Human-Agent Co-operation in Accessing and Communicating Knowledge Media – A Case in Medical Therapy Planning

Volker Dötsch[1], Kimihito Ito[2], and Klaus P. Jantke[3]

[1] Hochschule für Technik, Wirtschaft und Kultur Leipzig (FH),
Fachbereich Informatik, Mathematik und Naturwissenschaften
Postfach 301166, 04251 Leipzig, Germany
doetsch@imn.htwk-leipzig.de

[2] Hokkaido University Sapporo,
Meme Media Laboratory,
Nishi 8, Kita 13, Kita-ku, Sapporo, 060-8628 Japan
itok@meme.hokudai.ac.jp

[3] Deutsches Forschungszentrum für Künstliche Intelligenz,
Stuhlsatzenhausweg 3, 66123 Saarbrücken, Germany
jantke@dfki.de

Abstract. The concepts of memetics and the development of meme media implementations have set the stage for a new generation of knowledge processing systems in which knowledge evolution may take place. When the evolution of knowledge goes beyond the imagination of humans, we will arrive at a new quality of human-computer co-operation.

The evolution of knowledge is of a particular interest in domains where the future is clearly unforeseeable, but where a rapid growth of knowledge is highly desirable. Medical therapy is such an area of a particularly great importance.

IntelligentPad technologies as a form of meme media implementation are used to formulate and represent medical therapy knowledge, to set up meme pools and to allow for the evolution of knowledge beyond human expectations.

This paper is reporting about the authors' first steps and is intended to lay the cornerstone of a related research and development program.

1 Introduction and Motivation

Memetics is seen as outlined in the truly exciting books by Richard Dawkins [4], Susan Blackmore [3] and Yuzuru Tanaka [9]. Richard Dawkins has attracted the world's attention to the phenomena of cultural inheritance and has introduced his seminal concept named *meme*. Susan Blackmore has taken the initiative to discuss the relevance of Dawkins' perspective from a psychological and from a somehow philosophical point of view telling all of us that we are affected by Dawkins' work. It is, naturally, up to you whether or not you feel personally affected by memetics, and this might easily become a slightly esoteric discussion.

G. Grieser and Y. Tanaka (Eds.): Intuitive Human Interface 2004, LNAI 3359, pp. 68–87, 2004.

But Yuzuru Tanaka, fortunately, has seized Dawkins' suggestion and developed it towards concepts, implementations and applications in computer science. He has coined the key term *meme media*.

The present approach relies on Tanaka's trend-setting work taking Dawkins' and Blackmore's contributions seriously.

We want to contribute to the endeavour of enabling computer systems to foster true knowledge evolution – the benefit for humans will be paramount. This may be understood as a contribution to new knowledge media in action as envisaged by Mark Stefik [8].

There are two more or less independent starting points from which to undertake such an endeavour. One is technology development as an inevitable basis; see [6] for a recent publication of the second author. Another one is identifying existing meme pools and working on evolution-driving forces. For the present investigation, the first author has brought in the domain and currently undertakes corresponding work on knowledge processing and, in particular, knowledge generation as previously done by the third author; see [2], e.g.

In particular, the application domain chosen is medical therapy planning; Based on former work in planning for complex dynamic domains (see [1], [2]).

2 Peculiarities of the Application Domain

We briefly describe characteristics of the application domain to reveal it's appropriateness for a meme media based approach to foster knowledge evolution.

Planning is a traditional research area of Artificial Intelligence. Traditional planning was dominated by logical approaches and by procedures more or less deductive in spirit. Conventional actions in deductive planning have so-called preconditions and postconditions. Their executability can be logically verified, and they change the world in a deterministic manner such that corresponding effects can be logically recorded.

In contrast, many practically relevant application domains are complex and dynamic such that deductive planning is inadequate by nature.

Oksana Arnold [1] (see also [2]) has investigated planning in complex and dynamic environments like industrial process control. Chemical installations are characteristic application cases.

In those domains, automated reasoning, in general, and computer-assisted planning, in particular, is invoked when the underlying process is in trouble. Then, there is typically some lack of data and complete knowledge for deductive reasoning is not available. The most exciting planning tasks are inductive in spirit – planning is learning [2].

Reasoning over those application domains, especially computer-assisted or fully automated reasoning like diagnosis and control, is complicated by a number of peculiarities classified into three groups as follows [1]:

[i] fundamental peculiariries characterizing dynamics and complexity,
[ii] domain-specific dynamics,
[iii] derived, but essential peculiarities.

For every class of peculiarities, we are giving a few instances to illustrate the type of difficulties we are facing and going to attack by means of meme media technologies.

[i] Fundamental Peculiarities
- Several target parameters can not be controlled directly. For instance, a human's blood pressure can only be controlled through a number of indirect medications.
- There are several process parameters of which one can not regularly access current values. Repeated tests, though possible in principle, may by physically exhausting and mentally unacceptable to human patients.
- The execution of some actions may depend on external resources the availability of which may be locally undecidable. This is particularly true for actions which depend on environmental details like communication channels, transportation facilities and administrational customs.

[ii] Domain-Specific Dynamics
- Certain constraints underlying the executability of actions need to be satisfied throughout the whole execution period of some action. For instance, some medication may necessarily require the absence of fever.
- The execution of actions is time-consuming. The amount of time necessary to complete some actions can not be estimated, in general. Usually, so-called time-outs serve as an upper time limit for executing actions.
- Usually, there are alternative actions. Those actions may have advantages and disadvantages; there might be no clearly preferred decision.

[iii] Derived Peculiarities
- So far, the human body is only insufficiently understood. There are far too many process parameters to be taken into account. Data has to be dropped and, thus, all information is incomplete by nature.
- The state of the human body changes even in case no actions at all are executed. There is no assumption of persistency.
- There are many interacting processes, and even if a current plan is perfect, it may fail by an unexpectable interaction with some other process. For instance, a schedule of surgery may be perfect, but break down if some doctor falls ill. It may also surprisingly turn out that some therapy treatment is more exhausting for a patient than initially expected.

The domain of planning in complex, dynamic environments, in general, resp. medical therapy planning, in particular, is deemed highly appropriate for meme media applications. First, the building blocks of knowledge available in the area can be reasonably mapped to meme media objects. Second, a central type of knowledge manipulation in practice truly consists in plugging those building blocks together. Third, due to the complexity of the domain, the evolution of knowledge through successful assembly of building blocks and application of composite meme media objects is essentially unforeseeable. Fourth, the domain – especially the restricted medical case – is of great practical importance and does attract sufficient public interest such that there is hope to bridge the gap from academic cutting-edge investigations to sufficiently large application cases.

3 Therapy Plan Generation and Execution Scenarios

Before going into the details of memetics for medical therapy planning, we need a vision of the future of memetic knowledge processing and evolution.

In the medical domain, the importance of decisions to human beings, to their health and life, and the derived liability of decision makers form particular obstacles to automated knowledge processing. This is beyond the present paper, but it has to be taken into account.

The current state of affair is characterized by therapy planning completely done by humans who, at most, use computers as tools for typewriting, printing and documentation purposes.

In contrast to the current state of affair, information and communication technologies offer the prerequisites for automated knowledge processing ([1], e.g.). We are on the cusp of bringing memetics to work in this area, driven by our own interest to see knowledge evolution taking place.

How to proceed gradually? How to introduce memetic knowledge processing? How to embed science and technology under development into a quite complex environment?

The *plan generation* may be seen as a separated task depicted as a box in an IT infrastructure visualization. Other components are *constraint evaluation* having, so to call, a particular dynamic logic as plug-in, *plan execution*, and *monitoring* (see Figure 1).

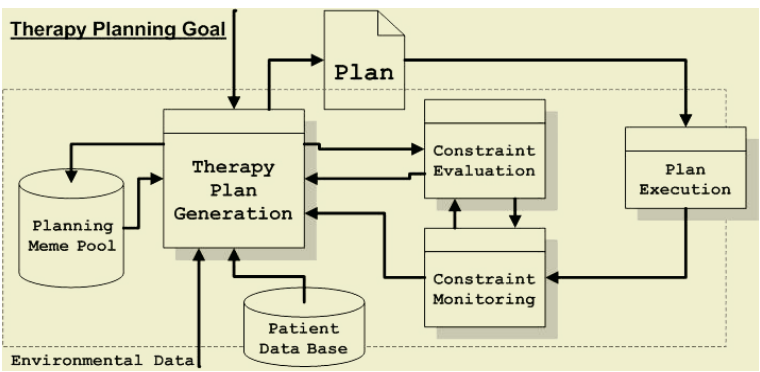

Fig. 1. Therapy Plan Generation in a Knowledge Processing Infrastructure

At a first stage, one may see `Therapy Plan Generation` as a task performed by hand. A meme media based planning tool may be invoked to support plan generation based on the `Planning Meme Pool`. If the employed logic is specified, `Constraint Evaluation` may be fully automated from the early beginning.

Re-planning is invoked when `Plan Monitoring` identifies invalid constraints.

On later stages of the implementation process, further components may be transformed into autonomous knowledge processing devices. But in contrast to process control [1], `Plan Execution` might never be completely automatic.

4 Medical Therapy Memes and Meme Pools

There is an enormous amount of knowledge in medical therapy which is currently more or less available. The authors are aware of the problem that making this knowledge available to computerized processing may take some time. Especially, a world-wide distribution of medical knowledge is not only facing technological obstacles, but also social, ethical, commercial, legal, and a large variety of other cultural difficulties. The authors refrain from going into these details and focus on scientific and technological investigations, exclusively.

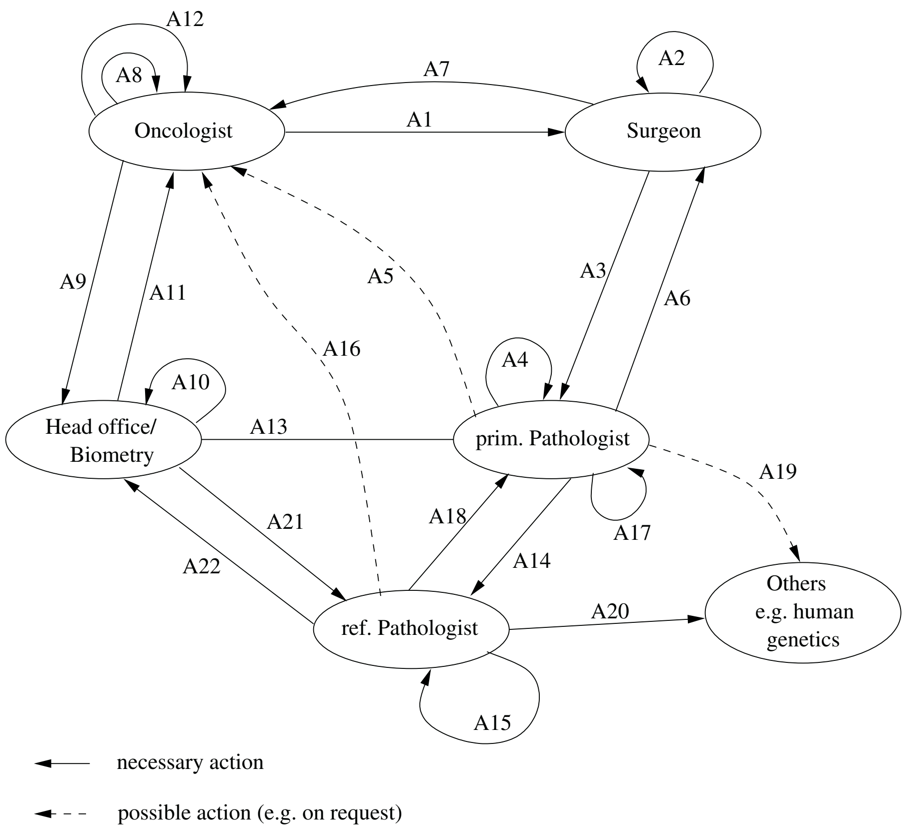

Fig. 2. Process of Patient Integration into a Randomized Clinical Trial in Germany

To avoid topical controversies on medical issues, for which the authors are not competent, the illustrations and examples stress more administrational activities in medical therapy than medication or direct physical treatment of patients. As a running example, we focus on the integration of patients into a randomized clinical trial as displayed in figure 2. Existing but finished protocols of German clinical trials on Non-Hodgkin-Lymphomas are taken as reference cases.

4.1 Basic Therapy Knowledge

There is a huge amount of elementary actions potentially taking place like the following which will be used subsequently:

- **[a 1]** referral
- **[a 2]** perform biopsy
- **[a 3]** send data and tissue
- **[a 4]** processing tissue, diagnostic findings
- **[a 5]** send findings (if requested)
- **[a 6]** send findings and data
- **[a 7]** selection of an appropriate clinical trial
- **[a 8]** patients registration (send data)
- **[a 9]** randomization
- **[a 10]** send data and the result of randomization
- **[a 11]** plan treatment in detail and start treatment
- **[a 12]** request to send data and tissue
- **[a 13]** reference findings
- **[a 14]** send reference findings (if requested)
- **[a 15]** send reference findings
- **[a 16]** compare findings and reference findings
- **[a 17]** request data and reference findings of integrated patients

These actions as described in natural language are still incompletely specified and clearly generic in the sense that some of their corresponding parameters can vary. For instance, sending actions have parameters like sender, addressee and material to be sent.

Consequently, actions form a conventional class hierarchy where classes may have subclasses and instances which form the leaves when the hierarchy is seen as a tree.

In addition, actions usually have conditions (named constaints) under which they can be executed. Conditions are classified into start constraints and interval constraints. In medication and physical treatment, interval constraints are of crucial importance. Administrational processes are usually a little simpler and, therefore, interval constraints are less important. For simplicity, we confine ourselves to start constraints, first.

Here are two sample start constraint for actions with reference to the listing above.

- **[sc 7]** The oncologist received diagnostic findings (and reference findings) and appropriate data.
- **[sc 11]** The oncologist received the result of the randomization.

As briefly mentioned, the actions **[a 1]** to **[a 17]** listed above are classes which may have subclasses. They become instances by substituting parameters. For example, action class **[a 3]** has the instances **[A 3]**, **[A 14]**, **[A 19]**, and **[A 20]** of sending tissue and data displayed in the process diagramm of figure 2.

Actions are represented as meme media objects which, in earlier days of AI, might have been called frames. They consist of a list of so-called slots containing all the relevant data ranging from an object's name over values for several parameters to constraints and certain flags like the one indicating whether execution has been started (cf. figure 3; for more details consult [9]).

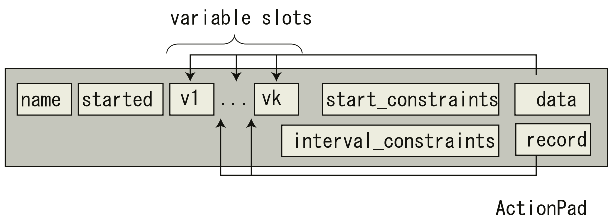

Fig. 3. A First Approach to Action Pads – The Slot Structure

Beyond former IntelligentPad applcations, the present one requires a specific type of relations between certain slots. There are some slots – those for start constraints and for interval constraints – that contain logical formulae. And these formulae contain variables which represent the content of some other slots.

4.2 Composite Plan Knowledge

Plan actions are building blocks which can be plugged together just like Lego toy blocks. A particular plan for the case under investigation is shown in figure 4. The naming of actions refers to the process model of figure 2.

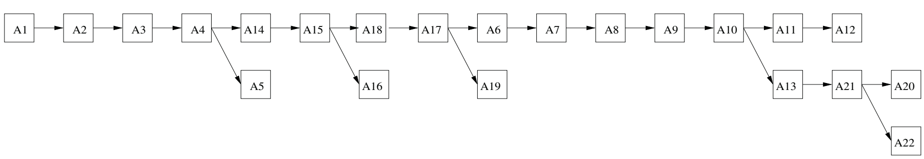

Fig. 4. A First Plan for Integration of a Patient into a Clinical Trial

Planning is plugging actions together in accordance with all the knowledge available. As explained in chapter 2, the ultimate goal is to come up with a plan consistent with everything known so far. Due to the vagueness and to the incompleteness of information, therapy planning is inductive in spirit and, thus, plan revision is frequently necessary (cf. [2]). Figure 4 displays an alternative.

The first plan is a more conservative one. Integrating a patient into a clinical trial is done only if reference findings are present. In contrast, the second alternative plan useful in seemingly urgent cases integrates a patient before reference findings did arrive. Especially in medication and physical treatment, there is usually a large number of alternatives.

Higher level actions usually consist of some structure according to subgoals to be reached either subsequently or potentially in parallel. They look like simple graphs. Planning is a stepwise refinement by substituting other actions for the nodes of those graphs. So, from a more structural perspective, planning is graph expansion. A planning step consists in plugging a graph into the node of a given graph under development.

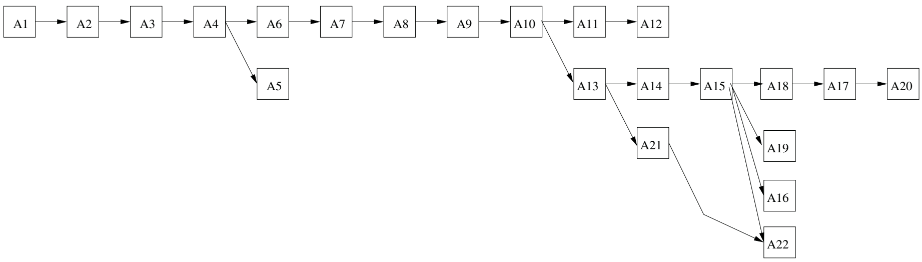

Fig. 5. An Alternative plan for Seemingly More Urgent Cases

Currently, all the work of planning and replanning is done by hand. Meme media technologies promise to gain efficieny by supporting the process of plugging actions together. The present paper aims at doing this automatically.

Plans as shown in figure 4 and 5 have to be represented as composite pads. Those pads result from elemtary pads or from other composite pads by plugging together, i.e. by establishing slot connections. The technicalities of this meme media manipulation process are subject of the following chapter.

Readers are advised to re-consult the communication architecture displayed in Figure 1. Therapy Plan Generation is the core activity yielding plans as output. Which plans are fit and suitable for execution depends on a rather dynamic environment including patients, doctors, medical stuff and environmental conditions like, for instance, available communication and transportation facilities. Even unforeseeable phenomena like weather conditions and, for instance, insect activities may interfere through impact on a patient's conditions.

Because of the complexity of the environment, we are far away from computer-controled plan execution, in general. The plans are executed by means of human activities. But constrained monitoring can be performed automatically to a large extent. To sum up, the evolutionary process through which plans are generated, are evaluated, may possibly succeed and may even survive for a longer period of time is only partially automated. Many agents – humans, computers, and others – work together.

The present approach focus the knowledge issues of this process with a certain emphasis on human-computer co-operation. The main agents are those that manipulate therapy planning knowledge. The authors' work aims at a shift in the division of labor such that a larger part of the knowledge processing is performed by computers.

Plan generation is meme media manipulation. Over time, it may result in a meme media evolution process.

5 Meme Media Manipulation and Development

When we think of meme media manipulation as taking meme media objects and plugging them together somehow like Lego toy building blocks – a methaphor frequently cited in [9] – we mostly focus on the manipulation of the core medical therapy knowledge.

But as pointed out above, we have to see the therapy knowledge in the light of recent data about the patient (ranging from the patient's history to recent measurements) and the environment (availability of equipment, e.g.).

When individual planning activities and composite plans are all represented as IPads, a plan under development, i.e. a composite pad under construction, has to be related to the underlying data mentioned.

We suggest a separation into two types of underlying knowledge:

– patient data,
– environmental information.

There are several arguments for such a distinction; a simple one is as follows. In different clinical environments there do exist conditions which are specific and not found in any other place. There are special communication connections to certain laboratories, there are special ways, distances, transportation facilities and the like to be respected, there are specific procedures of scheduling surgery and much more. All those peculiarities have to be taken into account. They can be bound directly to related action classes and will be inherited to individual actions upon instantiation time.

The separation of the underlying knowledge suggested above leads to different ways of accessing this knowledge during meme media manipulation.

In the present chapter, we go the step from *medical therapy memes* to *medical therapy meme media*. Knowledge will be encapsulated in meme media objects (knowledge media according to [8]) based on the IntelligentPad concept [9].

We have to discuss how patient data and environmental data can enter the meme media world and how knowledge about medical therapy, which is more procedural in spirit, can be formally represented to be put on top of the factual knowledge. These three aspects are shaping the present chapter.

5.1 IntelligentPad Technology for the Access to Patient Data

Increasingly, medical records of patients are being stored in databases. The main purpose of the computerization is to make treatment and accounting in hospitals efficient and, for legal reasons, comprehensible and undisputable.

To wrap such medical records into meme media objects, we follow the way that has been developed by Yuzuru Tanaka in his IntelligentPad research [9].

A MedicalRecordPad wraps any medical record stored in a medical database. It has slots for all medical data of a human according to a certain medical format. A MedicalRecordPad divides each medical record into a set of attribute-value pairs. Each attribute value is set to the slot with the same name as its attribute name. In its typical use on a MedicalRecordPad, it is pasted onto a

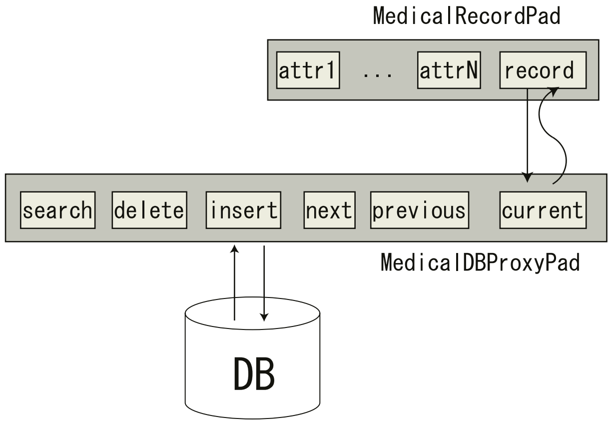

Fig. 6. Medical Record Access Architecture

MedicalDBProxyPad, which wraps basic database functions including search, insert, and delete.

Figure 6 displays a composite meme media object which is assembled of a MedicalRecordPad and a MedicalDBProxyPad.

Established meme media technology is invoked and adopted to the needs of the new domain.

5.2 IntelligentPad Technology for Environmental Embedding

Environmental data are very much like sensor information (see [1], e.g.) which is acquired somewhere outside.

Generally, there are two possible mechanisms of communication about environmental data – 'polling' and 'callbacks'. Polling is just periodical calling of a communication task for asking from time to time for the state of the environment by sending out requests for sensor data. On the other hand, callback is the way in which sensors notify the client that some event interesting to the client has been detected on the sensor. In response, the client invokes functions registered ahead of time. The use of callbacks provides event-driven planning. In this research, we adopt the mixed approach of polling and callbacks, according to the needs of clinical environments.

An EnvironmentalMonitorPad defines localized conditions and checks whether current and localized environmental information satisfies the conditions. In different clinical environments there do exist conditions which are specific and not found in any other place. Therefore, in different clinical environments different types of the EnvironmentalMonitorPad might be used.

An EnvironmentalMonitorPad has, at least, two slots named 'environmental data' and 'notification_constraints'. The slot 'environmental data' contains a set of attribute-value pairs of names of parameter and values of parameters. The slot 'notification_constraints' contains a formula that represents conditions to be satisfied for the notification. If the condition is satisfied, the pad notifies the environmental state change to its parent. It sends a set message with environmental data as its parameter to its parent.

Figure 7 shows a composite pad of a EnvironmentalMonitorPad and a BasePad.

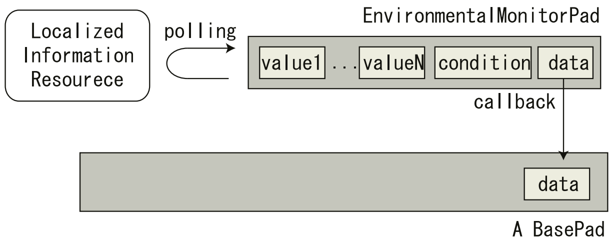

Fig. 7. Environmental Data Communication Architecture

Autonomous communication of a generated therapy plan, regardless whether the plan is handmade, automatically generated or resulting from an interactive session, is a highly interesting and novel issue. To the authors' very best knowledge, there is not much prior work (perhaps, nothing at all) on plans communicating autonomously with their environment. At least in the medical therapy planning domain, there are no related publications. A more detailed investigation will be postponed to a future publication.

5.3 IntelligentPad Technology to Represent Therapy Knowledge

In contrast to classical deductive planning, we never know about *executability* of a plan, because execution takes place in the future. All actions have constraints that must be satisfied at the beginning of an action or during the whole execution period. You never know whether or not all of the conditions will be true in the future. You can only check whether they are still possible from the perspective of the current knowledge. If there is nothing contradictory, the actions are considered *consistent*. If more information comes in, one may find out that consistency is lost.

Knowledge about medical therapy consists of an enormous amount of knowledge about elementary therapy actions and about ways to combine those actions into hopefully successful therapies. The latter is planning knowledge, from our present IT perspective. The core knowledge concerns individual actions and their related conditions of applicability. Knowledge of this type available in some location or, at least, in a human individual is forming a meme pool. Those meme pools are to be abstracted, formalized and represented as meme media pools.

The authors rely on the IntelligentPad concept by Yuzuru Tanaka [9], in general, and on the implementaional concepts of IntelligentPad, in particular. Therapy actions are mapped to so-called action pads.

An ActionPad (Figure 8) divides each medical record and environmental data into a set of attribute–value pairs in the same way as MedicalRecordPad descrived in the section 4.1.

Every ActionPad has the following four extra slots:

– action_name,
– action_started,
– start_constraint,
– interval_constraint.

The slot 'action_name' gives the name of this action, and the slot 'action_started' gives whether the action has already started or not. In the slots 'start_constraint' and 'interval_constraint', there are filled in logical formulae which may have variables. Those variables correspond to slot names of the action pad.

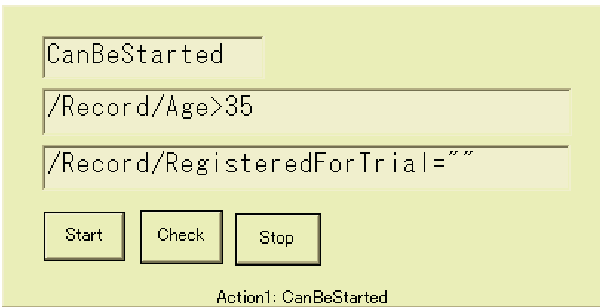

Fig. 8. Action Pads – Inhabitants of Medical Therapy Meme Media Pools

Every action pad has a method called 'Consistency Monitoring Method' (CMM, for short). Whenever an update of slot values takes place, this method is executed. It evaluates the formulae in the constraint slots. When the CMM computes for any constraint the value false, some activities are triggered. As a consequence, whenever data relevant to some action change, the action itself should check consistency.

Figure 9 shows the architecture of consistency checking with both patient data and environmental data. In its start_constraints slot, the ActionPad has the following constraint:

```
you_can_start(actionA) if constraints A, B, C hold.
```

And in its interval_constraints slot, the Pad has the following constraint:

```
you_can_continue(actionA) if constraints X,Y,Z hold.
```

ActionPad notifies that we can start actionA, if A, B, and C hold. This decision is made by replacing each variable v_i in A, B, C with the values of slot whose name is v_i. If actionA is started, the value of the slot 'action_started' will be turned into 'true'.

After the slightly simplified discussion of constraint monitoring above, the authors take the liberty of a short excursus on constraint monitoring in complex, dynamic, and non-persistent environments as laid out in [1] and [2].

Constraints are formulae that may express properties of process parameters to be satisfied in the future. Time plays an important role, and vagueness and incompleteness of knowledge about the future makes constraint evaluation somehow difficult. In fact, there is the need for a particular temporal logic.

Even more difficult, but also exciting, the choice of a suitable temporal logic has substantial impact on the whole knowledge processing under investigation.

To bring this to the point of memetics, the evolution of knowledge may deeply depend on the logics used in the planning process. To the authors' very best knowledge, there are no investigations at all about this problem area.

So, for the purpose of the present paper, we return to the simplified assumtion that there is any mechanism of reliable constraint evaluation.

During the actionA, whenever some value of slot in the ActionPad has been changed, CMM evaluates the formula in the 'interval_constraint' slots. If the formula does not hold any longer, the pad notifies the health care professionals involved.

Details of alert processing are beyond the limits of meme media technology and are important when putting the technology into application environments.

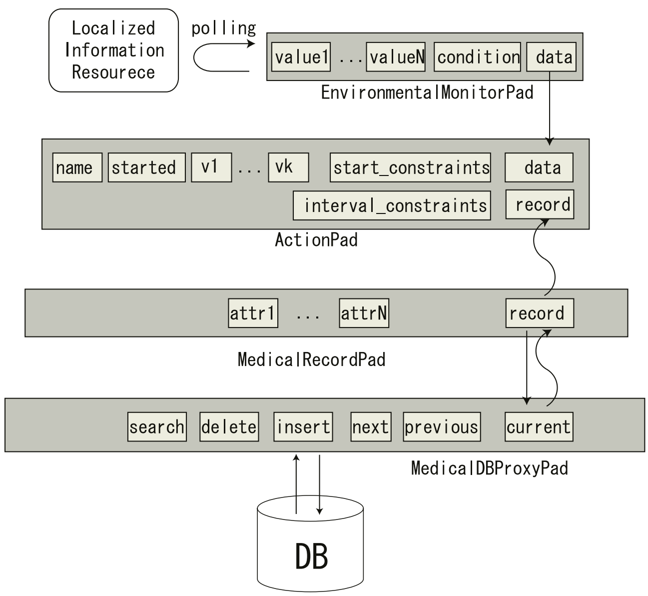

Fig. 9. Action Pad Communication Architecture

The present chapter's concepts are setting the stage for the representation of *medical therapy memes* truly as *medical therapy meme media*. On the basis of medical therapy meme media, computer-assisted or even automatic therapy plan generated can be introduced.

6 Automatic Planning and Knowledge Manipulation

This is a publication about memetics, about meme media technology and about knowledge evolution through meme media implementation and application. Planning in complex dynamic environments, especially medical therapy planning, serves as a testbed and demonstration case, only.

Therefore, the authors confine themselves to a treatment of the planning issue reduced to its essentials. More details can be found in [1], [2] and in some forthcoming publication about ongoing work of the first author.

Essentially, plan generation is seen as graph expansion. Medical therapy plans are hierarchically structured graphs.

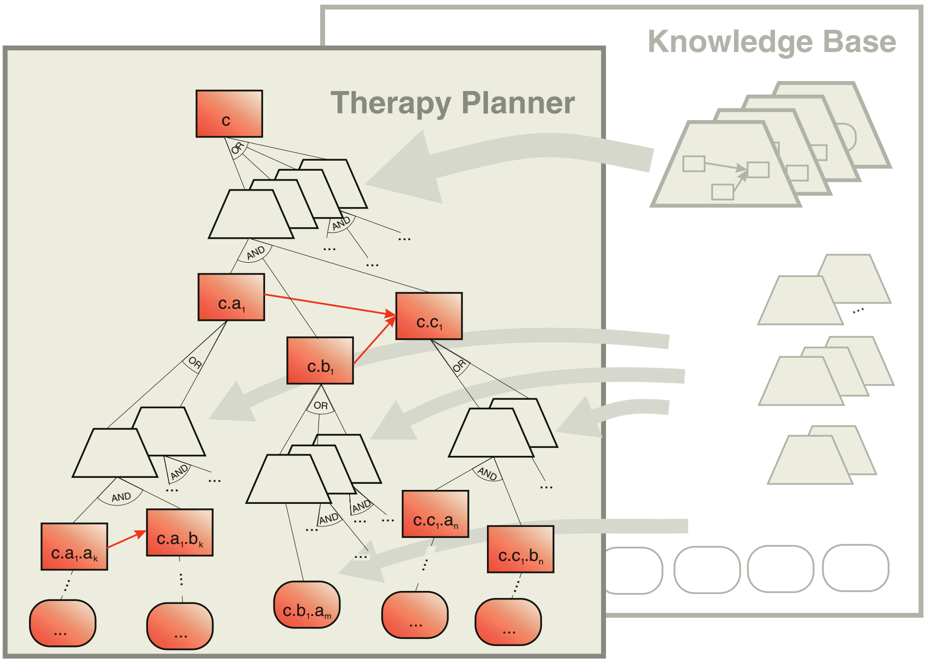

Fig. 10. Medical Therapy Planning as Hierarchical Graph Expansion

Planning starts with some highest level goal. For this goal, there are usually several alternative plan graphs which may be chosen and inserted. Alternatives can be ordered by preference values. As long as nodes of the stepwise extended graph are not yet atomic actions, they can be further expanded as illustrated in figure 10.

As discussed in chapter 2, complex dynamic environments do not provide sufficient information to decide plan executability. Planning is inductive by its very nature. But medical therapy actions depend on executability conditions expressed as constraints. Therefore, constraints need to be verified prior to execution.

Constraint supervision is necessary at planning time as well as during plan execution. Oksana Arnold has developed a sufficiently complete theoretical framework which can be carried over to medical therapy planning with a reasonably small amount of adaptation work.

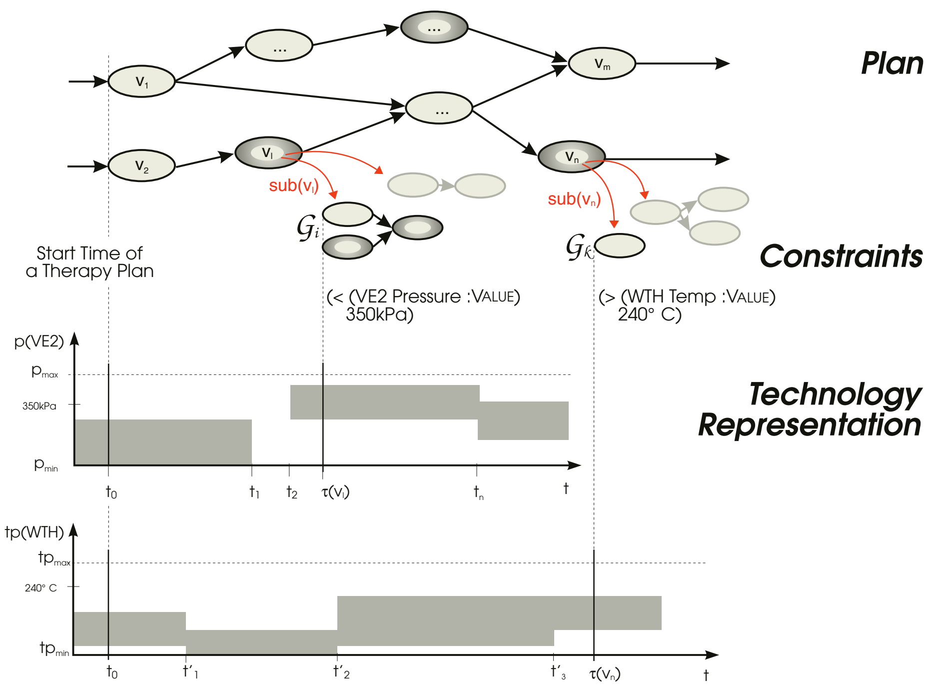

Fig. 11. The Problem of Constraint Satisfaction over Time (adopted from [1])

In medical therapy, we do not speak about technology representation, but about patient records and environmental data. Despite these differences, the constraint supervision problem appears as shown above. Actions inserted into a plan are consistent at planning time. Consistency has to be checked continuously. As soon as inconsistencies are detected, local plan revision is invoked. Action have to be replaced by other actions.

In our present approach, actions are pads and plans are more or less complex composite pads. Plan revision means a local reassembly of a composite pad.

The key message of the present chapter is that plan generation as sketched here can be performed automatically (cf. [1]). The first author is currently working on a variety of planning programs; details are to be published separately.

From the meme media point of view, planning programs do automatically assemble composite pads. According to the needs of a dynamic environment, plan revision my be frequently triggered. Ultimately, unforeseeable plans are generated and do succeed. Knowledge evolution takes place.

7 Meme Media Evolution Technologies

In the chapter 6 before, we have seen that planning can be done automatically as guarded graph expansion. But how does it work in detail that meme media evolution[1] is enabled?

Such as biological evolution has its internal mechanisms, so has meme media evolution. Though the present paper can not deal with all aspects of our quite complex research and development program, a brief look into the mechanisms of evolution seems desirable. A detailed discussion has to be postponed to a separate publication.

Automatic plan generation is a little like a robot playing with Lego building blocks. Under the peculiar conditions of a complex and dynamic domain and the absence of persistence, the outcome is usually unforeseeable, new and sometimes even innovative.

To stick with the metaphor of the playing robot for a moment, how does the robot find and choose its building blocks? How can the robot plug blocks together? In which way does the robot change a construction by unplugging blocks and inserting others instead?

Here, we do not ask for the strategy and for the supervision and control, but for the internal mechanisms of the present approach to meme media evolution.

Ultimately, a medical therapy plan is a tree-like graph outgrowing from a certain MedicalRecordPad as shown in figure 6. The leaves of the tree are actions to be executed at some future time point. From a conventional perspective, these leaves together with their partial ordering imposed by the underlying tree structure and with their immediate neighbouring relations inherited from the action hierarchy (cf. figure 10) form the plan.

Here, we are going to exemplify the meme media evolution technology by discussing three aspects in some more detail:

1. finding building blocks for planning,
2. finding consistent connections,
3. establishing new structures.

For finding suitable building blocks when generating a new plan, the planning algorithm needs to be able to search a meme pool for all available pads. In the authors' approach, pads have a lucid XML structure such that all relevant knowledge can be accessed.

The authors are aware of the requiorement for efficient implementations and for smooth embedding into given IT infrastructure. Nevertheless, the prefer a logical representation, in the sequel, for doubtless clarity.

[1] From the application perspective and at a first glance, this might seem to be a question of minor importance. Those interested in speeding up and advancing medical therapy planning might say that they don't care whether or not knowledge evolves, if only plan generation works well. But in the long run, if knowledge truly evolves, therapy planning will grow on an unforeseeably richer knowledge base. So, in its right perspective, the memetics' approach to knowledge evolution is substantial to medical therapy planning – and so is the present short chapter in this publication.

There is an implementation in Prolog which yields the desired behaviour with a predicate call as follows. For a given requirement `req`, we may search pad in a meme pool through the use of the following Prolog query:

```
?-pad_in_meme_pool(Pad),satisfy(Pad,req).
```

Let `p1` and `p2` be two pads. For a slot `s1` in `p1` and a slot `s2` in `p2`, the following Prolog query conncects `s1` to `s2`

```
?-connect(p1,s1,p2,s2).
```

where the predicate `connect` is defined accordingly. This is the point where automatic plan generation does really take place.

Connections established as above are not necessarily consistent. With an appropriate implementation of consistence checking – which is far beyond the present paper – one can step from just plugging pads together towards consistently establishing new structures.

```
?-satisfy(p1,s1,req1),satisfy(p2,s2,rec2),connect(p1,s1,p2,s2).
```

Whenever such a query succeeds the pad `p1` becomes a child of the pad `p2`, and `p2` becomes the parent of `p1`.

These steps of (automated) reasoning are some of the internal technicalities which take place when planning goes on and which, in the very end, are the establishing mechanisms of evolution.

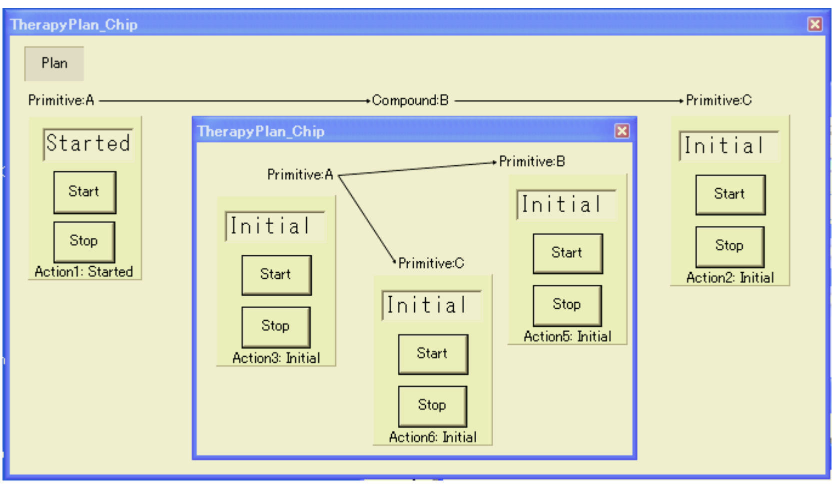

Fig. 12. Graph Expansion at Work Implemented in the CHIP Meme Media System

Last but not least, we should always be aware of the biotope issue [7] in meme media implementation and application:

We are the biotope or, at least, we belong to the biotope, together with all the regulations, formalisms and the like constraining social processes of health care.

8 Summary and Conclusions

Though our work is still in its earliest stage, it has brought up a number of insights and requirements which, in turn, lead to proper contributions to the current technology development. There are inventions in the large like, e.g., the introduction of sensor pads including the invention of corresponding architectural and communication concepts. And there are inventions in the small like, e.g., alternative approaches to communicate streams of patient data through pad hierarchies. One of them, the virtual slot approach resembles the view concept adopted from databases.

Another proposal recently brought up aims at the invention of agents that take care of constraint monitoring, for instance. A number of questions about agency do arise. Shall actions be seen as agents? If so, should these constraint monitoring agents sit inside of action agents? Or are they separately located and watching actions agents from a distance?

Talking about agency, we are obviously dealing with agents of a different time. As mentioned, there might be a hierarchical structure of nested agents. The modules of figure 1 may be seen as agents among which the planning agents plays a prominent role. The other distinguished agent is the one for plan execution. In the medical domain, this one may stay non-computerized for still some time. Thus, our general approach does inevitably rely on the communication between computer programs, between computerized agents and humans, and between humans, naturally.

Due to the lack of space and time, in particular, some quite exciting issues have not been discussed in sufficient detail. A crucial of them is the issue of communication among humans and therapy planning agents.

It is clear, especially in medical applications, that computer-generated results like medication, in detail, or like therapy, as a whole, do need to be given to humans for final confirmation. This, obviously, leads to the issue of communicating the computer's proposals to the human. Because mediacal therapy plans are generally hypothetic in nature, as discussed above in some detail, we are usually facing a lengthy sequence of communication acts. Understanding each other is a critical requirement which is, for the moment, postponed to a forthcoming investigation.

Acknowledgement

The authors are grateful to Yuzuru Tanaka for his influential work on meme media, in general, and for a manifold of activities supporting them in literally uncountably many ways. They also gratefully acknowledge the enormous help by Oksana Arnold beginning with her seminal work in planning and ranging to recent advice and kind permission to use material like some of the figures in the present publication. Oksana Arnold did a pioneering work on therapy planning in complex dynamic environments. Her achievements are still unsurpassed and have been serving as a rich source of inspiration.

References

1. Oksana Arnold. *Die Therapiesteuerungskomponente einer wissensbasierten Systemarchitektur für Aufgaben der Prozeßführung.* infix, DISKI, Bd. 130, 1996.
2. Oksana Arnold and Klaus P. Jantke. Inductive Program Synthesis for Therapie Plan Generation. *New Generation Computing* 15 (1997) 1, pp. 27–58.
3. Susan Blackmore. *The Meme Machine.* Oxford University Press, 1999.
4. Richard Dawkins. *The Selfish Gene.* Oxford University Press, 1976.
5. Kimihito Ito. *CHIP(Collaborating Host-Independent Pads).* accessible under http://km.meme.hokudai.ac.jp/people/itok/CHIP/
6. Kimihito Ito and Yuzuru Tanaka. A Visual Environment for Web Application Composition. In *Proc of 14th ACM Conference on Hypertext and Hypermedia* 2003, pp. 184–193.
7. Klaus P. Jantke. *The Biotope Issue in Meme Media Implementation.* this volume, 2004.
8. Mark Stefik. The Next Knowledge Medium. *AI Magazine* 7 (1986) 1, 34–46.
9. Yuzuru Tanaka. *Meme Media and Meme Market Architectures.* IEEE Press and Wiley-Interscience, 2003.

Appendix: The Integration Process of a Patient into a Randomized Clinical Trial in Detail

First some general information about the application-example: the integration of a patient into a randomized clinical trial. Often a randomized clinical trial consists of several therapy protocols. For example, the protocols can differ in the age of the target group (e.g. one protocol for younger people and another protocol for elderly people). Each therapy protocol consists of several therapy-arms. The therapy-arms for one protocol can differ for example in duration, in dose or in accompanying supportive therapies. Normally patients will be distributed to the therapy-arms by randomization, in order to make the success of several therapy-arms comparable statistically.

The following example explains the connection between figure 2 and the plans in figures 4 and 5. This should give a slightly better understanding of the application domain.

Fig. 2 does not show a plan but a scheme of actions, which should be carried out in the normal case of the integration into a clinical trial. Because of different requirements in the application domain (state of health, certainty of the diagnosis) the integration can be based on different plans. We will not discuss the different requirements nor the variants of plans in this paper. The figures 4 and 5 are two examples for different plans.

The plan in Figure 4 starts with Action A1. Before A1 the patient is suspected of having malignant lymphoma. Action A1 is the referral of the patient to the biopsy. The biopsy will be performed by the surgeon (A2). The tissue (the result of the biopsy) and all necessary data will be sent to the primary pathologist because of the diagnosis (A3). The pathologist processes the tissue (e.g. by using paraffin) and creates a diagnosis (A4) – in our example: high malignant Non-Hodgkin-Lymphoma (NHL).

The primary pathologist sends the diagnosis and all accompanying data back to the surgeon (A6), who gives these data to the oncologist. If the oncologist asked the primary pathologist for data and diagnosis (A5), he gets that from the primary pathologist directly (A7).

Based on the diagnosis and the latest patient data the oncologist selects a suitable clinical trial and a therapy protocol (A8). In our example, this is the protocol NHL-B, which is a protocol of a German clinical trial on therapy of malignant Non-Hodgkin-Lymphomas.

The selection of a suitable protocol often is very complicated. There exist a lot of clinical trials, each has several protocols. The number of available protocols is very dynamic. The number of clinical trials increases and protocols are valid only for ca. 5 years. In general the description of clinical trials and protocols is not available in electronic form. They are books of 100 or more pages. As decision support for the selection process each therapy protocol has two lists of criteria; one list to qualify and another list to disqualify a patient for this protocol. Each list consists of 10 to 20 items. It is allowed to integrate a patient into a protocol only if all qualifying criteria are true and no condition in the second list holds. So the selection of a suitable clinical trial and a protocol is tedious handwork.

After the selection of a suitable therapy protocol (A8) the oncologist registers the patient for that protocol in the corresponding head office (A9). To do so, the patient's data and some signed forms have to be sent to the head office. The head office performs the randomization (A10). The oncologist gets the result of the randomization and all protocol-specific data (A11). Only than the treatment and the detailed therapy planning can start (A12).

To avoid diagnosis errors, the primary diagnosis must be confirmed by a reference diagnosis. The head office asks the primary pathologist to give the tissue and all necessary data to the reference pathologist (A13). At the same time the head office asks the reference pathologist for the reference diagnosis (A21). After the primary pathologist has sent data and tissue to the reference pathologist (A14), the reference pathologist produces the reference diagnosis (A15). This diagnosis is sent to the head office and to the primary pathologist (A22 and A18). The oncologist gets this diagnosis directly from the primary pathologist only if he requests it. If not, he gets the diagnosis via primary pathologist and surgeon some days later. Now the rest of the tissue can be used by other institutions (A20).

Finally the oncologist, primary pathologist, surgeon, and head office compare both diagnoses. If they match, all is fine. Otherwise one must stop the started treatment, make a new diagnosis, etc. But this case is neither shown nor discussed in the running case study of our present paper.