Comput. Math. Applic. Vol. 14, No. 9–12, pp. 793–802, 1987
Printed in Great Britain

0097-4943/87 $3.00 + 0.00
Pergamon Journals Ltd

EXPERT SYSTEMS FOR CANCER CHEMOTHERAPY

S. GAGLIO
Institute of Electrical Engineering, University of Palermo, Italy

M. GENOVESI, C. RUGGIERO and G. SPINELLI
Department of Communications, Computer and System Science, Genoa University, Italy

C. NICOLINI
Chair of Biophysics, School of Medicine, Genoa University, Italy

G. BONADONNA and P. VALAGUSSA
National Cancer Institute, Milan, Italy

Abstract—In the present paper we describe BREASTCAN and NEWCHEM, two expert systems for the characterization of optimal adjuvant cancer therapies. The purpose of BREASTCAN is to support physicians in the postoperative breast cancer therapy, on the basis of currently used therapy protocols. It was developed in Prolog and positively validated, referring to the chemotherapies used by oncologists for some patients in the National Cancer Institute in Milan. NEWCHEM is a system oriented to the development of new cancer therapies, based on pharmaco-cell kinetic modeling and the newest molecular knowledge about neoplastic process. The system is being built and at first it will be validated by experiments on mice. Our aim with NEWCHEM is to extend our knowledge base and our rules to incorporate also all the most advanced knowledge at the molecular and cellular level, both theoretical and experimental, to make readily accessible to the health-community a system which will be really as expert as the present state of the art allows.

1. INTRODUCTION

The use of AI (artificial intelligence) techniques, with its logical-deductive mechanisms, that reproduce the human reasoning without the restriction of human capabilities in memory and association, is particularly suitable for those medical problems, in which you must regard a great number of prognoses and their effects and interactions. This is the case of cancer chemotherapy, used for 10 years with encouraging results.

In the present paper we describe BREASTCAN and NEWCHEM, two expert systems for the characterization of optimal adjuvant cancer therapies. They are both implemented on a VAX 11/750 at the Department of Communications, Computer and System Science of the University of Genoa, with the co-operation of the Chair of Biophysics of School of Medicine of the Genoa University and the National Cancer Institute in Milan.

The purpose of BREASTCAN is to support physicians in the postoperative breast cancer therapy, on the basis of currently used therapy protocols. It was developed in Prolog and positively validated, referring to the chemotherapies used by oncologists for some patients in the National Cancer Institute in Milan. In BREASTCAN the knowledge is represented using *production rules*, that are structures whose form is IF ... THEN, and *frames*, complex data structures for modeling stereotyped situations. Every frame describes a certain situation or hypothesis.

NEWCHEM is a system oriented to the development of new cancer therapies, based on pharmaco-cell kinetic modeling and the newest molecular knowledge about neoplastic process. The system is being built and at first it will be validated by experiments on mice. It is constituted by a set of *frames*, a *deduction rule* system, an *agenda-based* control system, an hierarchical *planning* system, a data base about drugs, a *qualitative reasoning* system, mathematical models for simulation and optimal control of phamaco-enzyme (Michaelis–Menten) and pharmaco-cell processes.

2. BREASTCAN: AN EXPERT SYSTEM FOR POSTOPERATIVE BREAST CANCER THERAPY

Chemotherapy as adjuvant treatment in patients with operable breast cancer and histologically positive axillary nodes has been used for about 10 years, achieving results that support the use of adjuvant chemotherapy in clinical practice [1–3].

Numerous clinical trials have been designed and activated, employing adjuvant chemotherapy alone or combined with endocrine or immunotherapy. Both study design and results of adjuvant protocols are influenced by a great number of prognostic variables (e.g. subgroups of patients with different size of primary tumor (T) and extent of axillary node (N) involvement, presence or absence of hormone receptors, quantities of drugs to be administered and intervals between each drug and treatment cycles, combinations of drugs, optimal treatment duration) whose significance and possible interactions need to be assessed, as well as by a great variety of clinical problems often originating from chemotherapy during the postoperative course of breast cancer (e.g, treatment discontinuation because of toxicity, continuous low drug dosage, different forms of salvage regimen, occurrence of general medical problems not related to adjuvant therapy).

It is therefore felt that AI techniques [4] could make valuable contributions in this field. In particular it would be worth while to set up an expert system to assist physicians in the selection and performance of optimal treatment—maximizing therapeutic effectiveness and limiting toxicity. The use of the expert system should facilitate an optimal decision making at each treatment cycle by taking into account all foreseeable medical variables.

An expert system for oncology protocol management, focused on the treatment of Hodgkin's Disease and the non-Hodgkin's lymphomas, has been recently set up [5]. The present paper describes an expert system for adjuvant breast cancer chemotherapy—based on few widely accepted clinical protocols [2]—which has been implemented to assist clinicians treating patients by standard chemotherapy.

2.1. Outline of BREASTCAN

BREASTCAN is written in Prolog [6]—a programming language that has been used in many areas of AI research—that has been chosen because it allows to quickly write clear, concise and readable programs, and provides an efficient built-in deduction system.

BREASTCAN is written employing *production rules*—that is general statements about objects and their relationships—and *frames*—that is data structures that model stereotype situations.

Production rules encode single pieces of expert knowledge and are written as Prolog clauses. They specify particular solutions to a given problem according to given conditions. As an example, the following BREASTCAN rule

```
                    rule(4,3,ctx,2,Time),300,Atruth,10):—
f(look_for(platelets,2,Time,Platelets,Truth1)),
  Platelets < 100000,
f(look_for(leukocytes,2,Time,Leukocytes,Truth2)),
  Leukocytes < 3800,
f(look_for(n_delays,2,Time,Ndelays,10)),
  Ndelays = 2,
min(Truth1,Truth2,Atruth).
```

establishes that the suggested dose of Ctx (cyclophosphamide) is 300 mg/m^2 if platelets count is less than 100,000, leukocytes count is less than 3800, and two consecutive delays have already occurred in drug administration. A rule is characterized by a truth value which is combined with truth values resulting from the satisfaction of antecedent conditions to provide a measure of evidence of the conclusion.

Production rules have been grouped, in BREASTCAN, into frames. Each frame corresponds to a given situation or hypothesis. One frame (named "initial") contain the rules establishing whether the patient is eligible for chemotherapy or not; another frame (named "registration") manages the description of the patient pathology prior to chemotherapy, to be used for adjuvent therapy choice and to be recorded for subsequent statistical studies that the user may want to carry out, and recommends one treatment protocol, each treatment protocol specifying one chemotherapy cycle. The user can then select one treatment protocol (accepting the system recommendation or not). At every stage the user can obtain explanations from the system, that can describe the pathway followed to reach any stage of consultation.

Frames are organized in a tree-like structure as shown in Fig. 1.

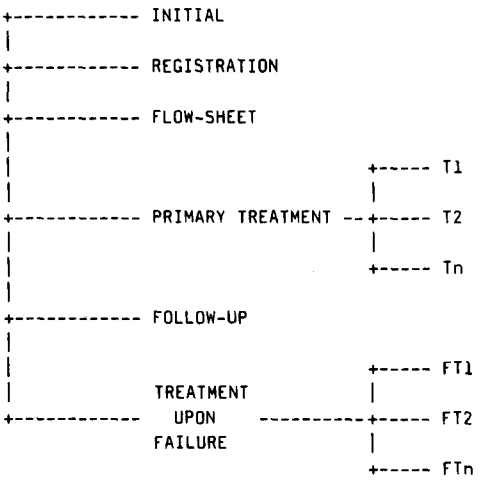

Fig. 1

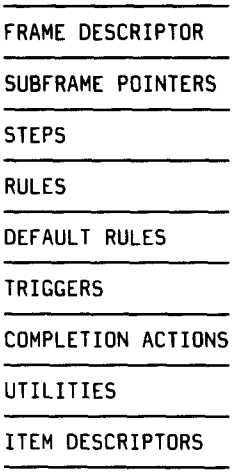

Fig. 2

The treatment protocols are subframes of one frame (named "treatment") that takes into account the stage of the treatment, the possible presence of pathologies developed during it, and suggests (activating one of the subframes) the drugs to be administered and their quantities, and the date of the next drug administration.

2.2. Control structure

The control system activates and executes frames. A frame can be in one of the following states:

—inactive;
—active (that is ready for execution);
—current (in execution);
—completed.

The control system provides a number of procedures not available directly in Prolog, designed to assist consultation. Frames are selected for instantiation and execution according to an "applicability value", that evalutes the agreement between the specific situation considered and the situation to which the frame relates; this is often accomplished comparing, for each frame, the postsurgical conditions characterizing the frame with the patient conditions.

2.3. Frame structure and execution

Frames consist of a sequence of steps, rules and facts. Steps consist of groups of statements that are executed sequentially, and steps are executed sequentially. Some steps are conditioned, that is their execution depends on a condition to be tested: if the condition is not true, the step is not executed and the control passes to the following step. The steps of a frame encode the procedural knowledge ("what to do") of the stereotype situation modelled by the frame. The rules of a frame encode the declarative knowledge (what is true and "how to do" things within the given situation). The facts are pairs item—value that encode what is known about particular situations (measures and information inferred by the system. Each value is provided with a descriptor containing the way the value has been determined (asked, deduced, default) and certainty factor, which expresses the extent of certainty associated with the value. Certainty factors are propagated and modified in the course of rule application as described in Buchanan and Shortliffe [7].

The structure of a frame is shown in Fig. 2.

The *frame descriptor* contains some general information about the frame, as its name, the number of steps, its possibility, and the condition for activation.

Subframe pointers indicate the subframes of the frame. *Triggers* are forward rules that perform particular actions when critical values of the parameters are added to the set of items.

Completion actions are performed when the frame enters the completed status, for instance the activation of another frame.

Utilities are auxiliary procedures used within the steps.

Item descriptors encode the constraints to be satisfied by the values of the items, the questions to be asked of the user concerning the items, and take care of data integrity (for instance, menopause age cannot be greater than the current age of the patient).

When, during the execution of a frame, the value of an item is needed, the system first consults its data base (the pairs item—value in the frames); if it does not find it, it tries to apply the rules of the frame to deduce it; finally, if it fails, it asks the user the value; when the answer is not even provided by the user, the system uses default rules and values.

2.4. Interfacing with the user

BREASTCAN offers the possibility to obtain indications about the line of reasoning that leads backwards from the current answer by the system. This capability is useful in many respects: realizing how the system comes to a conclusion enables users to make the most of the consultative advice and also helps users in difficult, "non-standard" situations, in which they may wish to violate a rule; moreover, it makes it easier for the user to make changes in the program.

When starting consultation about a patient for the first time, the user is asked questions that enable the system to decide whether chemotherapy is appropriate for that patient. If the answer is no the system gives indications about the reasons of this conclusion and some advice about that particular case. At the end of this stage the user can choose whether he wishes to carry on or not (irrespective of the conclusions of the system about the eligibility of the patient). The following stage, relating to the description of the patient pathology prior to chemotherapy, consists of six parts; at the end of each of them the data entered by the user in that section are displayed and the user is enabled to change values that he may have entered by mistake.

The system now suggests one or more treatments and the user can choose which treatment he wishes to perform (following the suggestions by the system or not).

2.5. Assessment and further developments

BREASTCAN has been proven capable to closely reproduce the actual decision taken over the period of 6 months by medical oncologists at the National Cancer Institute of Milan, Italy, in the treatment of a few patients, taken as test-cases in a preliminary retrospective confirmation of the system validity.

Our next step is then to extend this study to a large number of patients and to incorporate into the decision-making process all other possible strategies suggested by other alternative treatment protocols from worldwide clinical trials on the breast cancer.

3. NEWCHEM: AN EXPERT SYSTEM FOR THE TREATMENT OF DISSEMINATED CANCER

Since early 1974 a comprehensive program was developed at the Biophysics Division of Temple University in Philadelphia (U.S.A.), aiming to develop a rational basis for the chemotherapy of cancer [8]. This approach resulted in a widely interdisciplinary effort which, by a constant feedback between experimentation on animals and theoretical simulations, has been able to predict and explain multiple drug-action and interaction at molecular [9] and cellular [10] level for a variety of normal and cancer tissues.

Through the aid of optimal control theory [11] such a complex modeling can furthermore suggest the optimal treatment, with the sequence of rest-periods and drug administration at the proper timing and dosage. Recent preliminary results have indeed been quite encouraging on the treatment of lung metastates from melanoma B-16 in mice [12].

There appears to be however a basic premise for the successful utilization of the sophisticated pharmaco-enzyme and pharmaco-cell kinetic modeling, namely the detailed molecular knowledge of the drug-tissue metabolism and of the physico-chemical properties of both normal and cancer cell in the various functional state (cycling, non-cycling and in varying stage of differentiation). Fortunately, modern biophysical techniques [13] permit such a characterization at single cell level with high accuracy and frequently on real-time. The transfer of such a complex and multifold approach to routing medical practice appears however nearly prohibitive, for its highly analytical

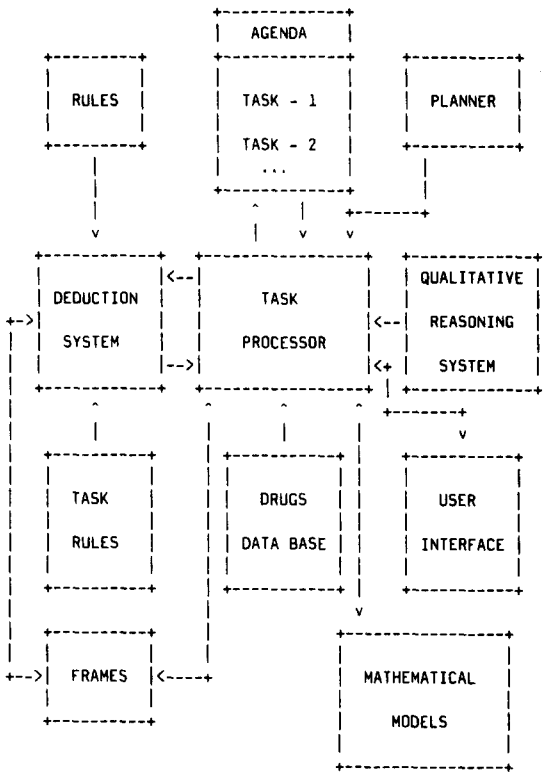

Fig. 3

nature and for the large number of parameters involved (constantly expanding due to the progress in cell and molecular biology).

We have then recently considered to use AI to bridge such a gap, by developing an *ad hoc* expert system named NEWCHEM to guide our present animal experimentation but with the aim to improve in the near future human cancer treatment.

3.1. The structure of NEWCHEM

The structure of NEWCHEM is shown in the diagram of Fig. 3. It is based on many paradigms of AI, since it puts together into one integrated system, a set of components, which are: a *production rule-based* system, a *frame* stucture, an *agenda-based* control, a *planner*, a data base about drugs, a *qualitative* reasoning system, a *blackboard* system. All those components cooperate to produce a plan (i.e. a sequence of actions) to be carried out by the user to achieve a given goal.

We adopt such a composite structure, because we believe that intelligent behaviour and accurate answers can be only obtained using specific techniques for the different kinds of knowledge and strategies to be embodied by the system. This is in agreement with current trends in expert system research. For instance, more recent systems like CENTAUR [14], ONCOCIN [5] and ABEL [15] use multiple representations of knowledge.

Furthermore, NEWCHEM uses mathematical models for pharmaco-cell kinetics [12], pharmaco-enzyme kinetics [9], and optimal control theory [11], written in FORTRAN, which are queried when the system needs precise answers concerning simulation and optimization of parameters at various levels of the decision-making process.

The system is written in Franz Lisp and uses the PEARL package [15] for efficient storage and retrieval, and the FLAVORS package for object oriented programming.

The interaction between NEWCHEM and the experimenter is as follows:

(a) The system asks the experimenter to provide some data describing the status of a subject to be treated.
(b) The system produces an optimal treatment plan together with a set of justifications for its choices and the expected results.

(c) The experimenter can propose changes in the plan and ask about possible consequences.
(d) A new treatment plan is proposed whenever the results are different from the expected ones.

In the following a short description of each component of the system is given.

3.2. Frames

Frames in NEWCHEM are associated with different points of view of the problem. They collect specific data and knowledge about inferences and actions concerning the specific domain. Data are related to different aspects of the case being examined. Frames are associated respectively with primary tumor, metastates, different tissues and organs (like bone marrow and small intestine). Different frames are also associated with discriptions respectively at the cellular level and the molecular level.

For each aspect, a frame acts as a blackboard in which all data are collected together with an indication about how they have been obtained and a measure of uncertainty. Consistency checking among data is also performed within the frame.

Frames also contain *action blocks*, which are sequences of tasks to be performed to fill the data entries and for different operations of the whole system. Tasks may refer, for instance, to the activation of other frames, to queries to the declarative knowledge base and to the mathematical model, and to planning activity. A particular *action block* is the *activation block*, to be executed at the activation of the frame.

The way a task is executed is based on a set of *task rules*, that for each task select the appropriate operations, depending on the data known so far (the known states of the subject). As an example, in the frame "treatment", a rule for the task "select a treatment plan" is the following:

```
(Task-rule
    (Task "select a treatment plan")
    (Context treatment)
    (Condition (isknown general_strategy))
    (Conclude_action '(progn (enter_date)
            (enter_new_data)
            (deduce general_strategy)
            (ask_user)
            (deduce strategy)
            (ask_user)
            (deduce treatment)
            (ask_user)
            (find_plan)
            (ask_user))))
```

Some data within a frame can be inferred from other data by means of backward production rules. For instance, the rule

```
(Rule
  (Ruleno 45)
  (Context treatment)
  (Condition
    (and ( > =subject performance_status 50)
         (for ?organ in (kidney,heart,liver,lungs)
         (same ?organ sufficiency satisfactory))
         (same_strategy kill_tumor)))
  (Conclude_value_treatment_type drug_combination))
```

gives treatment_type the value drug_combination if the performance status of the subject is $> = 50$, the sufficiency of kidney, heart, liver and lungs is satisfactory and the strategy is to kill the tumor.

With an item (or data entry) in a frame, one can associate a "trigger", which promotes appropriate actions when a critical value for the item is obtained.

Items can have multiple values, determined by means of different rules. Different values have, however, associated different scores representing their uncertainty.

At a given time a task can be active or inactive. When a frame enters the active status, the set of tasks, making its activation block, is loaded in the agenda of the system, waiting to be executed.

Groups of frames are organized in a tree-like fashion whenever some of them can be considered more specific points of view. In such a case a hierarchical relation "subframe" holds between pairs of frames.

3.3. The agenda and the task processor

The control of the system uses an agenda of waiting tasks and a task processor. The tasks in the agenda belong to active frames. For each of them the following data are maintained:

- The name of the task.
- The frame to which it belongs.
- A priority value.
- A justification.

At the end of the execution of a task, a new task is selected according to its priority.

The task processor is a problem solver that behaves in the following ways:

- For each task, it searches the associated task rules (within the frame the tasks belong to) to find a set of possible actions. Each action is a LISP code to which it is assigned a score, expressing how much it is appropriate for the current situation.
- It selects the action with the highest score
- It executes the selected action.

Actions may require to insert new tasks in the agenda. For each of them a priority value is computed and the current task is inserted in the list of its justifications.

3.4. The deduction system

The deduction system is invoked whenever in the system a value for an item is required. The deduction system first explores the frame to which the item belongs to see if a value is already known. If this is not the case, it looks for rules to deduce it. Rules are first searched for in the current frame (the frame of the current task), then in the frames which are ancestors of the current one in the subframe hierarchy, and, finally, in the frame to which the item belongs. If no rule is found, the value is asked of the user or a default value, if specified, is used.

The deduction system provides for an item, which is an attribute associated with an object, a list of PEARL structures of the following kind:

```
(Item
    (Object lisp)
    (Attribute lisp)
    (Context lisp)
    (Time lisp)
    (Value lisp)
    (Uncertainty integer))
```

Each structure expresses a value for the item at a given time with an associated uncertainty value.

3.5. Planner

The task of the planner is to produce a plan for the treatment of the subject, i.e. a partially ordered set of actions to be performed by the experimenter. The ordering of the actions refers to a temporal sequence, in the sense that some actions must be performed before others.

The planner is invoked by the task executor and receives as input an initial plan to refine. Refinement is performed by substituting a given action in the set with another partially ordered

```
;Task_plan R3 :
;IF the type of tumor is metastatized,
;   its progress istanzia reduced and
;   the state of the treatment istanzia in progress
;THEN change only one drug

(ci Task_plan R3
 (Name First_strategy) (Number 3)
 (Context General_strategy)
 (Type father) (Tax_father nil) (Tax_child nil)
 (Condition (prog (Date Listglob)
   (setq Date (daymonthyear))
   (setq Listglob (list 'cell_descr 'cell_descr 0))
   (return (fu_and
     (same 'type 'tumor 'metastatized Date Listglob)
     (same 'progress 'tumor 'reduced Date Listglob)
     (same 'status 'treatment 'in_progress
                          Date Listglob)))))
 (Action
   (conclude 'First_strategy 'change_only_drug 9.2))
 (Description (prog ()
   (printterpr " I find that the TYPE of TUMOR is
    METASTATIZED (CERTAINTY FACTOR : %f) "
    (same 'type 'tumor 'metastatized (daymonthyear)
   (list 'cell_descr 'cell_descr 0)) 5 1)
   (printterpr " the PROGRESS of TUMOR is REDUCED
    (CERTAINTY FACTOR : %f) "
    (same 'progress 'tumor 'reduced (daymonthyear)
   (list 'cell_descr 'cell_descr 0)) 5 2)
   (printterpr " and the STATUS of TREATMENT is
    IN PROGRESS (CERTAINTY FACTOR : %f) "
    (same 'status 'treatment 'in_progress
     (daymonthyear)
     (list 'cell_descr 'cell_descr 0)) 5 2)
   (cprinttab " THEN I recommend to CHANGE
          ONLY ONE DRUG " nil 5))))

(insertdb R3)
```

Fig. 4

set of more specific actions. This is accomplished using refinement rules, which are context-dependent production rules, that replace a given action in their left part with the set of actions specified in their right part. Consistency tests are also performed on the overall plan to resolve interactions and conflicts. This planning is similar to the hierarchical planning performed by NOAH [17].

The actions specified in a plan refer to different kinds of intervention on the subject, mainly surgery, radiology and administration of drugs. Some constraints, describing the particular modalities of the intervention (for instance, the properties of the drugs to be administered) are associated with these actions. These constraints are used to select, at the end of the planning activity, the most appropriate intervention through a query to the declarative knowledge base—for instance, the specific drug needed. This selection is made only when the plan is completely refined, in order to avoid an early choice that can compromise the whole plan. This is what is called a *least commitment strategy*. Constraints can also be propagated within the plan as in MOLGEN [18].

The plan is associated as a value with the item *plan* in the frame *treatment*.

An example of a rule of the Planner is shown in the Fig. 4.

3.6. The data base about drugs

The data base about drugs of the system contains data concerning all drugs that are to be used during the treatment considered by the system and their properties.

Drugs are represented as structured types of LISP objects, stored in an internal form as blocks of memory, regarding logical groupings of heterogeneous data as slots and slot fillers.

One main structure defines the way in which all drugs are characterized, while each drug is represented as an individual instance of this structure. The main structure can be regarded as a reference model of a drug in which the greatest number of attributes and properties are available,

and each individual drug is represented filling up all known attributes of the main structure and specifying properties to the greatest possible detail.

This knowledge representation is implemented in PEARL an AI language allowing to create hierarchically defined slot filler representations and to handle them efficiently, by easy inserting and fetching. PEARL is implemented in LISP, so the knowledge base can be directly addressed within a LISP system as the one described in the previous sections. Specific rules contained in the system directly address attributes of the drugs represented in PEARL in the declarative data base.

The main structure adopted is implemented in PEARL as follows:

```
(drug
  (name symbol)
  (type symbol)
  (main_attribute symbol)
  (class symbol)
  (toxicity symbol)
  (action struct))
```

According to this main structure, each specific drug considered by the system is represented as follows:

```
(drug
  (name cyclophosphamide)
  (type alchilant)
  (main_attribute polifunctional)
  (action (type inhibits)
          (object mitosis)))
```

3.7. The qualitative reasoning system

The qualitative reasoning system uses qualitative process theory, as developed by Forbus [19], to reason about processes at a cellular level. This component is used within the whole system to perform qualitative simulations in order to "roughly" predict the consequence of specific actions.

In qualitative process theory physical parameters are characterized by a quantity space in which only intervals of values are considered. Usually, what interests is if a quantity is positive, negative or null, and if it is increasing, decreasing or stationary. Furthermore, physical situations and processes are described symbolically, in a language based on predicate calculus. For a given process, the objects involved, the conditions in which the process takes place, the relations holding among the parameters, and the causes of change (influences) are specified.

The qualitative reasoning sytem is implemented in an object oriented style of programming which is provided by the "flavors package" of Franz Lisp.

As an example the inhibition of DNA replication by a substance can be described as follows:

```
(make_instance 'process
          ':individuals '(DNA cell ?S)
        ':preconditions
          '(and (substance ?S)
                (inhibitor_DNA_synthesis ?S))
          ':influences
          '(I- (Am (synthesis DNA))
               (Am (concentration ?S))))
```

The set of possible sequences of processes that follow a given action on a cell or on a cell population (for instance, an increase in concentration of a substance) is the result of a qualitative simulation.

This result is useful during the planning process in order to take care of the possible consequences of the actions that have been selected.

3.8. The mathematical model

The modeling concerns various levels of decision-making and attempts to provide an integrated and realistic approach to optimal cancer treatment. Each aspect of this modeling has been extensively dealt in previous publication [5, 8, 9–11, 18], and is based on numerous biophysical parameters, which are reviewed in details in a forthcoming book [13].

Each model acts synergistically within a proper multiple time scale, namely:

- — at the level of s–min–h, a model (DNAMET) for DNA enzymatic synthesis determines multiple DNA antimetabolites interaction to induce either differential synchrony or differential killing between normal and cancer tissues;
- — at the level of h–days, a cell kinetics model determines drug-tissues to suggest (SIVFIT) optimal short terms drug protocols (type, dosage and timing);
- — at the level of weeks–months, an optimal control-based model (SIVFIT 2) predicts long-term timing of treatment, where rest periods and changes in the type of drugs are determined taking into consideration long-term effects like drug resistance, due to gene amplification and/or cell mutation.

4. CONCLUSION

Two expert systems in Oncology have been described. The first of them, BREASTCAN, which follows a traditional approach, is based only on the traditional input derived from clinical experience with empirical trials, which are carried out on randomized patients with the prevailing chemotherapeutic protocols. Our aim with the second one, NEWCHEM is instead to extend our knowledge base and our rules to incorporate also all the most advanced knowledge at the molecular and cellular level, both theoretical and experimental, to make readily accessible to the health-community a system which will be really as expert as the present state of the art allows.

REFERENCES

1. G. Bonadonna and P. Valagussa, Chemotherapy of breast cancer: current views and results. *Int. J. Radiat. Oncol. Biol. Phys.* **9**, 279–297 (1983).
2. S. E. Jones and S. E. Salmon, *Adjuvant systemic therapy of cancer IV.* Grune & Stratton, New York (1984).
3. G. Bonadonna and P. Valagussa, Adjuvant systemic therapy for resectable breast cancer. *J. Clin. Oncol.* (in press).
4. J. C. Kunz, E. H. Shortliffe, B. G. Buchanan and E. A. Feigenbaum, Comparison of techniques of computer-assisted decision-making in medicine. *Modeling and Analysis in Biomedicine*, pp. 335–367. World Scientific, Singapore (1984).
5. E. H. Shortliffe, A. C. Scott, M. B. Bischoff, A. B. Campbell, W. Van Melle and C. D. Jacobs, ONCOCIN: An expert system for oncology protocol management. *Proc. 7th Int. Joint Conf. on Artificial Intelligence*, Vancouver, B.C., pp. 876–881 (1981).
6. W. F. Clocksin and C. S. Mellish, *Programming in Prolog.* Springer, New York (1981).
7. B. G. Buchanan and E. H. Shortliffe (Eds), *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project.* Addison-Wesley, Reading, Mass (1984).
8. C. Nicolini, The principles and methods of cell synchronization in cancer chemotherapy. *Biophys. Biochem. Acta* **458**, 243–282 (1976).
9. C. Nicolini, A. Belmont, M. Grattarola, C. Moore and E. Milgram, Pharmaco-enzyme kinetics simulations of experimental interactions among multiple antineopolastic drugs. *Biochem. Pharm.* **28**, 2891–2908 (1979).
10. S. Zietz, C. Desaive, S. Lessin and C. Nicolini, Modeling to determine dose dependence of drug and cell kinetic parameters. *Comput. Biomed. Res.* **13**, 297–305 (1980).
11. S. Zietz and C. Nicolini, Mathematical approaches to optimization of cancer chemotherapy. *Bull. Math. Biol.* **41**, 305–324 (1979).
12. C. Nicolini, S. Zietz, M. Tansy and S. Lessin, Modeling and analysis in cancer treatment. *Modeling and Analysis in Biomedicine*, (Ed. C. Nicolini) pp. 291–333. World Scientific, Singapore.
13. C. Nicolini, *Biophysics and Cancer.* Plenum, New York (1985).
14. J. S. Aikins, Prototipical knowledge for expert systems. *Artificial Intelligence* **20**, 163–210 (1983).
15. R. S. Patil, P. Szolovits and W. B. Schwartz, Modeling knowledge of the patient in acid-base and electrolyte disorders. *Artificial Intelligence in Medicine* (Ed. P. Szolovits) pp. 191–226. Westview Press, Boulder, Co. (1982).
16. M. Deering, J. Faletti and R. Wilensky, Using the PEARL AI package. Technical Report Department of EECS, University of California, Berkeley (1982).
17. E. D. Sacerdoti, *NOAH: Planning in a hierarchy of abstraction spaces. Artificial Intelligence* **5**, 115–135 (1974).
18. M. Stefik, Planning with constraints. *Artificial Intelligence* **16**, 111–140 (1981).
19. K. D. Forbus, Qualitative process theory. *Artificial Intelligence* **24**, 85–168 (1984).