both financial and ecological costs. The ecological cost is determined by the the increased chance of developing antibiotic-resistant strains that has been linked to the use of broad-spectrum antibiotics [18]. TREAT was designed to be used to identify hospital- and community-acquired infections in inpatients and suggest appropriate antibiotic treatment [19, 20, 21, 22, 23, 24].



Figure 2: Full treat network overview. Each infection site is highlighted blue, the bacteremia network red and the sepsis network green

TREAT aims to identify both the site of infection as well as the pathogen responsible for causing the infection. TREAT covers eleven sites of infection, with the basic units being the individual pathogens at each site. Figures 2 to 5 give an overview of some of the structures used within the TREAT CPN. Figure 2 shows how the network is organised into 11 infection sites, connected to networks for bacteremia and sepsis.

Each site is broken down further into the different pathogens that are likely to cause infections in that particular location. Each pathogen has an associated network for treatment/coverage which allows for a combination of up to three drugs. Figure 3 shows part of the pneumonia network structure. Four pathogens are shown (four groupings along the bottom) and the treatment networks for two of these (tree structures at the top).



Figure 3: Overview of the repeated structure used in each of the infection sites, in this case, the pneumonia network

Figures 4 and 5 show closer views of the treatment/coverage network and the individual pathogen, respectively. The pathogen network is connected directly to risk factors, for example, the pneumococcal pneumonia node shown is connected to nodes for copd, age, and alcohol abuse among others. Likelihoods for pathogens, and the severity of the infection they are likely to cause, are collected for each site at a point where site-specific symptoms are recognized. Each site is then connected to the sepsis and bacteremia networks, where likely severities are added together to give likelihoods for sepsis and whether bacteria will be found in the blood.



Figure 4: Closer view of the repeating structure shown in Figure 3 for coverage of pneumococcal pneumonia infections



Figure 5: Close up view of the network for pneumococcus pneumoniae

Following the identification of the most likely pathogen and the infection site, TREAT identifies the

most appropriate treatment based on a cost-benefit model. Benefits are derived from the coverage of the specific antibiotic in relation to the identified site and pathogen as well as increased or decreased efficacy when combined with other drug therapies among other factors. The cost component includes the cost of the drug, including administration and monitoring, cost of potential side-effects, the cost of resistance induction, and a penalty cost formed from a combination of local preferences and an additional cost for being a drug of last resort - one that after which, there is no broader-spectrum available. Before implementation in any site, TREAT is calibrated. This involves accounting for local pathogen prevalences as well as tailoring towards local preferences in terms of antibiotic usage, the latter of which is done by increasing the penalty cost on certain drugs.

TREAT has been studies in randomized comparative trials, described in greater detail in [23]. The results of the trial were that TREAT improved appropriate empirical antibiotic usage over clinicians (73% compared to 64%) when comparing wards where TREAT was used to control wards, with the figure for TREAT improving to 85% in all cases where the clinician followed the advice given by TREAT. In addition, the length of hospital stay, ecological costs due to the potential for developing resistance as well as the total antibiotic costs were also reduced in intervention versus control wards.

2.5 Previous Work

A previous semester project focused on identifying the pathophysiology of sepsis in greater detail. With the availability of a more complete picture of the relations between physiological mediators and sepsis symptoms, the causal network in TREAT could be examined for potential areas for improvement. The development of this causal diagram consisted of finding the missing links between the symptoms used in the current TREAT network, expanded to include a few common symptoms known to be missing, and the root cause of the septic condition - the microbial infection.

More specifically, the identification of sepsis pathophysiology focused on the existing symptoms and lab measurements qualified in the TREAT sepsis network; temperature, serum albumin, erythrocyte sedimentation rate (ESR), C-reactive protein (CRP), leukocyte count, serum creatinine, blood pressure, disseminated intravascular coagulation (DIC) and acute respiratory disease syndrome (ARDS). Essentially, sepsis is caused by an imbalance between the inflammatory response to an infection and the compensatory anti-inflammatory response. The extended network is shown in Figure 7.

The main inflammatory mediators are the interleukins, particularly IL-1 and IL-6, along with tumor necrosis factor alpha (TNF- α). There are then several pathways that are activated in different cases, involving different organ systems and other systemic markers. Temperature is affected by the adjustment of the thermoregulatory setpoint by prostaglandin E2 (PGE2) in the hypothalamus. The release of PGE2 is stimulated by $TNF\alpha$ and IL-1. Serum albumin can be modified through one of two pathways; the production of albumin can be modified, or changes in vascular permeability can lead to albumin leaking out into the interstitium. DIC and ESR are affected by the induction of the coagulation cascade by the complement system as well as the release of acute phase proteins. CRP is an acute phase protein, and thus is often combined/used interchangeably with ESR. Serum creatinine is a common marker of kidney function, with increased creatinine a result of a reduction in the glomerular filtration rate. Kidney damage in sepsis can be traced back to the release of vasoactive substances triggered by IL-1 and the activation of the complement system. Blood pressure (and as a direct consequence, heart rate) is modified on two main fronts, volemic issues are caused by increased vascular permeability leading to capillary leakage, and pressure is also affected by the release of vasoactive substances. Increased vascular permeability is the main cause of ARDS, with the fluid leakage causing edema, which when in the lungs leads to impaired oxygen perfusion and respiratory issues.



Figure 6: Diagram of the fever induction pathway in an infection [25] alongside the causal pathway extracted from Figure 7

One example of a pathway in the network is shown in Figure 6, which compares the pathophysiological process of fever induction as described in [25] to the causal pathway extracted from the extended causal network shown in Figure 7. The causal diagram is obviously somewhat simplified, and needs to be simplified further if it is to be used for diagnostic purposes. If a causal network is too complex, with too many unobservable nodes, it will not be useful.

During this project, the idea of replacing discrete chance nodes with continuous chance nodes was also addressed. Continuous chance nodes have potential to better reflect the changes between disease states by allowing for more fluid transitions. A 'playground' for development, consisting of just the sepsis and pneumonia networks was also created, and is again used in the present project. This playground allows a smaller subset of the network to come into focus, while also improving compilation and propagation speeds - less computing power required for a smaller network.



Figure 7: Causal diagram for sepsis, showing links between inflammatory mediators, physiological processes and observable symptoms and values

3 Literature Search

In the current TREAT sepsis network, viral infections are treated almost identically to non-septic patients, other than showing a mildly elevated temperature, see Figure 8. It is desirable to be able to distinguish between patients with bacterial and viral infections as only bacterial infections should be treated with antibiotics, however it is also important to be able to identify patients with a viral infection, which it is very difficult to do with the current setup. One part of the literature search focused on identifying ways in which viral and bacterial infections can be differentiated.



Figure 8: Characteristics of viral sepsis in the current TREAT network

Another issue identified where the literature would be helpful was the differentiation between the states of sepsis. The current sepsis network makes only very slight differentiations (discussed further in Section 4.2), and if this process can be improved, patients can be better stratified by risk, and receive treatment and tests accordingly.

3.1 Differentiation between bacterial and viral infections

One desirable feature of TREAT is the ability to differentiate between bacterial and viral infections. The literature search carried out here identified several studies in which inflammatory mediators/common laboratory measurements were compared for patients with confirmed bacterial, viral and/or mixed infections. The studies chosen focus specifically on LRT infections, which is reasonable given that most LRT infections are viral [26]. The measurements compare leukocyte count [27, 28, 29], procalcitonin (PCT) [26, 28], C-reactive protein (CRP) [28, 29] and creatinine [29].

Comparisons between pure viral, and pneumococcal infections are shown in Table 1. The results of the studies listed suggest that there are differences in leukocyte count as well as CRP and PCT levels in patients with viral and bacterial infections, although the differences in leukocyte count are not statistically significant.

Variable Pure Viral		Pneumococcal		Reference	
	n	Range	n	Range	
Leukocyte count	26	$13100\pm 6600 \text{ cells}/\mu \text{L}$	44	$16000\pm 6600 \text{ cells}/\mu \text{L}$	[29]
	34	$14100 \ [4350-36500]$	31	$19700 \ [8420-44380]$	[28]
CRP	26	$16{\pm}12~{ m mg/dL}$	44	$20{\pm}16~{\rm mg/dL}$	[29]
	34	37.35 [10.03-229.74]	31	$268 \ [9.62-575.8]$	[28]
Creatinine	26	$1.1{\pm}0.6~{\rm mg/dL}$	44	$1.1 {\pm} 0.5$	[29]
PCT	34	$0.854 \ [0.128-6.08]$	31	$9.42 \ [0.078-63.32]$	[28]

 Table 1: Comparison of septic marker levels for patients with pure viral and pneumococcal infections

Wright [27] also states that although at one point in time viral infections were associated with relative leukopenia, several studies contradict this idea, indicating that both etiologies result in leukocytosis and leukocyte count can not be used to differentiate them. At the very least, leukocyte count can be used to confirm some kind of infection, which it does not do in the current sepsis network, where for all factors other than temperature, viral sepsis and no sepsis are identical.

CRP and PCT appear to be much more useful in differentiating between viral and bacterial infections, with marked differences between the two. Although there is some overlap on the lower end, higher values of both are very strongly suggestive of a bacterial infection. PCT and CRP have also been evaluated as markers for differentiating between infectious and non-infectious SIRS, and the stages thereof, which is discussed further in Section 3.2.

3.2 Differentiation between sepsis and non-infectious SIRS, and the states of each

A problem often noted on the subject of sepsis diagnosis is being unable to distinguish between sepsis and non-infectious SIRS. As mentioned in Section 2.1, many of the markers used for sepsis are not specific to infectious diseases and are commonly markers of inflammation. However, several markers do show some differences between sepsis and non-infectious SIRS, and importantly, between the different stages of sepsis. Selected studies compare PCT, interleukin-6 (IL-6), CRP, leukocyte count, lactate, blood pressure, pulse rate, respiratory rate, oxygen saturation, body temperature and erythrocyte sedimentation rate (ESR) for non-infected and infected patients, with the further differentiation made between sepsis stages for PCT, CRP and lactate.

IL-6 is not measured commonly, nor is it included in the TREAT network, so it is excluded from this analysis. It should be noted, however, that IL-6 levels were significantly different between infected and non-infected patients as well as between non-infectious SIRS, sepsis, severe sepsis and septic shock (with the latter two compared to sepsis) [30].

Tables 2 and 3 show a collection of results from studies measuring lactate, CRP and PCT for patients diagnosed with different stages of sepsis, as well as non-infectious SIRS. Although lactate and PCT are not currently included in the sepsis network, nor have they been added during the course of this project, they are expected to be added in the near future. Lactate is routinely measured and PCT is becoming more common, although it is still a comparatively expensive test.

Diagnosis	n	Lactate (mmol/L)	Reference
Non-infectious SIRS	15	2.13(1.14-2.93)	[31]
Sepsis	34	1.37(1-2.61)	[31]
Severe sepsis	22	2.19(1.73-2.93)	[31]
Septic shock	15	3.7(2.6-6.4)	[31]

Table 2: Lactate concentrations for non-infectious SIRS and sepsis. All values are median (IQR)

Lactate levels appear to correlate well with disease state, with increases seen from sepsis to severe sepsis, and a further increase seen in septic shock patients. However, non-infectious SIRS patients appear to present with similar lactate levels to severe sepsis patients. This suggests that lactate is a useful marker of severity, but is not useful for distinguishing between infected and non-infected patients. The PCT and CRP values in Table 3 show broad differences within each categorization, that is, SIRS, sepsis etc, however, within each study there is a clear trend of increased values for both markers from non-infectious SIRS to sepsis, and then on with the severity of illness. This is discussed further in Section 4.3, where the distributions are designed to be implemented in continuous nodes in the new sepsis network.

Diagnosis	n	CRP (mg/L)	PCT (ng/mL)	Reference
Non-infectious	17	$113.4{\pm}23.9$	0.41 ± 3.04	a
SIRS	20 70 ± 48 3.8 ± 6.9		$3.8{\pm}6.9$	b
	18	119 ± 89	$0.6 \ (0-5.3)$	c
	211	79.9(52.9-103.4)	0.4 (0.2 - 0.7)	d
	11	N/A	3.0(0.7-29.5)	f
	29	N/A	5.45 ± 1.45	g
	89	15(3.3-54)	$0.07 \ (0.03-0.17)$	[30]
	15	51(19.5-80.5)	0.38(0.16-0.93)	[31]
Sepsis	61	104.8 ± 30.5	$0.53{\pm}2.89$	a
	24	$98{\pm}76$	$1.3{\pm}2.7$	b
	14	159 ± 51	3.5(0.4-6.7)	c
	132	125.6(79.4-174.6)	3.1(1.4-5.2)	d
	11	N/A	$7.29 {\pm} 3.16$	g
	202	82 (24-150)	0.19(0.06-0.95)	[30]
	34	150 (71-209)	1.58(0.41-3)	[31]
Severe sepsis	68	$138.1 {\pm} 25.5$	$6.91 {\pm} 3.87$	a
	27	145 ± 70	$9.1{\pm}18.2$	b
	21	$254{\pm}181$	6.2(2.2-85)	c
	25	73.6(60.9-148.9)	3.2(1.7-7.4)	d
	14	N/A	$19.1 \ (2.8-351.2)$	f
	17	N/A	$6.26 {\pm} 3.29$	g
	28	100 (67-152)	1.3(0.43-6.3)	[30]
	22	159(75-209)	5.58(1.84-32.93)	[31]
Septic Shock	39	$177.9 {\pm} 35.5$	$12.89 {\pm} 4.39$	a
	24	173 ± 98	38.5 ± 59.1	b
	16	$264{\pm}170$	89 ± 154	c
	25	228 ± 119	21.3(1.2-654)	d
	43	$108.0 \ (62.9-167.5)$	10.7(2.9-33.2)	e
	10	N/A	38.76 ± 14.75	g
	17	94 (45-188)	1.3(0.27-14.4)	[30]
	15	195(75-272)	$13.1 \ (6.1-42.2)$	[31]
a,b,c,d,e,f,g - different studies collected in review [32]				
				-

Table 3: Plasma CRP and PCT levels in non-infectious and infectious SIRS. Values given as mean±SD or median (IQR)

4 Sepsis Network Modifications

4.1 Strategy

In a previous semester project (described in Section 2.5), tasks were undertaken to convert some of the nodes in the existing TREAT sepsis network to continuous chance nodes to better represent the clinical variables they described and to provide smoother transitions between diagnoses. Suggestions were also made on the structure of the sepsis network, if it were to be redesigned based on sepsis pathophysiology. In this project, a hybrid approach is taken, where the initial approach is to replace discrete chance nodes with continuous chance nodes, then subsequently alter the states of the given nodes and make minor structural changes to create a more physically relevant picture. The current sepsis network is shown in Figure 9.



Figure 9: Former sepsis network in Hugin

Early in the project, it was decided that the links between the states of the leaf nodes and those of the sepsis node needed to be stronger. However, the problem exists that most of the clinical variables measured are not specific to sepsis; some are characteristic of SIRS, whether infection is involved or not, and some are purely indicative of organ failure. The fact that there is no specific marker of sepsis is why it is such a difficult condition to diagnose, and why decision support systems such as TREAT are so useful. A work-around for this problem, was to create a leak node to explain away some of the symptoms as being caused by non-infectious SIRS, or by other causes. A relative prevalence is set for this.

4.2 Modifications Made - Structural Changes

The modifications made to the network include additions of nodes, redesigns of parts of the structure and changes within the nodes themselves. In particular, the way in which the states in the factor nodes behave was completely redesigned, leak nodes were added to account for non-septic causes of variations in the sepsis network variables, and the ARDS network was redesigned. Nodes were also added for heart rate and age-dependence of blood pressure.

Addition of New States in Factor Nodes

The structure, in terms of the symptoms grouped beneath factor nodes, was retained, however, the states present in each factor were modified. Previously, each of the factors 'Fact_fever', 'Fact_albESR', 'Fact_leukocreat' consisted of three states: 'no', 'good', and 'bad', and the shock factor had the states 'yes' and 'no', as shown in Figure 10. In most cases, good and bad were almost identical, and degrees of sepsis were influenced by small shifts between their relative distributions. To allow for a more intuitive development of the sepsis network, these 'no', 'good' and 'bad' states were replaced with states that are more directly related to the physical situation to which the factor pertains. 'Fact_fever', for example, was given the states: 'no', 'mild', 'mod-sev' and 'hypothermia'. These states of fever are then related to the sepsis node.



Figure 10: States in the factor nodes and sepsis node in the current network

A similar practice is adopted for the other three factor nodes, with each node given a relevant set of state names. For example, the children of the shock factor are not involved in mild infections, so no mild state is included. The states of the factors are shown in Figure 11.

This introduction of extra states allows a more direct mapping from the states in the factor nodes to the sepsis diagnosis node. This is an ideal scenario, because not only does it simplify the conditional probability table, but it also gives the opportunity to give physically relevant distributions



Figure 11: States in the factor nodes and sepsis diagnosis node in the redesigned sepsis network

to the measurements. Figures 12 and 13 show the probability distributions for the fever factor and temperature nodes for given states of sepsis. The state names in the sepsis node have also been adjusted to match the common nomenclature of the literature, and map directly to the states of the previous sepsis node as shown in Table 4.

Old name	New name
no	no
moderate	sepsis
severe	severe sepsis
critical	septic shock
viral	mild*
	*includes cystitis

Table 4: Changes in state names in the sepsis diagnosis node

Expanding the number of states in the factor nodes also means making modifications to each of their respective child nodes. In several cases, this made it easier to set up continuous distributions, as distributions could be made to match a disease state as described in literature or the physical description now given in the factor node. In other cases where nodes were retained from the original sepsis network, distributions were adjusted to roughly match those given by instantiating the sepsis node in the old network. New nodes were introduced to perform a mapping task for leukocytes and ESR, as both require a very low count and a very high count to register as septic shock/critical sepsis evidence.

Figure 12 shows the lack of distinction between temperature distributions for the given states of sepsis. The only difference in distribution in seen for the 'no sepsis' case. One explanation for this is that disease state does not correlate well with the degree of fever. However, it is much more likely that a patient with a fever above 40 degrees is seriously ill when compared to a patient with a normal temperature. Currently every patient in the clinically defined fever range - temperature above 38.5 degrees celsius - is lumped together in one bin. It is also apparent that this version of

the temperature node does not take hypothermia into account. Septic patients can be considered to be at the severe end of the spectrum when they have both abnormally high and abnormally low temperatures. The new states created within the fever factor allow for these distributions.



Figure 12: Mapping of sepsis states to factor nodes and through to the leaf nodes in the current network



Figure 13: Mapping of sepsis states to factor nodes and through to the leaf nodes in the redesigned network

Other structural modifications include the split of ESR and CRP (Figure 14) so that they are not treated identically (Figure 15), addition of a heart rate node, and redesign of the ARDS network.



Figure 14: Previous structure of the nodes beneath the albumin-ESR factor. The node labeled ESR is an OR node that maps the three states in the crp node to those in the esr node.



Figure 15: New structure of the nodes beneath the albumin-ESR factor.

Leak Nodes

The initial strategy for implementing leak nodes was to have what was essentially a second factor node, linking one-to-one with the factor, and then to have a prevalence node as a parent of the second factor node. This allowed a relative prevalence to be set between non-infectious and infectious SIRS, while setting a distribution for the factor node states for infectious SIRS. In this case, the leak from each factor node could be adjusted independently. This first strategy is shown as a cut-out from the CPN in Figure 16.



Figure 16: Leak nodes with independent adjustment for each factor

Two further approaches, shown in Figure 17 and Figure 18 respectively, are combining the prevalence nodes into a shared 'other-SIRS' node, or removing the prevalence node entirely and integrating its effect into the secondary factor nodes in the a priori probabilities.

The chosen solution is the second method where there is a single leak node. This combines the reasoning from each factor and allows the node to act as an alternative diagnosis, with reasoning from the factors collected in much the same way as it is for the sepsis node.

Figure 17: Leak nodes with common 'other_SIRS' node

Figure 18: Leak nodes with integrated prevalences for non-infectious SIRS

ARDS Network Redesign

The ARDS network was also identified as an area in which improvements could be made. The effect of oxygen supplementation on the fraction of inspired oxygen was included, along with the introduction of continuous nodes for the respiratory rate, RR (previously 'tachypnea'), as well as oxygen saturation, SaO_2 (previously included in 'hypoxia'). A gradation is also introduced within the ARDS node to allow for the indentification of ALI. ARDS and ALI are defined by the following:

The previous ARDS network, shown in Figure 19 was also characterized by a small number of states. Tachypnea and hypoxia were effectively boolean nodes, with the split made 20 breaths per minute for the respiratory rate (≥ 20 is tachypnea), and at 60 mmHg for the partial pressure of oxygen (PaO₂ ≤ 60 mmHg gives hypoxia).

	Timing	Oxygenation	Chest Radio- graph	Pulmonary Artery Wedge Pressure
ALI cri- teria	acute onset	$Pa_{O_2}/Fi_{O_2} \leq 300mmHg$ (regardless of PEEP* level)	Bilateral infil- trates seen on frontal chest radiograph	$\leq 18mmHg$ when mea- sured or no clinical evi- dence of lef arterial hy- pertension
ARDS criteria	acute onset	$Pa_{O_2}/Fi_{O_2} \leq 200mmHg$ (regardless of PEEP* level)	Bilateral infil- trates seen on frontal chest radiograph	$\leq 18mmHg$ when mea- sured or no clinical evi- dence of lef arterial hy- pertension

Table 5: Recommended criteria for ALI and ARDS [33]

* PEEP: Positive End Expiratory Pressure

In the redesigned network, Figure 20, the P/F ratio and FiO₂ nodes are both split into eight states, and the PaO₂ node contains fifteen states to cover the combinations of its parents. Consequently, the oxygen saturation node, SaO₂ contains fifteen continuous distributions corresponding to the oxygen saturation states. Hypoxia is still essentially defined as an OR node, split at 60 mmHg. The supplementary oxygen contains states for flow rates of 0-10 L/min, in steps of 1 L/min.

Figure 19: ARDS network from the original sepsis network

Figure 20: New design for the ARDS section of the sepsis network

Addition of HR node and age-dependent BP

Other changes beneath the shock factor include the introduction of a node representing heart rate (HR). HR was a notable omission in the current network, with tachycardia (elevated HR) being one of the symptoms of SIRS. The age-dependence of blood pressure is also an important consideration; the general trend of blood pressure increasing with age means that a hard limit for shock is not likely to be appropriate for all patients. Previously set at 90 mmHg, the threshold for the blood pressure node may not have been reached by those patients with a normal blood pressure substantially above 120 mmHg, something that is common among older patients.

Figure 21 (from [34]) shows the relationship between age and blood pressure. In the sepsis network, three age ranges will be used; less than 40 years, 40-60 years and over 60 years. This will enable the capture of age-related blood pressure changes without excessive discretization. Distributions to be used in the network are described in Section 4.3

Figure 21: Variation in blood pressure with age

New Structure - Complete Sepsis Network

The complete structure of the modified sepsis network is shown in Figure 22. The nodes are coloured according to the kind of that variable they represent. The green node is the diagnosis. The observations are yellow and orange; orange if they are measured, yellow if they represent a clinical observation/opinion. The dark blue nodes are the factors, and light blue are nodes used to perform mapping. The red node is the leak node, effectively an alternative diagnosis, that is, a non-septic cause of the observed variables.

Figure 22: Redesigned sepsis network

4.3 Modifications Made - Continuous Distributions

The following nodes have been designed with continuous distributions:

- Temperature
- Serum Albumin
- Erythrocyte Sedimentation Rate (ESR)
- C-Reactive Protein (CRP)
- Leukocyte count
- Creatinine
- Systolic Blood Pressure
- Heart Rate
- Respiratory Rate
- Oxygen Saturation (SaO₂)

This section presents the distributions implemented (or set for future implementation) in the redesigned sepsis network, along with the rationale used in designing the distributions. Values were taken from literature where possible (referenced where applicable), otherwise distributions where designed to reflect physiological ranges and trends observed or noted in the literature.

Temperature

For each of the four states of the fever factor, the temperature node requires a continuous distribution. By defining the factor states with physically relevant terms, this process is easier. Distributions can then be designed after which they can be mapped to the appropriate states of sepsis. The 'The high and low temperatures are taken care of by the 'severe' and 'hypothermia' states, and normal (healthy) and mildly elevated temperatures are covered by 'no' and 'mild' respectively. The chosen distributions are shown graphically in Figure 23.

Figure 23: Continuous distributions within the temperature node in the new sepsis network

The normal values are taken from a large population study on normal body temperature [35], widened slightly to allow for variations in the measurement procedure. The mild distribution covers the mild viral infections which do not cause large temperature increases, while the severe and hypothermia distributions both map across the sepsis states. Disease severity is not correlated well with degree of fever, hence the single distribution to cover each of the high and low temperature regions.

Albumin

Little information was available in the literature regarding albumin level with respect to the various stages of sepsis, however information was available on the normal serum levels of albumin. It has also been reported in the literature that serum albumin level decreases as the severity of sepsis increases [refs]. Normal serum albumin is 41 ± 3 g/L [36], with hypoalbuminemia defined as levels less than 34 g/L, with severe cases falling below 30 g/L [37]. Distributions, as shown in Figure 24, were created to cover this range of serum albumin concentrations while ensuring that the distributions do not overlap excessively to a point where little distinction is made.

Figure 24: Continuous distributions within the albumin node in the new sepsis network

Erythrocyte Sedimentation Rate (ESR)

ESR in septic patients rises rapidly above normal levels, however in critically ill septic patients with DIC, it is possible to have an abnormally low ESR (i1mm/hr). There are also other causes of abnormally low ESR. ESR is dependent on gender and there are several factors that can cause artificially high ESR. Normal ESR is somewhere in the region 2-25 mm/hr [38, 39], with ranges for for sepsis and severe sepsis above this. The chosen distributions are shown in Figure 25.

C-Reactive Protein

CRP levels respond rapidly to infection [40] and have been studied extensively as a marker of disease state in sepsis. CRP correlates well with disease severity [30, 31, 32] and has often been used interchangeably with ESR [40]. CRP is also generally an indicator of bacterial infections, although there are some viruses that will cause elevated CRP levels [41].

Due to the skew noted in the distributions listed in Section 3.2, a log scale was chosen for the CRP distributions. The top panel of Figure 26 shows the literature values (on a linear scale) with the

Figure 25: Continuous distributions within the ESR node in the new sepsis network

horizontal lines showing those studies that gave a median and range, and normal distribution curves showing those that listed a mean and standard deviation. Curves are grouped into respective studies by line-style and into sepsis stages by colour. The bottom panel shows the log distributions chosen, with more weight given to the larger studies.

Figure 26: Continuous distributions within the CRP node in the new sepsis network

Leukocytes

The normal range for leukocytes is between 4000 and 10000 cells/mm³ with these limits also commonly used as cut-off values for leukocytosis and leukopenia in the SIRS criteria [42, 43]. In the literature, raised leukocyte levels are reported as high as ¿40000 cells/mm³. The degree of leukocytosis is correlated with the degree of the inflammatory response, which is usually proportional to the severity of the illness. Generally, leukopenia is associated with worse outcomes [43], as the compensatory anti-inflammatory response is stronger than the inflammatory response, weakening the hosts defenses against threats. Distributions for the leukocyte count are shown in Figure 27. The mild distribution covers viral infections as well as some bacterial infections, although most will fall into the mod-sev range. Septic shock is covered by both very high and very low leukocyte counts, which are mapped together onto the septic shock state.

Figure 27: Continuous distributions within the leukocyte node in the new sepsis network

Creatinine

Creatinine is produced at a fairly constant rate in muscle tissue and is cleared through the kidneys. Usually raised creatinine levels are a sign of kidney injury as they correlate with a reduction in the glomerular filtration rate (GFR) which shows kidney function or lack thereof. However, in sepsis, the production of creatinine is reduced and so septic patients may appear with mildly elevated or even normal creatinine levels [44]. Although levels may be normal, raised serum creatinine is a strong indicator of renal failure, and in this case will be related more to the critical stage of sepsis [45]. The normal levels in this case are derived from a combination of literature values [46], with mild levels taken as a slight right-shift from this, and mod-sev and septic shock levels taken from thresholds given in the current version of the TREAT sepsis network. The distributions used for the serum creatinine levels are shown in Figure 28.

Figure 28: Continuous distributions within the creatinine node in the new sepsis network

Systolic Blood Pressure

Blood pressure is dependent on age, as discussed in Section 4.2, and also on disease state. Blood pressure is reduced in septic patients by fluid loss due to increased vascular permeability, as well as the vasodilatory action of many inflammatory mediators. Shock is commonly defined as systolic blood pressure less than 90 mmHg, however, for patients who have a higher normal blood pressure, they may be in shock at levels above this. The distributions shown in Figure 29 contain the normal range, roughly corresponding to the values given in the literature [34], with septic levels staggered below this. Note the increase across the age ranges.

Heart Rate

The distributions for heart rate, shown in Figure 30, are based on opinion about heart rate elevation with sepsis severity. Normal heart rate in adults ranges between 60 and 100 beats per minute, with elite athletes slightly lower at 40-60 beats per minute [47]. The SIRS criteria define tachycardia

Figure 29: Continuous distributions within the blood pressure node in the new sepsis network

as heart rate greater than 90 beats per minute [9]. Distributions for severe sepsis and septic shock patients are included above this range.

Figure 30: Continuous distributions within the heart rate node in the new sepsis network

Respiratory Rate

In the current sepsis network, respiratory rate is addressed by a tachypnea node that has the option of either being greater than or less than 20 breaths/min. In the redesigned network, the parent of the respiratory rate node is the P/F ratio, which is used in determining the degree of lung injury in the patient, that is, P/F less than 300 is associated with ALI and P/F less than 200 is associated

with ARDS. As the P/F ratio decreases, the patient's use of oxygen is further limited, and so the respiratory rate will increase in an attempt to deliver sufficient oxygen. The distributions shown in Figure 31 have been designed to steadily increase with reducing P/F ratio, with the distribution for the normal respiratory rate not exceeding 20 breaths/min.

Figure 31: Continuous distributions within the respiratory rate node in the new sepsis network

Oxygen Saturation (SaO_2)

The oxygen saturation distributions are based on the oxygen saturation curve, where a particular partial pressure of oxygen (parent node) has a corresponding SaO_2 . The width of the distributions are adjusted to account for the relative slope of the curve. The curve is shown in Figure 32 (taken from [48]) and the distributions in Figure 33.

Figure 32: Oxygen saturation curve

Figure 33: Continuous distributions within the oxygen saturation node in the new sepsis network

Procalcitonin (PCT)

The PCT distributions shown in the lower panel of Figure 34 have been developed in a similar manner to the CRP distributions in Figure 26. Again, a log-normal distribution has been chosen to account for the skew shown by the data. The top panel shows the distributions reported in the literature, as listed in Table 3. Distributions shown as lines were those reported as median and range, and the curves show those reported as a mean and standard deviation. Line-style groups together individual studies, and the colors represent the sepsis stages.

Figure 34: Continuous distributions within the PCT node in the new sepsis network

Lactate

The distributions for the lactate levels are based off the literature values reported in Section 3.2 and are shown in Figure 35. Sepsis and severe sepsis produce only mildly elevated lactate levels, while septic shock is associated with greater increases.

Figure 35: Continuous distributions within the lactate node in the new sepsis network

5 EM Learning for Optimisation

The Hugin Researcher software package provides a useful tool for learning in causal probabilistic networks. Given a database containing a set of cases, the probability tables/continuous distributions of nodes within a CPN can be learned. This type of learning is known as batch-learning, and requires that all the data be available at the time of learning.

A case is defined as a set of values corresponding to some or all of the nodes in a given CPN, that is, the case can be complete or incomplete. Before learning can take place, experience tables are added to nodes for which learning is desired, and the database containing the set of cases is specified. If experience tables are left empty (containing all zeros), no restriction is based on the the form of the resulting tables, whereas if experience tables are modified, these values are used as 'prior counts' which are used to augment those found from the data set.

Hugin then uses the estimation-maximisation algorithm developed by Lauritzen in 1995 [49] to determine the optimal tables for each of the specified nodes. The process is iterative, with the log-likelihood computed after each iteration. Hugin attempts to maximize this value, that it, minimize the error, and will stop iterating either when the maximum number of iterations has been reached, or when the specified tolerance has been achieved. The tolerance is reached when the relative difference between the log-likelihood values falls below that specified. Providing the database is a representative sample of the population for which the CPN will be used, the new, learned values should give much greater accuracy when diagnosing patients.

5.1 Hugin EM Learning Tutorial

The use of and robustness of the EM learning process used in Hugin is demonstrated well by their EM learning tutorial. The situation involves a simple CPN with nine nodes, and a database of 10,000 cases. Hugin also includes a case generator which can be used to generate such a database which will contain a specified percentage of missing values. Conditional probability tables may be randomized and then learning performed with a given database.

The CPN used for this tutorial is the Chest Clinic CPN, pictured in Figure 36. This basic network contains three diseases; tuberculosis, lung cancer and bronchitis, two risk factors; visit to Asia and

whether the patient is a smoker, and two observations; positive x-ray and dyspnea.

Figure 36: Chest Clinic CPN, used for the EM learning tutorial in Hugin

The original probability distributions for the Chest Clinic example are shown in Figure 37.

Figure 37: Chest Clinic CPN, showing original probability distributions

Following a complete randomization of the probability tables, the EM learning process was run using a database of 10,000 cases, with 5% missing values, that had been generated from the original system using the built-in case generator function. Figure 38 shows the results of the learning process, the resulting probability distributions agree to within approximately 2% with the original.

This example gives an introduction to the concept of EM learning in the Hugin environment and

Figure 38: Chest Clinic CPN, showing learned probability distributions

allows the user to get a feel for how well the learning process works. Obviously, the process works very well for a large database defined across a small network with relatively few missing values. However, this is not likely to be the case when learning is attempted on the sepsis network. Medical databases are typically smaller, and are likely to have a greater number of missing values. Although the size of the database and the number of missing values detract, there are less nodes that need to be learned in the sepsis network, and it is expected that the initial guesses are much better than starting with random tables.

5.2 Patient Database used for Learning

The database supplied for an initial learning attempt was one from Hvidovre Hospital in Copenhagen. Data was collected while TREAT was in use at the location.

Data is recorded anonymously and grouped by patient. Patient data can then be split into episodes, each of which covers a hospital admission. Each episode consists of a set of encounters, which account for each time something new is entered into TREAT, for example, blood culture results. When analyzing the performance of TREAT, we are interested more in the first encounter when the patient is admitted to the hospital. However, for the purpose of building and optimizing TREAT, in this case the sepsis network within TREAT, it is desirable to have as much information at hand as possible. The problem lies in the fact that by the time some of the information is available, such as laboratory tests and blood cultures, the patient will have received preliminary treatment, and thus will not be the same as the patient who presented at the hospital.

Figure 39 shows the structure of the TREAT database, while Figure 40 shows the same diagram with important tables highlighted. These tables are the ones important for learning in the sepsis and pneumonia networks. The remaining tables can be removed as they contain information on other infection sites and decision related information, both sets of which are unnecessary in this case. Figure 41 shows the view created in SQL Server Manager for easy access to the relevant data.

5.3 Data Preparation and Mapping

In order to utilize the data available in the database, it first needs to be extracted and converted into an appropriate form for use with Hugin. When TREAT is in use, data exists in three different forms, one for the database, one for use in the CPN, and one for the user interface. Hugin expects the data in a specific format for use with the EM learning process, so the data must first be extracted from the database, irrelevant columns removed, and database entries mapped to the correct nodes and states of nodes in the network. Most of the cases do not map 1-1 and in certain places, information from several database columns is required to find the appropriate state for a node. A console application was written in C# to perform the mapping required, extracting only data to be used in the sepsis and pneumonia networks present in the 'playground' to be used for the learning process.

Figure 39: Database diagram for TREAT showing most tables

Figure 41: View created in SQL server manager so that unnecessary data is stripped away

6 Discussion

Although the changes made are thought to improve the performance of the sepsis network, the cases given are only examples and are not real patients. In order to give a more concrete statement on the effectiveness of the sepsis network redesign, a benchmark comparison between the two versions would need to be completed on a representative patient cohort. However, before this happens, further tailoring needs to be completed in terms of carrying out the learning process within the network, and making improvements to the pneumonia network. It is believed that the network is, at least, more physiologically relevant, with the new states added to nodes allow for increased transparency within the system. This transparency is particularly important for the developer, as it increases the understanding gained from the system, and makes further changes much easier.

The additional markers, PCT and lactate, discussed in this report have not been included, but should be in future versions of the sepsis network. The first step towards making this possible is to find out where they fit within the current structure, that is, which factor they should be added to, or whether they should be separated from the current factors. PCT in particular, although a less common measurement, has been shown to be particularly useful in sepsis diagnosis as well as differentiating between the states of the disease. Although lactate is less effective, it is a commonly measured marker, and should be included as long as it has some clinical utility. The learning process requires a much larger database than the Hvidovre database described in this report. A larger, more complete database will give much better results when learning is attempted. The data extraction methods also need to be improved, with recognition of certain features, such as encounters which are purely confirmations of information previously entered. It is important to avoid doubling-up on information as this has the potential to skew the learning. Actual learning from the database has not been attempted successfully in this case, however it is a necessary next step in the development of the sepsis network.

Only following the completion of learning will a reliable benchmark comparison be able to be made. Before this time, any comparisons between the current and redesigned sepsis networks must be made with caution. Although it may be sensible to compare some relative functions of the networks, for example, their behaviors in certain scenarios, it is unreasonable to make any quantitative comparisons as the tables in the network are still subject to change.

The pneumonia network, which has had minimal discussion in this report, is another area identified for improvement. Currently each pathogen is treated identically in terms of the severity of illness it causes, whereas this is not accurate. If pathogens can be at the very least grouped by the severity of the infection they cause, TREAT's ability to infer the causative bacteria and identify appropriate antibiotic treatment will be improved for patients with respiratory infections.

7 Conclusion

The TREAT sepsis network has been modified to give a more physiologically accurate representation of sepsis. Altering the representation of states within nodes has given increased transparency which allows for better understanding of the model and makes future improvements easier. Continuous nodes have also been implemented, using values from the literature where possible. Additionally, new nodes have been added to recognize septic characteristics not previously accounted for.

The EM learning process within Hugin was investigated, and seems to be a potentially powerful tool for refining existing CPNs. It is, however, necessary that the database to be used for learning is both sufficiently large and complete, as the success of the EM learning its very dependent on these factors.

References

- DC Angus, WT Linde-Zwirble, J Lidicker, G Clermont, J Carcillo, and MR Pinsky, "Epidemiology of severe sepsis in the united states: analysis of incidence, outcome, and associated costs of care," *Crit Care Med*, vol. 29, pages = 13031310, 2001.
- [2] GS Martin, DM Mannino, S Eaton, and M Moss, "The epidemiology of sepsis in the united states from 1979 through 2000," N Engl J Med, vol. 348, p. 15461554, 2003.
- [3] Richard S. Hotchkiss and Irene E. Karl, "The pathophysiology and treatment of sepsis," N Engl J Med, vol. 348, no. 2, pp. 138–150, 2003.
- [4] Shawn D. Carrigan, George Scott, and Maryam Tabrizian, "Toward resolving the challenges of sepsis diagnosis," *Clinical Chemistry*, vol. 50, no. 8, pp. 1301–1314, 2004.
- [5] Charalampos Pierrakos and Jean-Louis Vincent, "Sepsis biomarkers: a review," Critical Care, vol. 14, no. R15, 2010.
- [6] Niels C Riedermann, Ren-Feng Guo, and Peter A Ward, "The enigma of sepsis," J Clin Invest, vol. 112, pp. 460–467, 2003.
- Jean-Louis Vincent and Edward Abraham, "The last 100 years of sepsis," Am J Respir Crit Care Med, vol. 153, pp. 256-263, 2005.
- [8] Roger C Bone, "The pathogenesis of sepsis," Annals Int Med, vol. 115, pp. 457–469, 1991.
- [9] "American college of chest physicians/society of critical care medicine consensus conference: Definitions for sepsis and organ failure and guidelines for the use of innovative therapies in sepsis.," Crit Care Med, vol. 20, no. 6, pp. 864–874, 1992.
- [10] MM Levy, MP Fink, JC Marshall, E Abraham, D Angus, D Cook, J Cohen, SM Opal, JL Vincent, and G Ramsay, "2001 sccm/esicm/accp/ats/sis international sepsis definitions conference," *Crit Care Med*, vol. 31, pp. 1250–1256, 2003.
- [11] R. Phillip Dellinger, Mitchell M. Levy, Jean M Carlet, Julian Bion, Margaret M. Parker, Roman Jaeschke, Konrad Reinhart, Derek C. Angus, Christian Brun-buisson, Richard Beale, Thierry Calandra, Jean-Francois Dhainaut, Herwig Gerlach, Maurene Harvey, John J. Marini, John Marshall, Marco Ranieri, Graham Ramsay, Jonathan Sevransky, B. Taylor Thompson, Sean Townsend, Jeffrey S. Vender, Janice L. Zimmerman, and Jean-Louis Vincent, "Surviving sepsis campaign: International guidelines for management of severe sepsis and septic shock: 2008," Intensive Care Med, vol. 34, pp. 17–60, 2008.
- [12] A X Garg, N K J Adhikari, H McDonald, M P Rosas-Arellano, P J Devereaux, J Beyene, J Sam, and R B Haynes, "Effects of computerized clinical decision support systems on practitioner performance and patient outcomes," JAMA, vol. 293, no. 10, pp. 1223–1238, 2005.
- [13] Judea Pearl, Probabilistic reasoning in intelligent systems: networks of plausible inference. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 1988.
- [14] S L Lauritzen and D J Speigelhalter, "Local computations with probabilities on graphical structures and their application to expert systems," J R Statist Soc B, vol. 50, no. 2, pp. 157–224, 1988.
- [15] HUGIN EXPERT A/S, "Hugin expert about www.hugin.com/company-information/about," 2012.
- [16] L Leibovici, I Shraga, M Drucker, H Konigsberger, Z Samra, and S D Pitlik, "The benefit of appropriate empirical antibiotic treatment in patients with bloodstream infection," *Journal of Internal Medicine*, vol. 244, pp. 379–386, 1998.
- [17] J Garnacho-Montero, J L Garcia-Garmendia, A Barrero-Almodovar, F J Jimenez-Jimenez, C Perez-Paredes, and C Ortiz-Leyba, "Impact of adequate empirical antibiotic therapy on the outcome of patients admitted to the intensive care unit with sepsis," *Crit Care Med*, vol. 31, no. 12, pp. 2742–2751, 2003.
- [18] Mical Paul, Roberto Cauda, and Leonard Leibovici, "Electronic prescribing," IanM Gould and JosWM Meer, eds.), pp. 41–67, Springer US, 2008.
- [19] B Kristensen, S Andreassen, L Leibovici, C Riekehr, A G Kjaer, and H C Schoenheyder, "Empirical treatment of bacteraemic urinary tract infection. evaluation of a decision support system," *Dan Med Bull*, vol. 46, no. 4, pp. 349–353, 1999.
- [20] S Andreassen, L Leibovici, M Paul, A D Nielsen, A Zalounina, L E Kristensen, K Falborg, B Kristensen, U Frank, and H C Schoenheyder, "A probabilistic network for fusion of data and knowledge in clinical microbiology," D Husmeier, R Dybowski, and S Roberts, eds.), pp. 451–72, Springer-Verlag, 2004.
- [21] Steen Andreassen, Christian Riekehr, Brian Kristensen, Henrik C Schonheyder, and Leonard Leibovici, "Using probabilistic and decision-theoretic methods in treatment and prognosis modeling," Artificial Intelligence in Medicine, vol. 15, pp. 121–134, 1999.

- [22] Leonard Leibovici, Michal Fishman, Henrik C Schonheyder, Christian Riekehr, Brian Kristensen, Ilana Shraga, and Steen Andreassen, "A causal probabilistic network for optimal treatment of bacterial infections," *IEEE Transactions on Knowledge and Data Engineering*, vol. 12, no. 4, pp. 517–528, 2000.
- [23] Mical Paul, Steen Andreassen, Evelina Tacconelli, Anders D Nielsen, Nadja Almanasreh, Uwe Frank, Roberto Cauda, and Leonard Leibovici, "Improving empirical antibiotic treatment using treat, a computerized decision support system: cluster randomized trial," *Journal of Antimicrobial Chemotherapy*, vol. 58, pp. 1238–1245, 2006.
- [24] Mical Paul, Steen Andreassen, Anders D Nielsen, Evelina Tacconelli, Nadja Almanasreh, Abigail Fraser, Dafna Yahav, Ron Ram, and Leonard Leibovici, "Prediction of bacteremia using treat, a computerized decision-support system," *Clinical Infectious Diseases*, vol. 42, p. 12741282, 2006.
- [25] Charles A Dinarello, "Infection, fever, and exogeous and endogenous pyrogens: some concepts have changed," J Endotoxin Res, vol. 10, no. 4, pp. 201–222, 2004.
- [26] K L Becker, R Snider, and E S Nylen, "Procalcitonin in sepsis and systemic inflammation: a harmful biomarker and therapeutic target," *British Journal of Pharmacology*, pp. 253–254, 2010.
- [27] Harry T. Wright, "Acute respiratory infections," Current Problems in Pediatrics, vol. 5, pp. 3–63, 1974.
- [28] C Prat, J Dominguez, C Rodrigo, M Gimenez, M Azuara, O Jimenez, N Gali, and V Ausina, "Procalcitonin, c-reactive protein and leukocyte count in children with lower respiratory tract infection," *Pediatr Infect Dis J*, pp. 963–967, 2003.
- [29] A de Roux, M A Marcos, E Garcia, J Mensa, S Ewig, H Lode, and A Torres, "Viral community-acquired pneumonia in nonimmunocompromised adults," *Chest*, pp. 1353–1351, 2004.
- [30] E L Tsalik, Jaggers L B, S W Glickman, R J Langley, J C van Velkinburgh, L P Park, Fowler V G, C B Cairns, S F Kingsmore, and C W Woods, "Discriminative value of inflammatory biomarkers for suspected sepsis," *The Journal of Emergency Medicine*, 2011. In Press, corrected proof.
- [31] Gian Paolo Castelli, Claudio Pognani, Michael Meisner, Antonio Stuani, Daniela Bellomi, and Laura Sgarbi, "Procalcitonin and c-reactive protein during systemic inflammatory response sydrime, sepsis and organ dysfunction," *Crit Care*, vol. 8, no. 4, pp. R234–R242, 2004.
- [32] Chieko Mitaka, "Clinical laboratory differentiation of infectious versus non-infectious systemic inflammatory response syndrome," *Clinica Chimica Acta*, vol. 351, pp. 17–29, 2005.
- [33] George R Bernard, Antonio Artigas, Kenneth L Brigham, Jean Carlet, Konrad Falke, Leonard Hudson, Maurice Lamy, Jean Roger Legall, Alan Morris, and Roger Spragg, "The american-european consensus conference on ards: Definitions, mechanisms, relevant outcomes, and clinical trial coordination," Am J Respir Crit Care Med, vol. 149, pp. 818–824, 1994.
- [34] W E Miall and H G Lovell, "Relation between change of blood pressure and age," Brit med J, vol. 2, pp. 660–664, 1967.
- [35] H Eriksson, K Svardsudd, B Larsson, L Welin, LO Ohlson, and L Wilhelmsen, "Body temperature in general population samples. the study of men born in 1913 and 1923.," Acta Med Scand, vol. 217, no. 4, pp. 347–52, 1985.
- [36] Maria-Chiara Corti, Jack M. Guralnik, Marcel E. Salive, and John D. Sorkin, "Serum albumin level and physical disability as predictors of mortality in older persons," JAMA, vol. 272, pp. 1036–1042, 1994.
- [37] F R Herrmann, C Safran, S E Levkoff, and K L Minaker, "Serum albumin level on admission as a predictor of death, length of stay, and readmission," Arch Intern Med, vol. 152, no. 1, pp. 125–130, 1992.
- [38] L E Bottiger and C A Svedberg, "Normal erythrocyte sedimentation rate and age," Brit Med J, vol. 2, pp. 85–87, 1967.
- [39] S E Bedell and B T Bush, "Erythocyte sedimentation rate from folklore to facts," The American Journal of Medicine, vol. 78, pp. 1001–1009, 1985.
- [40] A Osei-Bimpong, J. H. Meek, and S. M. Lewis, "Esr or crp? a comparison of their clinical utility," *Hematology*, vol. 12, no. 4, pp. 353–357, 2007.
- [41] P Povoa, "C-reactive protein: a valuable marker of sepsis," Intensive Care Med, vol. 28, pp. 235–243, 2002.
- [42] NG Orfanakis, RE Ostlund, CR Bishop, and Athens JW, "Normal blood leukocyte concentration values," Am J Clin Path, vol. 53, no. 5, pp. 647–651, 1970.
- [43] A C Muller Kolbold, J E Tulleken, J G Zijlstra, W Sluiter, J Hermans, C G M Kallenberg, and J W Cogen Tervaert
- [44] Kent Doi, Peter S.T. Yuen, Christoph Eisner, Xuzhen Hu, Asada Leelahavanichkul, Jurgen Schnermann, and Robert A. Star, "Reduced production of creatinine limits its use as marker of kidney injury in sepsis," J Am Soc Nephrol, vol. 20, pp. 1217–1221, 2009.
- [45] Nestor Schor, "Acute renal failure and the sepsis syndrome," Kidney Int, vol. 61, pp. 764–776, 2002.

- [46] Ferruccio Ceriotti, James C Boyd, Gerhard Klein, Joseph Henny, Josep Queralto, Veli Kairisto, and Mauro Panteghini, "Reference intervals for serum creatinine concentrations: Assessment of available data for global application," *Clinical Chemistry*, vol. 54, no. 3, p. 559566, 2008.
- [47] A.D.A.M editorial team, "Pulse: Medline plus medical encycolpedia http://www.nlm.nih.gov/medlineplus/ency/article/003399.htm," January 2012.
- [48] S. Silbernagl and A. Despopoulos, Color Atlas of Physiology. Thieme flexibook, Thieme, 2009.
- [49] S L Lauritzen, "The em algorithm for graphical association models with missing data," Computational Statistics & Data Analysis, vol. 19, pp. 191–201, 1995.