

Ontology-Driven Hypothesis Generation to Explain Anomalous Patient Responses to Treatment

Laura Moss, Derek Sleeman, Malcolm Sim, Malcolm Booth, Malcolm Daniel, Lyndsay Donaldson, Charlotte Gilhooly, Martin Hughes, John Kinsella

Abstract Within the medical domain there are clear expectations as to how a patient should respond to treatments administered. When these responses are not observed it can be challenging for clinicians to understand the anomalous responses. The work reported here describes a tool which can detect anomalous patient responses to treatment and further suggest hypotheses to explain the anomaly. In order to develop this tool, we have undertaken a study to determine how Intensive Care Unit (ICU) clinicians identify anomalous patient responses; we then asked further clinicians to provide potential explanations for such anomalies. The high level reasoning deployed by the clinicians has been captured and generalised to form the procedural component of the ontology-driven tool. An evaluation has shown that the tool successfully reproduced the clinician's hypotheses in the majority of cases. Finally, the paper concludes by describing planned extensions to this work.

1 Introduction

It is widely acknowledged that anomalous scenarios provide a key role in knowledge discovery; an anomaly can indicate to an expert that their understanding of a domain may require further refinement which in turn may lead to the discovery of new knowledge [6]. However, in a complex domain such as medicine, it can be difficult, especially for junior clinicians, to generate the required hypotheses to resolve such anomalies. The resolution of anomalies is important as it can enhance patient treatment.

Laura Moss, Derek Sleeman^{1,2}

¹Department of Computing Science, University of Aberdeen, Aberdeen, AB24 3UE

Malcolm Sim, Malcolm Booth, Malcolm Daniel, Lyndsay Donaldson, Charlotte Gilhooly, Martin Hughes, John Kinsella

²Section of Anaesthesia, Glasgow Royal Infirmary, University of Glasgow, Glasgow, G31 2ER

The Intensive Care Unit (ICU) in a hospital provides treatment to patients who are often critically ill and possibly rapidly deteriorating. Such patients provide complex challenges for the attending clinician. To aid the clinician in monitoring and treating such patients, many modern ICUs are equipped with sophisticated patient management systems. These systems collect the patients' physiological measurements; record the infusions given to the patient, and note some interventions such as dialysis. The information recorded by the system is stored in a database and can be viewed at the patient's bedside or 'offline' at a later date. Access to the data stored in the patient management system at Glasgow Royal Infirmary's ICU has been provided for use in this study.

The focus of the work reported here is the development of a tool to assist clinicians in explaining why a patient has responded anomalously to treatment. To capture the process deployed by expert clinicians in such a scenario, detailed interviews were held with five ICU consultants from Glasgow Royal Infirmary (GRI) to identify anomalous patient responses based on their physiological data. Two further ICU clinicians from GRI were then asked to suggest potential explanations for the identified anomalies. Further, these interviews were examined by analysts to establish both how the hypotheses were generated and the types of hypotheses presented. The findings from this analysis form the basis of an ontology-driven hypothesis generation tool described below. A more detailed review of the interviews held with ICU clinicians can be found in Moss et al [13].

The rest of this paper is organized as follows: section 2 gives a brief literature review; section 3 discusses the interviews held with ICU clinicians; section 4 outlines the ontology-based tool, gives examples of use, and outlines an evaluation; section 5 discusses planned future work for this tool

2 Literature Review

Anomalous scenarios play a key role in knowledge discovery; Kuhn [6] defines an anomaly as a violation of the "paradigm-induced expectations that govern normal science". Such anomalies are of interest as they often point to the inadequacy of a currently held theory and require refinement of the related theory; consequently this can provide the impetus for the discovery of further domain knowledge. Analyses [16][12] have shown that the detection and explanation (using domain knowledge) of these anomalies force scientists to revise their knowledge in a number of ways; from a minor refinement of hypotheses to major changes of fundamental knowledge.

The generation of hypotheses as part of automated scientific discovery processes have been discussed widely in the literature [4][17][3], of most relevance to this work is Blum's work on the RX project which investigated the automatic generation and testing of hypotheses from a clinical database [14]. The work was successful in discovering a previously unknown correlation between prednisone and cholesterol when tested on a clinical database from rheumatology patients. However, it is not known whether this approach would be successful in other medical areas such as the

Intensive Care Unit (ICU) domain; much larger amounts of data can be retrieved in the ICU domain, and in particular the abstraction process that Blum deploys would have to become more sophisticated to handle temporal datasets.

Due to the complexity and amount of data in the ICU domain it is difficult for clinicians to be fully conversant with the data; further it is recognized that the ICU is a challenging domain in which to perform decision making [19]. Artificial Intelligence systems used in the ICU domain generally analyse real time data streams and provide sophisticated monitoring of the data streams, applying domain knowledge to assist in data interpretation. Some of these systems, such as those developed by the MIMIC II [1] project have been implemented ‘live’ in an ICU ward, whilst other systems use data ‘offline’ [7]. Despite the wide variety of decision-support systems developed for use in the ICU, such as RÉSUMÉ [20] and VIE-VENT [18], none have focused on providing support to clinicians when faced with anomalous patient behaviour.

3 Capturing the Identification & Explanation of Anomalous Responses to Treatment

Five ICU clinicians from Glasgow Royal Infirmary were presented with physiological data for 10 patients containing 1466 hourly sequences and asked to identify anomalous sequences in the data. Five of these patients had been previously identified as containing anomalous patient responses and five were randomly selected from the patient management system. A semi-structured interview was held with the clinicians and they were asked to ‘talk-aloud’ as they completed the task [10]. A grounded theory approach [8] was applied by the analyst (LM) to the protocols which resulted in the following coding categories (C to E describe anomalies):

- A** - Anticipated patient responses to treatment, possibly with minor relapses (default if clinician does not provide any other classification)
- B** - Anticipated patient responses to treatment, with significant relapses e.g., additional bouts of sepsis, cardiac or respiratory failure
- C** - Patient not responding as expected to treatment
- D** - Odd / unusual set of physiological parameters (or unusual rates of change)
- E** - Odd / unusual treatment

“..... the thing that I am puzzled by in all this is the fact that the cardiac output went up when we increased their vasoconstrictor, I wouldn't necessarily expect that.”

Taken from transcript of Clinician 4 discussing Patient 909

Fig. 1 Anomalous Example

The coding categories were successfully verified by a second coder (DS)¹. The clinicians identified a total of 65 anomalies. Figure 1 provides an example anomaly. The most interesting type of anomaly for the clinicians was category ‘C’, which describes scenarios where a patient is not responding as expected to treatment. It was decided that this category would be the focus of the second study.

The anomalies classified as ‘C’ (a total of 13) were used as the stimuli for a second set of interviews with two further ICU clinicians (not involved in the initial identification of anomalies). Separate interviews were held with these clinicians during which they were presented with a series of anomalies and given access to the patient data. The clinicians were asked if they could suggest potential explanations for the identified anomalies. An example of a hypothesis suggested for the anomaly shown in Figure 1 is given in Figure 2².

“The patient is clearly deteriorating over the course of the day though, the urine volumes are decreasing, the oxygen requirement has gone up, probably this patient is just getting considerably sicker and they are just not responding to the noradrenaline.”

Explanation provided by Clinician 6

Fig. 2 Hypothesis Example

Analysis [8] of the transcripts from these interviews led to the identification of a range of hypotheses which could be organised as the following categories: 1) clinical conditions, 2) hormone regulation, 3) progress of the patient’s condition, 4) treatment, 5) functioning of the patient’s organs and 6) errors in recordings.

The interviews were analysed further and a method of information selection and hypothesis generation used by the clinicians was identified. When generating a hypothesis each clinician began with an anomaly, for example the one shown in Figure 1, “*noradrenaline increased cardiac output and cardiac index*”³, which can be broken down into the treatment, ‘*noradrenaline*’ and the effect ‘*increase cardiac output and cardiac index*’. The clinician then proceeded to explain any combination of the treatment and anomalous effect. The clinicians appeared to use domain knowledge about treatment, medical conditions and the desired physiological state of the patient to suggest hypotheses that explain the treatment or effect. Further, the clinician appeared to use domain knowledge whilst examining the patient’s data to determine facts; for example, the patient is suffering from a myocardial infarction. The patient’s data can also be used to *eliminate* hypotheses. For example, one of the explanations for the anomaly detailed in Figure 1 was that the patient is getting worse, if the data does not show this, this hypothesis can be eliminated. The clinician repeated the process until they were satisfied that all viable hypotheses had been proposed.

¹ Further details on the analysis of protocols are provided in [13]

² Noradrenaline is considered as a vasoconstrictor

³ Noradrenaline at low doses should not increase a patient’s cardiac output and input, but at high doses this may be observed

4 Generating Hypotheses for Anomalous Responses to Treatment

A tool, named EIRA (Explaining, Inferring and Reasoning about Anomalies) has been developed based on the general model of hypothesis generation presented in Section 3. It is envisaged that EIRA will be used by clinicians as an ‘offline’ aid/tutoring tool when faced with an anomalous scenario. The work reported in this section describes the initial implementation of EIRA (planned further implementation of EIRA are discussed in Section 5). EIRA comprises: a knowledge base consisting of several OWL⁴ ontologies and a Java based program implementing strategies extracted from domain experts’ protocols.

4.1 Knowledge Base

As the analyst noted when analysing the interviews, the clinicians drew from a large knowledge base in order to create hypotheses. The knowledge base of EIRA comprises four ontologies which model the following sub-domains: the ICU domain, the patient data, human physiology, and time.

Although various ICU ontologies have been discussed in the literature, for example [15]), none were available to us at the time of development; further using a standard biomedical ontology (including [2, 11]) for this tool would not be appropriate, mainly due to their size. Instead we created our own (relatively small) ontologies for the task. These smaller ontologies, consisting solely of the knowledge required by the system provided many benefits when building and editing the ontologies with a standard ontology editor; inferencing using these ontologies is also quicker. We do, however, recognise that standardised biomedical ontologies have a role as an important reference point, particularly beneficial for interoperability, and should be used whenever possible. To support such usage, a simple alignment meta-ontology enables the definition of correspondences between concepts in our ontologies with concepts in these standard ontologies. We believe that the framework provided by these ontologies will enable their re-use in other medical domains.

The ICU domain ontology was developed in collaboration with an ICU clinician (MS) and it has been shown to be sufficient to support the clinical reasoning discussed in section 3. Four types of knowledge are described: disorders, treatments, disorder severity scores, and drugs.

The `Drug` class (a subclass of `Treatment`, visualised in figure 3) describes how drugs are used as treatments in the ICU domain. Features of a drug, such as `activeDrugName`, `alternativeDrugNames`, the anticipated length of time between the drug’s administration to a patient and its effect being observed (the drug’s `timeToReact`), and any contraindications of the drug (disorders or other treatments) have been modelled in the drug class; in addition, descriptions of a drug’s effects, inter-

⁴ <http://www.w3.org/2004/OWL/>

actions, and uses have been supported. Various types of drug effects are described: *expected* effects, *conditional* effects which occur under certain conditions, *rarely observed* effects which are rare but still theoretically possible, and (unwanted) *side* effects.

Drugs can interact with other drugs to produce anticipated and unanticipated physiological effects. For a particular drug, it is usually known which other drugs can interact with it and the associated effect. This is represented by the `Drug_Interaction` class (shown in figure 4). Each `Drug_Interaction` is associated with a drug (`hasDrug`), the interacting drug (`interactsWith`), and the physiological effects observed during the interaction (`interactionEffect`).

Various properties are used to represent different drug doses, as a particular drug may be given at different doses depending on the severity of the disorder for which the drug is being administered. Differences in how commonly (or not) a drug is used to treat particular disorders is also described.

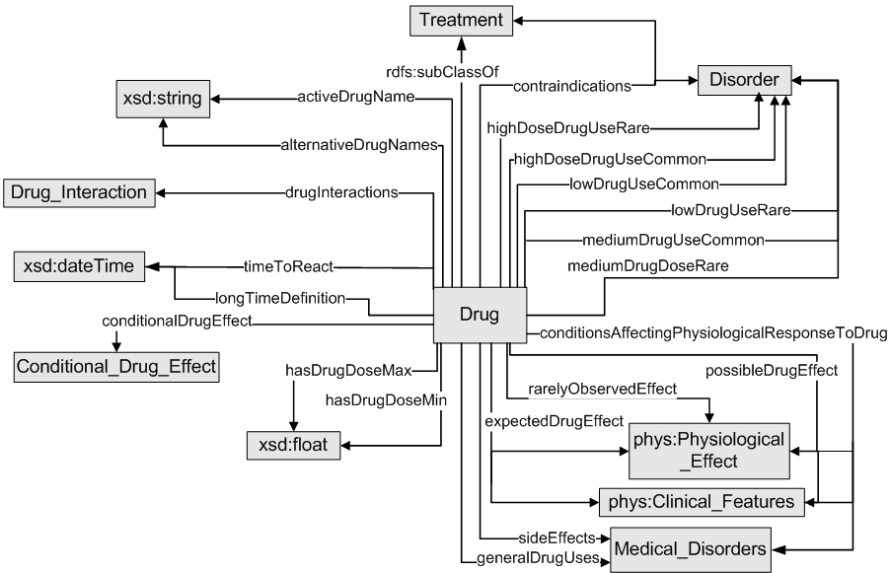


Fig. 3 Visualisation of the Drug Class from the ICU Ontology




Fig. 4 Visualisation of the Drug_Interaction Class from the ICU Ontology

The Patient Data ontology has been designed to model the time series data which is typically collected in medical domains. The ontology defines a *Patient_Data* class, which represents the patient's *Sessions* and *Location*. The *Session* class models a treatment session, which in turn links a series of *Timepoints*; the later describe the *Readings* for a particular *xsd:dateTime*. Each *Reading* has a *Parameter* and *value*.

The *Human_Physiology* ontology models, at a high level, knowledge regarding organs and organ systems, clinical features, and physiological effects. The *Organ_System* and *Organ* classes are used to represent basic human physiology, with the *hasOrgan* property associating the organ system with its primary organ. The *Clinical_Feature* class provides a template for describing physiological states such as *low MAP* and *high temperature*. The *Physiological_Effect* class represents different types of effects that occur in the human body; two types of effects have been defined: parameter changes and symptoms.

4.2 Reasoning about Anomalies

Section 3 suggests an anomaly can be defined as a *treatment* and an anomalous *effect* and can be considered to take the form:

For an anomaly at Time (T) \rightarrow Treatment (TR)_{T-1}, AND Anomalous Effect (E)_T

For example the anomalous statement shown in Figure 1 can be expanded to include relevant patient data as:

T = 12:00:00 Day 32 of Patient 583, TR = Noradrenaline, E = Increase Cardiac Output.

An anomaly is entered into EIRA in the format shown above (this is referred to as the 'original anomaly' in the rest of this document). EIRA then proceeds to determine if there are any other anomalous responses to treatment in addition to the original anomaly at this particular timepoint and then provides hypotheses for this set of anomalies.

4.2.1 Detecting Additional Anomalies

Anomalies, in addition to the anomaly identified by a clinician, can occur in a patient's dataset. EIRA can detect further anomalies at the anomaly time point(T). For example, in addition to the previous anomaly, the system may also detect that after administering noradrenaline, the patient's mean arterial pressure (MAP) reading is low. As discussed previously, an anomaly can be considered as a counterexample to domain knowledge. In this particular domain, this can be considered as:

Domain knowledge contains when *drug(D)* administered at T, *expected effect(E)* observed at *Time (T + 1)* AND patient data contains D at T AND *observed effect*, \neg E at T + 1.

Additionally, when attempting to detect anomalies, EIRA identifies the drugs given to the patient at the anomaly time point⁵ from the patient's data and retrieves the anticipated effects of administering each drug from the ICU ontology. When the anticipated response(s) have not occurred, the actual response observed in the data is noted. For example, the drug noradrenaline has been given to the patient at 12.00; further it is specified in the ICU ontology that noradrenaline is expected to increase systolic pressure, diastolic pressure, systemic vascular resistance (SVR) and MAP. These can be compared with the patient's readings and any discrepancies (for example, an observed decrease in MAP) are reported as additional anomalies.

4.2.2 Explaining Anomalous Responses to Treatment

To generate hypotheses for given anomalies, EIRA captures the various strategies (algorithms) and domain knowledge used by the two observed expert ICU clinicians and produce the types of hypotheses the clinicians proposed.

As identified in Section 3 the clinicians' hypothesis generation process focused on either providing an explanation for why the *treatment* in question did not work as anticipated, or concentrated on the reasons why the anomalous physiological *effect* might have occurred in the patient.

Although this system is implemented in the ICU domain, the algorithms listed below are believed to be generic and could be applied in other (medical) domains. Due to page limitations, we are unable to discuss each algorithm in detail, however, an example of the `Drug_Interaction` algorithm is provided in Algorithm 1. The following is the `Drug_Interaction` algorithm expressed as a SWRL⁶ rule to show how the concepts from the ontology are used:

$$\begin{aligned} & \text{Drug}(\text{?treatment}) \wedge \text{drugInteractions}(\text{?treatment}, \text{?interaction}) \wedge \text{Drug_Interaction}(\text{?interaction}) \wedge \\ & \quad \text{interactsWith}(\text{?interaction}, \text{?otherDrug}) \wedge \text{Timepoint}(\text{?t}) \wedge \text{hasTime}(\text{?t}, \text{?time}) \wedge \\ & \quad \text{hasReadings}(\text{?t}, \text{?r}) \wedge \text{Reading}(\text{?r}) \wedge \text{readingParameter}(\text{?r}, \text{?treatment}) \\ & \Rightarrow \text{DrugInteractionHypothesis}(\text{?hyp}) \wedge \text{hasDrug}(\text{?hyp}, \text{?treatment}) \wedge \\ & \quad \text{hasInteractingDrug}(\text{?hyp}, \text{otherDrug}) \end{aligned}$$

The algorithms are not dependent on each other, and all are executed in order to generate all possible hypotheses. Below we summarise the implemented algorithms:

1. `Conditional Drug Effects` - Under some known conditions a drug may have a different physiological effect on the patient than anticipated. For example, under high doses, noradrenaline may increase a patient's cardiac output, whereas at normal or low doses this would not be expected.
2. `Other Medical Conditions` - This method identifies whether the patient is suffering from another medical condition which has a symptom the same as the observed anomalous effect.

⁵ Although the example given here concerns the anomaly time point, any time point in the data can be used to detect an additional anomaly

⁶ <http://www.w3.org/Submission/SWRL/>

Algorithm 1 IdentifyDrugInteractions

```

1: Response - The anomalous response
2: Treatment - The anomalous treatment
3: Time - Time (T) in the patient's dataset at which the anomaly occurred
4: InteractingDrugs - Potential interacting drugs

5: Begin IdentifyDrugInteractions
6: Identify known drug interactions (DI) for Treatment
7: for Each DI do
8:   Identify the other drug ( $D_{int}$ ) involved
9:   Determine if the patient was being given  $D_{int}$  at T
10:  if  $D_{int}$  was given at T then
11:    Add  $D_{int}$  to InteractingDrugs
12:  for Each  $D_{int}$  in InteractingDrugs do
13:    return Hypothesis - Drug,  $D_{int}$ , may be interacting with Treatment

```

3. Other Medical Condition - Treatment - Identifies whether a treatment given to the patient for another medical condition is not working as anticipated and hence is responsible for the anomalous effect.
4. Drugs - Identifies whether another drug that the patient is receiving at the anomaly time could have the same effect (side effect or therapeutic effect) as the anomalous effect (and hence this further drug explains the observed anomaly).
5. Patient Improvement - An improvement in the patient may be the cause of the anomalous effect. This can be split into the following areas of improvement⁷:
 - a. Overall Patient Improvement - Identifies whether the anomalous effect may be explained by a patient's *overall* clinical condition improving. For example, in the ICU a severity score is often associated with a patient's overall condition. This score can be calculated from the patient's physiological data. If the severity score has shown an improvement at the anomaly time point, then it can be concluded that generally the patient is improving.
 - b. Improving Organ - Establishes whether one of the patient's organs could be spontaneously improving/recovering and hence could explain what otherwise would appear as an anomalous effect. For example, noradrenaline does not usually increase a patient's cardiac output, however, cardiac output is a measurement reflecting the condition of a patient's heart; if the majority of the other measurements of a patient's heart (e.g. heart rate) are also improving, it is possible that an improvement in a patient's cardiac performance may explain the anomalous effect (and not the noradrenaline).
 - c. Specific Condition Improvement - Ascertains if a patient could be recovering from a *specific* previous clinical event and the anomalous effect can be explained as a consequence of this improvement. For example, an observed increase in cardiac output seen when noradrenaline is administered could be

⁷ Improving Organ and Specific Condition Improvement can be considered as conceptually similar, however, they require distinct implementations and hence have been listed separately.

attributed to the patient recovering from sepsis. Unlike *Improving Organ* we cannot say that the heart has improved, only that one of the measurements of the heart has improved.

Complementary to the above analyses the following algorithms have been implemented to suggest hypotheses which may explain why the treatment has not worked as anticipated.

6. *Conditions Affecting Treatment* - Identifies whether the patient has a co-existing medical condition which can affect how well the treatment associated with the anomaly is working. For example, if a patient is an alcoholic, they can have a high cellular tolerance for the drug propofol i.e. they require a higher amount of the drug for an effect to be seen than with a non-alcoholic patient.
7. *Drug Interactions* - Establishes if the treatment associated with the anomalous effect has not worked as expected because of an interaction occurring with another treatment the patient is receiving. The secondary treatment may in some cases increase, decrease, or negate the expected effect of the primary treatment.
8. *Low Dose of Treatment* - Too low a dose of a drug may have been given to the patient and hence the intended effects have not been observed. This algorithm determines whether the dose of the drug associated with the anomalous effect is a 'low dose' (a low dose is taken to be the lower quartile (25%) of a drug's specified dose range).
9. *Overall Condition Deterioration* - Establishes if the patient's *overall* condition has deteriorated. If the patient's condition has deteriorated, the patient, in general, would require a higher dose of a drug for its intended effect to be observed. In some instances a patient can become too ill to respond to treatment. The overall condition of the patient is based on a severity score calculated from the patient's physiological data. If the severity score has worsened at the anomaly time point then it is inferred that the patient's overall condition has worsened.
10. *Resistance to Treatment* - Identifies if the patient could have become resistant to the treatment. When certain drugs are administered over a long period of time and/or at a high dose, a patient can become resistant to them and hence the drug has very little or no effect on the patient.

For each individual anomaly (from the set consisting of the original anomaly and any additionally identified anomalies), EIRA systematically works through each of the algorithms. At the end of the process the hypotheses for each anomaly are presented to the clinician. Figure 5 shows a segment of a sample output from EIRA. The numbers next to a hypothesis correspond to the numbered strategies above.

4.3 Evaluation

To evaluate the effectiveness of the hypothesis generation mechanisms, test cases were extracted from the interviews held with ICU clinicians. The test cases were

```
Patient ID: 119
Anomalous Treatment: Propofol
Anomalous Effect: HighMAP
Anomaly Time: 13/07/2006 06:00:00
*****
Anomalous Response added: Treatment: Propofol anomalous
effect: DecreaseHeartRate
Anomalous Response added: Treatment: Propofol anomalous
effect: DecreaseMAP
*****
Explaining why the anomalous effect may have occurred....
Anomaly being investigated:
Treatment: Propofol Response: HighMAP
(2) Identifying if the effect can be caused by another medical
(4) condition....
(5a) Identifying if any other treatments can explain the
effect....
(5b) Identifying if the patient is getting significantly
(5c) better....
(1) Too early in the session to calculate the trends
(3) Identifying if an organ has improved...
Identifying if the patient is recovering from an illness...
Identifying if a conditional drug effect has been met....
Identifying if the effect can be caused by treatment for a
medical condition....
The following conditions may explain the effect seen:
Hypertension
ST_Elevation_Myocardial_Infarct
Phaeocromocytoma
Explaining why the treatment may not have worked....
Anomaly being investigated:
Treatment: Propofol Response: HighMAP
(7) Identifying any drug interactions....
(6) Identifying any conditions which may affect response to
treatment...
The following conditions may have caused the treatment to
work differently:
(9) Cellular_Tolerance
Identifying if the patient's condition is getting worse....
(8) Identifying if the treatment dose was low...
The treatment, Propofol may not be working because a low
dose has been given
(10) Identifying if the patient could have become resistant to
treatment....
.....
```

Fig. 5 Example Output

formed from the anomalies identified from the first group of interviews and the subsequent hypotheses provided by the clinicians in the second group of interviews (as detailed in section 3). In total, 15 test cases were selected with 25 hypotheses (in some test cases clinicians suggested multiple hypotheses). Each test case comprised: the time and date of the identified anomaly, details of the anomalous effect observed, the associated treatment, the hypotheses given by a clinician to explain this anomaly, and the patient’s data for their complete stay (containing physiological readings e.g. Heart Rate and information about the drugs administered). During testing, each test case anomaly was entered into EIRA and we recorded whether the tool generated the same hypotheses as the ICU clinicians. A note was also made of whether the system detected the same additional anomalies that the ICU clinician had identified. Table 1 shows the comparison between the hypotheses suggest by the clinicians and by EIRA for the same anomaly.

	Test Cases															Total
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	
No. of Hypotheses (Interview)	4	2	3	1	2	2	1	1	1	1	2	1	1	2	1	25
No. of Hypotheses (EIRA)	2	2	2	1	0	1	1	0	0	0	1	1	1	1	0	13

Table 1 Test Case Results

A total of 13 hypotheses were matched by EIRA. The explanation for EIRA not identifying the remaining 12 hypotheses is that the clinician's hypothesis were not captured in the patient's data or the clinician's hypothesis was not reflected in the knowledge base. These points are discussed in some more detail below:

4.3.1 Patient Data

For 8 of the hypotheses, the patient's physiological data did not support the hypothesis given by a clinician. For example, in test case 15, the clinician suggested that a low heart rate observed in the patient in response to adrenaline and noradrenaline could be explained by the patient suffering from severe sepsis. The "severe sepsis" condition in the knowledge base requires a patient to have both a high MAP and either a low or high temperature. These conditions, however, were *not* observed in the patient's data and so the tool did not suggest that the patient may have severe sepsis.

Another hypothesis was not produced due to an error in the patient dataset (noradrenaline was recorded in mls/hr instead of mg/hr). This has been corrected.

Discrepancies between the clinician's hypothesis and the patient's data are not entirely unexpected. Studies [5][9] have shown that when considering potential explanations for an anomaly the expert's domain knowledge can influence hypothesis testing when dealing with inconsistent evidence. Chinn & Brewer[5] found that when the subject's background knowledge was not consistent with the evidence (in this case the patient's data), subjects largely discount the data. It is possible that this phenomena is being observed in some of the situations here. Further, the expert may also be reporting patient states which are *partially* supported by the data.

4.3.2 Knowledge Base Content

For 3 hypotheses, the knowledge base did not have the required facts to create the same hypothesis as the clinician. For example, in test case 3, in response to noradrenaline given, the patient's systemic vascular resistance (SVR) decreased. The clinician suggested that sepsis may be responsible for the decrease in SVR. EIRA did not suggest sepsis because the knowledge base did not have a *decrease in SVR* identified as a symptom, however, it did contain a *low SVR* as a symptom of sepsis. It is suggested that in these cases, the knowledge base does not accurately reflect

the clinicians' domain knowledge and clearly refinement of the knowledge base is required. This facility is suggested as an extension to EIRA in section 5.

Of the additional anomalies identified by the clinicians, 2 of the 3 were also identified by EIRA. The one case where the anomaly was not identified can be considered as a partial match; the clinician identified a *low* heart rate in response to the drug, propofol, as being anomalous whilst EIRA identified a *decrease* in heart rate as anomalous in response to propofol. EIRA did not use the term 'low heart rate' because the knowledge base only contained for that drug an 'increase in heart rate' as an expected effect relating to the patient's heart rate when administering propofol. If the knowledge base had also had a 'high heart rate' as an expected effect, EIRA would have identified the same anomaly as the clinician.

EIRA also produced a number of 'new' hypotheses for each test case (in addition to the hypotheses suggested by the clinician). These 'new' hypotheses haven't been evaluated but a planned future evaluation will involve an ICU clinician evaluating the 'new' hypotheses produced by EIRA for their clinical relevancy. In summary, the initial evaluation of EIRA is encouraging; in the majority of cases EIRA can adequately reproduce the precise hypotheses suggested by expert ICU clinicians.

5 Conclusions and Future Work

This paper outlines several interviews held with ICU clinicians to identify and provide hypotheses to explain anomalous patient responses to treatment. Further, a tool based on the processes captured from these interviews has been developed; additionally an evaluation of the system has produced promising results. Generally the hypotheses suggested by the clinicians are reproduced by the system; further the reasons why the remaining hypotheses were not reproduced have been adequately explained and will be addressed subsequently. Plans for future work include:

- **Further Evaluation** - The additional hypotheses produced by EIRA require evaluation by a domain expert for their clinical relevancy, further, upon completion of planned future stages of development, EIRA will be tested on a larger number of test cases.
- **Knowledge Base Refinement** - It would be beneficial to allow a clinician using EIRA to refine the domain ontologies. To enable this refinement, it is proposed that when a hypothesis is presented to the clinician, the clinician will be able to reject the hypothesis and provide a reason for its rejection. The relevant instances of the ontology will then be updated to reflect this 'feedback' from the clinician.
- **Inferring Information from the Knowledge Base** - As shown in the evaluation of the system, occasionally the domain knowledge has 'gaps' in it, which in some cases have prevented hypotheses from being suggested by the system. It would be useful to extend EIRA to make assumptions based on other instances in the ontology. For example, if property P has been observed for D_1 of drug class D then it is reasonable to infer that all sibling drugs of D_1 also have property P.

Acknowledgements Kathryn Henderson and Jennifer McCallum (CareVue Project) & the staff and patients of the ICU Unit, Glasgow Royal Infirmary. This work was an extension of the routine audit process in Glasgow Royal Infirmary's ICU; requirements for further Ethical Committee Approval have been waived. This work was supported under the EPSRC's grant number GR/N15764.

References

1. MIMIC II Project, Accessed August 2009. <http://mimic.mit.edu/index.html>.
2. Unified Medical Language System (UMLS), Accessed August 2009. <http://www.nlm.nih.gov/pubs/factsheets/umls.html>.
3. Blum, R L. Discovery, Confirmation, and Incorporation of Causal Relationships From a Large Time-Oriented Clinical Data Base: The RX Project. *COMP. AND BIOMED. RES.*, 15:164–187, 1982.
4. Bruce G. Buchanan and Edward A. Feigenbaum. Dendral and Meta-Dendral: Their Applications Dimension. *Artificial Intelligence*, 11:5–24, 1978.
5. Clark. A. Chinn, William F. Brewer. Factors that Influence How People Respond to Anomalies. In *Proc. of the Fifteenth Annual Conference of the Cognitive Science Society*, 1993.
6. D Kuhn. *The structure of scientific revolutions*. University of Chicago Press, 1962.
7. Jaap van den Heuvel, J. D. B. Stemerding, Ad J. J. C. Bogers and David S. Bree. GUUS an expert system in the intensive care unit. *International Journal of Clinical Monitoring and Computing*, 7:171–175, 1990.
8. Juliet Corbin, Anselm Strauss. Grounded Theory Research: Procedures, Canons and Evaluative Criteria. *Qualitative Sociology*, 13:3–21.
9. K. Dunbar. Concept Discovery in a Scientific Domain. *Cognitive Science*, 17:397–434, 1993.
10. K. Anders Ericsson and Herbert A. Simon. *Protocol analysis: verbal reports as data*. MIT Press, 1993.
11. Kent A. Spackman, Keith E. Campbell. SNOMED RT: A Reference Terminology for Health Care.
12. L Darden. Strategies for Anomaly Resolution. *Cognitive Models of Science, Minnesota Studies in the Philosophy of Science*, :, 1992.
13. Laura Moss, Derek Sleeman, Malcolm Booth, Malcolm Daniel, Lyndsay Donaldson, Charlotte Gilhooly, Martin Hughes, Malcolm Sim and John Kinsella. Explaining Anomalous Responses to Treatment in the Intensive Care Unit. In *Proceedings of the 12th Conference on Artificial Intelligence in Medicine, AIME09*, 2009.
14. Michael G. Walker and Gio Wiederhold. Acquisition and Validation of Knowledge from Data. *Intelligent Systems*, :415–428, 1990.
15. N.F de Keizer, A. Abu-Hanna, R. Cornet, J.H.M. Zwetsloot-Schonk, C.P. Stoutenbeek. Analysis and Design of an Ontology for Intensive Care Diagnoses. *Methods of Information in Medicine*, 38:102–112, 1999.
16. Paul Thagard. *Conceptual Revolutions*. Princeton University Press, 1992.
17. P. Langley, G.L. Bradshaw, and H.A. Simon. Rediscovering Chemistry With the Bacon System. *Machine Learning*, pages 307–329, 1983.
18. Silvia Miksch, Werner Horn, Chrisitan Popow, Franz Paky. VIE-VENT: knowledge-based monitoring and therapy planning of the artificial ventilation of newborn infants. *Artificial Intelligence in Medicine*, 10:218–229, 1993.
19. Vimla L. Patel, Jiajie Zhang, Nicole A. Yoskowitz, Robert Green, Osman R. Sayan. Translational cognition for decision support in critical care environments: A review. *Journal of Biomedical Informatics*, 41:413–431, 2008.
20. Yuval Shahar and Mark A. Musen. Knowledge-Based Temporal Abstraction in Clinical Domains. *Artificial Intelligence in Medicine*, 8:267–298, 1996.