

Modus Operandi and Means for Encouraging Innovation

In the following sections we shall consider THOUGHTSTICKER transactions in enough detail to bring out some points of epistemological interest, and to give an overall impression of the system. The discussion of the previous chapter is extended to indicate the main construction rules and to describe the transactions (based upon "epistemic symmetry" and the "extrapolation of principles") that are used as means to encourage many aim operation and innovation by the user.

Although THOUGHTSTICKER is a versatile system (the flamboyant phrase "epistemological laboratory" is not intentionally misleading), it has so far been used chiefly in connection with the environment of "Energy Conversion, Conservation and Regulation" (the subject matter for the examples in Chapter 7). To a lesser extent, THOUGHTSTICKER has been brought to bear upon an environment "Entrainment of Oscillators".

1. MODELLING FACILITIES FOR CONCRETE MODELS

The Lumped Modelling Facility for energy conversion is the standard modelling facility (on a par with STATLAB in this subject matter field) which is used for an ongoing tutorial project, together with patch-programmable analogue computing elements over and above those incorporated in the standard design (Fig. 8.1). The state of all analogue units (integrators, adders, multipliers) is traced by the LSI machine which acts as regulator. Similarly, all structural and patch-programmed connections in the

These ad hoc arrangements suffer from obvious and irritating defects. Ideally a user should construct and enlarge a lumped facility as required to accommodate the models he wants to manufacture, and his subsequent modelling operations in one component of the lumped facility should all be computer interpretable and constrained by the models already built. As it is, only the first of these requirements is fully satisfied. True, so long as the system is an experimental tool, these deficiencies are no more than a nuisance, on a par with the chore of copying out a revised and tidied version of the cognitive model (the mesh on the construction grid). But, in contemplating wider types of application, it is crucial to notice that the existing constraints are inessential.

Mechanically speaking, all the conditions for manufacturing "spare" modelling facilities can be implemented, and several slightly context dependent examples are in existence. Papert's (1970) LOGO was noted in the previous monograph as a paradigm mathematics laboratory, and the system could be modified slightly to accommodate the distinction of differently constrained universes. A further instance is a suite of interactive graphic manipulation programs originally designed for an art school and currently used for modelling in chemistry (at a plethora of different levels: molecular, atomic, quantum mechanical, etc.), which permits the user to make and retain "spare" modelling facilities (De Fanti 1975). One further example, is Negroponte's (1970) "Architecture Machine" which permits similar inventive liberties.

The issue of practical feasibility is very important, for without a means of giving users (who are *not* versed in programming) access to freely constructed "worlds," the system would remain no more than an experimental tool of limited value. The fact is that means exist, and though they are currently quite expensive, their cost is likely to decrease very rapidly as computer technologies come to fruition.

2. THE CONSTRUCTION GRID AND THE COGNITIVE MODEL

The arrangements for building up cognitive models and entailment meshes are currently implemented using a graphic display (Fig. 8.2) and a sketch pad input augmented by a keyboard. Previ-




Fig. 8.2. Current realisation of THOUGHTSTICKER using graphic display tube (ARDS terminal) and auxiliary equipment. A = ARDS Terminal, keyboard and display tube. B = Sylvania graphic tablet. C = Control equipment for graphics tablet. D = Minicomputer (LSI 2 with 24k core storage). E = Digital magnetic tape backup unit. F = Mini BOSS for aim validation. G = Display oscilloscope for modelling facility. H = Teletypewriter console. I = Backup display and modelling grids. J = Auxiliary display. K = Projector.

ously, the mesh construction was realised with certain limitations by using physical construction grids and physically placed electronic modules connected together by the user.

The previous arrangement gives a clear picture of processes which are now carried out automatically and as a result of which images are displayed. The system will be described in these terms and carries over into the current implementation, with the following caveats only.

(1) Node unit positioning refers to pointing operations; (2) connecting operations refer to key tagged link drawing operations; (3) displays, both of descriptor values (LEDs), and signal lamps (active node, and so on), are replaced by graphic conventions; (4) separate construction grids correspond to displaced tube locations; (5) regions are represented by a dashed line (quasi 3 dimensional) display.

The display tube can represent only a fairly small mesh (or part of a mesh at once) but can be augmented by concrete construction grids for representing relatively unchanging portions of a mesh. However, the mesh can be repruned under any head role (the heuristic of Section 3.2(f) is realisable), and the resorting of topic nodes according to the computer generated plan is automatic for all nodes displayed.

The programs governing the operation of THOUGHTSTICKER are under continual development: listings of the existing programs and their updated versions are available on request.

2.1. General Framework

The grids (one to each region as in Chapter 7) have modular cells associated with node positions (to be filled by the user), LEDs for exhibiting the values of semantic descriptors, and "attention lamps" via which the regulating heuristic B can bring the user's attention to one or a cluster of cells. Recall that the channels of the data bank are also associated with their own LED displays and are "tag name" labelled, but not ordered, under the (syntactic) depth descriptor.

2.2. Starting Set

The starting set of substructures is built up on the construction grid for region 0 (namely $CG(0)$) using modules (Fig. 8.3) identical with those employed in CASTE. Each module retains and displays the value of explore, aim, goal, and understand by means of flashing light codes based on three signal lamps. The data base (computer) inscription of the starting set modules is indexed by one family of descriptors (sufficient to access the topics): values being LED displayed on demand. Topics are accessed (as in CASTE, previous monograph) by specifying descriptor values via the interaction console. For the "conversion and conservation of energy" environment, the starting substructures are obtained by denuding the entailment structure in Fig. 7.1. At the outset, a user is faced with just these structures, and whilst he learns about the topics they adumbrate, his behaviour is regulated as it would be in a CASTE or INTUITION operating system.

However, the starting substructures do not delineate a full thesis




Fig. 8.3. A module used to represent a topic in starting set of substructures or a permanently instated node. Connections are specified in the computer and sketched as lines on display. As in an other-than-evolutionary operating system (CASTE or INTUITION) *explore*, *aim*, *goal* and *understand* markers are indicated by signal lamps. Sockets for inserting aim and explore probes are optional since topic may be accessed by descriptor values.

on the "conservation and conversion of energy," and the denuded fragments of the original entailment structure are deliberately truncated to secure this condition. As a result, the user can make *more* concrete explanatory models in the Lumped Modelling Facility than those attached to topics in the starting substructures. The possibility of constructing analogy relations is an obvious consequence of denuding the entailment structure. But it is practically important that topics other than analogy relations can also be invented.

2.3. Building Up the Cognitive Models

Apart from the starting set of substructures and the associated grids, the user has available an unlimited supply of electronic boxes and connecting links. As a matter of convenience and representational economy, the boxes are of several different kinds: (a) Units representing topics that are derived without analogy; (b) Units representing analogy relations, and representing topics of mutually exclusive and conditional hypotheses. Since all units stand for nodes in an entailment mesh, units are henceforward glossed as nodes: topic nodes, analogical nodes, and conditional nodes. Similarly, the links are classified as follows: (A) Unidirectional black links, representing an other-than-analogical derivation; (B) Bidirectional orange links, representing an analogical deriva-

tion; (C) White links, representing the "syntactic" component (isomorphism " \leftrightarrow " or a topic) which stipulates the similarity in an analogical topic relation; (D) Purple links, representing the names of semantic predicates, Dist or the difference in an analogical topic relation; (E) Brown links, representing a conditional derivation; and (F) Speckled black links, which have no functional distinction from black links but are useful in visually discriminating several derivation paths.

2.3.1. To instate a fresh topic T which is simply derived from existing topics P and Q, the user takes a topic node (Fig. 8.4), labels it with the name for T, and inserts it into a position on a grid. This operation illuminates the active lamp on node T (Fig. 8.4). The user next connects the output of P and Q, though black links to one of the input clusters (maximum of three) on node T. Each cluster is a kernel of T (first monograph), and it may have at most six members. If P and Q are sufficient entailment precursors (in




Fig. 8.4. Topic node. Each node is a "Box" with inputs (black lead) for a maximum of 3 kernels or conjunctive derivatives: each derivation being at most 6 sub-ordinates. The "active" lamp is illuminated if the node is positioned on the grid and is extinguished if the node is instated. The switches indicate that proposed derivation from more than 2 but not more than 6 other nodes is complete (the kernel in question is full). An insertion of a fresh derivation lead into any vacant kernel, reactivates the "active" lamp.

one kernel) of T, the user turns the switch (Fig. 8.4) on this input cluster. This operation signifies that node T is submitted for consideration by the regulatory heuristic. Amongst other things, a model for T must be constructed (in the processor associated with the grid on which T is mounted) before the submission can be accepted, and until this model has been successfully executed, T will remain active. However, the model could be, and commonly is, constructed and executed before any attempt is made to instate T.

If P and Q are interpreted in the same universe of compilation and interpretation (Fig. 8.5 on left), the account is complete. If P is interpreted in one universe X and Q in another Y (when P and Q are in separate substructures), then, in respect of the model for T, these universes are no longer independent. T unites X and Y; a priori independence is modified by the topic instated (Fig. 8.5 on right).

2.3.2. In order to instate a further derivation of an existing topic R from existing topics P and Q, the user connects black (or speckled black) links from the output of P and the output of Q into one of the unused input clusters of node R. The act of applying input connections to an unused cluster gives R the status of active. The user next presses the switch on the input cluster and submits his fresh derivation for scrutiny by the heuristic.

2.3.3. To instate an analogical relation between topics (either existing or due to be constructed), the user positions an analogical




Fig. 8.5. Derivation of *topic* T (at a topic node) from *topics* P and Q. On the left, the derivation is confined to one universe of interpretation (X); on the right P and Q are in distinct universes of interpretation (X, Y), which become related as a result of instating *topic* T.

node (Fig. 8.6) on one of the reserved grids. As a result, the node becomes active and remains so until certain inputs are furnished, though they may be furnished in any order whatsoever. First (though not necessarily in order of appearance), there must be orange connecting links from existing or yet to be instated topics in different universes $MF(X)$ and $MF(Y)$, which form the terms of the analogical relation. Next, there must be a white link from an existing or yet to be instated topic, which is the similarity of the analogy. The universe of compilation/interpretation of this (simi-

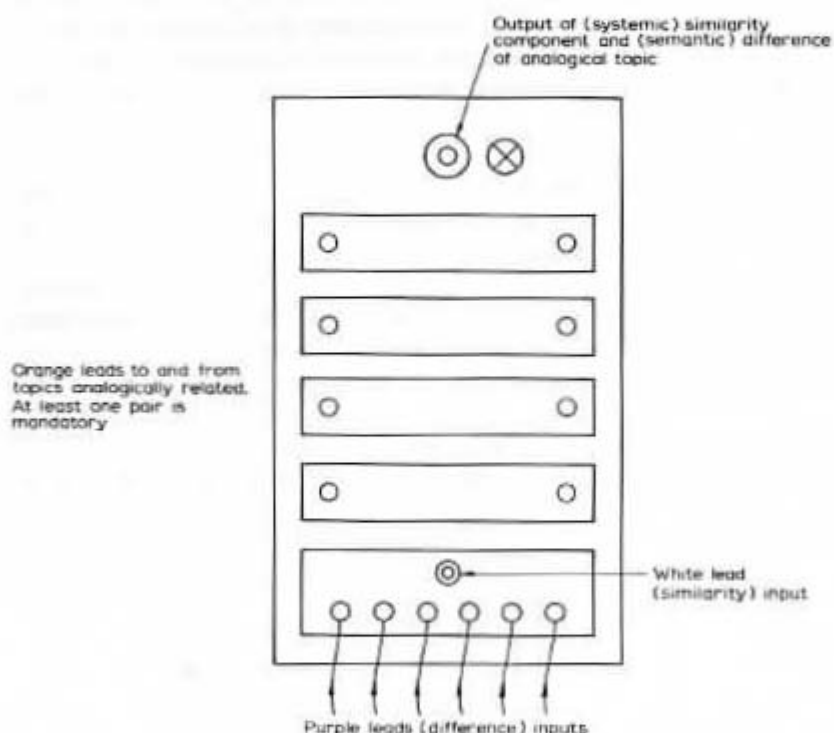


Fig. 8.6. An analogical node. This node must receive inputs (orange leads) from nodes in at least two different universes of interpretation X , Y (either partially, or completely, distinct derivation-linked substructures): An input (white lead) either from a topic indicating similarity of analogy, or from an isomorphism socket, and a (purple lead) input from either a topic or sockets labelled as semantic descriptors. Both orange and purple leads may be multiple (maximum of 4 orange and maximum of 6 purple).

larity) model may either be the union or the product of X and Y, or some distinct universe; it is a generalisation of the models for the analogically related topics.

In case the analogical relation is a strict and complete isomorphism, the white link may emerge from a special socket labelled *isomorphism operator* \cong . Finally, there must be a purple link from a topic or from one of the user labelled sockets representing free semantic descriptors, which are named as part of the description routine. The purple link thus signifies a so far unnamed difference Dist (x, y) upon which the analogical relation is based.

The user may press the submit switch whenever he has specified the collection of terms (topics he regards as somehow analogous), but the analogy relation is not adjudicated for legality until the various inputs are filled out. Fig. 8.7 shows typical completed analogy nodes, but at the risk of tedium, we stress that analogical nodes can exist (in an active state, of course) long before all the inputs are filled up.

2.3.4. The conventions built into the THOUGHTSTICKER system are deliberately pedantic. (The pedantries are justified insofar as THOUGHTSTICKER gives useful training in applied epistemology, as well as acting as a course assembly system.) According to these conventions, analogy relations hold between topics in distinct universes of interpretation (which is correct, though unduly fussy for ordinary purposes). Difficulties are thus encountered in dealing with analogies loosely said to hold between topics in the same universe.

For example, suppose it is desired to represent the isomorphism between graphs (or finite automata) F and G. As a general statement, there is one universe, \mathcal{U} , of graphs (or finite automata), a universe of the same kind of mathematical objects. However, the particular objects F and G cannot be simultaneously executed in the same independent and serial processor (as required if they are said to be analogous). They could, of course, be simultaneously *simulated*, but that is a very different matter; their *realisation* is actually called for. Hence, a user anxious to instate and model the F, G isomorphism must construct topic F as a node in one grid X and model it as M(F) in one a-priori-independent part MF(X) of the Lumped Modelling Facility; construct topic G as a node in an-

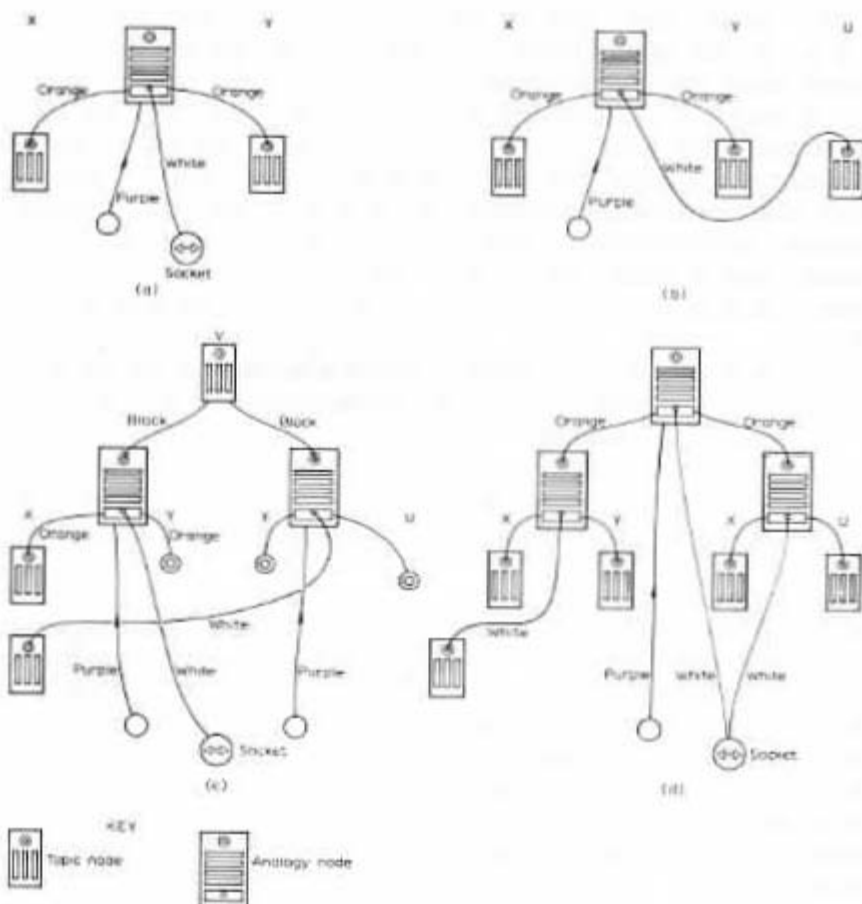


Fig. 8.7. (a) Simple isomorphic analogy; (b) Generalisation based analogy; (c) Derivation constructed from pair of analogy relations; and (d) An analogy between pair of analogy relations.

other grid (Y), and model it as $M(G)$ in another part $MF(Y)$ of the Lumped Modelling Facility. To complete his construction, he adds a white link to the isomorphism operation (the similarity) and seeks a difference between X and Y. But X and Y are equivalent so that $X \equiv Y$, which means that the universes could be represented as X, Y or the product $X \times X$. This possibility is accommodated by a special operator signified by a socket \equiv for "equivalent but dis-

tingent" (Fig. 8.8). The difference (between otherwise identical universes of interpretation) may be regarded either as spatial (X, X) or temporal (as in $X_{\text{now}} X_{\text{later}}$).

To press this important point home, consider a rather larger and more realistic example. The user wishes to model a finite ensemble of dynamic systems characterised by the same system equations and being replicas, but possibly differing in respect of initial conditions. Such formulations are ubiquitous in physics, genetics and numerous other disciplines; they underpin any application of statistical mechanics. The replica microsystems are analogous (not identical, but isomorphic). The similarity is the dynamic equation common to them all. The difference is equivalence with either spatial or temporal distinction, as capturing their a priori independence. The analogy relation is the ensemble of microsystems. Thus, the system equations are represented as a derivation structure copied in each analogous universe, X, Y, \dots . The statistical theory is a further derivation structure in a distinct (macrotheoretic) universe, say U . The head of this derivation structure in U is isomorphic to the analogy between the systems represented in X, Y, \dots .

2.3.5. To instate a topic representing mutually and perhaps conditionally exclusive hypotheses, the user positions a conditional node in a grid U . This node requires inputs from nodes of the




Fig. 8.8. Isomorphism between topic X and topic Y . Equivalence connection by purple link means that X and Y are regarded as coordinates of product set.

topics representing the hypotheses (which can be interpreted in universes X and Y), together with a further input through a brown link which (Fig. 8.9) either negates their conjunction in U (i.e., they cannot both be interpreted and correctly executed in U), or asserts their conditional tenure.

The topic represented by a conditional node is two or more alternative hypotheses T_1, T_2 that are purveyed or supported by different factions and are at loggerheads. In short, the topic represents a controversy between theses that are advanced or advocated by distinct P-Individuals. These P-Individuals may be as august as institutions, famous scientists, "the establishment," specific disciplines, or "schools of thought". They may be as miniscule as the different perspectives taken by one *person* (but two or more P-Individuals), as in the ambiguous figure example (Chapter 7, Section 3). Expert 1 and Expert 2 of Chapter 7, Section 2 would count as exponents of the rival theses T_1, T_2 if they failed to agree and their disagreement, the clash between T_1, T_2 , was inscribed in the network. In particular, a conditional is introduced if, and only if, there is a many aim resolution (B treats the user as $\langle A_1, \alpha \rangle, \langle A_2, \alpha \rangle$, in which A_1 and A_2 do not reach agreement).

2.3.6. A typical interpretation is as follows. Let T_1 hold and be modelled in universe X . For example, in elementary physics, T_1 is some prediction (the existence of sharp shadows) from the Newtonian corpuscular theory of light, and X is the universe proper to the geometry of this theory. By the same token, let T_2 (blurred




Fig. 8.9. A conditional analogy denying isomorphism between T_1 and T_2 .

shadows) hold in universe Y, proper to the geometry of a (Huygens, Fresnel-like) wave theory of light. If U is a further universe of experiments with shadow casting, then T_1 and T_2 are rival hypotheses in U, and this rivalry is expressed by the conditional node as a critical experiment between the theses of P-Individuals (a user's conception of Newton and his conception of Fresnel).

T_1 and T_2 are not formally contradictory. Further, both may be realised (in X, Y). But T_1 and T_2 are incompatible in some common (and accepted-to-be-standard) universe U. The conditional node denies a possible analogy relation.

If the experiment leads to falsification in Popper's (1959) sense, then one thesis or the other will be tentatively denied (until the issue is resolved by some more advanced discovery or theory). But there need be no such critical test (the rival claims may rest undecided, and the conditional node may represent only an open controversy and a fruitful research topic). As stressed repeatedly, we are not primarily or directly concerned with verification/falsification or absolute veridicality. However, such important notions must be representable in a body of knowables (as they are by conditional nodes), and it is essential to recognise that when conditional topics are manifest, they are invariably personalised: to Newton versus Huygens, Church versus State, or several distinct roles adopted qua P-Individuality by the user himself.

2.3.7. Whilst various node constructions are in progress the B heuristic detects any aim which it can identify. An aim may either be placed on a module, in which case it is identical with the aim of other operating systems (CASTE or INTUITION), or it may be an *active* node.

Many nodes may be simultaneously active; for example, in Fig. 8.10 there are five active nodes. B is programmed to interpret only some of these as candidate aim nodes; those that are superordinate and that have full kernels are submitted and accepted for submission. Thus, in Fig. 8.10, nodes S and T are the candidate head nodes; R is excluded because the construction, even if submitted, is incomplete.

After a period of construction, the user is able to submit nodes for instatement, in which case (as below in Section 2.4.), he must justify derivation of subordinates and the like. These transactions take place through the interaction console. Once an instated struc-




Fig. 8.10. Active nodes.

ture exists the user is impelled to state a *head* node and submit the structure. The planning routine is executed, and as a consequence, he must furnish a semantic description of the structure. For the most part, users are quite willing to choose heads; failing that, they are periodically forced to do so.

Next, if the B heuristic picks up a many aim configuration of the type shown in Fig. 8.10 (Section 2.3.7) and if it is also the case that at least a pair of aims have distinct descriptions (obtained by prior descriptor assignment under one aim), then B calls for resolution (placing the user in the position of A_1, A_2). In this case, the planning routine is executed, but description is replaced by comparing and updating the distinct descriptions of the aim nodes.

Thus, either B's requirement for resolution or the user's selection of a head belonging to an instated substructure initiates the planning and description routines of Chapter 7; the routines that tidy up the mesh and present it for description and/or resolution.

2.4. Instating a Node: Degree of Verification

The active lamp on a node is extinguished only if certain conditions are satisfied most strictly if the topics form a valid conversational domain, less strictly if the construction is agreed by another

user (including arbitrators and groups). In the strict case, the following conditions must hold:

(a) The model of any superordinate topic contains as constituents the models for all of the subordinate topics from which it is derived (that is, according to the user's derivation).

(b) The user's derivation of a superordinate topic from its subordinate topics loses no essential specificity and is cyclic, apart from its primitives, as a result.

An adequate, weaker form of this condition is summed up in a pair of injunctions that are to be obeyed by the user:

(I) If *topic k* is to be instated as derived (non-analogically) from *topic i* and *topic j*, then within the derivational structure the user must show (by a construction on the grid) how *topic i* and *topic j* are derived from *topic k* (perhaps using primitives) without loss of specificity. Further, the user must make (or assert that he can make) a model $M(k)$.

(II) If *topic k* is analogical, the user must show the reverse derivation (as above), given the waiver that the derivation depends upon the distinguishing predicates, Dist. Further, if the analogy is isomorphic, the user must show the one to one correspondence between *topic i* and *topic j* (directly, or by subordinate isomorphism), and if it is a generalisation, supported by topic ℓ he must make or assert that he can make a model $M(\ell)$.

Several degrees of rigidity are possible, depending upon the purpose in hand. At one extreme, the displayed network must be consistent and cyclic so that it (and the associated models) forms a conversational domain. If so, condition (a) and condition (b) are checked by applying the test routines of EXTEND, and these routines are also applied to isomorphic analogies between topics. This is a lengthy and rather expensive business.

At the other extreme, where THOUGHTSTICKER is used as an epistemological laboratory, we are only anxious to externalise an innovator's concepts and derivations. The B heuristic checks conditions (I) and (II). The user is required to state what he believes to model his beliefs and derive them from other topics. These statements are accepted without justification (or with only verbal justification), but there is no guarantee that the product is a conversational domain over which learnability and memorability are guar-

anted unless a model can be executed for each topic, or a verbal explanation exists. *

2.5. Description Methods

The description scheme evolves together with the topic network, and consequently, it is impossible to inscribe the values of descriptors as the fixed, maplike representation of an entailment structure, over positions on the grids. Moreover, in the interest of uniformity, the descriptors of all topics, whether in the grids or the disjoint substructures (with the exception of the syntactic depth descriptor with values superordinate/subordinate), are represented in the LED display.

Each position in the grid and each node in the disjoint substructures is equipped with an LED (light emitting diode) pair, able to shine red or green if the LEDs are illuminated. Consequently, the possible conditions of any *position* are red (which stands for the descriptor value +), green (the descriptor value -), and "off," the descriptor value * meaning "undetermined or irrelevant". † At any instant, it is possible to display all values of *one* descriptor or of one Boolean expression in the set of descriptors (all topics having P_1 and P_2 but not P_0 , for example). The user is able to obtain LED displays by typing the name of a descriptor or the form of an expression into the terminal. Conversely, the regulatory heuristic B can present an LED display to the user and identify it by printing out the name(s) of the descriptor(s) concerned.

New descriptors and their values are introduced by the descrip-

* The less rigid criterion may be based on the views of at least two users, $\langle A_1, \alpha \rangle \langle A_2, \beta \rangle$ and leads to an enhanced realisation of the "improved" operating system in Chapter 6, Section 5. For this purpose, instated nodes are temporarily replaced by modules and may thus be learned by a student in a ordinary operating system (CASTE and INTUITION) placing aim, goal, and understand markers on the topics.

Recall that the participating users $\langle A_1, \alpha \rangle \langle A_2, \beta \rangle$ have agreed to each node instated (the less rigid criterion). The tentatively transformed nodes are accepted permanently, as module based topics, if, and only if, $\langle A_1, \alpha \rangle$ can learn $\langle A_2, \beta \rangle$'s thesis when he addresses it under CASTE or INTUITION control as a student; similarly $\langle A_2, \beta \rangle$ in the role of a student can learn $\langle A_1, \alpha \rangle$'s thesis.

† This arrangement leaves open the possibility of representing the values of Fuzzy Predicates of the topics as intermediary shades of light.

tion routine of Fig. 2.8, executed if the user asserts a head topic, or B calls for a many aim resolution.

3. COOPERATIVE INTERACTION

If the user does nothing, he is bombarded with items of information from the data bank. At least, he must engage in explore transactions in order to stem this flux of data. Initially, he can only explore the data bank or the minimal topics in the starting set, and he receives in return items from the channel addressed by exploration.

As soon as some cognitive model has been constructed and the description routine has been executed, the user is able (and forced) to assign values of his own descriptors both to the topics or analogy relations he has instated, and to the data bank channels. True, in the limiting case when the data bank is deemed irrelevant, all descriptors have the value "*" on all channels. Otherwise, channels in the data bank act as information sources that back up topics or groups of topics.

Construction of analogies or topic nodes involves activity in the modelling facility and transactions instrumented through the interaction console and the construction grid display.

The whole process takes place under the following rules (recapitulated from Chapter 7): (a) If *topic k* is instated as derived (in a conjunctive substructure) from *topic i* and *topic j*, it is necessary to show how *topic i* and *topic j* are derived from *topic k* without loss of specificity overall. (b) Analogical derivations satisfy the same rule with the waiver that specificity may be lost (if replaced by the Dist predicates).

At the moment the user asserts a head (or the B heuristic detects a many aim configuration and demands resolution), the pruning and numbering routines come into play and provide a tidied up plan of the mesh (currently, on the Display Tube).

3.1. The Observer's Picture

We, the observers, see an exteriorised version of the user's mental operations. What does the user get in return for all his trouble? Part of the story has been told already. But there is a gap to fill:

namely, the transactions meant to encourage innovative action and many aim operation. In these transactions, the heuristic B acts as an innovative assistant to the user A. Succinctly, B promotes innovation on A's part by essaying innovation itself.

3.2. Promoting Innovation

(a) If more than one deductive scheme exists (as a separately headed or disjoint conjunctive derivation structure) and if the schemes (conjunctive structures) have analogous *parts* but are *not* identical, then B applies epistemic symmetry (Chapter 7, Section 2.5.2) to provoke the syntactic component (and a putative semantic component) of an analogy relation between topics of the existing scheme.

(b) If a principle exists (Chapter 7, Section 2.5.4.), then B applies extrapolation to provoke the development of any existing deductive scheme. *

(c) If an analogy is supported by a strict isomorphism, it stands. If there is an analogy k between *topic* i , *topic* j with $M(i)$, $M(j)$ in $MF(X)$, $MF(Y)$, and it is supported by a generalised *Topic* with $M(\ell)$ in $MF(U)$, then B asks the user to model a projection of $M(\ell)$ in $MF(X)$ or $MF(Y)$ or both. This operation ("Inversion") provokes innovation.

(d) If there are empty cells in the space of descriptors as there are (previous monograph) in an evolving entailment mesh, then B points to the empty cells and provokes the instatement of fresh topics to fill them.

(e) If there is a (suitable) many aim configuration, B requires resolution; if agreement is reached, B instates an analogy relation and, if not, a conditional node.

(f) Using the graphic facility, the mesh can be represented and displayed under any head node at the request of any user.

* The syntactic construction produced by extrapolation may not be interpretable in the existing universes, so, at the next stage, extrapolation leads to the construction of a novel universe in a spare modelling facility. For example, the information theoretic development of thermodynamics (Chapter 7) involves such a construction and is an innovative gambit. A further example is the invention of a (orthogonal) dimension to accommodate the mathematical extrapolation of "number" to "complex numbers". Goodstein (1962) and Polya (1954) give this example, as does Spencer Brown (1969), the latter author in terms that are precisely attuned to the present discussion.

4. ANOTHER VIEW OF AIM INITIATION

We argue that B acts as an innovative assistant to A because the aim initiating operations (a, b, c, d, e, f) have an interesting and equisignificant interpretation under the general title, "problem posing"; i.e. (given a network of topic relations) "form and pose problems that will generate further topics".

Von Foerster and Weston (1974) note, in their discussion of context oriented systems, that no problem exists without context. A relational specification on its own is insufficient to determine a problem, let alone an acceptable class of solutions to a problem. For example, under the relations x and $=$, the pseudo problem

$$2 \times 3 = ?$$

might be solved by 3×2 or by 6; or to cite a further example from Von Foerster and Weston, the curiously enigmatic pseudo problem posed by $6 = ?$ has any number of solutions depending upon the context in which this relation is embedded.

A fortiori, an uninterpreted network does not in itself determine a problem. But all of the procedures used to initiate or catalyse constructive activity are context proposing (hence, problem posing) operations. A few of the proposals may be as specific as the contextual resolution, " $6 = \text{some product of integers}$ ". Most are far less specific though possibly no less useful. The procedures are surely not complete and in that sense do not constitute an "Artificial Intelligence" (or, as we prefer, in the spirit of the context paper, a "General Intellect"). But they represent part of such a thing, and in combination with the other routines, yield a system in which it is impossible for an external observer to tell whether the innovation (if any) that takes place is due to the user A or to the heuristic B. As promised, B *encourages* innovation.

5. SPECIFIC PROCEDURES

The principles and operations of Chapter 7, Sections 2 and 3 are built into B as a number of "problem posing" or "innovation attempting" procedures.

5.1. B examines the network built up on the grid for analogies between a *topic i*, which is part of a subnet superordinate to node *i*, and a *topic j*, which exists in isolation. By *epistemic symmetry*, B infers that there may exist a subnet superordinate to *topic j* which is isomorphic to subnet *i* and is formed by copying the subnet *i* across the analogical distinction to form a hypothetical subnet *j*.

B displays this subnet by illuminating the *attention lamps*, of which there is one to each position on the grid. The display is intermittent since there may be, and quite commonly are, several topics with the status of *topic i*. A single display consists in illuminating the attention lamp on *topic i*, and whilst it is turned on, scan-illuminating the attention lamps in the hypothetical subnet *j*.

This operation is interpreted as a B question to the user, "Do you affirm or deny the existence of each topic on subnet *j*?" There are two equally productive ways of pursuing an answer: justifying affirmation, and justifying denial. So far, it has only been possible to implement the former method.

An affirmative reply from the user, in respect of an element *v* of subnet *j* consists in placing a node at the position on the grid occupied by *v*; this node being thereby given an *active* status (notice, however, that the node does not cover the attention lamp at this position).

Denial (which, in the current implementation, is not followed up) is achieved by pressing a key on the operating console at the moment when the denied element *v* is scan-illuminated coincidentally with *topic i*. As a result of denying that *v* is a topic in the thesis under construction, subsequent scan-illuminations of subnet *j* do not include *v*.

Once initiated, the display of subnet *j* in the context of *topic i* is repeated from time to time, unless

- (a) the tenure of all elements *v* in subnet *j* is denied, or
- (b) all affirmed elements (with nodes positioned) have been derived and instated, so that the corresponding nodes are no longer in an active state.

5.2. As soon as a fresh *topic i* is instated on the grid, B searches the entailment set of this topic for a node representing a *principle* (any *topic j* of the kind described in Chapter 7, Section 2.5.3.).

If such a topic exists, B infers from *Extrapolation of Principles*

that *topic j* might be applied to *topic i* as a means of obtaining some further *topic k*, and B thus displays the pair (node *i*, node *j*) coincidentally by illuminating the attention lamps at these positions on the grid.

The display is interpreted as a question to the user, "Can you obtain a further topic (which is part of your thesis) by applying principle *j* to *topic i*?"

An affirmative reply consists in placing a node *k* at a position superordinate on the grid to the node of *topic j*. This node becomes active, and the attention lamp display is repeated from time to time until *topic k* is derived and instated.

A negative reply is given by pressing a key on the console at the moment when the display is presented. As a result of denial, the attention lamps are extinguished, and the proposed application of the principle is deleted from B's repertoire.

5.3. B searches the *descriptor* space for any conjunction of descriptor values that specifies a unit set and is not occupied either by a node or a uniquely specified *channel* (recall that the descriptions cover the data base, as well as the topics). B prints out the description and asks if there is such a topic, which the user must affirm or deny. The procedure was exemplified in Chapter 7, Section 2.5.

5.4. In Chapter 7, Section 2.1.8, we discussed the construction of a generalised topic (GHWM) to represent the similarity in an analogy relation (HWC) between "heat engines" (HE) and "refrigerators or heat pumps" (RP) and noted that specialised forms of GHWM could be realised as isomorphic models (more general and more comprehensive than HE or RP) in the universes of compilation and interpretation proper to HE and RP, respectively.

Suppose that GHWM was, in fact, constructed in THOUGHT-STICKER. For this or any generalisation based on an other than isomorphic analogy relation (detected by the absence of the reserved isomorphism operator \Leftrightarrow), B asks the user to construct the specialised topics obtained by interpreting the freshly instated supporting generalisation (for example, GHWM) in the original universes of interpretation. The user A is required to "invert his generalisation". The request from B to A is a typed out question, "Is there a case of the generalisation supporting an X, Y, analogy ac-

tually realised in $MF(U)$ within the original universes of compilation and interpretation $MF(X)$ and $MF(Y)$?" An affirmative reply is evidenced by instating fresh nodes in X and Y , respectively, or in just one of them.

5.5. A completely negative reply, "the proposed construction is *impossible* according to my thesis," denies the validity of an analogy relation based upon the generalised topic. Such replies are stored by B and are the main evidence at B 's disposal for contradicting a mooted analogy relation (though not the generalised topic itself).

5.6. The last process, resolution of a many aim situation, is the most general weapon in B 's armoury. Notice that resolution of a many aim situation is always productive.

(a) It enlarges the set of semantic descriptors.

(b) If agreement is reached its syntactic component is inscribed in the mesh as the similarity part of an analogy relation (and usually a generalisation based analogy relation).

(c) If there is *disagreement*, the syntactic product is a conditional analogy, as the mark of rival theses.

Resolution is probably also the commonest transaction. We conjecture that all autonomously produced analogies and conditionals are due to "internal transactions" of this kind; only a few of them are captured as "official" and observable resolutions. To the extent that THOUGHTSTICKER *does* capture *some* of these internal transactions, it is able to exteriorise innovation.