COMPUTER SUPPORT FOR PROTOCOL-BASED TREATMENT OF CANCER[1]

PETER HAMMOND[2] AND MAREK SERGOT[3]

July 1994

Cancer treatment is often carried out within protocol-based clinical trials. An oncology clinic may take part in many trials each of which requires data to be collected for monitoring efficacy and toxicity of treatment. Subsequently, this data is analysed statistically to evaluate clinical objectives of the trial. To be scientifically valid, such analysis must be based on data that is both complete and correct. This is one motivating factor for introducing computer support for trial management. Further motivation is provided by concern that treatment is consistent with the protocol and the well-being of the patient. The complexity of many protocols, the life-threatening nature of cancer and the toxicity of treatment side-effects emphasise the safety-critical nature of oncology. The OaSiS system provides decision support for the protocol-based treatment of cancer patients with emphasis on the safety aspects of the advice it gives. It offers a highly graphical interface, employs integrity constraint checking techniques from logic databases to monitor compliance with a protocol and is implemented in PROLOG. The paper describes the main features of OaSiS and indicates work in progress and planned.

1. INTRODUCTION

OaSiS is a decision support system for the management of cancer patients based on the written protocols which govern clinical trials of therapies. Such computer support is necessary because the protocols are complex and there is a need for complete and correct accumulation of clinical data before scientifically valid statistical analysis can be carried out. This paper focuses on the current functionality of OaSiS and discusses the use of safety related knowledge identified from an extensive study of oncology protocols and from discussions with clinicians, pharmacists and medical informaticians. The safety critical nature of the domain imposes requirements on software designers and implementers to ensure that the translation from paper to computerised protocol is completed thoroughly and correctly [9]. OaSiS has been implemented within RED, a project funded by the UK DTI and SERC "Safety Critical Systems" research programme.

A major influence on OaSiS is the work at Stanford University on the ONCOCIN [32], EON [26] and OPAL [25] family of computer systems. In the OaSiS prototype, this is particularly evident in the user interface and its combined use of graphical, form-based

[1] An earlier version of this paper was presented at the 2nd International Conference on the Practical Application of PROLOG, London, 1994.

[2] *Address correspondence to:* Peter Hammond, Advanced Computation Laboratory, Imperial Cancer Research Fund, Lincoln's Inn Fields, London WC2A 3PX.

[3] Department of Computing, Imperial College of Science, Technology and Medicine, 180 Queen's Gate, London SW7 2BZ.

and spread-sheet like windows to present and manipulate clinical data. A second influence is that of ICRF's own work arising from the BOSS and DILEMMA projects which have affected the underlying protocol model, decision-making procedures and functionality. More specialised systems have been produced for radiotherapy planning [16, 27]; for breast cancer [10, 20]; for ovarian cancer [21] and for head and neck cancer [23]. The main differences between the work on OaSiS and that of other projects on cancer management systems is the emphasis on safety issues and the use of analysis and implementation methods based on logic based techniques.

Before illustrating what OaSiS can offer to a clinical user, we provide some background to the domain and explain in more detail why such software is necessary and potentially beneficial. The final section of the paper describes plans for future improvements.

1.1 Protocol-based Treatment in Cancer Trials

Variation in tumour behaviour and the variety of cancer therapies inhibit the selection of optimal treatments of many cancers. Instead, patients are entered in clinical trials involving combinations of surgery, hormone therapy, radiotherapy and chemotherapy with the aim of evaluating efficacy and safety. Clinical trials are governed by procedures, or protocols, which detail how they should be planned. A wide range of issues is presented in protocols which are detailed, complex documents. However, there is a common structure to the organisation of many cancer trials and those tasks amenable to computerisation are listed in figure 1.

eligibility
randomisation/trial registration
pre-treatment investigations
treatment planning (chemo-, hormone, radio-therapy; surgery)
monitoring response of disease (typically tumour size and incidence)
monitoring response of patient (typically side-effects)
modification of disease treatment as a result of the monitoring
treatment of side-effects
treatment follow-up and end-points

FIGURE 1. Tasks common to cancer trials and amenable to computerisation

Clinical trials typically have a number of "arms" specifying chemotherapies and other treatments to which patients are randomly assigned for subsequent comparison and for evaluation of trial objectives. Indeed, most protocols begin with an explicit statement of such objectives (figure 2).

(i) To determine whether treatment with the intensive regimen BOP/VIP-B is more effective than treatment with BEP/EP in the management of patients with poor prognosis metastatic teratoma with respect to complete response rate, progression-free and overall survival.

(ii) To determine the effect of r-metHug-CSF on the proportion of patients receiving full dose intensity of combination chemotherapy with either BOP/VIP-B or BEP/EP.

FIGURE 2: Objectives of protocol TE13 [24].

The high level tasks of a protocol and their temporal ordering are often presented diagrammatically (figure 3):

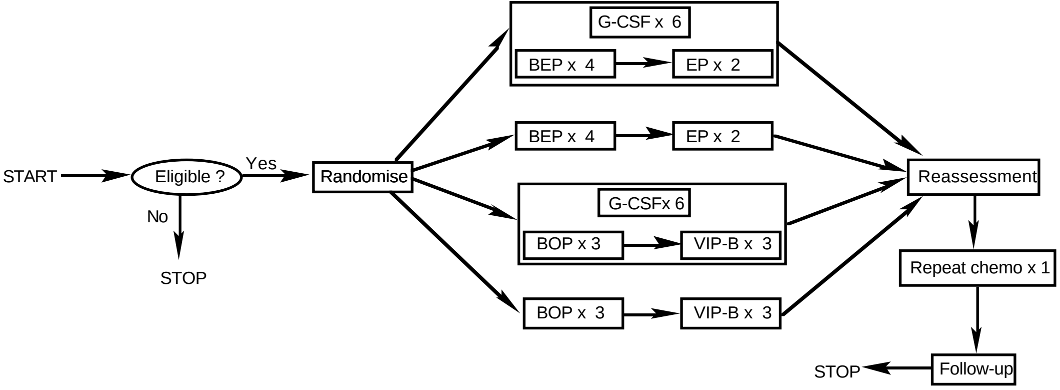

FIGURE 3. Overview of major tasks in a clinical trial of treatments for testicular cancer

The trial depicted [24] has four randomisation arms, each comprising different sequences of combinations of chemotherapeutic drugs. At the next level of detail, figure 3 illustrates the complexity of planning this particular chemotherapy. Drugs are given singly or in combination on particular days of cycles.

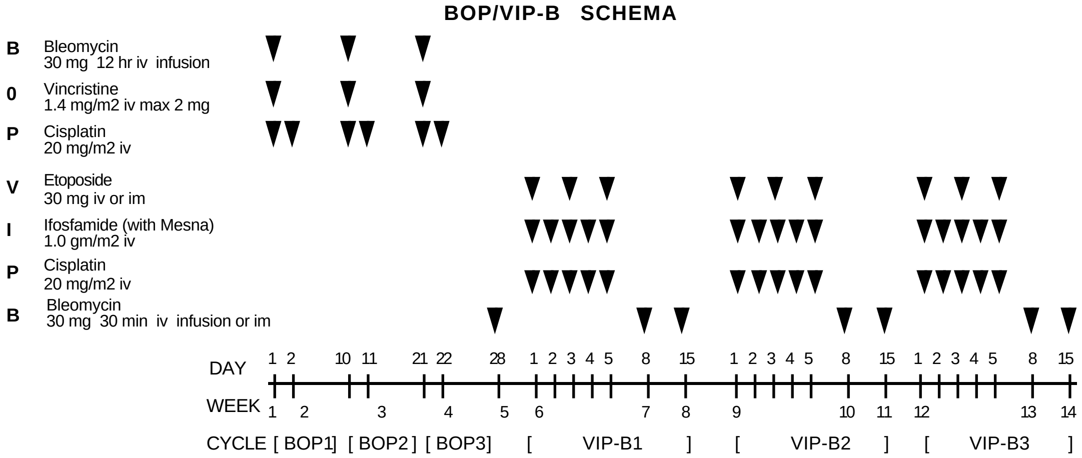
BOP/VIP-B SCHEMA
B Bleomycin
30 mg 12 hr iv infusion
0 Vincristine
1.4 mg/m2 iv max 2 mg
P Cisplatin
20 mg/m2 iv
V Etoposide
30 mg iv or im
I Ifosfamide (with Mesna)
1.0 gm/m2 iv
P Cisplatin
20 mg/m2 iv
B Bleomycin
30 mg 30 min iv infusion or im
DAY 1 2 10 11 21 22 28 1 2 3 4 5 8 15 1 2 3 4 5 8 15 1 2 3 4 5 8 15
WEEK 1 2 3 4 5 6 7 8 9 10 11 12 13 14
CYCLE [BOP1] [BOP2] [BOP3] [VIP-B1] [VIP-B2] [VIP-B3]

FIGURE 4. Part of a complex chemotherapy plan

1.2 The need for computer support in oncology

A written protocol can be between 20 and 60 pages in length and a major oncology clinic may use as many as 60 protocols concurrently [32]. The amount of clinical data collected (figure 5) and the complexity of some protocols call for computerised management of the patient's clinical details, the clinician's diagnosis, the treatment plan and the subsequent treatment modifications arising from therapeutic effects and toxic side-effects.

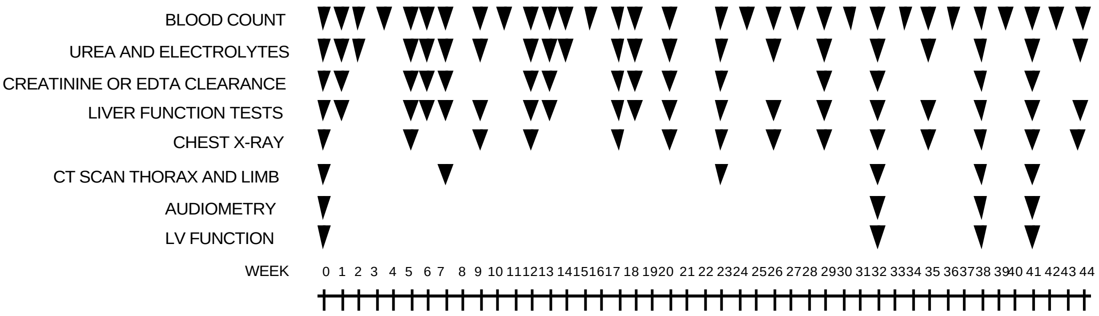

FIGURE 5. Data collected in 44 weeks for a bone cancer trial [3]

Patients do not always receive therapy in precisely the manner advocated by the protocol and data needed for formal analysis of treatment results are not always completely and accurately collected [32]. Inadequately collected data has invalidated a number of large-scale clinical trials[4]. Therefore, potential benefits of the use of decision support software are improved compliance with the protocol and improved data collection [17, 19]. With more accurate and more complete data, the statistical evaluation of trial objectives can be more effective with concomitant improvements in medical knowledge and in the treatment of patients.

1.3 The safety critical nature of decision support in oncology

Common side-effects of chemotherapy include bone marrow suppression and damage to the gastrointestinal mucosa. Bone marrow suppression lowers white cell count and makes a patient susceptible to serious infections. If the platelet count is low then severe bleeding may occur. Mucositis is unpleasant, causes severe weight loss and dehydration and provides a route for infection. Some unusual drug side-effects cause severe lung and heart problems. Thus, there is a narrow therapeutic window between giving sufficient drug for optimal anti-tumour effects and for life-threatening toxicity. Therefore, protocol guidelines for dose adjustment must be followed if the patient's life is not to be put at risk. Adverse toxic events are often presented in a protocol as a table of hazards, means for their detection and possible ameliorating actions (see table 1 below).

Some users of an oncology decision support system will be sufficiently expert to detect inaccuracies and potentially hazardous advice before it is acted upon. Less experienced junior doctors may fail to spot incorrect advice simply because they are unfamiliar with the protocol or the drug side-effects. Whether it is to avoid losing the confidence of experienced oncologists or to avoid unsafe recommendations being acted upon, extreme care must be taken to ensure that the transition from paper protocol to computerisation through abstract representation is as complete and correct as is possible. It is essential, therefore, that safety issues are thoroughly investigated before detailed software design

[4] The variation in quality control of data capture in multi-centre trials has been of particular concern [34].

and implementation begins. Preliminary results of our work on safety aspects of OaSiS are reported in detail elsewhere [13].

Hazard	Monitoring	Drug	Modification	Duration
Bone marrow suppression	wbc < 1.5 OR platelets < 50	all	suspend abort protocol	3 days if no improvement after 2 wks
	wbc > 3 AND platelets ∈ [50,75]	cisplatin	75% dose	This cycle of treatment
	wbc ∈ [2.1,3] AND platelets > 50	etoposide ifosfamide		
	wbc ∈ [1.5, 2] AND platelets > 75	etoposide ifosfamide		
	wbc ∈ [1.5, 2] AND platelets [50,75]	etoposide ifosfamide	50% dose	
Renal impairment	creatinine clearance < 40 mls/min	cisplatin bleomycin	75% dose withdraw	until recovery
Skin toxicity	Severe rash	bleomycin	withdraw	
Mucosal toxicity	Severe ulcers	all		
Dyspnoea	Shortness of breath	bleomycin		permanently
Anaphylaxis	Acute allergic reaction	cisplatin		permanently

TABLE 1. Hazards, monitoring conditions and corrective actions [3]

2 THE OASIS SOFTWARE

2.1 Background and overview

Following a study of ICRF's previous and existing projects on decision support systems in oncology [11, 29] and similar activities at Stanford University [19, 32], the development of OaSiS began in late 1992/early 1993. Primarily, the prototype has been produced as a demonstrator to oncologists with a view to recruiting their active contribution to the future design and implementation of a fully fledged system. In addition, the safety-critical nature of decision support in oncology could not be identified fully without clinical participation. Secondary aims have been to test methodological approaches to the associated decision-making and to their implementation in PROLOG on available hardware. A further aim has been to experiment with the user interface which needs to be highly graphical given the nature of the clinical data. The collaborating oncology clinic at the Churchill Hospital in Oxford uses Macintosh computers and so MacPROLOG [22], with its built-in facilities for dialogue and graphics management, was an obvious programming environment to select.

Currently, OaSiS has 7 major components as indicated in figure 6.

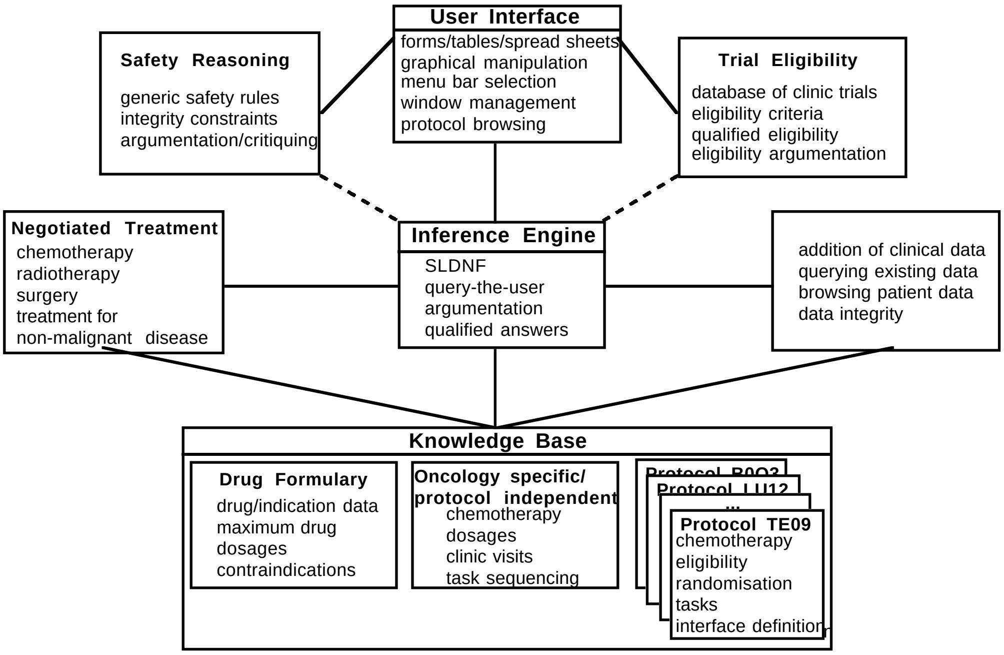

FIGURE 6. Major components and functionality of the OaSiS prototype

The main mode of interaction with OaSiS is analogous to that of a conventional spreadsheet. Typically, the user enters the relevant clinical data and, once vetted and recorded, OaSiS automatically generates advice for chemo- and other therapies according to the protocol recommendations for the appropriate stage of the patient's treatment. Should the user try to override these calculations by making some alteration then safety constraints are brought to bear in a style typical of logic databases. In subsequent sections of the paper, we illustrate how the clinician uses OaSiS to generate treatment recommendations and how alterations or suggestions by the user are checked for compliance with the computerised protocol. This will involve descriptions of patient eligibility criteria for trial entry, the user interface, some components of the knowledge base and the animation of the underlying protocol itself.

2.2 The representation of eligibility criteria

Eligibility criteria for clinical trial entry are typically represented as a collection of including and excluding conditions. Protocol TE09 [14], for example, has eligibility criteria:

(1) Histologically confirmed non-seminatomous germ cell tumour of the testis or combined seminoma/teratoma or seminoma with serum alpha feto protein concentration > 25 ku per litre;
(2) The following stage categories are eligible:-
Stage I with raised serum marker (see (3)).
Stage II with masses up to 10 cm in diameter.
Stage III with abdominal mass up to 10 cm in diameter or supraclavicular/mediastinal masses < 5 cm.
Stage IV with up to 20 lung masses. Patients with other sites (e.g., liver, bone or brain) are excluded.
(3) Patients with serum beta HCG >10000 iu per litre or serum AFP>1000 ku per litre are excluded.
(4) No previous chemotherapy given.
(5) Age greater than 15 years.
(6) Informed consent given for the proposed study.

FIGURE 7. Textual statement of eligibility criteria [15]

Eligibility is considered on a particular day, and so a straightforward representation[5] begins with:

```
patient_passes_eligibility_criteria_on_date(Patient, te09, Date) if
        has_satisfactory_diagnosis_for_protocol_on_date(Patient, te09, Diagnosis, Date) and
        satisfactory_staging_for_protocol_and_diagnosis(Patient, te09, Date) and
        not excluded_from_protocol_by_raised_serum_markers(Patient, te09, Date) and
        not already_received_chemotherapy(Patient, Date) and
        satisfactory_age_on_date(Patient, te09, Date) and
        satisfactory_consent_by_patient_before(Patient, te09, Date).
```

and, for example, the check that the diagnosis is acceptable is defined as

```
     has_satisfactory_diagnosis_for_protocol_on_date(Patient, te09, Diagnosis, Date) if
             patient_diagnosis(Patient, Diagnosis) and
             diagnosis_is_compatible_with_protocol(Diagnosis, te09, Patient, Date).

     diagnosis_is_compatible_with_protocol(['non-seminatomous germ cell tumour'],te09,Patient,Date).
     diagnosis_is_compatible_with_protocol([seminoma,teratoma], te09, Patient, Date).
     diagnosis_is_compatible_with_protocol([teratoma,seminoma], te09, Patient, Date).
     diagnosis_is_compatible_with_protocol([seminoma], te09, Patient, Date) if
        afp(Patient, Date, AFP) and
        AFP > 25.

askable(afp(Patient, Date, Result)).
```

The declaration *askable(Predicate)* means that all information regarding *Predicate* is to be obtained from the user (or an external source) in the Query-the-User fashion [30].

The representation of eligibility criteria also illustrates the current use in OaSiS of a temporal argument for those predicates involving time. This simplistic approach (which works reasonably well, it should be said) is being replaced in RED by an interval-valued temporal logic [5]. Protocols often include temporal references to treatments, particularly chemo- and hormone therapies, in a relative fashion in terms of cycles of treatment. Thus, it is necessary to have both relative and absolute descriptions of time points, with the former being used for drug calculation rules (see section 2.5) and the latter for eligibility, patient records and visit scheduling.

2.3 The user interface

The graphical user interface (figure 7) handles all interaction with the user through the manipulation of iconic, form and tabular based presentations of clinical data. All interactions are viewed either as data entry or query invocation. A graphics window (e.g., MASS/X-RAY POSITIONS in figure 7) is used for entering and presenting tumour incidence in an iconic format. Forms (eg PATIENT DATA) are generated from a form description language and are typically used to present and enter administrative data. Data that is tied to dates of clinic visits (see VISIT window) is presented spread-sheet

[5] Throughout the paper, we use the convention that strings beginning with an upper case letter denote variables.

style in a tabular format (examples include HAEM for blood test results and CHEMO for calculated drug dosages).

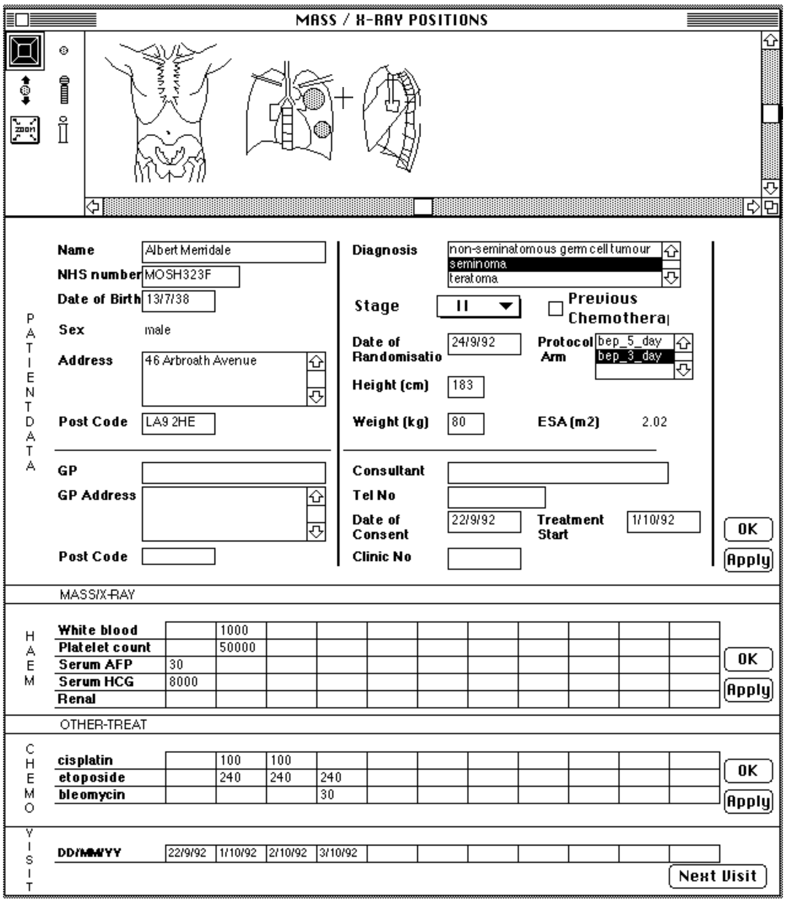

Name: Albert Merridale
NHS number: MOSH323F
Date of Birth: 13/7/38
Sex: male
Address: 46 Arbroath Avenue
Post Code: LA9 2HE
GP:
GP Address:
Post Code:

Diagnosis: non-seminatomous germ cell tumour / seminoma / teratoma
Stage: II
□ Previous Chemothera
Date of Randomisatio: 24/9/92
Protocol Arm: bep_5_day / bep_3_day
Height (cm): 183
Weight (kg): 80
ESA (m2): 2.02
Consultant:
Tel No:
Date of Consent: 22/9/92
Treatment Start: 1/10/92
Clinic No:

White blood		1000									
Platelet count		50000									
Serum AFP	30										
Serum HCG	8000										
Renal											

cisplatin		100	100								
etoposide		240	240	240							
bleomycin				30							

DD/MM/YY	22/9/92	1/10/92	2/10/92	3/10/92							

FIGURE 7. The OaSiS user interface for trial management

So as not to overwhelm a user with too much data, each window can be toggled to a closed state (by clicking on its label) whereby it is presented as a slim, labelled, horizontal band (for example MASS/X-RAY and OTHER-TREAT). These spread-sheet like windows have been implemented using lower level primitives since MacPROLOG does not support them directly.

2.4 The knowledge base and problem solving in OaSiS

A major part of the knowledge base deals with aspects of cancer management which are independent of particular protocols. Individual modules contain protocol specific knowledge for particular types of cancer as well as a small subcomponent describing the forms and "spreadsheets" for configuring the general user interface for that protocol. A third component contains drug information extracted from the protocols studied, supplemented by standard texts [1, 2, 8].

OaSiS can be thought of as a kind of "expert system shell" restricted to the domain of oncology. It can be used with any protocol provided the following details are provided:

a partial ordering defining the sequence of high-level tasks
a description of the randomisation arms
trial eligibility criteria

a description of chemotherapy or hormone therapy as cycles and days of treatment (for clinic visits)
a list of blood and other tests and their frequency for toxicity monitoring
default drug prescriptions (dose, units, route, formulation, period of administration etc)
dose modification rules according to toxicity monitoring

definition of the user interface for the particular form and tabular layout required

The remainder of the oncology knowledge base handles the sequencing of high level tasks, the scheduling of clinic visits, the computation of precise drug doses and their days for administration and the checking of the integrity of user input. The latter covers both elementary input errors (figure 8) as well as the application of specific safety knowledge (figure 9). Once the clinician accepts a computed treatment suggestion or OaSiS accepts the user's modification of its suggestion, the result of this negotiation is added to the record of the treatment plan. A more extensive discussion of safety related matters appears in section 2.6.

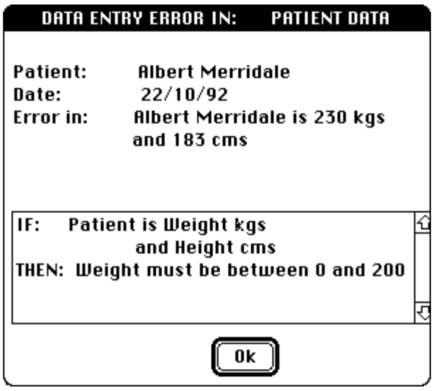
DATA ENTRY ERROR IN: PATIENT DATA

Patient: Albert Merridale
Date: 22/10/92
Error in: Albert Merridale is 230 kgs and 183 cms

IF: Patient is Weight kgs and Height cms
THEN: Weight must be between 0 and 200

FIGURE 8: violation of data integrity

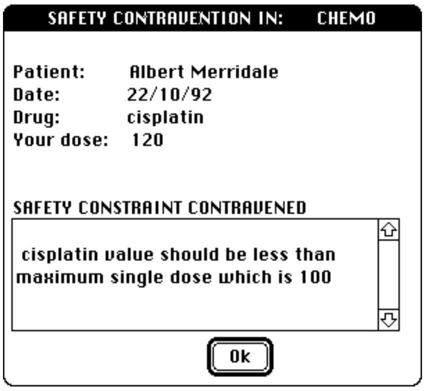
SAFETY CONTRAVENTION IN: CHEMO

Patient: Albert Merridale
Date: 22/10/92
Drug: cisplatin
Your dose: 120

SAFETY CONSTRAINT CONTRAVENED

cisplatin value should be less than maximum single dose which is 100

FIGURE 9: violation of safety

Currently, the knowledge base is represented in Horn clauses extended with negation as failure. Thus, most problem solving can be carried out directly by the underlying PROLOG interpreter. Interaction with the user, for obtaining clinical data about patients, is handled by a meta-interpreter in the Query-the-User style of Sergot [30].

2.5 Administration of therapies at clinic visits

Treatment is usually co-ordinated through clinic visits as indicated in a separate window 'VISIT' in the OaSiS interface (figure 7). By pressing the 'Next visit' button a new visit date can be generated according to the schedule specified in the protocol. A mouse click on a new date causes a summary of the purpose of the visit to be generated (figure 10).

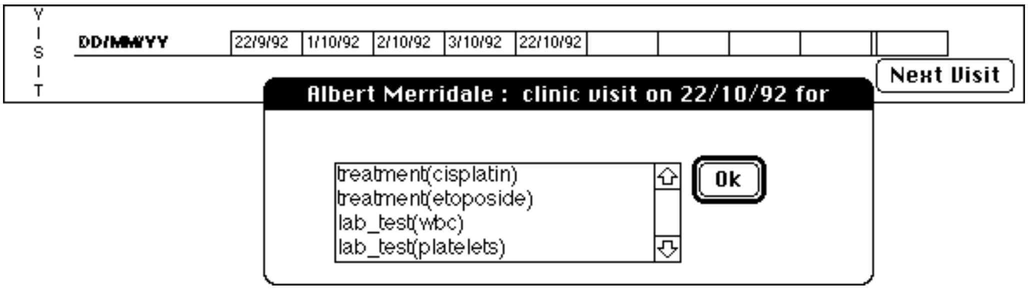

Albert Merridale : clinic visit on 22/10/92 for

treatment(cisplatin)
treatment(etoposide)
lab_test(wbc)
lab_test(platelets)

FIGURE 10. Querying purpose of next clinic visit

Figure 10, for example, indicates that the patient is due to receive two drugs as part of chemotherapy and that two blood test results are also required. The chemotherapy component of a trial will include rules about modifications to drug dosage according to the results of these blood tests and other factors. Unless the oncology clinic is connected directly to a laboratory computer, the data will need to be entered by hand. The haematological data is presented in a separate window indexed by the date of the clinic visit when blood specimens etc were acquired. New data entered in the HAEM spread-sheet style window is not recorded permanently until the user requests it. When a new clinic visit date has been generated and accepted, and the relevant laboratory results have been added, OaSiS automatically computes and displays drug dosages.

We use the predicate

```
dosage_modification_factor(<Chemotherapy>, <CycleDay>, <Drug>, <Patient>, <Factor>)
```

to define the drug modification <Factor> for <Patient> being treated on day <CycleDay> with <Drug> during a particular <Chemotherapy> regimen. For example, the TE09 protocol [14] requires the drug dose to be unchanged (i.e. <Factor> = 1) provided that toxicity (as measured by blood tests) remains within certain bounds:

```
dosage_modification_factor(regimen(te09, bep, Version), day(Cycle, Day), etoposide, Patient, 1) :-
        Cycle > 1,
        test_result_on_cycle_day(wbc, Patient,  te09, bep, day(Cycle, 1), WBC),
        WBC >= 2000,
        test_result_on_cycle_day(platelets, Patient, te09, bep, day(Cycle, 1), PlateletCount),
        PlateletCount >= 90000.
```

In other circumstances it needs to be reduced by 25%:

```
dosage_modification_factor(regimen(te09, bep, Version), day(Cycle, Day), etoposide, Patient, 0.75) :-
        Cycle > 1,
        test_result_on_cycle_day(platelets, Patient, te09, bep, day(Cycle, 1), PlateletCount),
        in_range(PlateletCount, (51000,89000)).
```

The results of these dose calculations are presented automatically to the user in the CHEMO spread-sheet style window. In Figure 11, we can see the result of applying these rules to the calculation of the latest modified dosage for the drug etoposide where the default value of 240 has been modified to 180 because of the results of blood tests shown in the HAEM spreadsheet:

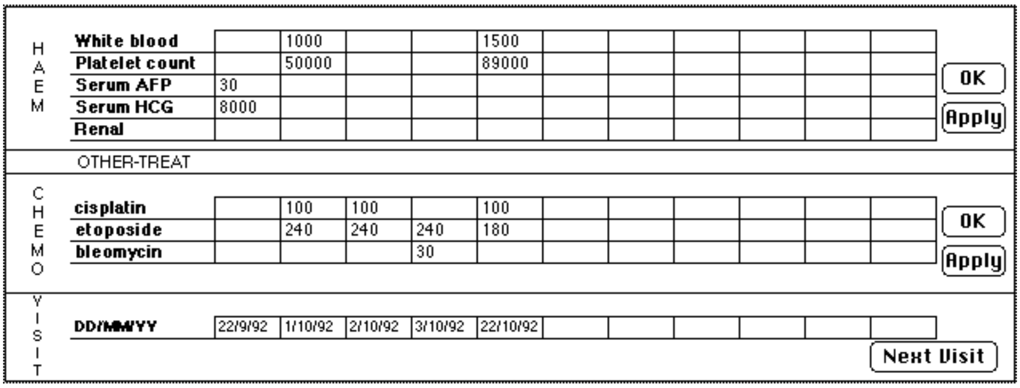

HAEM	White blood		1000			1500							
	Platelet count		50000			89000							
	Serum AFP	30											
	Serum HCG	8000											
	Renal												
	OTHER-TREAT												
CHEMO	cisplatin		100	100		100							
	etoposide		240	240	240	180							
	bleomycin				30								
VISIT	DD/MM/YY	22/9/92	1/10/92	2/10/92	3/10/92	22/10/92							

FIGURE 11. Dose modification of etoposide (240 to 180) on 22/10/92 due to toxicity.

Table 1 earlier illustrated how treatment modifications can be at various levels of severity and for different periods of duration. For example, individual drug administrations can be modified to a smaller or larger dosage, either singly or in groups, as well as being suspended temporarily or even aborted altogether. The duration of modifications can either be absolute (typically a week or so) or relative in terms of treatment cycles or until the patient's reaction to toxicity reaches some acceptable state.

2.6 Negotiating a safe treatment

OaSiS suggests treatments and their modification according to accepted clinical practice and the protocol governing the trial under which the patient is being treated. If the clinician disagrees then it must be possible to suggest an alternative. Of course, checks on such modifications will be made. The clinician may also need to record additional treatments, for example for illnesses unrelated to the malignancy or for side-effects. The proposed medication should obviously be validated against known contraindications, especially those specified in the protocol.

Next, we illustrate two of the safety principles mentioned earlier which have been extracted following an intensive study of many oncology protocols. An instance of the first safety rule is the following:

Ex R5 Nephrotoxic antibiotics such as Gentamycin should be avoided during the Cisplatin infusion. [3]

which is quoted verbatim as it appears in the protocol. Similar examples occur in many oncology protocols. Underlying them is a general principle that actions outside the recommended treatment plan are barred if they are likely to exacerbate an existing hazard arising from the treatment itself. It is such general principles that analysis of protocols has identified and which we seek to represent in OaSiS, in as general a fashion as possible. Informally, the principle underlying Ex R5 can be expressed as:

(R5) *Action1* should not be performed during *Action2* in *Plan* if
Action2 is necessary part of *Plan* and
Action2 produces *Effect* and
Effect is potentially hazardous and
Action1 aggravates or makes *Effect* more likely and
Action1 has alternative without *Effect*

The actual formulation (as an integrity constraint) will be shown in a moment. The specific information contained in Ex R5 is represented as a set of facts (some of which are implicit in its original formulation) to be used in conjunction with principle (R5), thus:

```
produces_effect(administration_of(cisplatin), impaired_renal_function).
hazardous(impaired_renal_function).
aggravates(administration_of(gentamycin), impaired_renal_function).
```

These safety restrictions are employed in OaSiS in two main ways: to generate a warning message immediately certain users consult OaSiS, and as integrity constraints to check any additional prescriptions that might be made during the actual consultation itself.

A second example of safety knowledge concerns the barring of actions that reduce efficacy of treatment. We have been able to formulate a general rule which, informally, can be stated as follows:

(R6) *Action1* should not be performed during *Action2* in *Plan* if
Action2 is necessary part of *Plan* and
Action1 reduces efficacy of *Action2* and
Action1 is unnecessary as part of *Plan* or has an alternative

Two instances of its use come from a trial involving the drug Interferon-a2α for the treatment of colorectal cancer [31]:

Ex R6a Aspirin may reduce the efficacy of interferon-a2α. Therefore, aspirin will not be used. Indomethacin does not have this effect.

Ex R6b Corticosteroids may reduce the efficacy of interferon-a2α and should not be prescribed for side-effects.

These specific instances are again represented as facts to be used in conjunction with (R6):

```
reduces_efficacy(administration_of(aspirin), administration_of(interferonα)).
reduces_efficacy(administration_of(corticosteroids), administration_of(interferonα)).
```

Rules such as (R5) and (R6), shown informally above, are represented as integrity constraints on the database recording the negotiated treatment plan, in the usual logic database style (see, e.g., [12]). For illustration, rules (R5) and (R6) are represented as follows:

```
(ICR5)   invalid(user_suggestion(perform(Action1), Plan) ) ←
                 part_of(Action2, Plan) and
                 produces_effect(Action2, Effect) and
                 hazardous(Effect) and
                 aggravates(Action1, Effect) and
                 is_avoidable(Action1, Plan).

(ICR6)   invalid(user_suggestion(perform(Action1), Plan) ) ←
                 part_of(Action2, Plan) and
                 reduces_efficacy(Action1, Action2) and
                 is_avoidable(Action1, Plan).
```

Although it is usual to write integrity constraints in the form of denials, we find it more convenient to employ the form shown above. The reason is simply that all conditions except for user_suggestion are static, in the sense that they are stored in parts of the OaSiS system which the user does not modify during a consultation with the system. Here invalid can be read as an alternative symbol to ¬ for (standard, truth-functional) negation. It also has an operational meaning, to signal that in the case of violation of the constraint it is this condition – i.e., the attempted input – which is to be rejected. Any integrity constraints that are violated can be presented in the form of a critique of the user's action as arguments for and against (see fig 13). Here the user has requested assistance with treating a fever and has then asked for an explanation as to why gentamycin is contraindicated. The argument against is generated from (ICR5).

We have illustrated only two of the half a dozen or so generic safety rules we have found [13]. The others are concerned with avoidance of hazardous side-effects by applying suitable prophylactic treatments; warnings about incorrectly performed treatments; recommendations to perform adequate monitoring of side-effects and actions to ameliorate toxic effects as soon as they have been observed.

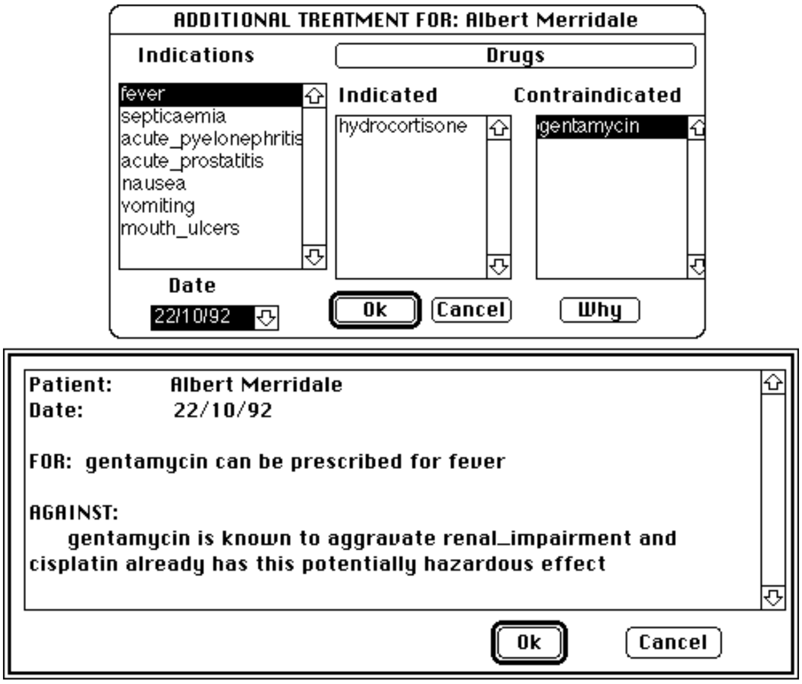

ADDITIONAL TREATMENT FOR: Albert Merridale

Indications: fever, septicaemia, acute_pyelonephritis, acute_prostatitis, nausea, vomiting, mouth_ulcers

Drugs

Indicated: hydrocortisone

Contraindicated: gentamycin

Date: 22/10/92

Patient: Albert Merridale
Date: 22/10/92

FOR: gentamycin can be prescribed for fever

AGAINST:
gentamycin is known to aggravate renal_impairment and cisplatin already has this potentially hazardous effect

FIGURE 13: Arguments for and against selection of a particular treatment

3 WORK IN PROGRESS AND PLANNED

3.1 Improvements to the user interface and graphical representation of data

Graphical presentation and entry of data is likely to be of considerable benefit to clinical users. With a suitable internal model of simple human anatomy, multiple views of tumour incidence should be displayed as well as more rapid or transient presentations of associated quantitative data. Alternatively, tumour incidence/dimensions could be presented graphically alongside haematological data or chemotherapy and indexed by clinic visit so that treatment progress can be rapidly reviewed. Graphical overviews of important clinical events are more usefully presented in terms of time lines with clear separation of treatment in the past, present and that planned for future clinic visits. Later, we shall also be considering the production of aids for clinicians to define their own protocols for in-house use and employing a graphical interface along the lines of the OPAL protocol editor implemented at Stanford University [25].

3.2 Representational formalism and the treatment of obligations

Further development of the representational formalism includes refinement of the temporal aspects; as identified earlier, these are currently done by explicit time-stamping. We are currently re-implementing the temporal behaviour of OaSiS in terms of an interval-valued temporal logic developed within the RED project [5]. This is intended to cover all areas where time needs to be explicitly manipulated:

scheduling of high level protocol activities and tasks in general
monitoring of toxicity and modifiable treatment planning
evaluation of treatment response criteria and planning of treatment follow-up

For the purposes of this paper, however, we wish to focus on the developments concerning the representation of safety and other constraints, especially as regards the treatment of notions of obligation that arise.

Although the current design of OaSiS separates constraints into two categories – data integrity constraints and those dealing with safety knowledge – the same treatment is given to both. They use the same underlying logic and generate the same behaviours; only the form of messages to the user differs.

We believe it is important to make finer distinctions according to the nature of the constraints to be supported in such systems, in particular, in order to deal more explicitly with the tension between compliance with the protocol and the need to allow clinicians some degree of flexibility or discretion. This tension is characteristic of any system where agents' behaviours and interactions are regulated by norms which specify how those agents ought to behave and how they are permitted to behave, whilst leaving open the possibility that actual behaviour may deviate from what is prescribed. The formal tool associated with the characterisation of such systems is deontic logic: for further development of OaSiS we will distinguish between database constraints that cannot be violated and a class of 'softer' *deontic* constraints whose violation can be tolerated. The application of deontic logic to integrity constraints in databases has been discussed in the literature (see [4, 15, 18, 36]) though proposals vary in scope and emphasis, and the techniques remain comparatively undeveloped; OaSiS provides an excellent example to drive these developments in a practical setting. The explicit representation of notions of obligation in safety-critical systems has also been addressed within the RED project [6, 9].

Although deontic logic is ordinarily described as the logic of obligation and permission, it can be misleading to rely on this description. The reason is simply that words such as 'obligatory', 'must', 'should' can have a variety of meanings. In [15] it is suggested that for the purpose of identifying practical applications, deontic logic is better regarded as the logic which deals with the distinction between what ought ideally to be the case on the one hand, and what actually is the case on the other.

Where a (safety) constraint is such that non-compliance is potentially life-threatening, it is represented appropriately by a standard 'hard' integrity constraint which allows no violation. Consider, for instance, the safety constraint

> If anaphylaxis (severe allergic reaction) occurs, Cisplatin must be stopped [24].

The 'must' is here adequately represented as the form of necessity already provided by the standard, 'hard' integrity constraint.

Obviously the possibility exists that such a constraint can be violated in the real world. One can easily imagine circumstances in which a clinician continues to administer a drug even though this should not, or must not, be done according to the protocol. However, in the context of the OaSiS application, the question is not whether such violations of constraints can occur in the real world, but whether we want to allow the possibility of violation into the database. Where safety constraints are represented as standard 'hard' database constraints the OaSiS system will not accept any attempt to input data which contravenes them, by the standard integrity checking mechanisms. Such a representation effectively builds compliance with safety constraints into the system (cf. the discussion of 'regimentation' in [15]). Of course, this mechanism does not *guarantee* compliance with safety constraints, because the actual behaviour of clinicians is outside the control of the system – it is possible that a clinician could enter one thing into the database but actually do something different. But subject to the assumption that clinicians act in good faith and allow themselves to be guided by the system, compliance with the safety constraint is guaranteed in a sense. And surely this is a realistic assumption, especially since attempted contraventions of safety constraints are accompanied in OaSiS by indications of the consequences of breach.

Compare now the next two examples, both taken from a protocol dealing with the treatment of testicular cancer [3]:

> Nephrotoxic antibiotics such as Gentamycin should be avoided during the Cisplatin infusion.
>
> The Bleomycin-Cyclophosphamide-Dactinomycin regimen will produce myelosuppression and it is **essential** to have a nadir count between days 7 and 12 and to warn the patient to report symptoms of infection or bleeding ...

We understand that for the first of these some degree of discretion would be desirable in practice. There are conceivable circumstances in which a clinician might choose to contravene this guideline and the system should then allow such non-prescribed treatments to be recorded. This constraint is appropriately represented as a deontic integrity constraint. In the second example, we are unsure about the status of 'essential'; this would need to be checked with a clinical oncologist and the constraint represented as a 'hard' or 'soft' integrity constraint accordingly.

Informally, the operational behaviour and the implementation of the deontic constraints (of the kind sketched here) is straightforward. Attempted violations are detected in identical fashion to standard integrity constraints; a warning message is displayed to the user, but freedom to over-ride the recommendation is available, subject to satisfaction of other constraints. In practice it will usually be necessary to maintain an audit trail so that treatment which deviates from the protocol can be monitored and subjected to later analysis, but this feature can be provided without difficulty.

The point of introducing a special deontic-logical component into the representation formalism is not just that it provides a cleaner conceptual framework for the constraints arising in the application. The interactions between constraints, safety and other, can be quite complex in practice, and it becomes more and more difficult to keep track of these interactions as the set of constraints in a given application grows. It is obviously desirable to check for consistency among constraints, to simplify them if possible, and generally to subject them to other meaning-preserving transformations, independently of their application to a specific database. For instance, it seems clear that the sets of 'hard' and 'soft' integrity constraints should both be internally consistent, but it is also necessary to ensure that the two sets do not conflict: if there is a 'hard' constraint requiring that A must be in the database, then there should be no 'soft' constraint requiring that A should not be in the database; any set of constraints not satisfying this property would ordinarily be regarded as ill-formed at least. This property is an instance of the kind of 'ought implies can' principle common in the formulation of many deontic logics.

There are good reasons to think that a simple form of deontic logic will prove adequate for dealing with the type of constraints encountered in OaSiS, although this statement is still subject to some investigation. In particular, once 'soft' constraints and the possibility of violations are admitted, then it is natural to consider what other constraints come into effect in those circumstances. The proper formalisation of 'contrary-to-duty structures' – roughly, the situation where there is a primary obligation and a secondary obligation which comes into effect when the primary obligation is violated – has proved notoriously elusive and is the subject of much current research in the field of deontic logic. (See [28] for further discussion and references to the standard works in the deontic logic literature; this paper also makes some comparisons between contrary-to-duty and the more familiar exception structures studied in default and non-monotonic reasoning.) It remains to be seen to what extent contrary-to-duty constraints need to be treated in the OaSiS applications.

We have motivated the use of deontic constraints by reference to safety constraints, but we do not want to give the impression that deontic constraints are applicable only to the representation of safety knowledge. Deontic logic can also be applied – less commonly and less urgently perhaps – to what we have called data integrity in this paper. In this respect we disagree with the position expressed in [36] where it is proposed that database constraints can be classified (roughly) according to whether they correspond to statements which can be violated in the real world or to truths of the real world, which cannot be violated. Much of the discussion of this section has been concerned with indicating that norms which can be violated in the real world are often appropriately represented in a database by 'hard' integrity constraints which cannot be violated; and, conversely, we are in agreement with, e.g. [4], who point out that truths of the real world need not always be represented as necessary ('hard') constraints on a database since we may allow, for instance, that our representation is incomplete.

It may also be desirable to pick out other categories of constraints, such as those concerned with ensuring the scientific validity of the data collected during the course of treatment. We mean by this that some constraints are included in protocols, not because contravention would constitute a life-threatening hazard or affect the efficacy of treatment, but because it would compromise the intended contribution to the objectives of the clinical trial being conducted. We can make a case then for three categories of constraints – for data integrity, for safety knowledge, for ensuring scientific validity of trial results; within each category, any given constraint would be represented either as a hard or a soft database constraint according to the nature of the necessity – or the 'strength' of the constraint – to be represented.

3.3 Consideration of clinical trial recruitment

It is theoretically possible for OaSiS to be used to screen clinical databases for patients who might be eligible for entry into particular trials. With completely automated screening, situations would arise where patients fail criteria even though satisfaction could be ensured if small changes to existing treatment or if delays in consideration of trial entry could be made. For example, a patient may be receiving a drug for a non-malignant condition that is explicitly contraindicated in the trial protocol. There may be an alternative treatment that is not contraindicated or it might be possible to delay consideration of trial entry until the non-malignant condition has cleared up - provided the clinical judgement is that such a delay is consistent with the patient's best interests. Similarly, there will often be a lack of up to date haematological or biochemical data needed for eligibility consideration. In some situations it might be possible to reason about the persistence of previously obtained data and estimate the likelihood that the patient could be eligible and that such considerations need to be pursued further.

In some medical domains, such as in trials of new drugs in AIDS therapy, successful trial recruitment is very important. The Stanford group have recently been applying statistical and related techniques to the consideration of trial recruitment for this very reason in a project called T-Helper [33], a follow-on activity from work on the ONCOCIN system. We shall investigate the same problem but armed with different tools of analysis which exploit more obviously symbolic approaches to identify the qualifications [35, 37] necessary for eligibility. This approach can be viewed as the generation of incomplete arguments which need to be "repaired" and subsequently structured in some framework of logical uncertainty [7] so that eligibility can be suitably qualified.

ACKNOWLEDGEMENTS

The development of OaSiS has been undertaken in the RED project which is funded under the UK DTI and SERC safety-critical systems research programme. Colleagues at ICRF and collaborators in the RED project have been very helpful in providing comments and suggestions for improving the content of this paper and earlier drafts of the material presented. In particular, thanks are due to Subrata Das, John Fox, Colin Gordon, Paul Krause and Jeremy Wyatt. The authors are grateful to Adrian Harris of the ICRF Clinical Oncology Unit at the Churchill Hospital in Oxford for comments and suggestions in relation to safety and clinical issues in oncology.

REFERENCES

1. ABPI Data Sheet Compendium 1993-94, Datapharm Publications Ltd, London.

2. British National Formulary, British Medical Association and Royal Pharmaceutical Society, London, 1992.

3. BO03 (1986) A randomised trial of two chemotherapy regimens in the treatment of operable osteosarcoma, MRC Cancer Trials Office, 1 Brooklands Ave, Cambridge CB2 2BB.

4. Carmo, J. and Jones, A., Deontic Database Constraints and the Characterisation of Recovery, in: *Proc. Second International Workshop on Deontic Logic in Computer Science (DEON'94)*, Oslo, Tano Publishers, Norway, 1994.

5. Das, S. K. and Hammond, P., An interval-valued temporal modal logic., in: M. A. Bramer and A. L. Macintosh (eds.), *Research and development in Expert Systems X* Technical proceedings of Expert Systems'93, Cambridge; BHR Group Ltd and BCS SGES, 1993.

6. Das, S. K. and Fox, J., A logic for reasoning about safety in decision support systems, 2nd European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty, 8-10 November, 1993, Granada, Spain.

7. Elvang-Gørenson, M., Krause, P. J. and Fox, J., Acceptability of arguments as "logical uncertainty". 2nd European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty, 8-10 November, 1993, Granada, Spain.

8. Fischer, D. S., Tish Knobf, M. and Durivage, H. J. The cancer chemotherapy handbook. Mosby, St Louis, 1993.

9. Fox, J., Engineering safety into expert systems, in: F. Redmill and T. Anderson (eds.) *Directions in Safety-Critical Systems*, Chapman and Hall, 1993.

10. Gaglio, S., Ruggiero, C., Spinelli, G., Bonadonna, G., Valagussa,. P and Nicolini, C., BREASTSCAN: An expert system for postoperative breast cancer therapy, *Computers and Biomedical Research* 19: 5 (1986).

11. Gordon, C., Herbert, I., Renaud-Salis, J-L. and Jackson-Smale, A., *Care Protocols and Healthcare Informatics*. (1993).

12. Grant, J. and Minker, J., Integrity constraints in knowledge-based systems, in: H. Adeli (ed.), *Knowledge Engineering Applications Vol III* , McGraw Hill, 1990.

13. Hammond. P., Harris, A. L., Das, S. K. and Wyatt, J. C., Safety and decision support in oncology. *Meth Inf Med*, in press.

14. Horwich, A., A prospective randomised comparison of Carboplatin and Cisplatin in combination chemotherapy for good prognosis metastatic non-seminomatous germ cell testis tumours. MRC Cancer Trials Office, 7 Green Street, Cambridge CB2 3JU, UK, 1989.

15. Jones, A.J.I., Sergot, M.J. "On the Characterisation of Law and Computer Systems: The Normative Systems Perspective." Ch 12 in Deontic Logic in Computer Science: Normative System Specification (Meyer, J.-J.Ch., Wieringa, R.J., eds), Wiley (1993).

16. Kalet, I. J., Jacky, J. P., Knowledge-Based Computer Simulation for Radiation Therapy Planning. Proc of 9th Int. Conf. on the Use of Computers in Radiation Therapy. The Hague, The Netherlands, 22-25 June 1987. North-Holland, Amsterdam, 553-556.

17. Kent, D., Shortliffe, E. H., Carlson, R., Bischoff, M. and Jacobs, C. Improvements in data collection through physician use of a computer-based chemotherapy treatment consultant. *J Clin Oncol* 3: 1409-1417 (1985) .

18. Kwast, K.L., A Deontic Operator for Database Integrity, in: *Proc. First International Workshop on Deontic Logic in Computer Science (DEON'91)*, 262-280, 1991.

19. Lane, C. D., Walton, J. D. and Shortliffe, E. H., Graphical access to expert systems:II. Design of an interface for physicians, *Meth Inform Med* 25: 143-150 (1986).

20. Leaning, M. S., Ng, K. E. H. and Cramp, D. G., Decision support for patient management in oncology. *Med Inform* 17(1): 35-46 (1992).

21 Lhommé, C., Giron, A., Jan, P., Wartelle, M., Michel, G., Prade, M., Duvillard, P., Castaigne, D., Rochard,F.,Haie-Meder,C.,Gerbaulet,A.,George,M., Le Beux,P. and Flamant,R., Pénélope, an expert system for diagnosis and therapy of ovarian adenocarcinoma, *Bull Cancer* 79:1055-1070 (1992).

22. LPA Ltd (1992) MacPROLOG 4.5, Logic Programming Associates Ltd, Studio 4, The Royal Victoria Patriotic Building, Trinity Road, London SW18 3SX.

23. Mamelle, G., Richard, J. M., Wartelle, M. and Jan, P., Decision support with expert system in head and neck cancer (abstract). AIME’91, University of Limberg, Maastricht, the Netherlands, 1991.

24. Mead, G. M. and Kaye, S. B. A Medical Research Council/EORTC collaborative randomised trial of BEP/EP vs BOP/VIP-B with or without G-CSF in the management of poor prognosis metastatic and extragonadal teratoma. MRC Cancer Trials Office, 1 Brooklands Ave, Cambridge CB2 2BB (1991).

25. Musen, M. A., Combs, D. M., Walton, J. D., Shortliffe, E. H. and Fagan, L.M., OPAL: Toward the computer-aided design of oncology advice systems. Knowledge Systems Laboratory Memo KSL-86-49, Medical Computer Science, Stanford University, Stanford, California 94305-5479, (1986). Also Proceedings of Computer Applications in Medical Care. Washington, 43-52, 1986.

26. Musen, M. A., Tu, S. W. and Shahar, Y. A problem-solving model for protocol-based care: from e-ONCOCIN to EON. Knowledge Systems Laboratory Memo KSL-91-61, Medical Computer Science, Stanford University, Stanford, California 94305-5479, (1991). Also proceedings of MEDINFO'92 Geneva, Switzerland, September 1992.

27. Nariainen, K., Ojala, A., Rantanen, J. and Saranummi, N. CARTES - a knowledge system for decision support, therapy design and quality assurance in cancer therapy [Preprint] Tampere, Finland. Technical Research Center of Finland, (1987).

28. Prakken, H. and Sergot, M. J., Contrary-to-duty imperatives, defeasibility and violability, in: *Proc. Second International Workshop on Deontic Logic in Computer Science (DEON'94)*, Oslo, Tano Publishers, Norway, 1994.

29. Renaud-Salis, J-L. and Taylor, P., The Bordeaux Oncology Support System: Knowledge representation and prototype. ICRF Internal Report No 97, (1990).

30. Sergot, M. J., A Query-the-user facility for logic programming, Integrated Interactive Computer Systems, in: P Degano and E Sandwell) (eds.), North Holland, 1983.

31. Seymour, M. T. and Slevin, M. L., A randomised trial of 5FU and leucovorin versus 5FU, leucovorin and interferon-a2α for advanced colorectal adenocarcinoma. MRC Cancer Trials Office, 1 Brooklands Ave, Cambridge CB2 2BB, (1991) .

32. Shortliffe, E. H., Scott, A. C., Bischoff, M. B., Campbell, A. B., Van Melle, W., Jacobs, C. D., ONCOCIN: An expert system for oncology protocol management. Proc IJCAI'81, 876-881, (1981).

33. Tu, S. W., Kemper, C. A., Lane, N. M., Carlson, R. W. and Musen, M. A., A methodology for determining patient's eligibility for clinical trials, *Meth Inform Med* 32 (4): 317-325 (1993).

34. Vantongelen, K., Steward, W., Blackledge, G., Verweij, J. and Van Osterom, A., EORTC Joint Ventures in Quality Control: Treatment-related Variables and Data Acquisition in Chemotherapy Trials. *Eur J Cancer* 27 (2): 201-207, (1991).

35. Vasey, P., Qualified answers and their application to transformation. Proc 3rd International Conference on Logic Programming, July 1986, London, UK. (Ed E Shapiro), Springer-Verlag.

36. Wieringa, R. J., Meyer, J.-J. Ch. and Weigand, H., Specifying Dynamic and Deontic Integrity Constraints, *Data and Knowledge Engineering* 4: 157-189, (1989).

37. Wolstenholme, D. E., (1990) External data in logic-based advice systems. PhD Thesis. Department of Computing, Imperial College, 180 Queen's Gate, London SW7 2BZ.