ISA : THE BUSINESS LINK.

Rory K. Woodward

November 11, 1987

Cognitive Science Research Paper

Serial No. CSRP 093

School Of Cognitive Sciences,
The University Of Sussex,
Brighton, BN1 9QN

ISA : THE BUSINESS LINK.

Rory Woodward.

ABSTRACT

Business simulation models have been used for many years as an important element of the corporate planning process. We suggest that many business simulations are flawed, because they expect exact quantitative values for all data. We are developing an architecture for an intelligent simulation assistant (ISA), which will aid a decision maker in using a large number of existing business simulations.* This paper outlines the architecture, concentrating upon input of non-precise data, simulation using both qualitative and quantitative information, and qualitatively determining causality to aid in debugging of the simulation and construction of explanations.

CONTENTS

ISA - The Business Link.

1. INTRODUCTION

Business simulation models have been used for many years as an important element of the corporate planning process. They help in forecasting the evolving state of a business and its environment, and assist evaluation of future strategy. A company's performance is affected by many factors and all must be given a quantitative interpretation if they are to be included in a simulation. Many of these factors, however, are based on subjective and qualitative information, where a bad judgement can lead to spurious model output on the first, and all subsequent, runs. We suggest, therefore, that many business simulations are flawed. They require input in precise quantitative terms, and produce results in the same form. The discrete nature of such results (right or wrong) can lead management to either disbelieve the ability of the model, or accept its predictions too eagerly.

We are developing an architecture for an intelligent simulation assistant (ISA), which will aid a decision maker in using a large number of existing business simulations. A basic requirement is that ISA accepts a formal description of a model as data, and is therefore not tied to a specific simulation. Consequently, it is a general purpose tool easing the interface between a quantitative simulation and the decision maker. This paper outlines the architecture, concentrating upon input of non-precise data, simulation using both qualitative and quantitative information, and qualitatively determining causality to aid in debugging of the simulation and construction of explanations.

ISA is currently being designed and implemented, and this paper concentrates on one of the central modules of its architecture. Section 2 gives a brief introduction to some of the goals and design principles involved in ISA. Section 3 describes the use of qualitative reasoning to assist in simulation with incomplete knowledge, and indicates some of the current research directions, intended to allow qualitative reasoning to be used alongside traditional quantitative simulation.

2. STRUCTURE AND HJRPOSE OF ISA

2.1 Goals Of The Project

ISA is designed to be a general purpose tool to assist in the lose of a large range of business models. A given model, and associated knowledge, is treated as data to the system, and it is therefore necessary to reason about the structure of the model, the ways in **iich its parts interact, deficiencies in the logic, etc. We identify three main functions to vrtiich ISA should provide assistance:

1. Model calibration.

   This goes beyond maintenance of those parameters (e.g., tax rate, market response to supply), \dhich can be deemed constant over a number of model runs, and includes adjustment of the model structure (i.e., the equations / processes / etc, involved) to better reflect the simulated environment in the presence of new data.

2- Forecasting.

   Forecasting includes the use of heuristics to aid the sinaalation or simplify the ccatputations, the use of qualitative and quantitative reasoning to allcw simulation to continue if seme of the data is imprecise, sensitivity analysis (e.g., the degree of sensitivity of the results to the input data), and analysis of the results to infer meaningful explanations.

3. Diagnosis.

   Diagnosis includes identification of the factors vjhich have caused a disparity between the simulations predictions and either actual events or the user's goals (postdiction), suggestion of potential policy options vdiich might achieve the goals, and inference about the type of external (to the model) events that could have caused the difference.

For a system to assist intelligently in these areas, domain Itootfledge is required - knowledge of the cenpany and its o m H vrmirerrfc. erf tfhA simulation structure and rarroose. of

2.1.1 Acquisition Of Qualitative Data

One of the major aims of the research is to develop techniques which reduce the initial requirement on the user to give strict quantitative values to all input parameters. Instead, the user should be allowed to express uncertainty over the value which a parameter may take. The user may, therefore, wish to express the value of a parameter as being within some range (e.g. market share is between 10% and 15%), or in a qualitative way (e.g. profits are "low", sales are "reasonable"). The danger is that, as the input data becomes less exact, the chance of being able to calculate a precise answer to a precise question diminishes. This can easily lead to a situation in which such a large range of answers are possible, that no useful information can be derived from them. In this way, a model can easily become under-constrained, producing an answer which covers a large range of possibilities, and is therefore useless. However, given knowledge of the degree of precision with which values were entered, and heuristic knowledge about the parameters and the model, it will be possible to determine a subset of parameters which, when constrained further, would significantly affect the degree of precision of the final output. This suggests the potential to prompt the user in an intelligent way for further refinements of the initial input data, concentrating upon those critical parts of the model which would produce the requisite degree of accuracy.

The idea of reasoning about a business in purely qualitative terms is appealing [1] [2] [3] [4] [5], however, there are a number of problems which arise when qualitative reasoning is applied to business simulations, rather than physical systems, attributable largely to the lack of constraint involved when the model is interpreted qualitatively. We suggest that when only a subset of the parameters are given in qualitative terms it will be possible, with the addition of heuristic knowledge, for simulation to produce an answer which is constrained sufficiently to be useful.[6]

2.1.2 Causal Reasoning For Debugging And Explanations

Qualitative reasoning performs another important function within ISA - determination of causality. In the diagnosis stage of model usage, it is necessary to discover the reasons for differences between predicted and actual events. By applying a qualitative interpretation to the algorithms of the model, the parameters likely to have caused the divergence can be identified, and these can be compared to a catalogue of uncontrollable events (described in section 4), allowing guided questioning of the user.

Causality is also important in the forecasting module. For a manager to accept and believe a model and its predictions, he/she must understand how and why it has come to certain conclusions - a reiteration of the equations involved is not sufficient. What is required is an explanation module which uses qualitative reasoning to provide causal knowledge, allowing a comprehensive description of the important factors which contributed to the output.

3. QUALITATIVE REASONING

This section provides an introduction to qualitative reasoning, describing some of the problems, and suggesting a number of partial solutions. It also introduces the idea of linking qualitative and quantitative reasoning - one of the most important abilities of ISA.

3.1 The Goals Of Classical Qualitative Reasoning

Qualitative reasoning* is a relatively new sub-field within Artificial Intelligence that is fast gaining recognition. Although the goals of the topic vary with different researchers, we will use the following definition: "derivation of a description of the behaviour of a mechanism from a qualitative description of it's structure"[7]. Although there are a number of tasks to which qualitative reasoning can be put [8], there are two that are usually credited with particular importance:

1. Prediction.

   Determining the future behaviour of some device or system (e.g., will sales fall if supply is cut, can the delivery facilities accommodate a major order from a foreign customer).

2. Postdiction.

   Determining how a known state might have been caused (e.g., why has turnover fallen, why is the production line under-utilised).

The theory of qualitative reasoning is relatively straightforward. A model is described by a set of qualitative equations (e.g., if the number of customers increase, then sales increase, that is, sales monotonically increase with the number of customers). To perform a simulation, the parameters of the model are set to initial qualitative values (i.e., sales are positive but falling, work in progress is low and steady), and the system

---

* The field of qualitative reasoning comes under a variety

propagates the values based upon the equations (if sales are falling, the number of customers must be falling, etc), until an equilibrium situation is achieved. The trace of the values of the parameters is the qualitative behaviour of the system over time, given the initial situation.

Unfortunately, the reality of qualitative reasoning is more complex, because of the imprecision of qualitative arithmetic. For example, if the increase in trade and other debtors are both positive, then their sum is positive, however, if one of them is negative, then their sum is undefined, because we may not have sufficient knowledge of the magnitude of the two values. In these situations, we have multiple possible future states, creating a tree of possible behaviours, and the system must investigate each of them. The normal domain of qualitative reasoning is physical devices, which are generally well constrained (the tree can be pruned), however, this is less frequently true of business models, and can result in such a large number of possible behaviours (assuming it remains computationally tractable) that "anything could happen".

3.1.1 Qualitative Process Theory

One of the better known approaches to qualitative reasoning is Qualitative Process Theory (QP), developed by K. Forbus [8]. The central idea behind the theory is that physical processes are the mechanisms which cause change in the physical world, and a language for describing a process is presented. From this, it is possible to determine when a process will begin and end (limit analysis), and the combined effect of several processes (achieved by resolution of inferences).

The variables upon which the processes act are represented as "quantities", which consist of an amount and a derivative. Both amounts and derivatives are numbers, which have a sign and a magnitude. Numbers and magnitudes are described by reference to a "quantity space" (ordered set of landmark values), which is a partially ordered set of numbers or symbols, including the value '0' allowing the sign of the number to be determined.

Because the properties of objects can change, they are described by a set of "individual views", i.e.:

```
Individual View Products-On-Shelf(p)
  Individuals:
    sf a shelf
    pd a product
    ro a retail-outlet
  Preconditions
    In-Sales_Area(ro, sf)
    Space-For_Product(sf, pd)
  QuantityConditions
    A[number-in-stock(ro, pd)] > ZERO
    A[number-on-shelf(sf, pd)] > ZERO
    A[number-on-shelf(sf, pd)] <=
      A[number-in-stock(ro, pd)]
  Relations
    There is p, set-of-object
    amount-of(p) = number-on-shelf(sf, pd)
    made-of(p) = pd
    on(p) = sf
    in(p) = ro
```

This individual view is one way of describing a quantity of a product available on a shelf in a retail outlet. The individuals are objects that must exist before the view can become active, the quantity conditions are inequalities referring to the quantity space, and the relations are statements that are true whenever the individual view is active.

Hayes' notion of a "history" [9], is used to represent how changes occur over time, with a history composed of a series of episodes and events. Events last for an instant, and serve as the start and end points of episodes. The term "situation" is used to represent a "slice" through time for a number of objects (this can be either an instant or an interval), which is necessary for processes which act between several objects. It is mentioned that it would be useful if the objects in a model could be split into an optimal number of situations, which could then be considered semi-independently, and this is called the "local evaluation problem".

The theories above allow the description of physical situations at points in time, but provide no means to change from one situation to another. Processes are used to describe the ways in which the parameters of objects change over time, i.e.:

```
process salejofjproduct
  Individuals:
    prod a prociucts-on-shelf
    cons a consumer-group
    dp a process-instance,
      process(dp) = desire-product
  Preconditions:
    Status(dp, active)
  Quar±ityOonditions:
    A[desire(cons, prod)] >
      A[quantity(cons, prod)]
  Relations:
    let required be a quantity
    required = A[desire(cons, prod)] -
               A[quantity(cons, prod)]
    required > ZERO
  Influences:
    I-(quantity(prod), required)
    1+(quantity(cons, prod), required)
    I-(cash(cons), required)
```

The (simplified) process above describes a consumer group purchasing a desired quantity of a product (a separate process would describe the case where quantity available is less than the quantity desired). The process will become active if a consumer group requires a product that is available, and when active, enforces constraints on the individual views involved: that the amount the consumers require is related to the amount they desire less the amount they possess (note that the syntax ᶠA[Quantity]• extracts the magnitude of that quantity); that as their requirement increases, the amount of the product available for sale will fall; etc.

A process is similar to an individual view, *except* that it has influences upon the parameters of the individuals. When the preconditions and quantity conditions of a process hold, an instance of the process is active. The influences in a process describe direct effects *upon* the parameters of the individuals, and can be either negative, positive, or unspecified. Additionally, indirect influences occur when two individuals are related by the "qualitatively-proportional" function, and one of them is changing (due to a direct or indirect influence). The examination of all influences active \jpon an individual can be vised to determine its derivat.ve.

When this occurs, a limit point is reached, and the individual vie^ becomes inactive (and, in most situations, a new view becomes active to take its place - i.e., consumer-sated) • Similarly, processes may became inactive when the conditions for their existence no longer hold, and new processes become active (i.e., re-order-product).

Because any changes in the system can only be caused directly or indirectly by processes (the "sole mechanism assumption"), it is necessary to maintain a vocabulary of all the processes that can occur in a given domain. With this process vocabulary and a collection of irxiividuals, it is possible to make a number of basic deductions. By examining the individuals involved and the conditions of the processes, the set of processes that are active in a situation can be determined (this is called the "process structure" of the situation). Once the active processes are known, the changes occurring to the parameters of the individuals can be determined from the direct or indirect influences, and these are represented as the derivatives of the values involved. Deterirdning the derivative of a value is known as "resolving its influences", and is achieved by collecting the direct and indirect influences together, and cxaribining their effects (this cannot always be achieved unambiguously) • Another operation that can foe performed is limit analysis. By considering the neighbouring points (landmarks) to each changing quantity, and determining whether these are limit points, and would therefore cause a change to the active processes and individual views, the possible ways in which the processes could change from the current situation can be inferred.

The techniques described above can be examined to achieve qualitative reasoning. The use of QP theory is appealing, because the definition of causality is intuitively simple - a change in an object is caused by the processes that are active or the propagation of their effects through active constraints (constraints exist because processes are active) • The use of histories also enables the system to describe when objects are created or destroyed, and how their parameters change, in relation to other objects.

3.1.2 The Difficulties Of Qualitative Reasoning About Business Models

In fill a model of the dynamic flow of funds within a that Stockholders Funds are diminished by the movement of Share Capital into the organisation, and that these funds, together with other sources, form the input to the Cash Reservoir. This is not qualitative reasoning - it is merely describing a portion of the model.

Attempting to use the model, as given, for prediction in a qualitative reasoning system is not productive. Because the illustration only provides flows of funds from and to various "tanks", and the ways in which they combine, the number of potential next states is large. Even assuming that this were computationally tractable, the output would be a huge number of possible behaviours, because, quite literally, almost anything could happen. The problem is the lack of constraint. In general, the model uses addition and multiplication as a means of propagating the parameters, but qualitatively, these operations are not well defined (see section 3.1).

3.1.3 Some Partial Solutions

The constraints in [1] are merely what the model shows us for any business. If used for a specific business, the company strategy on the flow of funds will become relevant, that is, additional constraints will be suggested because the model will be used within the context of the current business strategy. These additional constraints might be, for example, that investments may increase, but only to a limited point (this may be difficult or even impossible to do qualitatively), that the raw materials inventory is to be maintained at a constant level, or that credit sales may no longer be given.

However, not every element of company strategy, that is relevant to the model can be implemented as qualitative constraints. Similarly, many of the heuristics that have been learned about the nature of the business environment cannot be implemented in this way. For example, the difference between credit sales and cash sales (i.e., the level of credit which is given) may depend on many factors other than those which can be modeled qualitatively, for example, repayment history of certain customers, suspicions (market research, etc) regarding their future intentions (that is, strategy intentions, e.g., minimise repayment of trade credit, or unplanned intentions, e.g., impending bankruptcy), etc. These are examples of the types of suggesting, however, that qualitative simulation could include an "expert system[11] as a part of its filtering mechanism, to avoid consideration of future states which, although feasible according to the simulation, are not relevant to the current line of *enquiry.*

This type of "embedded expertise" could also help avoid seme of the problems inherent in qualitative simulation. For example, the points or intervals of time at which states occur are purely qualitative, that is, the only determinable relationship is their ordinal value, and we cannot, for example, assume any relation on the interval between each one (i.e., that they are equi-distant from each other in real time). However, we can impose sane constraints using heuristics or rules, for example, the raw materials inventory may only be increased when the work in process inventory readies a specific level (i.e., only maintain sufficient raw materials to supply a certain level of work in process) • This type of filtering is not available in standard qualitative reasoning, because it was designed to reason about physical mechanisms, where all possible behaviours which the physical (qualitative) constraints allow are legal.

3.2 Qualitative And Quantitative Seasoning

Using both qualitative and quantitative knowledge during the same simulation is an appealing, yet difficult, goal. It is appealing, because in many cases, scare of the parameters involved will be knewn in precise quantitative terms, yet converting them to qualitative values (in order to perform qualitative reasoning) loses significant information. It is difficult because the nature of a qualitative value means that it nay not have any direct quantitative analogue - it is defined by its ordinal position to other qualitative values. Very few researchers have attempted to solve this problem, yet a limited degree of success has been achieved by a few. The approaches can be split into postdictive and predictive reasoning, and these will be discussed in the following two sections.

3.2*1 Postdiction

One of the few attenpts to use qualitative and quantitative values for postdiction, is that of Simmons and Davis [10], who are interested in the problem of geological diagnosis -

process of "imagining" - given a set of activities (geological events) and a goal state (cross-section of the earth) simulate the effects of the activities. The set of activities can be determined by "scenario matching" (backward reasoning from the effect of a process to it's cause), and this hypothesised sequence of activities is then tested by imagining (that the events are taking place).

The basic approach is as follows. Given a start state (i.e., just bedrock) and a goal state, apply scenario matching by reasoning backwards to determine a possible sequence of activities, using a set of rules which provide explanations (causes) for the local effect of a process. The activities determined are merely a hypothesis because:

1. The rules provide causes for local effects which need not be globally consistent (i.e., the hypothesised cause of a local effect might be 'tilting', but this might not be globally true across all rock formations).

2. The evidence of a process might no longer be explicit in the diagram (i.e., a deposition may have been eroded).

Given a start state, a goal state, and a hypothesised sequence of events, reason forward from the start to the goal state by imagining that the events are taking place. This is done in two stages - qualitative and quantitative. The events are first simulated qualitatively using the process descriptions (see below) to reason forwards, creating objects, and producing a sequence of changes to the attributes of those objects. The next stage is to determine the quantitative values of the attributes, by reasoning back from the known goal state. The value of an attribute is measured from the goal state, and corrected for the changes that have occurred to it over time. It will not always be possible to determine an exact value for an attribute, in which case, a range is inferred. Finally, forward reasoning is used to simulate the model quantitatively, using the quantitative values just determined, and replacing numeric ranges with a value arbitrarily chosen from within that range. Attribute values are stored in a quantity lattice, so that the effects of choosing a value or range correctly propagate through the system. Consequently, the effect of choosing a value from a range does not affect the final

Ttoo events during the iitagining phase can cause the system to return to the scenario matching stage. The first occurs when the final state of the sinplation does not match the goal state* The scenario itatxhing module will then attempt to create an alternative hypothesis, presumably by applying a different selection strategy, backtracking, etc. The second situation occurs when an event is not applicable to the current situation (because first, events are only hypothesised due to local phenomena, and second, the events were hypothesised by reasoning backwards, however simulating them forwards might prove them inappropriate). In this case, the imaginer produces an explanation of why the simulation cannot continue (the event that could not be simulated, and the difference between the current state and the state required). Scenario matching vises a form of means-end analysis to infer a process (or sequence of processes) that could eliminate this difference.

There is significant similarity between the qualitative aspects of this theory, and QP theory (see section 3.1.1). Firstly, qualitative reasoning is performed by creating and destroying objects, or altering their attributes, due to processes that are determined to be active at a point in time. Histories are used to represent objects (that is, an ordered set of values for each parameter, which may be temporally related to the histories of other objects depending on the constraints present), process descriptions have a similar form (preconditions, existing objects, objects created, effects, and relations), however, objects and their instantiations (individual views in QP theory) do not have such a complex description, because of the specific domain within which the Simmons and Davis system operates (i.e., it only needs to know about rock units, boundaries, and geologic points).

3.2.2 Prediction

Using qualitative and quantitative knowledge for prediction has been attempted, with various degrees of success, e.g., by [6], [11], [12], and [13]. The techniques used can be suramerised as follows: perform qualitative reasoning on the model, and where ambiguities are found, attempt to resolve by quantitative analysis. Quantitative reasoning, therefore, is used to disambiguate the qualitative analysis, while qualitative simulation is ur^ed to minimise the quantitative arithmetic necessary.

irrelevant to the present discussion). In their domain, quantitative values could be determined (by measurement) for a variety of the parameters involved, however, each measurement has a precision (the degree to which the measured value can be believed), and a cost (of taking the measurement, i.e., due to installation of measurement equipment). The system, when an ambiguity is discovered, must determine the different quantitative computations that could be expected to remove the ambiguity, and choose the option which minimises cost and maximises precision of the parameters which require measurement.

This has similarities to our domain. Quantitative values of parameters in a business model have a precision (a subjective estimate of copy efficiency would be less precise than the cost of a given advertising run) and a cost (copy efficiency could be estimated by the user, while cost of advertising might require interaction with advertising agencies, designers, etc). The work of [6] is, unfortunately, limited. Their aim is to spot a fault in the condenser, which they achieve by simulation, and table lookup. The table lists five faults that could occur, and the qualitative values of various parameters which would indicate it (i.e., pump malfunction is indicated by decreasing power of the pump, and constant thermal transmission coefficient, etc). The problem is two-fold. First, the program must have knowledge of every possible fault that could occur, and a means to unambiguously identify it from the qualitative values of the parameters. Second, the program cannot spot a situation in which more than one fault has occurred. If we read "fault" as "unexpected event" (i.e., a competitor suddenly reduced his price), then our system must be capable of hypothesising that more than one unexpected event has taken place, and further, we cannot be expected to create a complete and finite list of these.

4. STRUCTURE OF ISA

The proposed structure of ISA is based around the qualitative and quantitative reasoning components, with support modules which use (generally) heuristic knowledge to assist and interpret the simulation. Five knowledge bases and three main reasoning components are involved in this process. The knowledge bases required for the task are:

1. Object Knowledge Base.

   The object knowledge base defines the types of objects that can be used by the various models in the system. An object is defined as some entity (physical or conceptual) which may exist and can be manipulated by the model. Objects have parameters, sub-objects, conditions for existence, and relations to other objects. The types of objects needed will be determined by the model being used. For example, if modelling the transportation of goods to customers, we may need a <lorry> object (with parameters <fuel> and <position>, sub-objects <driver> and <load>, conditions for existence <not_in_use> and <fuel.quantity = +required>, and relations between <position> and the <route> object, etc), or if modelling the competitive effects of a new advertising strategy, we might require objects such as <advertising_medium>, <customer>, <retail_outlet>, <generic_competitor>, <complementary_goods_producer>, etc. Objects define a general type, and when required, a model will create one or more "instances" of them to represent real items which the model can manipulate.

   The definition of a parameter of an object includes knowledge of its type, and heuristic knowledge to assist in determining its value. The system should also have knowledge of how to determine the quantitative value of a parameter (ask the user, ask someone else, interrogate another on-line system), and the cost and likely precision of that value. For example, it might know that last years sales (if an attempt to locate this value in some database fails) is likely to be available to the user in exact quantitative terms, while copy effectiveness for the coming period is unlikely to be known with such precision.

2. Process Knowledge Base.

   During a simulation, objects will be created (a new competitor enters the market), their parameters changed (sales increase), and destroyed (a promotion ends). This takes place because one or more processes are active. When a process is active, it can affect objects (and processes), and by this method, the simulation is performed. Processes might be defined for <market_penetration>, <conveyor_belt_in_motion>, or for <diffusion_of_innovation>. The definition of a process includes the objects it affects, the conditions for its existence (e.g., which objects and processes must be active, what the parameters of an object should be), and the effects that it has (e.g., on the object parameters).

3. Model Knowledge Base

The model knowledge base contains the descriptions of the models that are available to the system. For a given model, it identifies the objects and processes that can be created (i.e., points to the object and process knowledge bases, respectively), the ways in which the model can be used, and the various initial configurations of objects and processes that can be set up before simulation. It can also contain specialised heuristics which are applicable to a single model, to reduce the number of possible next states, and simplify the computations - this will be be termed "meta-knowledge", because it is knowledge about how to perform a specific simulation. Meta-knowledge is necessary, because a model uses a general set of objects and processes, which are designed to be applicable across a wide range of models, and therefore contain sufficient detail for any model which might make use of them. However, depending on the purpose of a model, this level of detail may not be needed. A model, for example, may need to know that goods are delivered to a customer, but will not require knowledge of the different ways the lorry can be loaded, etc - merely that transportation exists, and therefore the goods will be delivered.

4. Marketing Activities Knowledge Base.

Catalogue of potential marketing activities, over which the company has control, that could alter the to "understand" the reasoning behind some of the users input values.

5. Uncontrollable Events Knowledge Base.

   Catalogue of classes of uncontrollable events (e.g., strikes), and the effects that these might have on the input data, together with probabilities of each of them occurring, to interact with the user, suggesting possible causes for results which do not agree with the simulation's predictions, and which can be checked and accepted / denied. The user might also tell the system which parameters are definitely unchanged, resulting in a more intelligent search of possible events. An unpredictable event would be indexed by the effects that it has on the model parameters, with sample descriptions of the events that could have resulted in such a change. For example, a strike affecting distribution, and bad weather conditions, might have a similar effect on the parameters - these two would therefore be instances of a general class of unpredictable event which limits availability of the commodity (a specific instance might slightly modify the parameter effects of it's class).

The marketing activities and uncontrollable events knowledge bases seem, at first sight, rather ambitious. Their main purpose is to suggest explanations which might account for an observed behaviour, and possible actions which might achieve a given result. Surely the explanation process is at the heart of intelligence? In [14], Schank suggests that explanations are merely the use and mis-use of stored explanation patterns. It is the mis-use which is most interesting - take an explanation pattern which does not fit the anomaly observed, and alter ("tweak") it until it provides a valid explanation, and if successful, generalise the new explanation pattern, and add it to memory. This process, Schank claims, is an algorithm for creativity.

For our purposes, an uncontrollable event can be regarded as an explanation pattern - it is a stored, general, pattern, which explains an anomaly which might be observed between predicted behaviour, and the activities which actually occurred. Using these patterns, the system should be able to infer possible explanations which are not completely backed up by the known facts (i.e., to borrow an example from [14], their inventory at the same time. This type of explanation is of the "delta agency[11] type - the actor who produced the anonaly may be acting under orders frora a higher authority). With a hierarchical explanation (as the pattern becomes more specific), the system should be able to recognise the general class of an event (i.e., lew production due to lack of personnel), even if the specific instance has never been described (i.e., secondary picketing).

Ihe preceding description of the knowledge bases has largely described the functioning of the three main reasoning modules, but to summarise, these are:

1. Quantitative Reasoning Module.

   This module performs standard quantitative calculations, and range arithmetic.

2. Qualitative Reasoning Module.

   This module performs qualitative reasoning, as described in section 3.

3. Control Module.

   This is the core module of ISA, which interacts with the user, determines the question to be asked, the model to be used, etc, uses the qualitative reasoning module to perform the simulation, and attempts to resolve ambiguities using the quantitative reasoning nodule. When new information is received, it compares this to previous predictions, and if anomalies exist, attempts to provide explanations, updating the knowledge bases as required.

Obviously, the control module is a complex entity, with access to all sub-modules of the system. It's main purpose is to monitor a simulation as it is performed, maximising the quality of the data available.

5. OONCLUSIONS

Simulation models typically require exact figures for every relevant input, before they can project the current world into the future, however, if the input data is dubious, the output cannot be regarded with much credibility (GIGO). Mthough a company would be expected to knew accurate quantitative data for .it's accounts, sales figures, stock levels, etc, it would be unlikely to know such exact information for it's competitors, suppliers, the economy, and so on. For example, it may be known that a competitors product line is profit making, but the degree of profit would be harder to determine. However, the fact that a competitor is making a prof it is still useful knowledge, and should be usable in assessing that competitors business, and its effects upon our own.

In this paper, we have suggested that the theory of qualitative reasoning can be used to address these problems, and that by using both qualitative and quantitative reasoning together, significant iitprovements can be achieved in the performance of simulation inodels.

To date, the research has cxancentrated on development of the arcMtecture, and identification of the problems to be resolved. By inplernenting a sub-set of the QSIM algorithm [7], we have tested qualitative versions of various models (e.g. [1]), to determine the difficulties involved (see section 3.1.2). We have experimented with range arithmetic, as a useful addition to the quantitative reasoning module (see section 2.1.1). We have studied various approaches to causal reasoning, as a means of improving the explanation capability (see section 2.1.2). The initial stages of iroplementation of ISA are currently being performed.

References

1. Hart, P. E., Barzilay, A* and Duda, R. O. (1986). Qualitative Reasoning For Financial Assessments: A Prospectus. The AI Magazine.

2. T&ngen, K. and Wretling, U. (1986). Intelligent Front Ends To Numerical Simulation Programs. Research And Development In Expert Systems III, Bramer, M. A. (ed).

3. Fjellheim, R. (1986). A Knowledge Based Interface Td Process Simulation. AI Applied To Simulation, Simulation Series Vol. 18, No.l.

4. Rajagcpalan, R. (1986). Qualitative Modeling And Simulation: A Survey. AI Applied To Simulation, Simulation Series Vol. 18, No.l.

5. Reddy, R. (1987). Epistemology Of Knowledge Based Simulation. Simulation Vol. 48 No. 4.

6. Gallantir M.f Gilardoni, L., Guida, G., and Stefanini, A. (1986). Exploiting Physical And Design Knowledge In The Diagnosis Of Complex Industrial Systems. The 7th European Conference On Artificial Intelligence.

7. Ruipers, B. J. (1985). Qualitative Simulation Of Mechanisms. MET laboratory For Computer Science, TM-274.

8. Fortous, K. D., (1984). Qualitative Process Theory. Artificial Intelligence 24.

9. Hayes, P. J., (1985). Naive Physics 1: Ontology Of Liquids. In Formal Theories Of The Ccramonsense World, Hobbs, J. R., & Moore, C. R. (Eds), Ablex Publishing Corporation.

10. Simmons, R. G. & Davis R. (1983). Representations For Reasoning About Change. AI Memo No 702, Massachusetts Institute Of Technology.

11. De KLeer, J. (1977). Multiple Representations Of Knowledge In A Mechanics Problem-Solver. 5th

University Of Edinburgh.

13. Cross, S. (1983). Qualitative Reasoning In An Expert System Framework. FhD Dissertation, University Of Illinois at Urbana-<harapaign.

14. Sdharik, R. C. (1986). Explanation Patterns: Understanding Mechanically And Creatively. Lawrence Erlbaum Associates, Inc.