Instructional Science **24:** 259–293, 1996.

Modelling students' construction of energy models in physics

ROSHNI DEVI[1], ANDRÉE TIBERGHIEN[2], MICHAEL BAKER[2] & PAUL BRNA[3]
[1] *Department of Mathematics and Computing Science, University of the South Pacific, Suva, Fiji;* [2] *C.N.R.S.-GRIC, COAST Research Team, Lyon, France;* [3] *Department of Computing, University of Lancaster, Lancaster, UK*

Abstract. One of the main differences between novice and expert problem solving in physics is that novices mostly construct problem representations from objects and events in the experimental situation, whereas experts construct representations closer to theoretical terms and entities. A main difficulty in physics is in interrelating these two levels, i.e. in *modelling*. Relatively little research has been done on this problem, most work in AI, psychology and physics education having concentrated on how students use representations in problem solving, rather than on the complex process of how they construct them. We present a study that aims to explore how students construct models for energy storage, transformation and transfers in simple experimental situations involving electricity and mechanics. The study involved detailed analysis of problem solving dialogues produced by pairs of students, and AI modelling of these processes. We present successively more refined models that are capable of generating ideal solutions, solutions for individual students for a single task, then models for individuals across different tasks. The students' construction of energy models can be modelled in terms of the simplest process of modelling – establishing term to term relations between elements of the object/event 'world' and the theory/model world, with underlying linear causal reasoning. Nevertheless, our model is unable to take into account more sophisticated modelling processes in students. In conclusion we therefore describe future work on the development of a new model that could take such processes into account.

Introduction

Modelling is a fundamental activity in all physical sciences, and particularly in physics itself. However, this activity is generally reserved for advanced students and scientists themselves. If we are to help students to one day gain a deeper understanding in science, and to bridge the gap between a naïve phenomenology on one hand, and on the other effective manipulation of equations, we believe that students need to learn modelling during their time at school. However, our knowledge of how to assist students in this enterprise is relatively poor since we do not have very reliable models of the relationship between student's knowledge of the domain and the models that they build.

Research in cognitive psychology and artificial intelligence over the last thirty or so years has given us a reasonable level of understanding of problem solving processes. Two main approaches have been pursued: Piaget's basic approach (Inhelder & Piaget, 1979) and the information processing approach developed principally by Newell and Simon (1972). In the first case, starting from a problem solver's knowledge, the researcher looks for how knowledge produces actions in given situations. Applying knowledge in a specific situation is done 'through' representation which means that knowledge needs to be activated in order to produce representations that result in actions (Richard, 1990: 251). In the second case, a search space is defined by constructing a representation of the initial state of the situation, of the goal to be achieved and of the available problem-solving operators. The researcher then looks for search heuristics that show how the actions observed in the situation result from the application of the operators to the current representation of the situation in the search space (Richard, 1990).

Instead of seeing these two points of view as opposed, we consider that the learners' representations of problems are constructed both from their own knowledge and from their interpretation of the situation. In this paper we focus our analysis on *the processes by which students construct representations*, in relation to the different kinds of knowledge that they already possess (in long term memory), to information given in the problem-solving situation, and to the heuristics used to solve the problem.

Most research within these cognitive science paradigms has, indisputably, been concerned with modelling students' problem-solving processes, *given an initial problem representation*, usually input by the researcher. Some research has been done, however, on how students construct problem representations themselves, but usually from textual or verbal statements of the problem (e.g. Lewis, 1989).

The research of Larkin (1983), on the way in which experts and novices construct representations for physics problems, has shown that the main difference between novices and experts is that the novices' representations are constructed mainly in terms of the real entities given in the problem, whereas experts' representations are more directly in terms of conceptual entities which are needed to solve the problem. Novices do also frequently construct their solutions in terms of formulae and variables. However, in this case novices often attempt to work back from formulae presented in the problem statement (means-ends resolution), whereas in the case of experts such formulae are derived directly from their representations of conceptual entities. These results are also confirmed by work on students' conceptions in physics (Trowbridge & McDermott, 1981; McDermott, 1984; Tiberghien, 1984; Viennot, 1993).

Our research question is somewhat different, and relatively unexplored: we want to understand **how students construct representations** themselves,

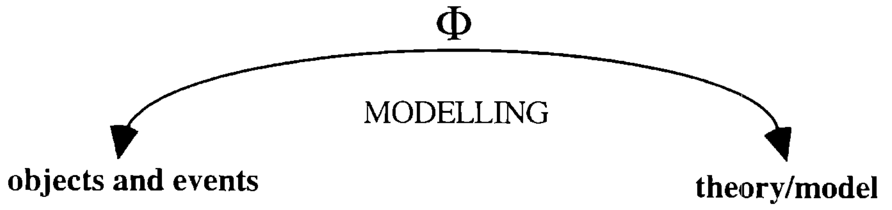

Figure 1. Modelling as establishing the semantic relation between model/theory and objects/events.

for **real experimental situations**, and with respect to **abstract conceptual entities** in a **specific domain** (energy in physics). In the case of energy, the construction of representations presents special difficulties. The students must *select* the objects and events in the experimental situation that are relevant, and establish links with abstract entities (model and theory of energy). In the other direction, clearly not all theory and model terms correspond to specific objects or events.

This may be represented as in Figure 1 above (simplified from Tiberghien, 1994). Greeno (1989) has analysed the situation in a related way but our emphasis is more on the role of knowledge and its interrelationship with learning.

So our interest is in studying how students create the relation Φ in Figure 1 in a way that is coherent to them. This relation is constructed for the experimental situation in terms of the relevant theory/model. This is, of course, a simplification of the problem, to be discussed in the next section.

We have adopted two complementary approaches for investigating the above research questions: detailed analyses of students' problem solving processes, working in pairs, and computational modelling based on our analyses. In this paper we concentrate on the latter. The students' task was to construct "energy chain" diagrams, consisting of energy reservoirs, transformers and transfers, that "translated" the objects and events of the experimental situation. For this they were provided with a specially designed computer-based learning environment ("CHENE",[1] developed in the COAST research team – see e.g., Bental et al., 1995), that enabled us to automatically record their actions on the interface. This task was conceived as an integral part of a larger teaching sequence on energy.

The paper is structured as follows. We begin by presenting hypotheses underlying design of the teaching sequence, of which the specific energy chain construction tasks studied were a part, followed by a description of the tasks themselves and the experimental situations. In the second main part of the paper, we concentrate on AI modelling of the protocol data. Beginning from a model for an "abstract" (or "typical") student for one experimental task (battery-wires-bulb circuit), we compare the model behaviour with analyses of dialogue data. Then we discuss modifications to the abstract student model to account for specific protocols, and for developing student expertise. Finally, we discuss limitations to the work described and further work to be done.

Hypotheses underlying design of the problem solving tasks

All of our research on the teaching of energy has been carried out on students at the level of the beginning of "Lycée" (16–17 years old). Here, we concentrate on analysis of the modelling[2] process, rather than on the teaching content itself (see Tiberghien, 1994).

The modelling process in physics has been analysed (Tiberghien, 1994) in terms of three main 'worlds': **theory**, **model** and **experimental field**. In the rest of this paper, we group '**theory-and-model**' into a single 'world', since this is sufficient for analysing the tasks described here. This includes physics paradigms and principles (e.g. conservation of energy), the properties of energy, and relations between energetical physical quantities and the associated symbolic representations. The experimental field corresponds to all the objects and events involved in the experiments.

We consider these two "worlds" as corresponding to different semantic fields: the words introduced in the theory/model that relate to energy have a distinctly different meaning from that found in everyday language. Moreover the language of the theory/model has to be consistent with the symbolic representation of the energy chain.

In order to study the relations students establish between the two 'worlds', we need to design a task where one of the main problems in physics learning is eliminated: the **radical gap** (or "incommensurability") **between the students' theories and the theory of physics as it is taught**, even if – obviously – we do not claim that they are the same. In this perspective, we draw on the results of work on students conceptions, particularly on electrokinetics, and also on a synthesis of different areas of this research work. In electrokinetics (Shipstone, 1985; Shipstone et al., 1987; Closset, 1983; Johsua et al., 1986), it is well known that students use "sequential reasoning", which is associated with simple causal reasoning. In these terms, the "generator" is the agent which provides a constant flux whose magnitude does not depend on the whole circuit. We know that even young children (9–12 years old) use causal reasoning to interpret a battery-bulb experiment: for them the battery is the agent which contains electricity and which gives it to the bulb, which then shines. In addition, research on students' conceptions (Driver, Tiberghien & Guesne, 1984) shows that most of the students preferentially use only *one* variable (when in physics several are necessary to predict events correctly) as well as causal reasoning between an agent and a patient.

Two final related hypotheses with respect to learning underlie our design of the energy teaching sequence:

(1) when students are presented with new material to be learned, they need to understand how to relate it to what they already know;

(2) in order for learning of new material to take place, a necessary condition is that the students recognise an *intellectual need* for this material.

With respect to the second hypothesis, our claim is that one way of inducing an intellectual need in students, is to present them with problems to solve which are just beyond their current knowledge, and where they will thus be led to recognise for themselves that they cannot solve the problem without new knowledge, i.e they recognise an *intellectual need* for the new knowledge. We then require that the student is provided with information that will act as a source for learning -we term this process one of (metaphorically) *providing the seed* for learning the necessary theoretical knowledge.

We therefore designed the following teaching sequence.

1 – We started from this conception to construct the "seed" of the energy theory/model which is initially presented to the students (see Table 1). In this case, there is only one physical quantity (energy), and a simple linear causality between a first reservoir and the following transformer is compatible with the presented theory/model.

2 – We chose a battery – bulb experiment as the first task for which the students have to perform a modelling$_{SP}$ activity based on energy. In particular, two aspects are important:

- the role of the battery as a reservoir giving "something" to the bulb in a causal relation between the battery (agent) and the bulb (patient). This is close to the way in which the energy chain can be interpreted. At this introductory point in the teaching the energy chain can be considered as a linear causal chain.
- the fact that "the battery is being used up" is relevant from the energy point of view but not from that of electrokinetics. Thus a difficulty can appear for the students who have learnt electrokinetics as a conflict between their knowledge of electrokinetics and the rules given for the construction of energy chains.

In order that the students' performance of this task should allow us to analyse their knowledge (from their long term memory or from the available information) and the heuristics that they use to establish relations between the two worlds (theory/model and objects/events) we need to be able to separate these worlds in our data. This is why we control the information available on

Table 1. Information and rules for energy chain construction provided to students

Energy can be characterised • by its **properties:** • **Storage** a reservoir stores energy • **Transformation** a transformer transforms energy • **Transfer** Between a reservoir and a transformer, or between two reservoirs, or between two transformers, there is energy transfer. The different *modes of energy transfer* of a system to another one are: work, heat and radiation. • by a **fundamental principle of conservation** Energy is conserved whatever its transformations, its transfers or its forms of storage	To build an energy chain you have to use these symbols and take into account these rules: Reservoir Transfer 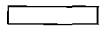  Transformer 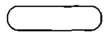 • a complete energy chain starts and ends with a reservoir • under each rectangle indicate the corresponding object (or the objects) in the experiment • under each arrow indicate the mode of transfer. If there are several modes of transfer use one arrow for each mode of transfer (between two rectangles).

the theory/model on energy; we give a real experiment to the students and we ask for an answer which only belongs to one world (theory/model).

The computer-based learning environment CHENE was developed as a research tool for studying students' modelling$_{SP}$ in terms of energy, and for assisting them in the performance of these tasks.

The problem solving tasks

The tasks are conceived within a larger scale teaching sequence for energy in physics. This sequence is designed for students who have only an informal knowledge of energy, but who have previously received teaching on electrokinetics.

Teaching sequence

The teaching sequence implies an organisation in time of the knowledge to be taught in which these tasks are included. It consists of four parts:

(1) Producing a text that describes an experiment, and analysing the text into three categories: "objects and events", "electrokinetics" and "other". This analysis has the aim that *the learner senses an intellectual need* for a new theory/model – the "energetic model" – since students usually spontaneously identify text in the "other" category as corresponding to "energy".

(2) a modelling activity for different experiments. The "seed" of the theory/model on energy is provided for the learner; his or her task being to construct the model of a given experiment. The aim is that *the learner constructs a meaning of the theory/model when establishing relations between the theoretical construction and the experiment.* We consider that the theory/model provided is only a 'seed', since the students will have to refine the model in the subsequent teaching sessions;
(3) development of the model on the initiative of the students to interpret experiments for which the knowledge to be taught uses two physical quantities (energy and power) and quantitative relationships (the chosen experiment consists of a coil heater with an electricity meter, a thermometer and a chronometer);
(4) enlargement of the field of applicability of the quantitative model with construction of a new parameter (efficiency), and the qualitative idea of the "quality" of energy.

At present, CHENE is used in part 1 (analysing a text) and in part 2 (modelling$_{SP}$). Here we attempt to model$_{AI}$ only part 2: energy chain construction.

Modelling$_{SP}$ tasks

We now focus on aspects of the sequence for which CHENE is used. These tasks provide the elements with which learners may themselves *construct meaning* for the theory/model given as such to them. The information on the theory/model is given to the learner between task 1 and task 2 (see Table 1). It is available to the learner during all these tasks. These tasks consist of drawing a symbolic representation in terms of the model (energetic chain) for a given experiment.

The interface of CHENE has two roles here. Firstly, it provides the learner with the possibility of accessing information of different types in order to help with problem solving (e.g. the information on the energy model shown in Table 1). Secondly, it provides a means for drawing (graphical construction) the solution, in terms of a specific form of expression: rectangles (two types) for reservoirs and transformers of energy, arrows between rectangles to represent transfers of energy and labels on the arrows to represent modes of energy transfer. On each rectangle, the student can indicate the corresponding object (or objects) in the experiment.

First experiment: As we have already discussed, we chose "battery – bulb – wires" as the first experimental situation presented to the students because we assume that its energy chain point of view is not too far from the students' current point of view. It is important to remark that for this first experiment,

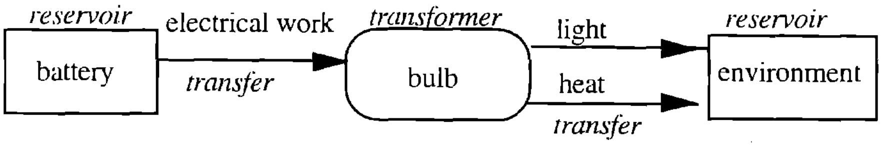

Figure 2. Example energy chain – the "ideal" solution for the battery-bulb-wires experiment.

the students have never seen an energy chain before. The only information they have on this chain is given in Table 1, their only knowledge about energy coming from informal sources. After the construction of their chain (without the teacher's help) the students are given a solution (see Figure 2) and discuss this with their teacher. They then have to construct other chains for the following experimental situations. Between each chain they are not given the solution.

Second experiment: An object is attached to a string wrapped round the axel of a dynamo (in fact it is a motor working as a dynamo) which is connected to a bulb. When the object falls the bulb shines.

Third experiment: A battery, a motor, a pulley and an object attached to a string. At the beginning the object is down and when the circuit is closed, the object rises.

This order of tasks was chosen because we assumed that the second experiment is easier to solve than the third one. In fact, the difficult point in the second one concerns the elements of the experimental devices which correspond to the first reservoir. Given that the students use linear causal reasoning, they may identify the object falling as being the cause of the following events – the motor (as a dynamo) runs and the bulb shines. The students then match the falling object as the first reservoir. In this experiment we have a completely different first reservoir from the first one, it is not associated with a single object as in the case of the battery. For the physicist, the reservoir is the system object – earth, but for the knowledge to be taught at this step of teaching (introductory sequence), the first reservoir can be considered as the falling object. Therefore it is a property of an object (falling and giving movement to the motor) and the the object itself which is associated to the reservoir. In the third experiment, the difficult point corresponds to the last reservoir which is the object rising. This association mainly comes from a rule of the model and not from the students' very frequent simple causal reasoning.

These tasks are *generic* in the sense that they play a (proto)typical role in modelling activities: creating a need for a new theory/model, constructing a meaning of the theory/model by establishing relations between these levels and that of the objects/events. They are not specific to a given situation of the experimental field, nor to the age of the students (such tasks can be relevant at middle school level or at university level).

We collected data for 6 pairs of students, 3 from two Lycées in Lyon, working together on the CHENE environment in the classroom during normal lesson time, and 3 from data obtained with students who came into our laboratory. Their dialogues were recorded and transcribed for analysis. Other data collected included an automatic trace of the students' actions on the interface (together with timing in seconds), and videos of their manipulation of the experiment. This data enables us to build up an accurate picture of their problem-solving activity.

Analysis of students' problem solving using "psCHENE"

Choice of a modelling$_{AI}$ approach

Given the nature of modelling$_{SP}$, and the fact that little research in cognitive science has so far attempted to construct AI models for it, the choice of an initial modelling$_{AI}$ approach was problematic. As mentioned in the introduction, modelling$_{SP}$ involves a complex process of selecting relevant elements from, and establishing correspondences between, at least two representational 'levels', each of which has its own rules for combination of elements, via additional task-related knowledge. In the case of students, the latter knowledge is to be (co-)constructed during task execution.

One possibility for modelling$_{AI}$ such a task would thus be to define declarative representations of the two 'levels' (e.g. using semantic networks, or schemas), and then to define complex matching mechanisms for establishing relations between the two. However, for the task studied here, relations between elements *at a given level* are not *fixed* initially, since the experimental field has to be constructed, as does the energy chain model, from a set of 'syntactic' rules. The way in which knowledge is grouped into schemas (even if hierarchically structured) in an initially fixed manner by the system designer does not therefore seem to lend itself 'naturally' to modelling$_{AI}$ this aspect. Secondly, any schema, or analogical matching system, must in any case contain *rules* at some level, for deciding what elements can match or not, whether these are 'embedded' in the matcher itself or made explicit. Thus, for example, Falkenhainer, Forbus and Gentner's (1989) AI model for analogical matching was defined in terms of a rule-based system. Finally, even if we were to choose a schema-based system, for example, Anderson (1983, pp. 36 et seq.) has described how such systems may in fact be simulated using rule-based systems: it is therefore possible to develop such AI model without making too many firm architectural commitments.

Given these considerations, we have initially chosen a well-established approach for modelling$_{AI}$ of modelling$_{SP}$: *production-rule systems* (see below for detailed discussion). This may appear paradoxical given that such systems have usually been used for modelling so-called procedural skills (see below), rather than for the process of constructing problem representations. However, on the one hand we do not necessarily subscribe to the *learning mechanisms* that are usually associated with production-rule systems, and on the other hand, at this exploratory stage of our research we do not claim that the process of modelling$_{SP}$ is in some sense a 'procedural' cognitive process. Rather, this first stage of our research is to be viewed as a *knowledge elicitation* exercise. By this we mean that our initial research goal was to produce a formal (rule based) model for modelling$_{SP}$, that incorporated knowledge that could reasonably be attributed to the students on the basis of analysis of their protocols. The process of developing such an AI model obliged us to render the students' knowledge explicit, and to make hypotheses concerning the additional knowledge required in order to enable it to reproduce the students' solutions in a particular series of intermediary stages. The fact that the model is able to produce these solutions, with this ordering of stages, shows at least that it is *logically* possible to do so on the basis of this knowledge.

In terms of the correspondence between model and data, therefore, this is to be found at the levels of 'input-output', intermediary problem-solving stages, and the knowedge used, rather than at the level of a particular cognitive architecture. Our model would have been deemed to be *invalidated*, therefore, to the extent that it: (a) was not capable of reproducing the students' solutions, (b) did not produce these solutions in the macro-level order (see below) that the students' produced them in, and/or (c) it relied on ad-hoc facts and rules to an extent that surpassed the small number of 'house-keeping' rules that are required to make any real production-system run.

Clearly, although many further cognitive constraints (such as on memory) would need to be imposed on the model in order for us to be able to claim 'cognitive' validity, we view the research presented here as a first step along the way to such a model. We discuss in conclusion the ways in which the present AI model is in fact insufficient to take into account more complex modelling$_{SP}$ processes, and how the knowledge incorporated within it could be recast into a more suitable framework in further work.

The production-rule modelling$_{AI}$ approach

Our AI modelling approach was to begin with a production rule model ("psCHENE" = "production system" CHENE) that produced "ideal" performance for a given population of students on the problem-solving tasks studied. This was based on our existing knowledge of students' conceptions. We term this the model for an **abstract student**. Then we aimed to refine the model so that it could model *individual students' performance*, beginning with *one task*. Subsequently, these individual models were refined further, so that they could account for *individual performance* across *different tasks*.

On this point, our use of data consisting of dialogues between two students as the basis for models of individual students requires comment. We are interested in studying students working in pairs for a number of reasons. Firstly, because the role of dialogue itself in collaborative problem-solving is a subject of interest (Baker, 1994), but secondly for the methodological reason that in this context the students' externalisation of their problem solving processes will be intrinsic to their task (O'Malley, Draper & Riley, 1984) since the students were required to reach agreement on a common solution. In future work we intend to extend the individual models to a model for collaborating dyads.

Our modelling$_{AI}$ work has been elaborated in the framework of a production rule system psCHENE that we present below. We initially chose to implement our own system to retain full control over architectural features rather than use an existing production system such as GRAPES, PRISM (Langley, 1983) or OPS5 (Brownston et al., 1985). The modelling$_{AI}$ described in this paper did not require a very complex architecture. It was therefore felt that, as ideas developed (which often involve domain-specific needs), it would be an advantage to be able to modify our own interpreter. In the longer term, we require a more structured architecture in order to provide greater assistance for the representation of student activities of different kinds. For example, to permit us to model those parts of a protocol in which a student discusses his/her state of progress.

Modelling using the psCHENE production rule system

psCHENE is a forward chaining production rule system (Brownston et al., 1985) implemented in Mac Common Lisp v2.0. For its conflict resolution strategy, a rule instantiation is selected by matching with the working memory item that has the greatest strength. In this case rules which match with those items in the working memory that have been asserted more recently are selected over those matching older working memory items. The final element of the strategy is to use the rule which is found first in the rule database.

The system is designed to model the "generic tasks" that we discussed in the previous section. We present this system in terms of modelling an "abstract" student working on the battery-bulb experiment. In so doing, we also introduce some of the key domain specific concepts.

In the initial task, there is an electrical circuit featuring a battery, two wires and a bulb in front of the students. The problem itself is based on a set of objects featured in an experimental situation. The system maintains a representation of the structure of the experimental setting.

Modelling and the "Abstract" student: the battery-bulb experiment

For the purposes of modelling students, we chose the level of granularity of knowledge and of the rules used for problem solving to be consistent with the level of our protocol analysis as a check on the validity of our modelling. This level of analysis is approximately equivalent to statements made by the student. This is also consistent with the level of granularity at which we represent the text for the theory/model on energy (Table 1). This makes checking the behaviour of the model against the protocols much more straightforward. A further important point is that additional rules needed to represent implicit knowledge in an explicit manner are at the same level of granularity.

In this model, we adopt an hypothesis relating to students' reasoning that we have made in designing the teaching-learning tasks: students develop energy chains by selecting objects in the experimental situation in a *standard sequence*, according to *linear causal reasoning*, i.e. the causal agent (battery) first and then the causal patient (bulb).

Clearly, deciding what order to consider the objects, and which ones to select to model, is an important problem for modelling$_{SP}$. However, our studies have revealed that the objects *are* always considered in this order: the battery first and then the bulb or the wires. Consequently, in this initial system, we have 'hard wired' this standard sequence into the system. Thus the "abstract" student is considered to use linear causal reasoning.[3]

Next we illustrate what we regard as an efficient solution to the problem as performed by the "abstract" student.

Initially, the abstract student may have in mind several different representations of the battery-bulb experiment building from his/her electrokinetics knowledge and/or his/her everyday knowledge – e.g. the electrical current flows all round the circuit, the battery is a generator and the bulb a resistor. In contrast to this descriptive knowledge he/she may also be able to use sequential reasoning (knowledge with explanatory power) such as that the battery gives current to the bulb and allows it to light. Such reasoning is embedded in a causal relation. The "abstract" student also knows from everyday knowledge that the battery has energy.

Then (s)he takes information from the text on the theory/model (germ) and interprets it using his/her everyday knowledge. (S)he learns that there are reservoirs which store energy, transformers which transform energy, and three different modes of energy transfer: heat, work and radiation (i.e. in the form of light).

So (s)he can start by relating the battery which has energy with a reservoir which stores energy and concluding that the battery is a reservoir. The student might then remark that, as current flows from the battery to the bulb then electrical work must be done. Since work is a mode of energy transfer there is a transfer of energy from the battery to the bulb. Another line of reasoning results in the deduction that the bulb gives out both heat and light to its environment. From this, the student might infer that the environment is a reservoir. Since the bulb receives energy from the battery in the form of electrical work, and gives energy to the environment in the form of heat and light then the bulb is a transformer. These steps in reasoning result in the building of an energy chain such as that found in Figure 2.

Working memory

Working memory is represented by a global variable, and is initialised to a set of facts that represent a student's previous knowledge. As the interpreter performs the actions associated with the rules, items are added to the working memory. The working memory is a list containing its elements in embedded lists.

We considered that an "abstract student" A:

- using his/her everyday knowledge, takes into account the objects and events of the experiment: *bulb lights up environment, bulb heats environment, wire between battery and bulb*
- uses his/her everyday knowledge associated to the fact that s/he is studying energy: *battery has energy*
- uses his/her knowledge about electrokinetics: *battery is a generator, there is current flow in an electrical circuit.*

The following is an example of the structure of working memory:

```
( ... (has energy battery 10)
(receives current bulb 0)
(gives light bulb 0))
```

The numbers at the end of each item represent the strength associated with the item. In the following we present details on the function of the strength values and how they are calculated. Working memory is updated by adding items at the front, that is the left-hand end of the list.

Production rules

A rule is represented in terms of four parts:

i) The name of the rule.
ii) a list containing one or more patterns, where each pattern is a 'condition'.
iii) a list containing one or more patterns to be added to the working memory, where each pattern is an 'action';
iv) the "strength" of the rule, which determines order of firing.

The following is the syntactic representation of a production rule:

(M2 ((transforms =A to =B =X =St)) $\rightarrow$ ((is =X transformer 0)))

which represents 'if an object transforms something to something else, then it is a transformer' (=St and 0 are concerned with strengths). A set of such rules are formulated to represent steps in problem solving.[4] Conditions can have negated items. Apart from negated clauses, the conditions are in such a form that testing against working memory items involves simple pattern matching. The actions contain clauses which are like those in the conditions (except for negations). These clauses are added to the working memory directly, with the variables instantiated. A negated condition succeeds if its non-negated condition does not have a matching item in the working memory.

Turning to the rules, we consider that the abstract student uses the information concerning the theory/model given to them directly. This information includes the statements "the reservoir stores energy" or "the transformer transforms energy".

This class of rules arranges for actions such as transfers, transformations, and storing to be associated with the corresponding object roles transformer and reservoir. For example:

IF X transforms energy,
THEN X is a transformer

One way in which students infer that an object of the experiment is a transformer is that it "receives" one form of energy (or mode of transfer) and "gives" out another.[5] We believe that the use of the words "give" and "receive" are suggestive of the students building a consistent explanation of how the energy chain works. This information is turned into the rule:

IF Y receives S from X and Y gives T to Z,
THEN Y transforms S to T

The rule is specific to this domain and is not intended to be general for all possible values of the variables 'S' and 'T'. For example, if the rule had been derived from everyday knowledge, where 'S = water', 'T = wine', 'Y = John', 'X = Mary', 'Z = Julie', then one possible instantiation would be: "IF John receives water from Mary and John gives wine to Julie, then John transforms water into wine" (!).[6] The rule is domain dependent to the extent that S and T are both assumed to be *forms of energy*. What is 'transformed', therefore, is the *form* of energy. An everyday example would therefore be rather where John receives water in the form of ice and gives it in the form of a liquid. However, it is *linear causal reasoning* (see our previous hypothesis) – rather than more general everyday knowledge – that underlies this rule, since the object identified as the transformer is identified as the direct *cause* of the difference between input and output. A further problem, that is not addressed in the present model, could arise in the case where an object has more than one input and/or output.

There are a few rules that are closer to 'housekeeping' rules: ones that facilitate problem solving but ones for which there is very little evidence. For example, there is a rule for deducing that if a transfer takes place from one object to another then the receiving object is either a transformer or reservoir. While this is a logical deduction it is by no means certain that students engage in this kind of reasoning. The protocols provide little evidence here.

Finally, in the initial model for the 'battery-bulb' experiment, there is a special rule for the "environment" (rule 15), which simply says, in effect, that *the environment is a reservoir,* i.e. we assume that the students simply 'know' this (in fact, they usually only possess this knowledge when *told* by the teacher, in a 'correction session' after their problem-solving). It is reasonably plausible to model students in terms of whether this rule is known or not since it is hard to see how students might modify their view of the experimental situation to include the environment. We present the rules in Table 2.

The rules found in Table 2 often feature actions that relate to the theory/model but some actions relate more closely to everyday knowledge. However the status of each rule is in question since the rules do not make clear the intended meaning of variables. These variables could, on the surface, be any kind of entity whatsoever. We note that students also fail to discriminate between kinds of entity – however here we do not claim that the use of variables in the rules in any way represents the student's use of rules. These rules therefore have to be interpreted very carefully. When such care is taken they have a utility that relates to the level of granularity of knowledge adopted.

Rule 9 associates two types of knowledge: electrokinetics and linear causal reasoning. The first two clauses of the antecedent state that there is electrical current in the circuit between the battery and bulb, and the third that the

Table 2. Rules in psCHENE

No	*Rule*
1	IF X stores energy, THEN X is a reservoir
2	IF X transforms from A to B, THEN X is a transformer
3	IF X gives energy to Y, THEN [energy goes from X to Y] & [X cannot receive energy from Y] & [there is a transfer of energy from X to Y]
4	IF [Y receives S from X] & [Y gives T to Z], THEN Y transforms S to T
5	IF there is electrical work from X to Y THEN the transfer from X to Y is electrical work
6	IF there is a transfer of Z from X to Y, THEN Y is a transformer or a reservoir
7	IF X is transferred from Y to Z, THEN Y gives X to Z
8	IF X can not receive Y from Z, THEN there is no transfer of Y from Z to X
9	IF [electrical current from X to Y] & [electrical current from Y to X] & [X is a generator of energy] THEN there is electrical work from X to Y
10	IF X gives A to Y, THEN Y receives A from X
11	IF X receives A from Y, THEN Y gives A to X
12	IF X has Y, THEN X stores Y
13	IF X heats Y, THEN there is a heat transfer from X to Y
14	IF X lights Y, THEN there is a light transfer from X to Y
15	IF the environment is a transformer or a reservoir THEN the environment is a reservoir
16	IF there are wires between X and Y THEN [electrical current from X to Y] & [electrical current from Y to X]

battery is the 'first cause'. Given that "work" is defined as a form of energy transfer, on the information provided to the students (see Table 1), these two types of knowledge together with this information imply that the work is *electrical*, and that it is transferred *from the battery to the bulb.*

The Recognise-Act cycle

The three main functions of the interpreter are matching, conflict resolution and rule execution (the so-called recognise-act cycle). In the process of solving a given problem, the interpreter performs the cycle of matching rules against working memory items, selecting a rule and 'firing' it, until either the problem has been solved, or there are no more rules to 'fire'. A problem is solved if the following two criteria are satisfied: (i) the energy chain begins and ends with a reservoir and (ii) each item in the energy chain has an object associated with it from the experimental situation. Perhaps the most important thing to note about the matching process is that different variables in the production rules cannot be instantiated to the same constant.

Conflict resolution makes use of some information from the matching process along with the strengths associated with different rules. Initially, strengths of working memory items are calculated based on their order in the variable *assembly* (strengths of each working memory item are currently either 0 or 10).[7] At the beginning, all working memory items that contain the first element of *assembly* are updated to have a strength of 10. At each cycle, the interpreter checks whether some link has been made between the partial solution to the problem and the experimental situation being modelled. If so, then the strengths of all working memory items that contain the next item in the *assembly* are increased to 10, and the strengths of all other working memory items reduced to 0. This assignment of increased working memory strength to the 'next item' in the predetermined structure of the experimental field therefore corresponds to one of the ways in which we have incorporated linear causal reasoning into our model (other aspects are incorporated in certain rules shown in Table 2).

Solution path

In this section, we present the way psCHENE solves the battery-bulb task. Given an initial model of the world represented as follows,

((has energy battery 0)
(battery is generator 0)
(lights environment bulb 0)
(heats environment bulb 0)
(wires between battery to bulb 0))

together with the set of rules which we have discussed in the previous section, the production-rule system generates the solution path given in Table 3.

This corresponds to the final "ideal" solution is as previously shown previously in Figure 2.

The run leads to success with the role of reservoir being assigned to the battery and to the environment, and the role of transformer being assigned to the bulb. Note that the energy chain begins and ends with a reservoir which is consistent with the requirements of the theory/model. Also note that the wire is not assigned a role in the energy chain – consistent with the solution sketched previously.

This table shows the importance of everyday knowledge not only in the interpretation of perceptions. It turns out that this knowledge is fundamental to the interpretation of the theory/model on energy. For example, everyday

Table 3. Solution path in relation to the types of knowledge used. The left hand column features the solution path as produced by psCHENE. The right hand column features a simple analysis of the ways in which different sources of knowledge are associated to produce the final result. We have highlighted the energy-chain solution elements indented and in **bold**. Also note that "K" stands for Knowledge (e.g. Everyday K stands for everyday knowledge) and "En." stands for energy

Solution path	*Origin*
(stores energy battery)	Association
(is battery reservoir)	Everyday K – En. theory/model K
(electrical current from bulb to battery)	
(electrical current from battery to bulb)	
(electrical-work from battery to bulb)	Electrokinetics K
(is electrical-work transfer from battery to bulb)	En. theory/model K
(gives electrical-work battery to bulb)	
(receives electrical-work bulb from battery)	Everyday K – En. theory/model K
(bulb is transformer or reservoir)	Association
(is light transfer from bulb to environment)	
(environment is transformer or reservoir)	Association
(is environment reservoir)	
(gives light bulb to environment)	Association
(receives light environment from bulb)	Everyday K – En. theory/model K
(is heat transfer from bulb to environment)	
(gives heat bulb to environment)	Association
(receives heat environment from bulb)	Everyday K
(transforms electrical-work to light bulb)	
(is bulb transformer)	

knowledge is used to construct a meaning for reservoir and storage, transformer and transformation. Moreover a meaning for the association between reservoir, transformer and transfer is constructed through a linear causality associated with the action of giving and receiving and which is used frequently in everyday life.

Making the association between different types of knowledge explicit allows us to analyse how students establish relations between the conceptual world and the object/event world. The two associations battery-reservoir and bulb-transformer are made term to term between these two worlds. The same property of action: storage or transformation are given to both elements of these two worlds. In fact in these two worlds the same causality is used which allow this term to term correspondance.

This statement is confirmed by the results obtained in real classrooms where the written chain is obtained with a similar protocol (the teacher does not give any help to the pupils). 99% of 143 pupils associate battery and reservoir and 95% associate transformer and bulb.

Concerning the notion of transfer, term to term correspondance is not possible between a "real object" and an element of the model. There are several steps:

- first, the actions of heating or lighting have to be transformed into "transfer of heat (or light)";
- second, recognise that this transfer is an energy transfer (this step is implicit in our rules)
- then consider that this transfer comes from from "something" to "something else" which is translated in our rules by to give and to receive;
- then the recognition of these "somethings" has to be done but they can be done at the model level or at the object/events level. In our modelling, we choose the level of object/event: "heat or light transfer from bulb to environment".

In the case of the transfer between the battery and the bulb, we used electrokinetics knowledge, if is a current and there is electrical work.

This analysis illustrates how much an apparently obvious perception (the bulb shines or heats) has to be transformed by cognitive treatment using the theory/model knowledge in order to be used in physics modelling. It is not everyday knowledge which allows the student to go from these perceptual facts to "heat" as a mode of energy transfer.

In psCHENE the sequence of actions of the rules that fire for solving the problem is shown in Figure 3 [the rules fired, as shown in Table 2, are marked "R1", and so on].

As mentioned earlier, the initial items in working memory that represent the experimental field are fired in the order 'battery-wires-bulb-environment', with a final 'return' to the inputs and outputs of the bulb in order to determine that it is a transformer. This gives the following order for production of the energy chain solution (items shown enclosed in square boxes in Figure 3):

1. matching battery – reservoir;
2. matching electrical-work-transfer from battery to bulb;
3. matching transfer-light, transfer-heat from bulb to environment;
4. matching environment-reservoir;
5. matching bulb-transformer.

Comparison between psCHENE and students' problem solving

We describe our modelling in terms of the construction of a succession of models which are based on our detailed analysis of problem solving protocols. The initial model is sufficient to describe some, but not all, of the important

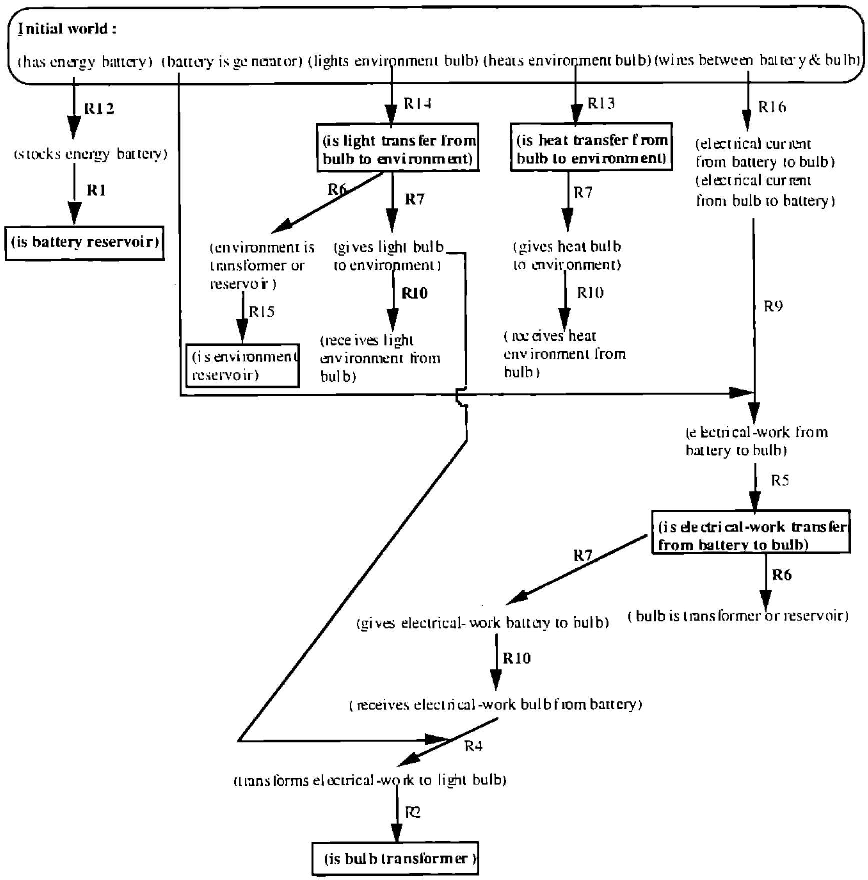

Figure 3. Rules fired in producing the ideal solution for battery-bulb experiment.

characteristics of a specific protocol of two student's attempting to solve the 'battery-bulb' task. This deliberately "minimal model" is then extended to account for a further protocol with the same students solving the 'object-motor-bulb' task. We explain how significant limitations of the initial model are parallel to problems faced by the students themselves. The model is further extended to account for a third protocol based on the 'battery-motor-object' task. Again, limitations of the model parallel problems faced by the students.

Common characteristics of performance between all students and psCHENE "abstract student"

There are two main characteristics that all the pairs of students studied, and the performance of psCHENE have in common for the battery-bulb experiment:

(1) all students associate **battery** and **reservoir** in the first minutes of their construction of the chain and, psCHENE makes this matching first.
(2) all students associate **bulb** to a **transformer** (sometimes initially to a reservoir, which they subsequently modify to be a transformer).

Moreover, the protocols validate an important hypothesis stated earlier: the use, at least, at one point of the problem solving, of **linear causal reasoning** by all students. Thus, all students use this form of reasoning at some point in their problem-solving, believing that *the battery is the cause of the lighting-up of the bulb*. For example:

"When you have a transfer which goes from the bulb to the battery, it is impossible because the light energy produced by the bulb, it cannot produce energy in the battery"
(Stéphanie, line 121, Group 1 – equivalent to rule n°: 3).

Followed by:

"...because it is that in fact which receives, it is the filament which transforms, which transforms the energy produced by the battery into light energy."
(line 156)

"This arrow, it means that the energy, it goes from the transformer to the reservoir. That has been never seen, if not the light would never go out."
(Group 2, Fulvia, line 224; contradicting proposal of other student that there should be an arrow from the transformer "bulb" to the reservoir "battery"; equivalent to rule n°: 3)

"It's the chemical reaction in the battery which makes the bulb light up"
(Group 3, Céline, line 426).

"It's [the battery] used up ... because it gives energy to the bulb."
(Group 4, Fabien, line 830).

The model is not, however, capable of taking into account 'trial and error' and impasse driven problem-solving. These limitations are discussed in the conclusion.

Differences between performance of students and psCHENE

The main differences concern the identification of the wires as objects to be modelled in the experiment, and identification of the *environment* as the last reservoir. This may be summarised as follows:

(1) Nearly all students identify the two wires between the battery and bulb as corresponding to *two transfers* very early on in problem solving (often just after the battery as a reservoir). Although psCHENE uses the wires in its resolution, it identifies (correctly) only one transfer (electrical work) which none of the students do.
(2) None of the students identify the "environment" as the final reservoir (necessary to respect the rule of the model that a complete energy chain starts and ends with a reservoir), whereas psCHENE does so. One group of students approached this, by identifying their eyes as a reservoir to which the light is transferred and their body as one to which heat is transferred.

These differences are illustrated by six example energy chains actually produced by students for the battery-bulb experiment, shown in Figure 4 below.

Now, we go on to describe the relationship between the initial psCHENE model for an abstract student and a specific problem solving protocol (the one where the second solution in Figure 4 above, "Fulvia-Lionel" was produced). We discuss how the initial model was modified to account for this protocol. This is followed by a short description of how this initial model relates to two further tasks in order of their increasing complexity.

Modelling problem failure for individual students: the battery-bulb experiment

In one protocol ("Fulvia-Lionel"), Fulvia fails to find a satisfactory solution. In the protocol, there is no evidence that she ever considers the role of the environment. Given that she has only one potential reservoir she could solve the problem by allowing the battery to be at both ends of the energy chain (the energy would then start and end at a reservoir, and so the model rule would be satisfied). However, she correctly rejects this possibility, preferring to *reject the model rule*:

> Lionel: "... a reservoir to begin with and a reservoir to end with."
>
> [repeats model rule]
>
> < ... >
>
> Fulvia: "I don't give a damn. For me that's the way it is, it's not because they've written that we have to do that"
>
> [rejects model rule]

She is therefore left with no way of solving the problem (it is the solution that *she* prefers that the students finally write down above).

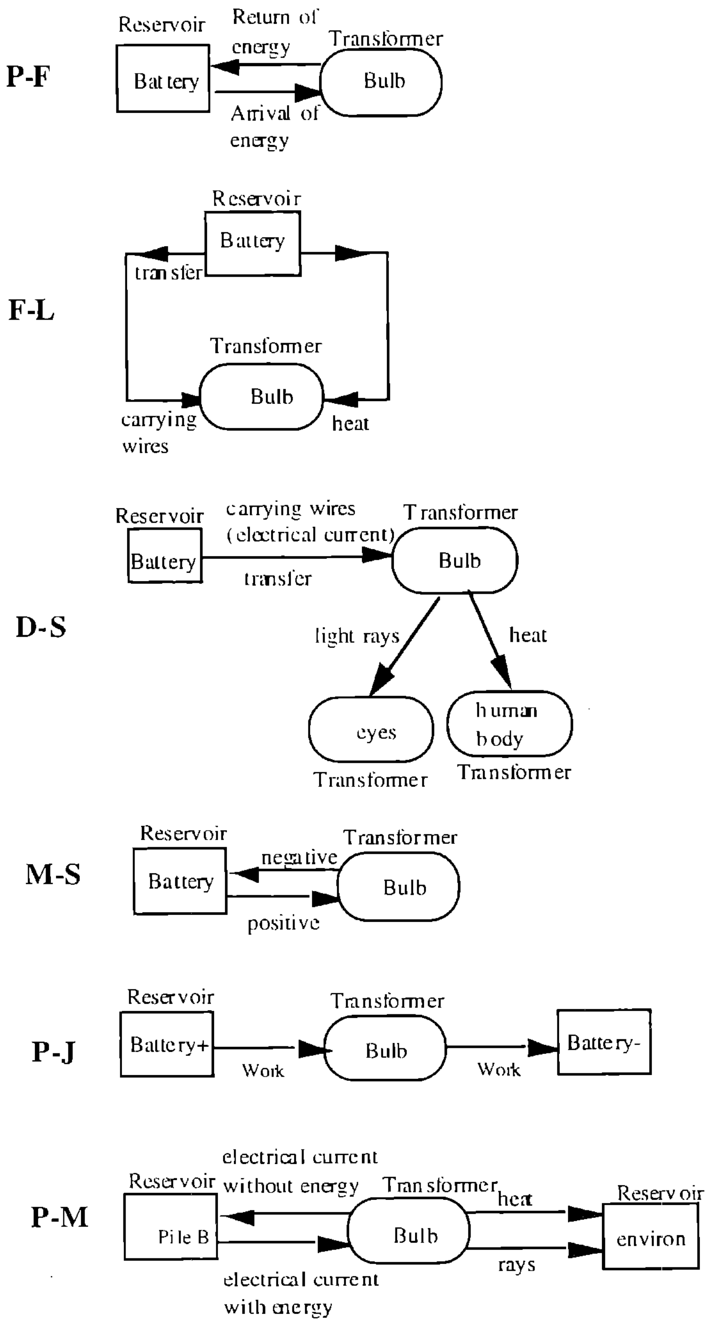

Figure 4. Example energy chains produced by students for the battery-bulb experiment.

To model Fulvia's problem solving, we require some small changes to both the initial state of her knowledge and the rules that are available to her. All the changes are based on an analysis of the protocols but some are due to characteristics of her problem solving which we will not discuss here. The important point for now is that the model has access to a rule that permits *the assignment of the role of transfer to the wire*, and that no access is provided to a key rule for assigning a role to the environment. In fact, both students agree early on that since there are two wires from the battery to the bulb, there must be two transfers – they differ in the *direction* of transfer given (Fulvia = both from battery to bulb; Lionel = one in each direction).

The additional rule is one given to the students in the model of energy:

(M3 ((transfer energy =X =St)) → ((is =X transfer 0)))
[IF X transfers energy, THEN X is a transfer]

Fulvia, in the protocol, makes the association (one to one mapping) between a wire (an experimental item) and transfer (a model item), hence one of her conclusions is "is wire transfer". The rule selected for deletion is:

(N3 ((environment is transformer or reservoir =St)) →
((is environment reservoir 0)))
[IF the environment is a transformer of a reservoir, THEN the environment is a reservoir]

Except for these two changes to the rules, Fulvia's problem solving is similar to the example we presented previously. In her case, the stopping condition, a constraint provided by the model of energy, "an energy chain begins and ends with a reservoir", is not satisfied (because she does not reach the conclusion that the "environment is the final reservoir"), and hence the interpreter stops when there are no more rules that are applicable. Fulvia believes that if an object gives something to something (in this case, the battery gives energy to the bulb), then it cannot receive it (the battery cannot receive energy from the bulb). Her reasoning can be seen in the following:

"I say, if the arrow goes in this direction [from bulb to battery], *it means that the energy returns to the battery all the time and it's not possible; otherwise, it will never go flat"*

Lionel (the student who worked on the problem with Fulvia), has a model closely related to the electrokinetic model. His protocol suggests that he sees transfer of energy as transfer of current. For him, transfer begins at one object and ends by returning to the starting object. Lionel's divergent behaviour can be interpreted as an ingenious application of the model rule "A complete energy chain must start and end with a reservoir", since in his case

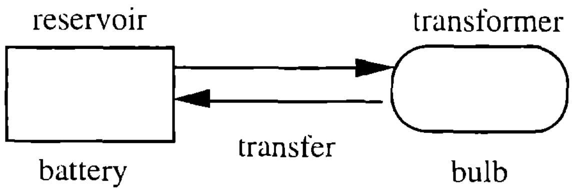

Figure 5. Lionel's energy-chain for the battery-bulb experiment.

the initial and final reservoirs correspond to the same object (the battery). His understanding of energy is not yet sufficiently developed to enable him to distinguish it from the concept of electrical current. In addition he is not faced with a constraint that is sufficiently strong to oblige him to modify/develop his existing knowledge acquired from previous teaching. This can be modelled by using a rule such as:

(T2 ((electrical current from =X to =Y =St1) (=X is generator =St2)) → ((electrical work from =X to =Y 0)))

[IF there is electrical current from X to Y and X is a generator,
THEN there is electrical work from X to Y]

by dropping the condition (=X is generator =St2). In this case the modified rule T2 applies to X being a battery and Y being a bulb to reach the conclusion 'electrical work from battery to bulb'. It also applies to X being a bulb and Y being a battery to reach the conclusion 'electrical work from bulb to battery'. Figure 5 shows the energy chain that Lionel produced when solving the battery-bulb problem on paper.

Modelling developing student expertise across different tasks

We now describe two further stages in learning to build energy chains, within our teaching sequence: the object-motor-bulb problem, then the battery-motor-object problem.

Object-motor-bulb problem

As previously discussed, in this experiment an object is attached to a piece of string that is wrapped around a motor axle. The motor is attached to the terminals of a bulb. The students wind the string round the axle and let the object fall, which lights up the bulb.

A solution is illustrated in Figure 6. This problem follows on from the first, and requires an extension to the model.

A major problem with the model we have described is one of generality. Previously, we could finesse the problem of how to choose the first reservoir

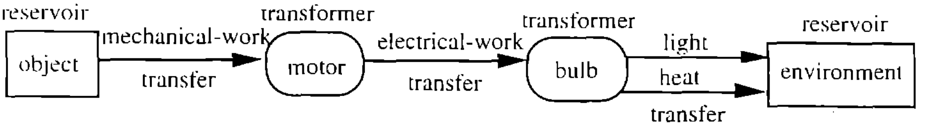

Figure 6. Energy chain for the object-motor (as a dynamo)-bulb experiment.

as it was always taken to be the battery – mainly, because students appeared to have little or no difficulty in assigning the role of reservoir to it. In the situation we now describe, this assumption is not valid so we have to generalise the rules. The rule New-M1 is the result: the model will now determine that some object is a reservoir if it can determine what object is associated with the first item (i.e. 'first cause') of the energy chain.

(New-M1 ((=X is first-item =St)) → ((is =X reservoir 0)))
[IF X is the first item in the energy chain
THEN X is a reservoir]

The following rules are also included for the relationship between movement and mechanical work: the object falls, which makes the motor move (the object is the first reservoir as it provides energy to the motor which is a transformer of energy):

(A1 ((=X falls =St1) (connected =X to =Y =St2)) →
((=X makes =Y move 0)))
(A2 ((=X makes =Y move =St)) →
((gives mechanical-work =X to =Y 0)))
(A3 ((receives mechanical-work =X from =Y =St)) →
((=X is transformer or reservoir 0)))
(A4 ((gives mechanical-work =X to =Y =St)) →
((is mechanical-work transfer from =X to =Y 0)))

These rules, together with those that are included unchanged from the battery-bulb experiment, and an appropriate representation of the experimental field, give the solution shown in Figure 6.

Battery-motor-object problem

We now turn to the third experimental situation which has also been investigated. In this experiment a motor has a string attached to its turning axle, with an object attached to the other end of the string. When the motor is connected to the terminals of a battery, the motor axle turns, the string winds round it, and the object rises.

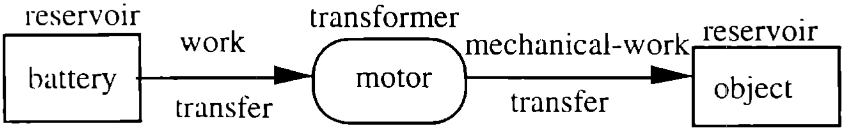

Figure 7. Energy chain produced for battery-motor-object task (Fulvia & Lionel).

Figure 7 illustrates the final solution produced.

The interpreter is provided with some additional rules to handle the new case. The key rule for this problem is:

(New-N4 ((receives =L =E from =B =St1) (not (gives =T =E to =A =St2)) (=E is transformer or reservoir =St3) (final-reservoir =E =St4)) $\rightarrow$ ((is =E reservoir 0)))

["IF E receives L from B and E does not give T to A, and E is a transformer or a reservoir and E is the final reservoir,
THEN E is a reservoir"]

This new rule corresponds to the student knowing, or inferring, that some object is the end reservoir in the energy chain. This need to generalise the model is paralleled by the student's need to determine which entity in the experimental situation is the final recipient of the energy in the chain. How students achieve this step is unclear: they might generalise in much the same way as we have done in developing the model; they might perform some domain-based inference; or they might use a coping strategy and argue that the end of the chain must be the moving object since this entity is the only one left which has not been assigned a significant role.[8]

The rules used in the previous experiment for the relationship between movement and mechanical work are included. The only further exception is New-A1 below, that is modified to reflect the fact that in this case the motor makes the weight rise (rather than the falling weight making the motor move):

(New-A1 ((=X makes =Y rise =St)) $\rightarrow$ ((=X makes =Y move 0)))

The solution path produced by this modelAI fits the protocols quite well. For example, Fulvia has little problem in building an energy chain for this task. An extract from her protocol indicates the nature of her solution.

"The motor pssst (moves) and the weight rises"

[$\Rightarrow$ it is a transformer]

"The object goes down ..."

[$\Rightarrow$ it is a reservoir]

"The object, when it falls, it is like if you throw it if you throw an object, you don't do anything about the energy in it."

[⇒ the last object is a reservoir]

"The object is falling and it produces energy, and it goes through the motor, which arrives to the bulb and the bulb lights. Perfect!!"

In future we will nevertheless need to address the issue of how the student determines which reservoir is the final repository of energy.

Discussion

In our modelling we have sought to make clear distinctions between the different formal aspects of the situation. The student using CHENE has to build a model of the energy chain for a specific experimental situation and using a specific interface (the representation of the model, the means of building it). The student has been given specific information about the modelling process as well as some information about the underlying physics.

The contribution of work with psCHENE

Work with psCHENE highlights both the role of different types of information/knowledge in problem solving, and the information that is missing that needs to be derived from various sources. Different students use different kinds of knowledge. For example, Fulvia places more emphasis on everyday constraints (energy can not return to the battery otherwise the bulb would never go out), while Lionel gives priority to rules of the energy model (energy coming back to the battery satisfies the model constraint that an energy chain must start and end with a reservoir).[9] This work has demonstrated the simple aspects and those more difficult to relate the underlying physics to physical objects and events. The work has allowed us to further assess the potential of our approach to teaching energy.

Problem solving stages and their sequencing

All students construct energy chains in basically the same sequence. They begin by one to one matching of the main *objects* in the experimental situation (battery, wires, bulb, etc.). Moreover, they do this in a *specific order*, relating to linear causal reasoning ('the battery causes energy to flow along the wires to the bulb'). psCHENE is quite good at modelling this behaviour.

This work shows two main difficulties for students:

1. invention of the last reservoir in certain cases – the "environment";
2. elaborating the notions of the modes of energy transfer, work, heat or light.

In the first case, this invention requires a "leap of the imagination" that is difficult to model. It seems that either students think of it or they do not. The model rules are designed to oblige them to invent some "final reservoir" of this kind (in one case this was identified with the perceiver of the experiment, which is an interesting result in its own right).

In the second case, we saw that a transformation of perceptual knowledge to the form of theoretical knowledge is not at all straightforward. For example, students see that the bulb is shining then have to turn that perception into an explicit statement that "the bulb gives light to its surroundings".

In both cases, we may hypothesise that the reason for these difficulties is because they have a belief in a bidirectional rule of the kind 'all physical objects in the experiment correspond to elements in the model, and all elements in the model correspond to a physical object in the experiment'.

Since the environment is not a physical object in the sense that a battery is (it is the general framework of the experiment), they are precluded from thinking of it. One way of modelling this would be to allow our model to create unnamed or abstract reservoirs and transformers (at present they must be given a specific name), which the students could name 'last reservoir' in order to satisfy the model rule, and for which they could subsequently search their memory or the problem situation for other entities that have not already been represented.

There are two final and related ways in which our modelling approach needs to be extended. The first concerns the *unilinearity* of our model, which predicts that students solve the energy chain problems by a single chain of reasoning, whereas (fairly obviously) they sometimes revise their solutions during production. In order to model this we are currently using a "belief system" (or "Truth Maintenance System" – Doyle, 1979; de Kleer, 1986; Gardenförs, 1988) which will eventually be linked to the production system, that records dependencies (positive and negative justifications for beliefs derived). Thus when students detect contradictions in their beliefs, they usually try to resolve them (not always with success, as the Fulvia-Lionel protocol illustrates), leading to revision of their solution in progress.

The second way in which our model needs to be extended relates to the fact that we are studying pairs of students, engaged in dialogue in order to collaborate in problem-solving. At present we only take into account the students' individual long-term memory information, and their individual representations of the problem situation and the information provided to them in modelling their problem solving. However, in the collaborative problem-solving situation, each student may be (and usually is) influenced in their problem-solving by information communicated by the other student. In future work we therefore aim to move towards modelling problem-solving in dyads by combining two individual problem solvers and belief systems with a dialogue model.

Some unresolved questions

Students also do not differentiate the notion of current from that of energy, so that energy and current have shared properties. This kind of problem is a major difficulty for the teaching of physics since students have great difficulties distinguishing between certain concepts such as energy and power, velocity and acceleration and so on. We argue that modelling student behaviour in using CHENE is a suitable vehicle for exploring the ways in which this problem might be mitigated.

We are also able to model some of the ways in which students go wrong. For example, some students want to represent each physical object as a specific entity within the energy chain constructed. So, if there are two wires in the experimental field then there must be two corresponding entities in the energy chain. Additionally, if the constraints imposed are not respected, some students might try and model wires in terms of box-like 'transferors' instead of in terms of links between boxes. Either way, we can interpret the behaviour of such students in terms of an attempt to interpret energy flow as being essentially identical to current flow. For example, Lionel sees wires as being the means to carry electrical current and therefore transport electrical work. Hence he appears to hold a faulty relationship between aspects of the experimental field and the model. Fulvia, on the other hand, has a more accurate understanding of the distinction between electrical current and energy flow.

We also need to take into account the ways in which students recover from mistakes. Analysis of the protocols indicate that students sometimes arrive at conclusions that are contradictory to their earlier conclusions. One of our next steps is to analyse and to implement ways in which students resolve such conflicts. The solution to this problem is of general interest in that it cannot be restricted to operations only at the level at which the impasse occurs – unlike the BUGGY family of systems and VanLehn's repair theory which is based on the notion that students have no real understanding beyond that found in the ability to perform mainly syntactic operations (Brown & VanLehn, 1980).

Moving to a new modelling system

While the way in which we 'slice' the world up is open to debate we do have to select at least one way of so doing. The student may not see the CHENE situation in terms of experimental field, interface, physics theory etc. However, these are some of the important didactic variables. So modelling the behaviour of students in terms of their utilisation of physics theory, knowledge of the interface etc. is of significance to the exploration of the relationship between student learning and these didactic variables.

The model upon which psCHENE is based is moderately complex but psCHENE utilises a simple representation and is therefore not able to capture all the phenomena relating to how students use CHENE and how, and what, they learn. The simplicity of psCHENE is adequate enough to represent specific incidents in the protocols upon which it has been used but psCHENE will need considerable improvement before it can be used to provide a more in-depth study of student learning. It is not difficult to develop a more complex system: the issue is to do so in a principled way.

At the implementation level, the distinctions between interface, experimental field, physics theory etc. are not as clear in psCHENE as we would like. All the knowledge is represented as a single set of rules. Modelling can be improved by distinguishing between the different components of problem solving. Examples of the components are model rules, which are rules that are given to the students for their problem solving, rules that deal with concrete objects found in the experimental situation, rules that are concerned with terms of the model and with abstract terms or concepts only, and rules that are comprised of a mixture of concrete and model level terms. This process will help us with further analysis of students' problem solving strategies and also with the analysis and design of appropriate feedback to the students.

While this may not affect the specific results we have so far obtained, it is both methodologically desirable and practically useful to devise a cleaner representation than that provided by a flat production rule system. Especially if we wish to move towards modelling the collaborations between students which requires that the model has an explicit representation of the process of problem solving and the choices that are available. Any lack of such meta-knowledge can lead to considerable difficulties – as Clancey (1987) pointed out in describing the development of GUIDON to utilise the knowledge found in MYCIN.

We need metarules for additional reasons. Namely, for choosing rules that are applicable when modelling a particular student. These meta-rules will need to be based mainly on our empirical data. These will also need to exploit general problem solving mechanisms. For example, some of the heuristics that are important to Fulvia's problem solving are: everyday knowledge is most important, and hence, such knowledge in her working memory should always have more strength; respecting physics constraints is one of the last things she is worried about; use of analogy with previous problems; she is more concerned about what makes sense to her – the teacher's goal, other students' solution, etc. are not important to her. Lionel, on the other hand, gives more attention to knowledge of physics (and hence, such knowledge in his working memory should have high strength values). The constraints that are most important to him are physics constraints or constraints imposed on them by the given model of energy. It is also important for him to satisfy the teacher's goal.

Our future plan is to develop a much more structured architecture that permits us to explore some of the theoretically important aspects of the situation, and to separate out some of the factors at work. For example, we are aware that the student is confronted with both a computer interface and an experimental situation. We wish to be able to capture the interplay between these aspects of the student's modelling – as well as aspects relating to the theoretical issues and modelling knowledge needed to solve the energy chain problem.

Future work

The protocols indicate a fairly smooth transition between tasks. Our work with psCHENE suggests a line of reasoning as to why this is the case. Very little change is needed to the ruleset even though the surface appearance of the three situations suggests that they are not that closely related. Although we have been able to generate models to match our data, a further possibility for the future is to confirm that new protocols can be modelled using the same set of rules. This would provide an evaluation of the system, and of our representation of the rules for their adequacy in modelling 'new' cases.

Further possible work includes investigating mechanisms for generating various learning errors such as failure to distinguish important bodies of knowledge (e.g. electricity and energy). This would entail both a more complex modelling system and further empirical work in the laboratory and classroom. Another possibility includes a more detailed study of the relation between implicit and explicit learning: though much of the previous literature has studied this in terms of the control of dynamic systems (e.g. Sanderson, 1989). However, a future plan entails the extension of the set of three tasks outlined in this paper to include a further task which begins to address issues connected with dynamical systems in a more direct way.

Another possibility for future work includes taking into account some of the perceptual/linguistic/conceptual relationships at play. For example, there are connections between students' descriptions of certain situations (rise/fall) and event-based physics descriptions (move). There is a direct link between physics descriptions of events and the physics model (e.g. the battery-motor-object task). Instances of the utility of such links can be found in the protocols and within our modelling. The implication is that teaching schemes should take into account intermediate descriptions (such as the object-event level) between students' descriptions and the physics model.

Another example is the use of notions of giving and receiving: notions that suggest that students are in the process of constructing a coherent underlying mechanism to explain how energy chains work. The increasing use of such notions might be a good indicator that genuine and effective learning is taking place. We have begun to look at this possibility empirically but there is a need to provide a clear mechanism for how such an effect might materialise.

Conclusions

The main result of our preliminary modelling activity is that we have developed a succession of models of the knowledge needed by specific students in order to construct energy models for experimental situations within a specific sequence. Though these models are fairly simple they have forced us to consider a number of issues connected with the relationships that hold between the information given to the student, the nature of the task, the nature of the experimental situation, and the student's growing understanding of energy.

We have provided some further examples of the kinds of difficulty faced by students in many domains. In a sequence of three simple, specific tasks we have observed some instances of these difficulties, and we have been able to model these instances.

The work has emphasised the value of relating specific teaching sequences to student's actual problem solving behaviour through the process of developing computational symbolic models. It assists us to validate our own solutions to the tasks; to be explicit in our assumptions; and to be clear about the events which are important from a theoretical perspective.

In sum, the research reported in this paper demonstrates the feasibility and fruitfulness of implementing computer models of certain aspects of students' modelling processes in one area of learning physics. The iterative method of collecting and analysing empirical think aloud protocols together with the requirement to develop computational models provides a promising approach towards the development of models of students' cognitive processes.

Further, we believe that we can now move on to develop a more powerful modelling scheme which can provide us with the necessary platform to begin to integrate issues connected with the social construction of meaning with the individual student's cognitive processes.

Acknowledgements

We would like to thank the teachers and students who participated in the protocol collections.

Notes

1. "CHENE" = "*CH*aîne *ENE*rgetique", or "Energy Chain". (In French "Chêne" also means "oak").
2. Throughout the rest of the paper we use the following simple notation in order to avoid possible confusion between "modelling" as a process performed by the students, in the domain of physics, and "AI modelling" of the former modelling process: students' modelling in physics = modelling$_{SP}$; AI modelling (of modelling$_{SP}$) = modelling$_{AI}$.
3. It has not been necessary to use more sophisticated strategies at the stage of our work reported here. The next system, modelCHENE, will directly address this issue.
4. Note that what we refer to in this context as "problem solving" may in another context be viewed as construction of a qualitative representation for subsequent quantitative problem solving.
5. The protocols indicate that students use this as a kind of anchor in their reasoning – some students returning to it in order to resolve impasses. As the current problem solver provides no mechanism for handling impasses we cannot model the reuse of information in any meaningful way.
6. We are grateful to an anonymous reviewer for this example.
7. This provides the minimum distinction for our initial needs.
8. Though there is a difficult step in deciding that the moving object really is the last unassigned object -since that requires setting aside any need to assign roles to, for example, connecting strings etc.
9. At the moment we have to provide psCHENE with slightly different rulesets: providing rules with priorities would be more convenient.

References

Anderson, J.R. (1983). *The Architecture of Cognition*. Cambridge Mass.: Harvard University Press.

Baker, M.J. (1994). A model for negotiation in teaching-learning dialogues. *Journal of Artificial Intelligence in Education* 5(2): 199–254.

Bental, D., Tiberghien, A., Baker, M.J. & Megalakaki, O. (1995). Analyse et modélisation de l'apprentissage des notions de l'énergie dans l'environnement CHENE, in D. Guin, J-F. Nicaud & D. Py, eds., *Environnements Interactifs d'Apprentissage avec Ordinateur*, Tôme 2., 137–148. Paris: Eyrolles.

Brown, J.S.& VanLehn, K. (1980). Repair theory: A generative theory of bugs in procedural skills. *Cognitive Science* 4: 379–426.

Brownston, L. Farrell, R. Kant, E. & Martin, N. (1985). *Programming Expert Systems in OPS5: An Introduction to Rule-Based Programming*. Massachusetts: Addison-Wesley.

Clancey, W. (1987). *Knowledge-Based Tutoring: The GUIDON Program*. Cambridge MA: MIT Press.

Closset, J.L. (1983) *Le raisonnement séquentiel en électrocinétique*. Thèse 3ème cycle, Université Paris 7.

De Kleer, J. (1986). An assumption based TMS. *Artificial Intelligence* 28(2): 128–162.

Doyle, J. (1979). A truth maintenance system, *Artificial Intelligence* 12: 231–272.

Driver, R., Guesne, E. & Tiberghien, A., eds. (1984). *Children's Ideas in Science*. Milton Keynes, England: Open University Press.

Falkenhainer, B., Forbus, K.D. & Gentner, D. (1989). The structure-mapping engine: algorithm and examples. *Artificial Intelligence* 41: 1–63.

Gardenförs, P. (1988). *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Cambridge MA: MIT Press.

Greeno, J. (1989). Situations, mental models and generative knowledge, in D. Klahr & K. Votovsky, eds., *Complex Information Processing: The Impact of H.A. Simon* (pp. 285–308). Hillsdale, NJ: Lawrence Elrbaum.

Inhelder, B. and Piaget, J. (1979). Procedures et structures. *Archives de Psychologie* XLVII(181): 165–176.

Johsua, S. & Dupin, J.J. (1986) L'électrocinétique du collège à l'université. *Bulletin de l'Union des Physiciens*, n° 683: 779–799.

Langley, P. (1983). Exploring the space of cognitive architectures. *Behaviour Research Methods and Instrumentation* 15(2): 289–299.

Larkin, J. (1983) The role of problem representation in physics, in D. Gentner & A.L. Stevens, eds., *Mental Models* (pp. 75–98). Hillsdale, NJ: Lawrence Erlbaum Assoc.

Lewis, A.B. (1989). Training students to represent arithmetic word problems. *Journal of Educational Psychology* 81: 521–531.

McDermott, L. (1984). Critical review of the research in the domain of mechanics, in *Research on Physics Education: Proceedings of the First International Workshop* (pp. 139–182). Paris: CNRS.

Newell, A. & Simon, H.A. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

O'Malley, C., Draper, S.W. & Riley, M.S. (1984). Constructive interaction: a method for studying human-computer interaction, in B. Shackel, ed., *Proceedings of INTERACT'84* (pp. 269–274). Elsevier Science Publishers B.V. (North-Holland): The Netherlands.

Richard, J-F. (1990). *Les Activités Mentales*. Paris: Armand Colin.

Sanderson, P.M. (1989) Verbalizable knowledge and skilled task performance: association, dissociation, and mental models. *Journal of Experimental Psychology: Learning Memory and Cognition* 15(4): 729–747.

Shipstone, D.M. (1985) On children's use of conceptual models in reasoning about current electricity, in R. Duit, W. Jung & C. v. Rhöneck, eds., *Aspects of Understanding Electricity. Proceeding of an International Workshop* (pp. 73–93). Ludwisburg 1984 (IPN, Kiel).

Shipstone, D.M., von Rhöneck, C., Jung, W., Kärrqvist, C. Dupin, J.J., Johsua, S. & Licht, P. (1987). A study of students' understanding of electricity in five European countries. *International Journal of Science Education* 3: 303–316.

Tiberghien, A. (1984). Critical review of the research aimed at elucidating the sense that notions of temperature and heat have for the students aged 10 to 16 years, in *Research on Physics Education: Proceedings of the First International Workshop* (pp. 75–90). Paris: CNRS.

Tiberghien, A. (1994). Modeling as a basis for analyzing teaching-learning situations. *Learning and Instruction* 4: 71–87.

Trowbridge, D.E. & McDermott, L.C. (1981). Investigation of student understanding of the concept of acceleration in one dimension. *American Journal of Physics* 49: 242.

Viennot, L. (1993). Temps et causalité dans les raisonnements des étudiants en physique. *Didaskalia* 1: 13–27.