

A Conceptualization and General Architecture of Intelligent Decision Support Systems

P. Struss^a

^a*Technical University of Munich*
Email: struss@in.tum.de

Abstract: The paper presents an attempt to conceptualize decision support and various generic subtasks and to develop a general architecture of intelligent decision support systems. We decompose the task of decision support into subtasks whose input, output, and function are characterized. This is based on a small number of concepts: besides “decision”, the essential ones are “observation”, “situation”, “goal”, “action”, and “process”, which are in turn defined using elementary concepts for characterizing the system under consideration, or our model thereof.

This is not an academic exercise aiming at providing definitions, but a prerequisite for a generic architecture of decision support systems with interfaces for certain generic functions, the comparison of basic modules implementing these functions, and the configuration of systems from a set of such modules.

The primary subtasks whose (intelligent) solution is heavily dependent on domain knowledge are situation assessment, i.e. inferring what is happening in a system from a set of observations, and therapy proposal, i.e. developing plans for interventions to achieve certain goals starting from the current situation.

Secondary tasks are situation and plan evaluation (checking whether and to what extent a situation or plan satisfies or violates goals), prediction (forecasting the future development starting from a situation with or without interventions), and observation/experiment proposal (designing activities to collect information, possibly after stimulating the system in a particular way, useful to disambiguate situation assessment and also situation evaluation).

Keywords: *Decision Support System (DSS), Model-based reasoning, Process-oriented modeling*

1. INTRODUCTION

1.1. What is a Decision Support System?

A survey of definitions or characterizations of decision support systems (DSS) and of proposals for general architectures or (sub-) functions of DSS delivers rather disappointing results. Many of the offered definitions boil down to “A DSS is a computer system (or set of tools) that supports making decisions”, which turns any data base, MATLAB, excel, even Google into a DSS. When architectures are proposed, they are often presented as a huge set of tools and computational steps embedded in a confusing web of interconnections, often heavily emphasizing data-driven techniques. The components are mainly characterized as various alternative or complementary techniques, rather than by the function they implement.

A systematic analysis and a conceptualization of DSS seem to be missing. This is not an academic question aiming at delivering a set of definitions. It is a practical necessity, if we are interested in the systematic re-use and integration of different tools and methods. And it is a prerequisite for establishing requirements on DSS and its components, especially when we want to build “intelligent DSS”.

The conceptualization and architecture proposed in the following does not claim to offer a general account for all existing DSS. Its objective is to provide the foundation for the design and comparison of **knowledge-based** DSS and the identification of subtasks and modules and their principled interfaces. This leads to a second disclaimer: although we are fully aware that observations, knowledge, and inferences can be subject to a significant degree of uncertainty, we do not explicitly represent this at this stage of the formalization. The consideration behind this is that uncertainty (e.g. in terms of probabilities) can be **added** to the presented concepts, potentially “softening” results and introducing ambiguity and alternatives. We start with the fundamentals, the concept of decision and what it means to make a decision.

1.2. Decision Making

A decision has to be taken only if there is a **choice**. It means picking one from a set of alternative options. These options are different ways to **act**. If the possible ways to act are complex (rather than only a single action), we call them a plan, which is a set of actions, possibly with a particular order or temporal extent. Hence, a decision has a set of plan options as an input and results in choosing one of them. The necessity to act and the criteria for the selection depend on two crucial concepts: the **goals** that should be achieved (or maintained) and the current (or an assumed) **situation**. Hence, the basic concepts are related as depicted in a static structure diagram of the Unified Modeling Language (UML) of Figure 1.




Figure 1 A decision selects a plan from a set of options to achieve goals in a particular situation. The line with the diamond indicates that a Plan is a collection (set) of Actions. Other straight lines are simple associations.

The resulting characterization of the activity as “selecting the best plan for achieving or pursuing goals in a given situation” is decision making in the narrow sense. It assumes the main tasks that require knowledge, domain expertise, and reasoning to be done:

- achieving a sufficiently detailed understanding of the situation that allows a good choice, and
- developing a set of alternative plans that promise to be applicable and to satisfy the goals.

If we want to describe the tasks and subtasks of intelligent DSS, we need to consider a decision process that starts from observations that provide only a partial description of the state of the system, goals, and possible actions and generates a plan to be executed, i.e. the outer box in Figure 2, which illustrates the information flow, as the following two figures do.




Figure 2 A decision support system and its breakdown

In the next section, we identify subtasks and their required inputs and outputs in an informal way. Section 3 is dedicated to associating a precise meaning with the involved concepts.

2. STRUCTURING INTELLIGENT DECISION SUPPORT SYSTEMS

2.1. High-level Decomposition

Following the definition of the concept “decision”, a DSS should support or automate the selection of a plan for achieving certain goals based on a set of observations. Doing this in a way that deserves the attribute “intelligent” certainly implies determining whether or not and to what extent carrying out a plan will effectively transform the given situation to one that satisfies the goals. This, in turn, requires the step to interpret the observations and in order to derive a more complete picture of the current situation, which includes, in particular, a representation of the present causal interactions in the system. We call this subtask, which certainly requires deep domain knowledge, **situation assessment**. Another subtask, which is necessary, unless the different options to choose from are given *a priori*, is the generation of plans that promise to achieve the goals. The inputs to this **therapy proposal** step are the situation, the goals, and the potential (types of) actions that can be taken. Together with the **decision making** step, i.e. the choice of one plan from the options, we obtain the high-level decomposition of the DSS shown in Figure 2. We will further decompose the subtasks.

2.2. Therapy Proposal

The initial step in therapy proposal is to evaluate the situation with respect to the given goals, i.e. determining which goals are violated and, perhaps, in which way (e.g. a particular threshold exceeded or undershot, relevant objects missing, or unwanted ones being present). This **situation evaluation** (see Figure 3) may not only be necessary for the current situation, but also for future situations with or without applying actions. This requires a **prediction** module, which produces the description of such forecast situations, e.g. through simulation. The discrepancies between situations and goals are an input to **goal generation**. This is needed, although we already have a set of explicit goals, because, usually, we need to determine some intermediate goals that guide the search for a plan. If the iron concentration in drinking water is above the threshold, even after a remedial action, the respective goal will be violated for some time, while a feasible intermediate goal will be reducing the iron concentration. This set of goals, together with the situation description and the available actions is an input to the **plan generation** module.




Figure 3 Therapy proposal and sub-tasks

2.3. Decision Making

The key step leading to a decision is **plan evaluation**, by accumulating **situation evaluation**, i.e. checking whether and to what extent they establish situations that fulfill the (intermediate or general) goals. Again, this analysis may not only be carried out for a single situation, but for a sequence of situations created by executing a plan with several steps or situations in the evolution of the system after the intervention has been finished. Therefore, **prediction** is involved in this subtask, too. This is, usually, a knowledge-intensive task, since determining the impact of actions on the system and its evolution requires an understanding of the driving forces and interactions in the system. In contrast, the step of picking one plan from a set of suitable ones based on the plane evaluation is a more or less formal task based on the plan evaluation, which is guided by weights and priorities of goals and cost of plan execution, factors that are external to the domain model.




Figure 4 Decision making and sub-tasks

So far, we characterized the inputs and outputs of the DSS and its various modules by informal concepts. We continue by defining these concepts more precisely, which then allows us to specify the required functionality of the modules more formally.

3. CENTRAL CONCEPTS

3.1. Foundations

The main concepts we used in the description of the interfaces and their intuitive meanings were:

- **Observation:** “What is the case” (information)
- **Situation:** “What is going on” (applied knowledge)
- **Goal:** “What should be the case” (objectives)
- **Action:** “What can be done” (potential)
- **Plan:** “What should be done” (intention)

We will define these concepts in terms of more elementary concepts and, thus, establish the basis for formal representations and automated reasoning. The above concepts are all assertions about the world under different modalities, which means we need to fix a basic ontology for statements about the considered domain, a class of real or potential systems. We decide to categorize such statements as propositions about

- **Structure:** in terms of a set of existing **objects** and **object relations** between them
- **Behavior:** by constraining quantities or properties of the objects to a particular range.

Figure 5 introduces the unifying concept of an **assertion** in a UML static structure diagram, depicting that (non-)existence is assigned to structural elements, which are objects and object relations. Objects have quantities, and assertions can associate a certain range to them.

Although this may seem restrictive, we believe it is appropriate and expressive enough for ecological and environmental systems and also social, commercial or organizational systems, at least if the concepts of object and quantity are taken in their broad meaning and not restricted to solid physical objects and their parameters. Objects can be species, populations, individuals, chemical substances, locations, social groups, companies, currencies, etc. Their relations include “living-in” (a population in an area), “feeding-on” (a link in a food web), “competitor” (populations or companies), “applied” (a certain policy by farmers).

Relevant quantities might be usual measurable or observable properties, such as size of a population, concentration of a dissolved substance, pH of a water body, amount of crops, but also more abstract characteristics: intensity of exploitation, degree of degradation of a mangrove forest, and even environmental awareness of people or social attitudes. Existence statements may be negative, e.g. if certain species are not present in an area.

While this represents a purely descriptive perspective, the dynamics of a system, or, rather, our understanding of it, is captured by the

- **Causal structure:** by a set of **processes**, i.e. the active mechanisms and “laws” that we consider to be the drivers of the evolution of the observed system.

Processes emerge from configurations of related objects and modify objects and relations, including their creation and elimination and changes in continuous quantities. Examples are chemical reactions, reproduction of a population, energy consumption, accumulation of capital, and even social interactions. We will provide a formal definition of this concept in section 3.6.

In the following, we define the more complex concepts in terms of these elementary ones.

3.2. Observation

Observations are the starting point of the entire process of decision making, in providing a description of a certain section of the world, but restricted to its visible or accessible surface. This superficial and incomplete picture of the system is represented as a subclass of the concept of assignment introduced above (Figure 6).

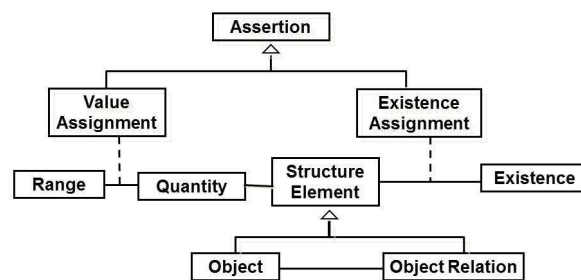


Figure 5 Assertion. The triangle expresses a subclass relation (E.g. Object and ObjectRelation are subclasses of StructureElement). The dotted line is meant to specify objects as binary relations, e.g. ValueAssignments as pairs of Quantity and Range.

The objects and relations in this description will be tangible ones, usually objects visible or perceivable by a user (who may not be a domain expert), and their observable or measurable features. For instance, there is information about the existence and size of populations of different species in a particular area, drinking water and a measurement of its iron concentration, or the crops in certain agricultural area and their market prices. One may aggregate assertions to represent structured observations, such as “high concentration of iron in drinking water”, which contains existence assignments to the objects drinking water and dissolved iron as well as to the relation dissolved-in and the value assignment “high” to the concentration quantity.




Figure 6 Observation, goal, and situation

Observations may be temporally indexed and represent time series and system evolution, but they do so in an “objective” and descriptive way only. This is to say that, in our purist view, they do **not** contain any representation of the causes for the existence, state, and evolution of the involved objects or relations: processes are **not** part of the observations, because they are never visible themselves – although their results may be – but are mental reflections of a system capturing our understanding of a domain (the regularities and “laws” that have been discovered, so far, i.e. models).

Observations may not always be the result of actual measurements, but unverified assumptions about some conditions that normally hold. This may be helpful, or even necessary, to enable a reasonable situation assessment, but inferences based on them have to be defeasible. And even the reliability of actual observations may be limited, which motivates us to allow for the association of some uncertainty to an observation, without restrictions on how to represent it and use it in inference processes. While we use only a binary categorization into facts and assumptions, others may consider probabilities, certainty factors, etc.

Finally, we emphasize the distinction between observations in our sense and the raw data that may be available. We are aware that bridging this gap and transforming the latter into an appropriate input to the core of the knowledge-based DSS may require substantial efforts pre-processing of the data, which is actually a focal point of many existing DSS, but considered here as a pre-processing step and not discussed in detail.

3.3. Situation

A situation or, more precisely, a representation of a situation is meant to capture not only a completed description of the observed system, but also reflect our understanding of the underlying causal interdependencies. Hence, we represent it as a set of assertions and a collection of processes (Figure 6). The activity of certain processes has to be inferred from the observations. The situation description has to comprise the original observations (at least the factual ones). However, the inferred or hypothesized processes may entail the existence of objects that are not directly observed and also impose further restrictions on quantities. For instance, a situation description that causally explains the increase of dissolved iron concentration by a process of re-dissolving of solid iron from the sediment of the reservoir with a pH in the acid range includes the existence of solid iron, which may not have been directly observed or mentioned before, as well as the value assignment to the pH.

The more complete picture of the system, primarily its causal interpretation, is a prerequisite for designing interventions, because they will only work if they can eliminate the existing negative factors or generate new counteracting processes or strengthen existing ones – without creating new negative influences. Such a situation description is helpful for decision support only if it is stated in fairly technical or scientific terms which allow drawing rigorous conclusions, e.g. about the future evolution of the system and the effects of certain interventions. We have to be aware of the fact that the given observations may reflect the user perspective and be stated in different terms than the ones used in the scientific model. A practical solution may therefore require some translation between assertions in observations and those in situation descriptions, an aspect which we do not explicitly address at the abstract level of our description.

3.4. Goal

While observations and situation refer to the actual or hypothetical state of a system, goals characterize its desired properties. Goals are not intrinsic to the system, but reflect the perspective of some stakeholder(s). There is no natural limit for iron concentration in water, but only one enforced by regulations and laws. And even the prevention of extinction of a species is a highly legitimate objective of people, but does not follow from any “laws of nature”.

Again, we represent goals as a specialization of assertion (Figure 6), which may have an associated degree of relevance, whose appropriate representation will depend on the domain and on the preferences of the system developers. One may just use a classification into mandatory and optional, or decide for a qualitative or numerical ranking. This will be used for a refined solution for the modules that evaluate situations and plans and, thus, influence plan selection. Aggregation of assertions in order to represent complex goals may be handy in practical applications. Note, that we exclude the activity of a process to be a goal – reflecting that only the overall result in terms of objects and their properties matters –, but if one wants to be able to formulate constraints on how the result is achieved, one may consider to modify this design decision.

As discussed in section 2.3, initial goals may be general, absolute and long-term ones, such as “dissolved iron concentration c_{iron} in drinking water below $\text{conc}_{\text{limit}}$ ”; but for deciding on remedial or supportive actions, one needs to establish intermediate goals, in the above case e.g. “reduce c_{iron} ”.

Goals may refer to directly observable features (like population size), but often they do not, which is one reason why the completed situation description is necessary to determine whether or not goals are fulfilled in a particular situation. Even if the fulfillment of goals is observable in principle, the task of plan proposal requires powerful situation descriptions and predictive inferences, because it requires projection into the future. Goals may relate directly to terms in the situation description “preserve the population of an endangered species” (which could be translated to “population size $> p_{\text{limit}}$ ”), but since they reflect the stakeholder perspective, they may not be stated in the same language as the situation description. Often, they will be more abstract (“preserve biodiversity”) and require a translation and breakdown into elements expressible in the situation description. Given the level of this presentation, we do not address this translation step, without denying that it will often be important and challenging.

3.5. Action, Plan

Actions represent possible human interventions affecting the ecological, environmental, social, and/or economic system. While carrying out an action in reality may involve complex interactions of people, the use of tools and machines, etc., its relevant aspect in the context of DSS is its impact on the system under consideration. The DSS will have to determine the overall impact based on the situation description and a description of the direct effects of the action if carried out. In a first approximation, these **effects** are assertions, since the action may create or destroy objects and object relations and modify properties of objects. It is important to only represent the **direct** effect of an action, rather than all the changes due to the internal dynamics of the system that are triggered by this immediate effect. For instance, an action may create an increased flow of chlorine into a treatment plant, while the (intended) effect of a reduced concentration of dissolved iron is caused by the oxidation process, which is part of the system and not part of what is directly controlled by the human action. This “fringe” of the system, which is amenable to human intervention may shift and depend on the granularity of the representation. For instance, the action may not directly cause a flow of chlorine, but only increasing the opening of a valve.

Furthermore, one may need to express restrictions on the situation as prerequisites for being able to carry out the action. Also these preconditions can be represented as assertions. It is important to emphasize that they express only what needs to be satisfied by the situation to make the action possible, not what makes it useful or promising. For instance, for increasing flow of chlorine, the precondition might be stated as “availability of chlorine”, but not as “high iron concentration” as condition which suggests the action. The interventions of humans can be considered as becoming part of the system, and under this perspective, actions appear as a special kind of processes. We will, therefore, indeed specify actions as a subclass of processes later on.

3.6. Process

Processes represent the mechanisms that are the driving forces of the system’s dynamics (or, rather, our model thereof). Although the term “law” is highly inappropriate for denoting what we extracted as an understanding of how natural systems or artifacts behave and change from empirical data and theoretical inferences, processes are considered to dictate these changes in the real world. They are, in turn, considered to be generated or triggered by certain conditions of a particular situation. Accordingly, one part of a process description is a collection of **preconditions**, which are assertions. Whenever the




Figure 7 Process

mentioned objects and object relations do exist and the value assignment hold, the process is inferred to be active. For instance, whenever the existence of solid iron, sediment, and a water layer as well as the contained-in relation between iron and sediment and the above relation between water layer and sediment are established, and the value of the pH is restricted to an acid range, re-dissolving of iron occurs.

The second element of a process description is the descriptions of the **effects** this process imposes on the system. This is the creation of iron contained in the water layer above the sediment, which can be represented as the existence assignments to dissolved iron and the contained-in relation between the iron and the water layer.

However, there is an impact of the process that cannot be represented as an assertion: the re-dissolving process will transport iron into the water layer above the sediment. Why should this be beyond a value assignment? Why not represent it as a restriction of the derivative of the iron concentration in the water to a positive value? The reason is that the change in the iron concentration may not only depend on the re-dissolving process, but also by other, potentially counteracting processes, such as ascending of iron to a higher water layer and oxidation of iron, which may override the positive influence and finally even lead to a decrease of the iron concentration. Hence, we need a new concept, which actually goes beyond the expressiveness of differential equations. What we need to express is not a positive derivative of the concentration as an effect of re-dissolving, but its positive **contribution** on the concentration. Only after combining influences from all processes that impact the concentration, the actual change in the concentration can be determined. Following the Qualitative Process Theory by Forbus (1984), we call these contributions **influences** and have to represent them as separate concepts. This completes the representation of processes as depicted in Figure 7.

3.7. Actions as Processes

If humans interfere with a system, the actions taken can be represented as a special kind of processes. There can be preconditions in terms of the respective situation to enable the effect of the action. But actions are special because activating them depends on yet another precondition, namely the **decision** to carry out the action (or the assumption thereof): for instance, the opening of the valve is not caused by system-internal processes. The decision to carry them out must be explicitly represented as a decision. We can introduce a new concept, or we can simply turn such a decision into a virtual object in our representation and make it subject to existence assignments. Thus, actions become a special subclass of processes.

Please note, that if human activities are considered as deterministic reactions completely controlled by factors in the system, like the migration of farmers to cities if their income drops below a threshold, this corresponds to a process in the situation description.

4. CURRENT STATE AND FUTURE WORK

In this paper, we propose a conceptual foundations and general structure of knowledge-based DSS, but did not present details on particular theories and techniques for realizing the different subtasks. A restricted version, based on a rigorous logical specification of the involved subtasks, has been implemented (see Heller and Struss (2002)). It is based on a logical reconstruction of Qualitative Process theory introduced by Forbus (1984) and exploits theories and techniques from model-based problem solving (Struss (2008)) Situation assessment is performed by constructing minimal situations that are consistent with the observations based on a domain library of processes. Therapy proposal operates in the same way, adding actions to the domain library and seeking to establish consistency with the goals. This solution does not represent evolution over time and is restricted to finding sets of simultaneous actions (rather than sequences) that achieve (intermediate) goals. Spatial aspects are represented by relations between specific spatial locators and other objects. In the future, more work will be dedicated to the development of heuristics for a best-first search for situation assessments and plans. Long-term research will address the integration of temporal representations and planning.

REFERENCES

- Forbus, K. (1984). Qualitative Process Theory. *Artificial Intelligence* 24 (1-3)
- Struss, P. (2008). Model-based Problem Solving In: van Harmelen, F., Lifschitz, V., and Porter, B. (eds.). Handbook of Knowledge Representation, Elsevier
- Heller, U. and Struss, P. (2002). Consistency-Based Problem Solving for Environmental Decision Support. *Computer-Aided Civil and Infrastructure Engineering* 17