



Diagnostic reasoning in action

Rasmussen, Jens

Published in:
IEEE Transactions on Systems, Man and Cybernetics

Link to article, DOI:
[10.1109/21.247883](https://doi.org/10.1109/21.247883)

Publication date:
1993

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Rasmussen, J. (1993). Diagnostic reasoning in action. *IEEE Transactions on Systems, Man and Cybernetics*, 23(4), 981-993. <https://doi.org/10.1109/21.247883>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

DIAGNOSTIC REASONING IN ACTION

Jens Rasmussen¹

Abstract. The task of diagnosis is a very important topic in many different contexts. In highly complex technical installations involving high hazards, such as process plants, diagnosis is a crucial part of disturbance control; in technical maintenance, diagnosis is necessary to locate the root cause of system failures; and in medicine, diagnosis is the basis for any patient treatment. The paper presents a discussion of the basic nature of causal reasoning as applied for diagnosis and the mental strategies applied when diagnosis is viewed as an integrated part of "natural decision making" for interaction with the environment. A typology is suggested to characterise diagnosis in different domains such as process control, maintenance and medicine. In addition, an attempt is made to distinguish between the features of diagnosis depending on the ultimate aim, whether it is explanation, compensation, repair, or punishment and the difference in the context of the task, "the causal field," related to the mental model involved in the different cases is outlined.

Keywords: Diagnosis, diagnostic strategies, medical diagnosis, technical diagnosis, causal reasoning, natural decision making, mental strategies, mental models, accident analysis.

I. INTRODUCTION

Diagnostic reasoning has been studied from many academic points of view, such as artificial intelligence, cognitive science, psychology, social judgement, and decision theory. Within each of the disciplines diagnostic reasoning has been chosen as a cognitive task well suited to explore some basic psychological mechanisms in focus of the particular discipline and the findings are, quite naturally, biased by their fundamental research interest and methodology. It is not the aim of the present paper to review these different approaches. Instead, the diagnostic task is approached from a problem oriented point of view: How is diagnostic reasoning shaped in the actual, complex work context, and how can diagnostic reasoning be supported? The presentation, therefore, cuts across the topics studied in detail by the individual academic schools and will, very likely, raise objections as being incomplete and even wrong from the different specialists. However, we have met difficulties in applying the results of most academic research on diagnostic reasoning for design of support systems and, hopefully, the *structure* of the problem as presented by the paper will serve to create a better interface between the selective research efforts and system design.

According to Webster, the basic meaning of the term 'diagnosis' is "the act or process of deciding the nature of a diseased condition by examination", or: "a careful investigation of facts to determine the nature of a thing." This definition is focused on diagnosis as an analytical process, separate from the planning of the ultimate act based on the outcome of the diagnostic effort and its definition is

¹ In: IEEE Trans. SMC. Vol. 23, No. 4, pp. 981-993.

closely related to most academic approach to the study of diagnosis. In the present paper, diagnosis is considered in a wider context including its role in action. Consequently, it appears that a more appropriate definition is the original connotation which, also according to Webster, originates from Greek: "diagignoskein," i.e., "knowing the difference." That is, diagnosis involves the act of distinguishing one case from another, of separating a relevant item from the general context. In that sense, diagnosis is closely related to categorisation, to the act of labelling. From this it is immediately clear, that in order to describe the act of diagnostic judgement, it is essential to determine, *what* is subject to categorisation and, as well, *why* is the categorisation made.

Even in an "objective" attempt to classify and to give a name, the diagnostic context is essential. The decomposition of the world into elements to classify and the choice of attributes for classification depend on the particular point of view applied by the categoriser. A biologist and a cook classifying plants and animals will apply quite different classes, based on completely different attributes, as already pointed out by Linneaus [1]. In other words, the purpose and the related point of view define the classes and their attributes or, in the terms chosen below, the relevant diagnostic field.

In conclusion, to be able to describe the diagnostic process, it is necessary to consider carefully the object world in which the categorisation will take place and, as important, the context in which an actor finds himself, i.e., the goals and the repertoire of alternative actions, relevant in the diagnostic situation.

II. BASIC ISSUES

To set the stage for a detailed discussion, some basic issues will be reviewed in the following paragraphs. First, the difference is considered between diagnosis viewed as a separate decision task and as an integrated part of the cognitive control of goal directed actions.

A. DIAGNOSTIC JUDGEMENT IN THE LABORATORY

It is not the aim here to present a review of the large number of laboratory studies of diagnosis in various work domains. To illustrate the problem of generalising from well controlled laboratory studies to an actual work context only one comparison will be given to compare with our experiences from field studies, that is the study of diagnosis from the point of view of social judgement theory. Diagnostic behaviour has been studied extensively within this paradigm, normally by an analysis of the utilisation of the available cues in laboratory judgement tasks. This approach has been used to study diagnostic judgement in several professional activities such as stockbrokers, clinical psychologists, and physicians, see e.g., [2]. In experiments, cues identified as diagnostically relevant by expert judges are used to prepare trial cases to present to subjects, generally in the form

of cards with sets of attributes. From this evidence, a statistical model describing diagnostic behaviour is developed. The general result has been that linear statistical models, such as multiple regression analysis, have been adequate. Four characteristics of expert judgement are typically found by such experiments. First, the judgement process tends to be very simple. Even though experts identify up to 10 attributes or cues to be relevant to diagnosis, they actually use very few, usually only two or three, and the process tends to be purely additive. Second, the process tends to be inconsistent. Subjects do not use the same rule from case to case, and judgement in a second presentation of a case may differ considerably from the first time. Third, there are wide individual differences even among subjects with years of experience. They differ with respect to the cues used and the weights they apply. The fourth general result is that people are not very good at describing how they make judgements. [2]

A very similar approach to the study of diagnostic reasoning has been taken in medical philosophy [3]. Diagnosis has been defined as the act to seek, isolated from the application context, the attributes necessary for classification of a "case." The theoretical basis has been influenced by the causal theory of Mackie [4], as discussed in a subsequent section.

B. MULTI-ATTRIBUTE JUDGEMENT AND DIAGNOSIS IN ACTION

However, analyses of the diagnostic judgement in an actual work context tend to paint a different picture. Comparing the results of laboratory studies with our analyses of diagnostic tasks in hospitals and repair shops, we can identify some important differences which will signal great caution for transfer of the laboratory results to the actual professional work context and call for a wider definition of a "diagnosis." This statement does not imply that the results of laboratory experiments are not valid for multiple attribute judgement tasks, but rather that isolated multiple attribute judgement is not always a characteristic of a real-life diagnostic judgement.

First, the experimental design suggests that decision-makers be subject to an information input, which they have to process. The task is isolated from its normal context and, therefore, the 'tacit knowledge' of the subject has no opportunity to be synchronised. In actual work, subjects are emerged in the situational context and they are, therefore, tuned to ask focused questions to the environment, rather than to process multiple attributes sets. The various features of the context through time according to Gibson serve to up-date the 'attunement' of the organism [5] or, in our terms, to synchronise the "internal, dynamic world model" [6].

Second, in actual work, a diagnostic judgement is not a separate decision task but intimately connected with the subsequent choice of action. Diagnostic judgement for action is not a theoretical categorisation of the observed data, but a search for information to select, among the perceived alternatives for action,

the one matching the case in question. Models of decision making are normally structured such as to be a sequence including situation analysis, goal formulation and priority judgement, and planning. This normative sequence is the basis of the decision ladder in figure 1. Experts in action, however, have a repertoire of heuristic short cuts by-passing the higher levels of the ladder. In any familiar situation, they only perceive a small number of alternative plans - this is the core of expertise - and they only need information enough to resolve the choice among those plans. Therefore diagnosis and action is intimately connected.

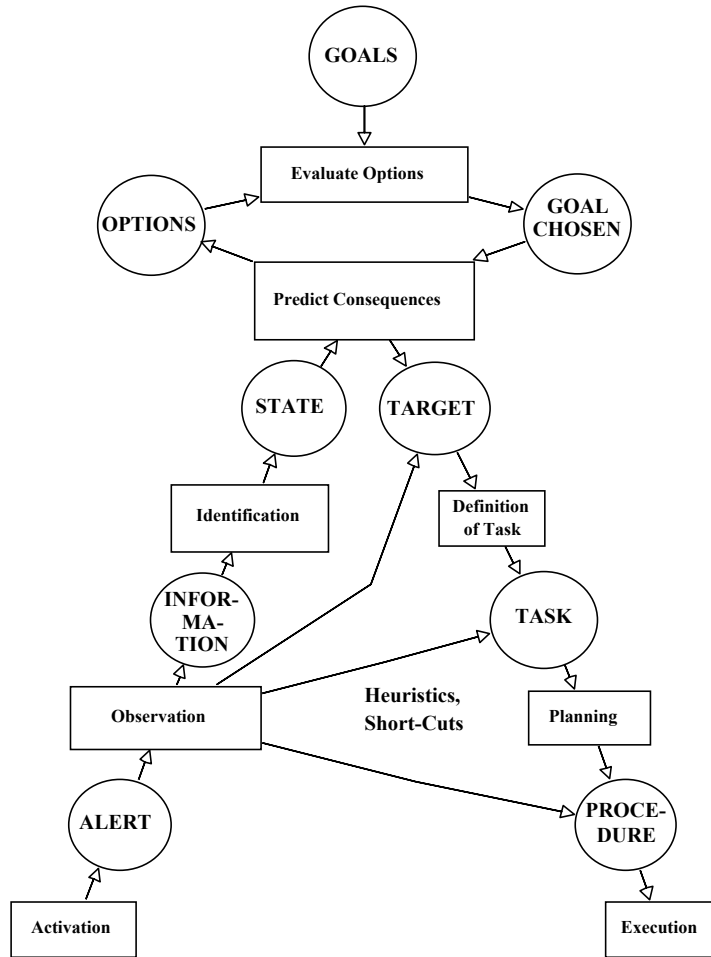


Figure 1. The figure illustrates the sequence of basic information processes in a decision task together with a number of heuristic short-cut paths. It serves to identify a number of basically different decision functions, which are used to connect different "states of knowledge" with respect to the activity in the work domain. The figure is used in our field studies as a sketchpad for representation of the interaction of situation analysis, goal evaluation, planning and action, and for indication of "recognition primed" shortcuts.

of expertise - and they only need information enough to resolve the choice among those plans. Therefore diagnosis and action is intimately connected.

If the categories to consider are defined by the context in terms of the relevant action alternatives, then the attributes to consult, in addition to be dependent on the state of the object of diagnosis, they also to a large degree depend on the aim of the diagnostician and the context, that is, the nature of the diagnostic field. The lesson to be learned from this discussion is that an analysis of the diagnostic process must be based on a wider definition of the diagnostic task, it must consider the *actual context* of decision making, and it must take into consideration the different *modes of decision making*.

C. DIAGNOSTIC JUDGEMENT IN NATURAL CONTEXT

Diagnostic judgement in the present context is taken in its wider meaning of categorising the state of affairs in the environment with respect to selection of the proper action towards some target or goal. The concept of diagnostic judgement

in this way implies the first phase of decision making in a natural context, which, in the normative case, includes situation analysis, goal selection and action planning. In a familiar context, decision making and planning degenerate into cue-action chaining. It is, therefore, essential to consider the fact, that "decision making" and, consequently, "diagnostic judgement" will take place with reference to different levels of cognitive control of action in a dynamic environment and to different domains of action.

In familiar context, skilled behaviour unfolds as an integrated, continuous flow with no need for interrupt and discrete decisions and relies on data-driven chaining of movement patterns generated by an internal, dynamic representation of the context of behaviour, i.e., of the environment. However, conscious mental operations, 'decisions,' can play an important role by preconditioning the required dynamic representation. Such preplanning can be done by recall of previous, similar activities and situations in advance and rehearsal of likely, useful action scenarios together with preview of the expected points in time when choice between action alternatives will be required. This will serve to introduce landmarks in the simulation scenario of the internal world model and to prepare it for the proper cues for choice. In this way, intuition can be prepared for the events to come, and the actual decision-making then becomes "recognition primed decision" [7], [8] for which no action alternatives are considered *at the time of action*.

This kind of high-skill decision making depends on the conditioning in advance of the internal world model, which is required to generate automatically the proper behavioural pattern on occasion. When the conditioning as described has not been effective, mismatch can be experienced by the person, between the state of affairs in the environment and the predictions by the internal world model. In this case, a number of alternatives for action may be perceived, and the environment will be consulted to read a sign, which can resolve the ambiguity. For diagnostic judgement, this means that the information sought, the attributes consulted, will depend strongly upon the action alternatives perceived to be available for adjustment of the state of affairs in the light of the current objective. This perception may also depend on the perceived consequences of the alternative actions available. E.g., the perception of the involved risk of failure and punishment.

If an acceptable set of action alternatives is not available, recall of prior similar cases can, as mentioned, assist in the identification of the relevant actions and their activation cues. If no resolution is found in this way, and only in this case, resort will be taken to the analytical mode of diagnosis and knowledge-based decision making which has been the focus of most academic research.

III. CATEGORISATION AND CAUSAL REASONING.

It follows from this discussion that diagnostic judgement in the present context is taken to be an identification of the state of affairs in the environment with reference to the actions relevant for the immediate objective. Diagnostic judgement

implies the perception of a causal relation between a state, an action, and the ultimate effect, as related to the current objective. Causal reasoning, therefore, is an important issue in the diagnostic task and a discussion of the nature of causality will be useful. Bertrand Russell [9] discussed the characteristics of causal reasoning in contrast to scientific reasoning.

A classical *scientific analysis* is based on mathematical equations relating physical, measurable variables. This approach depends on a selection of *relationships among variables*, which, in Russell's terms, can be 'practically isolated.' This separation is possible if the relations are isolated by nature (e.g., as they are found in the planetary system) or because a system is designed so as to isolate the relationship of interest (e.g., by a scientific experiment or in a machine designed to support a physical process in a controlled way). In this kind of representation, material objects are only implicitly represented by sets of parameters of mathematical equations. The representation is particularly well suited for the analysis of the optimal conditions and theoretical limits of physical processes in a technical system, which, by its very design, carefully separates physical processes from the complexity of the outside world.

Causal reasoning, on the other hand, depends on *regular connections of events* in time. Causal representations are found in terms of the propagation of events in the environment, i.e., changes of the states or configurations of objects. Russell emphasises the ambiguity of the terms used to define causality: the necessary connection of events in time sequences. The concept of an 'event,' for instance, is elusive: the more accurate the definition of an event, the less it is likely that it is ever repeated. Completeness removes regularity. The solution is not, however, to give up causal explanations. Representation of the behaviour of the physical world in causal terms is very effective for describing complex phenomena, such as e.g., accidents, because the objects of the real world are explicitly mapped by the model and changes, such as faults, are easily modelled. This is the case because causal reasoning is related to changes in the normal context as experienced by the analyst. Therefore, rather than to give up causal explanations, as Russell requests, or to seek objective definitions of events, it must be realised that regularity in terms of causal relations is found between kinds of events, between *types*, not between particulars, i.e., individually defined events or *tokens*.

Russell's distinction is based on a definition of cause and effect as being consecutive events in a conditioned environment. This point of view is focused on the relationship among events. "Causes" are discrete antecedents of other events. In consequence, laws of nature, such as e.g., gravitation, cannot be a 'cause.' To talk of the gravitational force as the cause of the movement of a grandfather clock is a category mistake, mixing concepts from Russell's two classes: causal and deterministic models. The gravitational force does not 'cause' but it 'determines' the movement of the clock mechanism. Measuring time is the 'reason' for the presence of the clock. The 'cause' of the motion of its parts is the owner's push of the pendulum after winding the clock.

The present point of view is that causal reasoning is important for scientific as well as practical reasoning because of its direct mapping of objects and because it maintains a unified representation of the properties of objects and, therefore, is well suited to represent changes of the properties of objects, e.g., as an effect of human actions.

When events and objects found in causal representations cannot be defined by an exhaustive list of attributes, they can only be understood as being *prototypes* representing classes defined with reference to the shared context as defined by the tacit knowledge of the community within which the causal model is the accepted basis of communication.

The behaviour of the complex, real world is a continuous, dynamic flow, which can only be explained, in causal terms after decomposition into discrete events. The concept of a causal interaction of events and objects depends on a categorisation of human observations and experiences. Perception of occurrences as events in causal connection does *not* depend on categories which are defined by lists of objective attributes but on categories which are identified by typical examples, prototypes (as defined by Rosch [10]). This is the case for objects as well as for events. Everybody knows perfectly well what 'a cup' is. To define it objectively by a list of attributes that separates cups from jars, vases and bowls is no trivial problem and it has been met in many attempts to design computer programs for picture analysis. The problem is that the property to be 'a cup' is not a feature of an isolated object but depends on the context of human needs and experience. The identification of events in the same way depends on the relationship in which they appear in a causal statement. An objective definition, therefore, will be circular.

Mackie as discussed below has given a classical example. His statement, that "the short-circuit caused the fire in the house" is a record of a particular case. In the general sense, its significance is to interrelate two prototypes: the kind of short-circuit that can cause a fire in a particular kind of house. The explanation that the short-circuit caused a fire may be immediately accepted by an audience from a region where open wiring and wooden houses are commonplace, but not in a region where brick houses and piped wiring are the more usual reality. If the explanation is not accepted, a search for more information is necessary. Short-circuits normally blow fuses, therefore further analysis of the conditions present in the electric circuit is necessary, together with more information on the path of the fire from the wiring to the inflammable elements of the house. A path of unusually inflammable material was probably present. In addition, an explanation of the short-circuit - its cause - may be needed. The explanation depends on a decomposition and search for unusual conditions and events. The normal and usual conditions will be taken for granted, i.e., implicitly given by the intuitive frame of reference. Therefore, in a causal explanation, the level of decomposition needed to make it understood and accepted, depends entirely on the intuitive background of the intended audience. If a causal statement is not accepted, for-

mal logical analysis and deduction will not help because it will be easy to give counter-examples, which are not easily falsified. Instead, further search and decomposition are necessary until a level is found where the prototypes and relations match intuition. (The reason that nuclear power opponents do not accept risk analysis may be that they have an intuition very different from the risk analyst's intuition, rather than a lack of understanding of risk and probability).

A. RULES FOR TERMINATION OF DECOMPOSITION AND SEARCH

The conclusion of this discussion is that the very nature of causal explanations shapes backtracking in causal explanations such as diagnosis, e.g., of a particular accident or the disease of a particular patient. Decomposition of the dynamic flow of changes will normally terminate when a sequence is found including events, which match the prototypes familiar to the analyst. The resulting explanation will take for granted his frame of reference and in general, only what he finds to be unusual will be included: the less familiar the context, the more detailed the decomposition. By means of the analysis, a causal path is found upstream from the ultimate effect. This path will be shaped by resident conditions, which are latent effects of prior events or acts. Also these resident conditions can be explained by causal backtracking and in this case branches in the path are found. To explain a particular case, such branches are also traced backward until all conditions are explained by abnormal, but familiar events or acts. The point is: how does the degree of decomposition of the causal explanation and the selection of the side-branches depend on the circumstances of the analysis? Another question is: What is the stop-rule applied for termination of the search for causes? Ambiguous and implicit stop rules will make the results of analyses very sensitive to the topics discussed in the community involved at any given point in time. There is a tendency to accept as explanation what you expect to find. For example, during one period of industrial safety concern, technical faults were in focus as causes of accidents, then human errors were predominant, presently the focus is moving up-stream towards errors of designers and managers [11]. For medical diagnosis, similar fashions have been identified by Burnum [12]. This points to the question whether accidents and diseases are related to higher level functional structures and feedback mechanisms rather than to local causal connections. In that case, traditional causal attribution turns out to be fighting symptoms rather than the structural origin of breakdown. This raises the question of generalisation, which is discussed in a subsequent section.

Stop-rules controlling termination of search are not usually formulated explicitly. The search will typically be terminated pragmatically in one of the following ways: (a) An event will be accepted as a cause and the search terminated if the causal path can no longer be followed because information is missing; (b) A familiar, abnormal event is found to be a reasonable explanation; or (c) A cure is available. The dependence of the stop rule upon familiarity and the availability of

a proper action by the analyst makes the judgement very dependent upon the role in which a judge finds himself. The implicit nature of the stop rule frequently influences studies of the causes of accidents. In analysis of anaesthetic mortality, for instance, two concepts are used to categorise causes, i.e., the blame concept (something has been done incorrectly) and the event concept (something has happened) and studies have frequently used mixed categories [13].

To summarise: identification of the cause of a particular case is controlled by pragmatic, subjective stop-rules. These rules depend on the aim of the analysis, i.e., whether the aim is scientific, that is to explain the course of events; legal, that is to allocate responsibility and blame; or therapeutic, that is to identify possible improvements in order to avoid similar future cases.

IV. DIAGNOSTIC STRATEGIES IN ACTION.

On this background, diagnostic reasoning applied as a basis of action in different work situations is discussed in the following sections, as we have seen it in our field studies.

A clarification of the concept of strategies is important here. In the present context, a particular "strategy" is taken to be one idealised, formal category of cognitive processes used for a diagnostic task. All particular implementations of a strategy will be different, but will share a particular kind of mental model, a certain kind of interpretation of the observed evidence, and a coherent set of tactical planning rules. In consequence of this definition, different strategies require very different resource profiles of an actor with respect to mental models, a priori knowledge, empirical evidence, etc. From this point of view, a "strategy" is an abstraction in terms of an idealised, normative diagnostic inference procedure. Performance in an actual work situation involves frequent shifts among the various relevant formal strategies in order to resolve local demand-resource conflicts. The formulation of a set of coherent strategies is necessary for a formal description of the complex task of the diagnostic reasoning in an actual work context.

There are many possible ways to characterise diagnostic strategies. In the present context the focus of interest is the implication of the nature of causal reasoning and, consequently, a particular approach to a typology is discussed in the subsequent sections.

The first distinction is related to the *direction of inference* involved in the diagnostic reasoning. Here three categories are found to be relevant:

1. Going from the *normal* functionality to the actual *abnormal* case: This is the analytical variationist approach, by which judgement is made with reference to a normal or an as-designed state of affairs independent of general, empirical evidence from prior cases. This category includes diagnostic strategies typically found in technical diagnosis.

2. Going from the *general* picture to a *specific* case: Judgement is based on instantiation of empirically established, general causal relations based on evidence

from prior cases. This category includes strategies typically found in medical diagnosis.

3. Going from a *specific* case to *general* statements: Judgement is concerned with generalisation, which involves going from evidence about a particular case to statements about general causal relations. This category includes strategies used to identify possible improvements of future system performance from analysis of particular cases (e.g., accidents).

In the following sections, a number of diagnostic strategies are reviewed with reference to the features of causal reasoning discussed in the previous sections.

A. THE VARIATIONIST APPROACH: FROM NORMAL CONDITION TO ACTUAL CASE

This set of strategies involves judgement of an observed disturbed or abnormal case with reference to the normal state of affairs within the object of diagnosis or to the intended (e.g., as designed) normal functional structure. Inferences in the diagnostic task will typically be made independently of general, empirical evidence from prior cases of faults, accidents, or diseases. This class of strategies is very effective for man-made, technical systems, which serve well-defined purposes by well-known functions. The class is, however, also effective for biological systems, e.g., in medicine when reliable models of the normal physiological functions and states have been established (e.g., in case of physical injuries).

Different idealised strategies are possible for an analytical variationist approach depending on the *purpose* of the diagnosis.

The Analyst's Perspective: Explanation

In an analysis to explain a particular event such as e.g., a break-down in a technical installation, the course of events will be followed backwards from the ultimate effect. As mentioned above, the nature of causal reasoning will require the backtracking to be continued until a cause is found which is familiar to the analyst. If a technical component fails, a component fault will only be accepted as the prime cause if the failure of the particular type of component appears to be 'as usual.' Further search will probably be made, if the consequences of the fault make the designer's choice of component quality unreasonable, or if a reasonable operator could have terminated the effect, had he been more alert or been better trained. Therefore, depending on the perceptions of the analyst, a design error, a manufacturing error, or an operator error can be accepted as an explanation of the same case.

In most recent reviews of larger industrial accidents, it has been found that human errors are playing an important role in the course of events. Very frequently, errors are attributed to operators involved in the dynamic flow of events. This can be an effect of the very nature of the causal explanation. Human error is familiar to an analyst: "To err is human." However, the high skill and efficiency of professional people normally depend on their ability to optimise their work prac-

tise, that is, to depart from instructed procedures. To work according to rules has been an effective replacement for formal strikes among civil servants. It is therefore very likely that an analyst after the fact will find departures from instructed procedures and, thus, identify human errors as part of the stream of events.

The *diagnostic field* in which the course of events will be traced is the as-designed physical anatomy of the system, including normal, causal input-output relations for technical components and equipment, and the operating instructions given to the staff. Understanding the abnormal state involves identification of a change with respect to the normal functioning of a particular piece of equipment, which is familiar to the diagnostician. The course of events is identified as the propagation of aberrations with reference to the intended, designed processes of equipment and components. The search will, very likely, be supported by the use of a functional diagram, by calculated performance plots, by instruction manuals, and by other design blueprints. The diagnostic field underlying this class of strategies is independent of empirical evidence from prior cases. One does not need empirical evidence from prior failures to distinguish failed components from those working properly. The *stop rules* controlling the termination of decomposition and search, however, depend on the subjective experience and level of expertise of the analyst. The more experienced, the less detailed will be the functional decomposition of the object of diagnosis required for the search.

This kind of diagnosis will also be possible in the medical case; no evidence from prior cases will be necessary for the identification of a broken leg or a physical injury of a human circulatory system.

The Operator's Perspective: Compensation

The first concern of an actor faced with an abnormal system will, very likely, be a compensation of the immediate influence on some vital performance parameter of the observed abnormal state. This is the case when a failure disturbs the operation of an industrial process system and thus indicates the possible advent of an accident. In this case, the task is to protect the system and immediately to bring it into a safe state before repair is considered. Similarly, in the medical case, it can be important to stabilise the conditions of an injured victim of a traffic accident, before therapy is considered.

In this case, the aim of the diagnostic search is to identify the endangered system functions and to locate the target of a proper compensating action. This can be done analytically in terms of a causal backtracking in a representation of the internal functional structure of the system from the observed effect until a parameter is found, sensitive to corrective action. The source of the disturbance is only of concern later. In case of fire, you will first look for a bucket of water, not for the possible short circuit.

This kind of causal search is, in particular, possible in case of well-structured technical systems. In such systems, compensatory actions are important for protection of the system against major accidents and damage to the environment.

Major accidental damage can only emerge from the loss of control of the major flows of mass or energy and compensatory diagnosis will be focused on identification of means for regaining control with the major mass- and energy- balances of the system. The *diagnostic field* of the search will represent those major flow structures, which can cause accidents. Means for control then can be identified from an analysis of the physical process involved in the disturbed balance, irrespective of the cause of the particular disturbance and independent of any empirical evidence from prior cases.

Similarly, in the medical case, stabilisation of the state of a patient depends on the control of vital circulatory systems, e.g., the stabilisation of blood circulation or the flow of oxygen, irrespective of the prehistory or cause of the actual condition. That is, the diagnostic field represents the basic circulatory systems of the patient, not empirical evidence from prior cases.

The Repair Man's Perspective: Correction

For the ultimate correction of a faulty condition, the objective is to restore the normal physical state of the system, e.g., by replacing a failed pump or mending a broken leg. Also with this objective, an analytical diagnosis is often possible simply with reference to normal function and without empirical evidence from prior cases. The causal field is defined by the normal, physical anatomy of the object of diagnosis.

In the medical case, correction depends on introduction of a change, e.g., by medication, having a corrective effect. In this case, the diagnostic field may be the same as for compensatory search, i.e., the functional structure of the system in question. Search is aimed at finding an element or parameter in the functional structure, which is sensitive to one of the means available. That is, the causal field is strongly influenced by the diagnostician's medical "tool box."

In the typical case of technical repair or medical surgery in response to physical damage, the diagnostic target is to locate the faulty component or organ and, locally, to replace or mend it. This means that reference to the *location* of the root-cause of a disturbance is the aim of the diagnostic search, i.e., the location of the particular faulty organ or component which is the origin of the abnormal functional condition. In this case, topographic reference to the location of the disturbance must be drawn from the observations. Such topographic reference can be drawn from observations in basically two different ways:

One is to draw a reference from the location of the source of the observations (topographic diagnosis). Another is to infer a topographic reference from the functional significance of the content of the observation, i.e., the pattern of symptoms found (symptomatic diagnosis). The repertoire of diagnostic strategies applied by skilled technicians for locating faults in technical systems have been identified from analysis of verbal protocols. They illustrate the basically different structure and resource requirements of strategies applicable to one particular diagnostic task [14].

Topographic diagnosis. The search for the origin of a fault can be performed directly in the system itself. That is, observations are made at various locations in the system, the observations are judged to be good or bad with reference to a template representing the normal functional state and the possible location of the fault is judged from the *location* of the observations. The topographic search is a kind of good-bad mapping of the system. The diagnostic field in this case is a topographic map of the location of functional elements within the physical anatomy of the diagnostic object together with a set of "normal state" reference templates given for suitably located observation points. The region in which the fault is located can be systematically narrowed down by suitable tactical search rules. For instance, a frequently used fault finding strategy in electronics is the "half-split" heuristic. If the signal is normal at the input of a path, but missing at the output, then the most information economic search tactic will be a progressive "split-in-half" and test of the path, i.e., to zoom-in on the location by always making the next observation at the mid point of that part of the path in which the signal disappears). Also in this case, the diagnostic field is independent of empirical evidence from prior cases. The stop-rule for termination of decomposition and search is very pragmatically determined. There is no need to seek beyond the level of decomposition at which parts can be replaced as standard units, i.e., the stop rule is given by the spare parts and/or the tools available.

Search by hypothesis and test. While the topographic strategy derives reference to the location of the fault from the location of the observations, search by hypothesis and test derives such reference from the *information content* in a set of observations. This involves functional inference based on a symptom pattern. The search is based on deduction of the effect of a hypothetical fault by means of a model of the normal functional constitution of the system. Diagnosis involves generation of suitable hypothesis (frequently found by the initial use of one of the empirically based strategies) and a verification by deduction of the propagation of the functional aberrations from the hypothetically faulty part to the observed features.

The causal field for deduction of symptoms in this case is a representation of the normal functional and causal structure of the system, which can be used to trace the propagation of the effects of hypothetical causes of failure.

The Attorney's Perspective: Penalty

In order to allocate responsibility and judge the need for retaliation, a variationist approach will normally be used to locate a responsible person. The search will take place up-stream along the unusual course of events including erroneous human activities. That is, aberrations will be identified with reference to the usual or intended state of affairs. The search will be very similar to that applied for understanding the case, but a particular stop-rule will be used to terminate search, i.e., the search will stop when a person is found who made an error and at the

same time, 'was in power of control' of his acts. The very nature of the causal explanation will focus attention on people directly and dynamically involved in the flow of abnormal events. This is unfortunate because they can very well be in a situation where they do not have the 'power of control.' Traditionally, a person is not considered in power of control if physically forced by another person or when subject to disorders such as e.g., epileptic attacks. In such cases, acts are involuntary [15] [16], from a judgement based on physical or physiological factors. It is, however, a question as to whether psychological factors also should be taken into account when judging 'power of control.' Inadequate response of operators to unfamiliar events depends very much on the conditioning taking place during normal work. This problem also raises the question of the nature of human error. The behaviour of operators is conditioned by the conscious decisions made by work planners or managers who will be more 'in power of control' than an operator in the dynamic flow of events. This means that the causal field of diagnosis for charging people should be influenced by or include the features of the cognitive control of the people involved in the actual situation. In fact, a tendency is presently seen to include normal management functions in the analysis of accidents although the blame generally is for the people in the stream of events (see for instance the analysis of the Zeebrügge and the Clapham Junction accidents [17], [18]).

In conclusion, the diagnostic strategy and the context in which it unfolds in the legal perspective are very similar to the strategies applied for understanding how a situation evolved, but the stop rule used depends on criteria defined by legislation.

B. THE EMPIRIST APPROACH: FROM GENERAL PICTURE TO ACTUAL CASE

Another extensive category of diagnostic strategies is based on a causal context defined by empirical evidence from prior cases. When the internal functional structure is poorly known, the variationist approach described in the previous sections cannot be applied and diagnosis must then be based on a search in a body of empirical evidence from prior cases. This applies for the typical medical case. Medical diagnosis is normally based on evidence organised in hierarchically structured categories of diseases. Diseases, symptoms, causes and likely therapies are connected by epidemiological and statistical analysis of past cases and controlled experiments with new therapies. The search space will be represented by a kind of decision tree which is a hierarchical structure of nested empirical categories, see figure 2. How this empirical evidence is used depends on the expertise of the diagnostician.

This kind of diagnosis depends on a symptomatic search where the abnormal state of the system is represented by a set of observations - a symptom pattern - which is used as a search template to find a matching set in a library of symptoms. In such a body of empirical evidence, the labelling of categories and relations is based on correlation and the labels can relate to very different conceptual domains - e.g., the related cause, the failed part or function, its location or an available cure. In effect, the strategies used for different diagnostic purposes such as compensation or correction are similar and only the search label is different.

The Textbook Perspective: Decision Tree

The body of empirical evidence for diagnosis in medicine is a pool of experience

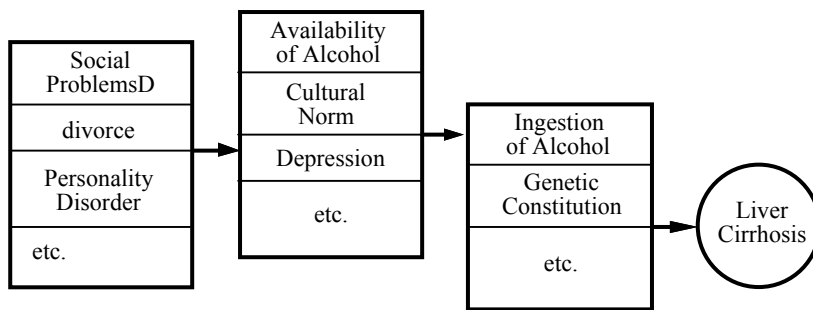


Figure 2. Causal complexes in medicine. The doctor who treated the patient was in no doubt that the liver disease was caused by alcoholism, as the ingestion of so much alcohol can be taken as a sufficient cause of liver damage. But it is also known that the extent of the damage and the cause of the disease are to some extent determined by the patient's genetic constitution. If the patient's ingestion of alcohol had been less extreme or if he had a stronger genetic constitution, he might not have developed cirrhosis of the liver. Furthermore, many other processes can lead to cirrhosis of the liver, as, for instance, various kinds of virus infection, metabolic defects, etc. (Adopted from [24]).

from prior cases with no systematic structure related to the functionality of the "system". However it is possible to choose an idealised, normative order, which can aid in teaching diagnosis to novices and such a normative strategy is found in medical textbooks. The categories are defined empirically with reference to disease categories related to patterns of symptoms and available therapies. The complete tree for even one category will be very complex and have a

structure as illustrated schematically in figure 3. A botanical field guide can be taken to be an example of a consistent and complete decision tree for a diagnostic task. The decision tree is defined by empirically established category memberships and the causality implicit in the structure of the representation reflects logical necessity, not temporally ordered cause-effect relations.

The diagnostic strategy using this body of knowledge is a kind of decision table look-up having the symptom patterns as entries and guided by heuristics derived from the therapeutic options available to the diagnostician. Similar diagnostic strategies are applied for 'symptom based operating procedures' used in certain kinds of hazardous process industries. Operators are guided through a decision tree through a specification of the symptoms to consult (instruments to read) at

each branching point of a decision tree for searching in a library of operating procedures where each entry corresponds to a particular critical situation.

The Expert's Perspective: Recognition Based Diagnosis

Experts have a large body of knowledge from prior cases immediately available and will not enter a comprehensive search of a decision tree but use heuristic guesses.

Direct recognition.

Frequent cases can be recognised directly by an expert when their symptom patterns are met in a familiar context. The categories of abnormal conditions, represented by particular symptom patterns, can be labelled directly empirically in terms of its cause, the location of disturbed function or its proper cure. No co-




Figure 3 serves to illustrate that in a particular diagnostic session, the causal tree considered emerges from a more complex tree of possible branches which are, however, not activated in that particular case. The tacit knowledge of the analyst focuses the attention on the relevant part of the total causal network. This selective focus of a first guess is the hallmark of an expert diagnostician. (Adopted from [24])

herent or explicit diagnostic context can be identified; memory search depends on templates with no guidance from the symbolic connotation of the pattern. The diagnostic context in this case is the subjective pool of episodic evidence of the particular diagnostician.

According to Pejtersen [19], a variant of this recognition strategy is the empirical strategy applied by skilled librarians. The population of users is divided into stereotypical classes associated with certain categories of book con-

tents. For instance, elderly ladies with grey hair and glasses are very likely to be offered family novels without any actual interrogation of their needs. Similarly, it is very likely that medical doctors link classes of patients with treatments and their initial associative guess will determine the diagnostic context activated for a more conscious search and verification by other strategies.

Recognition-guided decision making. In less familiar cases, even experts have to search through a decision tree. The decision tree actually considered for a given case cannot include all possible branches of the knowledge base, and normally the sub-tree used will evolve along with the diagnostic process. This selective use is guided by the tacit knowledge of the analyst which, for the expert diagnostician, results in a very effective "first guess". This in turn serves to limit the necessary search space; see figure 3.

In conclusion, different variants of empirical diagnostic strategies are found, based on empirically established decision trees. Depending on the work domain and the expertise of the diagnostician, the context will take the form of a formal, hierarchical representation of category relationships, an ad-hoc generated search tree including subjective experience or a pool of subjective episodic evidence. Likewise, the search process can be an orderly consideration of the individual branches of a tree; it can be based on recognition guided short cuts or it can be limited to a single direct recognition.

C. GENERALISATION FROM A PARTICULAR CASE

In order to improve a work system in the face of an earlier unacceptable event or happening, another kind of diagnostic judgement is required. It becomes necessary to infer from the particular case a change in the system or its environmental conditions, which will serve to decrease the likelihood of similar occurrences in the future. For this purpose, the causal context has to be analysed to identify a change, which will prevent or break a similar accidental flow in the future. This implies an identification of the actual causal chain in the proper temporal order, as shown in figure 4.




Figure 4 illustrates the variationist approach to analysis of an accidental course of events. Nearly all the events included are labelled with reference to a change of a normal state of affairs or a normal work practice. (Adopted from [23]).

A search for remedies for earlier cases in a causal representation is problematical. Frequently, prevention will be associated with removal of 'root causes'. However, in general, a particular accidental course of events would have been prevented if any link in the causal tree or its conditioning side branches had been broken or blocked prior to the particular occurrence. Therefore it is very easy to suggest many possible counter measures from the analysis of a single case. The danger is that most of these will be special ad-hoc measures and will only be ef-

fective in the specific case. A careful generalisation from such an analysis is therefore necessary. However the opportunities for generalisation depend on the nature of the system considered. The difficulties in generalisation from epidemiological studies, that is, to conclude that because a certain event occurs in many places that have X in common, then X is a causal factor, are well known from the unsettled discussions about the relationship between food and lifestyle and heart diseases.

The Engineer's Perspective: Functional Generalisation

Functional generalisation can be used for systems having a stable functional structure such as technical systems, which are designed to serve a particular purpose. A process plant is a particular exemplar of a class characterised by its conceptual design. In such cases, generalisation involves simply referring back from a particular accidental chain of events to the general design concept. Prevention of a repetition of the particular (and any similar) accident can be planned from an analysis of the course of accidental events through the functional structure. A consistent generalisation can then be based on an analysis of the effects of potential improvements, such as an improved component quality, introduction of redundancy, protective functions against the loss of control of major mass and energy flows, etc.

In other words, the causal analysis serves to identify the propagation of accidental changes in the normal functional structure of the system, as designed, and generalisation involves a reconsideration of the conceptual design.

The Manager's Perspective: Variationist Generalisation

An important class of systems for which generalisation from unacceptable cases is important includes organisations for work system management. Such organisations have a kind of "designed" functional structure and an established pseudo-stable procedural practice. Causal analysis of the propagation of accidental events in an organisation is typically done in terms of the propagation of effects from violations of established practice, that is, by a variationist strategy. One basic restriction of generalisation from a traditional causal analysis is that it presupposes a stable causal and intentional system structure. However, the causal structure as found by an accident analysis (see figure 4) is only a record of one singular case. For intentionally structured systems, such a representation is not a model of its internal, relational structure. It does not take into account closed loops of interaction created by individual and organisational adaptation to changes. Therefore, in systems including feedback loops within the intentional control structure, the identification of "improvements" in causal terms may be unreliable due to inherent feedback compensations. As an example, when radar was introduced to increase safety at sea, the result was not increased safety but more efficient transportation under bad weather conditions [20].

Feedback paths depending on the intentional control structure will change dynamically and thus obscure post event analyses. Moreover it will be difficult, even if such paths are recognised, to determine whether causal arguments in a closed loop constitute a convergent or divergent series. The instability of Watt's steam engine regulator was only understood when Maxwell [21] replaced causal analysis with a solution using differential equations.

Such problems make reliable generalisations in a social organisation by means of a variationist strategy difficult. Therefore, a functional approach to generalisation will ultimately be necessary also for socio-technical systems. If anything, this demonstrates the need for the development of a framework as suggested in the present text.

The Safety Council Perspective: Empirical Generalisation

Many systems have very unstable functional structures (e.g., work conditions at a construction site) and/or they are not transparent to a functional analysis (e.g., complex, dynamic organisations). In such cases, generalisations for improvement have to be based on an aggregation of cases across systems and/or over time as well as comparisons with 'healthy systems' by statistical and epidemiological analyses in order to identify general causal factors. This type of generalisation serves to identify cures, which will be effective across the systems supplying the data. The problem then is that the improvements identified can be rather general - for instance, improved training and/or motivation. A more effective use of a variationist approach also in this case has, therefore, been suggested [22].

The causal context is, like general medical diagnosis, a body of evidence ordered in a hierarchical decision tree, which constitutes the reference case from which the particular "causal complex" used for decision making in the particular case is drawn. As described under the textbook perspective, this empirical strategy relies on causality defined by class membership and logical necessity in contrast to an analytical strategy, which is based on causality defined by the regular, temporal connection of events.

To conclude, the various diagnostic strategies serving as a basis for "system" improvements have also very different structures - including functional reasoning, a search in decision trees and a search for class membership by statistical methods. In addition, they are based on searches in different contexts ranging from well defined functional structures for technical systems to established practice and rules of conduct for social organisations to collections of general empirical evidence for loosely coupled systems.

V. DIAGNOSTIC FIELDS

Formal Strategies

To sum up, for a representation of diagnostic behaviour, a description of the context in which the strategies are used is equally as important as the description of

the structure of the strategies themselves. The previous discussion of the various idealised strategies demonstrates that the context in which diagnostic reasoning takes place varies considerably.

In search by recognition, the context is only implicitly defined in terms of the pool of episodic experience of the diagnostician.

In decision table search, the diagnostic context is a hierarchically ordered set of categories which can be determined inductively (typically for medicine) or deduced from a model of the functional structure of the system (typically for operating instructions for technical systems). However for an expert in technical process control as well as in patient treatment, the context will be based on subjective experience and the relevant decision tree will unfold ad-hoc during the search. That is, the potentially interesting branches together with the attributes to look for are determined by the immediate context.

In diagnosis by hypothesis and test, the diagnostic context is formed by the functional structure of the system serving as the basis for deducing symptoms from a postulated cause.

In topographic search for the location of a fault, the diagnostic context is a representation of the physical anatomy of the system.

This discussion shows that even for one particular decision task, such as diagnosis, the manifestation of the properties of a particular system in terms of the behaviour-shaping features to be considered for modelling behaviour and for system design will vary significantly - not only with the objectives brought to bear by the actors, but also with the perspectives taken and the levels of expertise of the actors. This problem is complicated even further because the strategies, their context and the perspectives taken will change dynamically during a decision task.

Diagnostic Reasoning in Work

The description of the idealised, formal strategies serves to formalise their processes, their information requirements and their cognitive loading. The diagnostic reasoning applied in actual work situations will for several reasons involve frequent shifts among these formal strategies and among the different perspectives - according to very subjective and situation-dependent factors. An example of shifting strategies is shown in figure 5, which shows the trajectory in the workspace taken by a computer maintenance engineer. He starts (1) by guessing a familiar fault from mere recognition. When proved wrong, he gets some hints from a passing colleague who reported on an odd experience with the system the day before and he (5) continues from that episodic evidence without success. Finally, he (13) enters a sequence composed of pieces of strategies from topographic and hypothetical search. In this way, each particular footprint of a strategy becomes unique and complex.




Figure 5 shows the trajectory in the workspace taken by a computer maintenance engineer. He starts (1) by guessing a familiar fault from mere recognition. When proved wrong (2), he gets some hints from a colleague's experience the day before (5) and continues from episodic evidence without success (12). Finally, he (13) enters a sequence composed of pieces of strategies from topographic and hypothetical search.

One reason to shift between strategies is because of their very different resource requirements with respect to time taken, information needs and necessary background, etc. Shifts in strategy are a very effective way to circumvent local difficulties along the decision-making path.

Another reason for shifting strategy in a particular work scenario is that the diagnostic objective can change during a session. Initially a medical doctor or a process operator will be concerned with the question of whether he is confronted with a need for a rapid compensatory action. That is, he is concerned with the potential consequences of the present state (the prophet's perspective). Next, he will be concerned with the choice of a function to stabilise (the operator's perspective). Then he will be concerned with the correction of the present abnormality (the repairman's perspective) and, finally, he may be concerned with preventing a repetition (the manager's or safety council's perspective). This means that a diagnostician can shift perspectives and strategies several times in response to the changing priorities of different objectives. He will, of course, not completely restart a diagnosis for each objective and a complex transfer of results will have to take place between the phases applying the different strategies. Figure 6 illus-

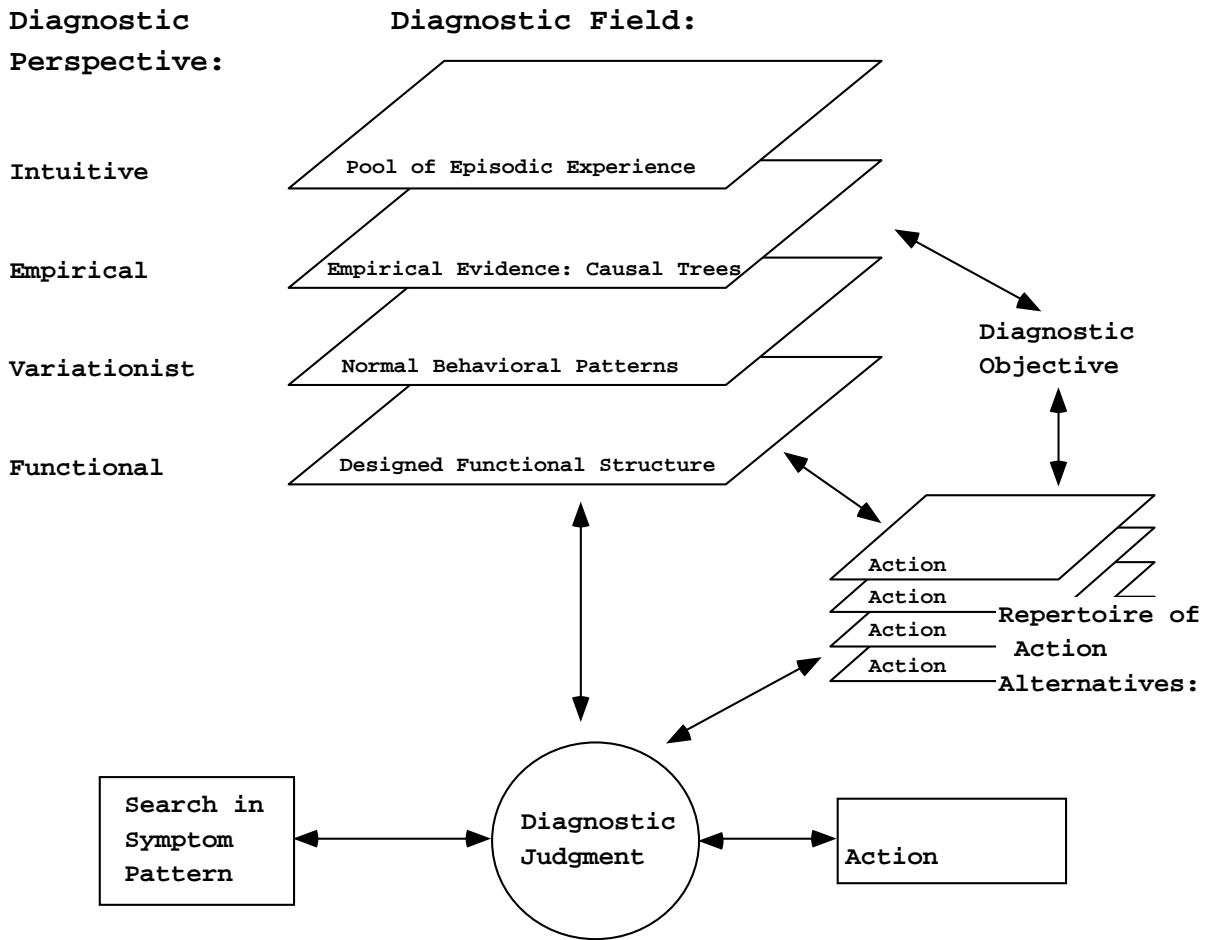


Figure 6 illustrates the elements of diagnostic reasoning during an actual case. The approach taken and the diagnostic field activated depend on the immediate circumstances, including aspects such as demand/resource conflicts, cues from prior experience, the perceived action alternatives, etc. Consequently, diagnostic fields of a very different nature will be used corresponding to the frequent shifts in the diagnostic strategy applied. A basic circularity is found between the diagnostic objective, the interpretation of the observations depending on the activated diagnostic field and the perceived repertoire of action alternatives for the particular case. The process, therefore, will have the character of zooming-in from an intuitive initial guess.

trates the complexity of a natural diagnostic process caused by the frequent shifts among strategies and objectives during a diagnostic session.

IMPLICATIONS FOR DESIGN

Several important implications about the nature of natural strategies should be mentioned here. One is that modelling and simulation of diagnostic performance in actual work should take into account a repertoire of possible strategies. A second important modelling problem is to identify the performance criteria and the (subtle) cues in the work setting which control the transitions from one strategy to another.

Another point is that an information system should support all the strategies relevant for a task in order to allow users to shed mental workload by shifting strategies. Forcing users to work through problems using only one (designer-preferred) strategy instead of enabling them to try other strategies will strongly influence their acceptance of a system.

Finally, the role of the context underlying mental strategies in order to identify the behaviour-shaping constraints and represent them in interfaces should be studied. The means-ends representation gives a description of the structure of the work system and the conceptual levels at which the functional and intentional constraints can be formulated (more or less) objectively, as seen from the system point of view. The analysis of activities which are characteristic for the work system serves (a) to select those constraints which are active at certain times and for certain work functions and (b) to transfer them to the information processing domain. The analysis of strategies which are effective in connecting states of knowledge in a decision task will reveal an abstract structure of the mental processes. In an actual case, these are guided by the behaviour-shaping constraints as represented in a given 'reasoning context', which often is embedded in a person's 'tacit knowledge'. This implies that the 'objective' behaviour-shaping constraints as defined by the work system are transformed by the interpretations of the individual actor. These transformations depend not only on the level of representation of the work system the actor selects, but also on the implicit representation of the constraints in the actor's repertoire of heuristics derived from previous encounters. That is, the perception of the work environment changes as constraints become implicit in work practice and end up being embedded in a tacit and subjective representation of the work context.

VII. CONCLUSION

A couple of important implications of this discussion of diagnosis in action should be mentioned in conclusion. One is, that several "diagnostic fields" very likely will be accessed during a particular diagnostic scenario. An expert diagnostician will work on the background of a "multi-dimensional intuition," switching between different strategies by transitions which are cued by observed evidence which presents "affordance" [5] with respect to one of the latent, familiar diagnostic contexts.

Another important observation is that modelling and simulation of diagnostic performance outside controlled laboratory environments will have to take into account all the strategies that are relevant to the actual work context. An important modelling problem will, in particular, be to identify the performance criteria and the cues in the subtle work setting which control the transitions among strategies.

Finally, it must be realised in design of decision support systems, that such systems should support all the effective idealised strategies, or else the diagnos-

tician can be severely constrained in shifts among strategies which could serve the resolution of local demand-resource conflicts.

VII. REFERENCES

- [1] Linnaeus, C. (1735): *Systema Naturae*. Edition: Engel-Ledeboer & Engel. Amsterdam.
- [2] Brehmer, B. (1981). Models of diagnostic judgements, in J. Rasmussen and W. B. Rouse, eds. *Human Detection and Diagnosis of System Failures*. New York: Plenum Press, 231-241.
- [3] Wulff, H. R., Pedersen, S. A. and Rosenberg, R. (1986): *Philosophy of Medicine: An Introduction*. Oxford: Blackwell Scientific Publications.
- [4] Mackie, J. L. (1975): "Causes and Conditions." *American Philosophical Quarterly*, Vol. 2.4 pp. 245-255 & 261-264 Reprinted in: E. Sosa (Ed.): *Causation and Conditionals*, Oxford University Press.
- [5] Gibson, J. J. (1966). *The Senses Considered as Perceptual Systems*. Boston, Ma: Houghton, Miffling
- [6] Rasmussen, J. (1986): *Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering*. North Holland, 1986.
- [7] Klein, G. A. (1989): Recognition-Primed Decisions. In Rouse W.B. (Ed.): *Advances in Man-Machine System Research*, 5, 47-92. Greenwich, CT: JAI Press.
- [8] Rasmussen, J. (1990): Deciding and Doing: Decision Making in Natural Context. In: Klein, G. and Calderwood, C. (Eds.): *Decision Making in Action: Models and Methods*. Ablex Publishing. In Press.
- [9] Russell, B. (1913): "On the Notion of Cause". *Proc. Aristotelean Society*, Vol. 13, pp. 1-25.
- [10] Rosch, E. (1977): Human Categorisation. In: N. Warren (Ed.): *Advances in Cross-Cultural Psychology*. New York: Halsted Press.
- [11] Rasmussen, J. (1990): Learning from Experience? How? Some Research Issues in Industrial Risk Management. Invited contribution to: Leplat, J. and G. de Terssac (Eds.): *Les Facteurs Humains de la Fiabilité dans les Systemes Complexes*. Marseilles: Octares. 1990.
- [12] Burnum, J. F. (1987): Medical Practice A La Mode: How Medical Fashions Determine Medical Care. *The New England Journal of Medicine*. Vol. 317, No. 19, pp. 1220-1222.
- [13] Duberman, S. M. and Bendixen, H. H. (1986): Mortality, Morbidity, And Risk Studies in Anaesthesia. In: (Ed.): *Epidemiology in Anaesthesia*. New York: Edward Arnold.
- [14] Rasmussen, J. and A. Jensen (1974): "Mental Procedures in Real-Life Tasks: A Case of Electronic Trouble Shooting", *Ergonomics*, vol. 17, 1974, pp. 293-307.
- [15] Feinberg, F. (1965): Action and Responsibility. In: M. Black(Ed.): *Philosophy in America*. Allen and Unwin. Reprinted in: A.R. White (Ed.): *The Philosophy of Action*. Oxford Univ. Press
- [16] Fitzgerald, P.J. (1961): Voluntary and Involuntary Acts. In: A.C.Guest (Ed.): *Oxford Essays in Jurisprudence*, Clarendon Press. Reprinted in: A. R. White (Ed.): *The Philosophy of Action*. Oxford Univ. Press
- [17] HMSO (1989): *Investigation into the Clapham Junction Railway Accident*. The Department of Transport. London: Her Majesty's Stationary Office, 1989.
- [18] HMSO (1987): *M V Herald of Free Enterprise*. Report of Court No. 8074, ISBN. 011 550828 7, Department of Transport, November 1987. London: Her Majesty's Stationary Office
- [19] Pejtersen, A. M. (1979): Investigation of Search Strategies Based on an Analysis of 134 User-Librarian Conversations, in T.Henriksen (ed), *Third International Research Forum in Information Research*, Oslo.
- [20] Rasmussen, J. (1990): Human Error and the Problem of Causality in Analysis of Accidents. *Phil. Trans. R. Soc. Lond.* B 327, 449-462.
- [21] Maxwell, J. C. (1867): On Governors. *Proc. R. Soc. Lond*, 16: 270-283
- [22] Leplat, J. and Ramussen, J.(1984): Analysis of Human Errors in Industrial Incidents and Accidents for Improvement of Work Safety. *Accid. Anal. and Prev.* Vol. 16, No. 2, pp.77-88
- [23] Pedersen, S. A. and Rasmussen, J. (1991): Causal And Diagnostic Reasoning in Medicine and Engineering. Contribution to Mohwac workshop, May 1991, Stresa, Italy. To be published.

