Contextual Control Model (COCOM)

The information processing approach to the modelling of cognition uses a set of elements or structures of the human information processing system as the basic building blocks for models and theories. These may also be called process genotypes. Thus elements such as short-term memory, long-term memory (knowledge base), attention capacity, etc., are common primary components. Such structural approaches are attractive because they provide an apparently objective frame of reference where human information processing is a reflection of information processing in the machine. Their disadvantage is that they refer to an information processing mechanism in near splendid isolation that is triggered by events in the external world. Cognition thereby becomes a kind of higher level information processing that occurs entirely within the human mind, and the holy grail of cognitive science is to unravel the mechanisms of pure cognition. Although information processing in some disciplines can be described as an abstract process, it makes little sense to speak of basic human information processes in the same manner. While the information processing metaphor is useful to understand some fundamental features of human thinking, it does not mean that the mind is an information processor or that cognition is computation.

The alternative to a structural approach is to describe the regularities of performance rather than the details of human information processing, i.e., on performance phenotypes. This functional approach is driven by the requisite variety of human performance rather than by hypothetical conceptual constructs (process genotypes). The observed regularities of human behaviour by definition only exist in a given context, and actions occur in anticipation of events as well as in response to them. Functional approaches avoid the problems associated with the notion of pure mental processes, and in particular do not reduce cognition as an epiphenomenon of information processing. In a structural approach it is necessary to account for the context separately from the processes of the mind; in a functional approach this problem disappears. The advantages of that should be obvious.

Human performance is determined, largely, by the situation. People can do many things and achieve their objectives in many different ways. The selection among the possible actions is not determined by normative characteristics of the action elements (as components), but by the current needs and constraints - that is, by the demand characteristics of the situation. Due to the regularity of the environment there may be frequently recurring patterns or configurations of actions, but this is not evidence for procedural prototypes. The challenge of cognitive systems engineering is to provide a reasonable account of how this regularity can occur without making too many assumptions about human cognition or about the capabilities of an internal information processing system. A contextual control model is based on three main concepts: competence, control, and constructs.

- Competence represents the set of possible actions or responses that a system can apply to a situation according to the recognised needs and demands. The extent of this set depends on the level of detail or the granularity of the analysis, and it is not necessarily denumerable. Furthermore, in terms of the model, the system cannot do something that either is not available as a possible action or which cannot be constructed or aggregated from the available possible actions.
- Control characterises the orderliness of performance and the way in which competence is applied. The COCOM deliberately simplifies the description of control to a set of control modes: scrambled, opportunistic, tactical, and strategic.

These four control modes are, however, only regions on a continuum, which ranges from no control at all to completely deterministic performance. One issue of control has to do with the conditions under which it changes from one mode to another; another has to do with the characteristic performance in a given mode – i.e., what determines how actions are chosen and carried out. Both issues are addressed by the COCOM, and define the requirements to the internal functions of the model (Hollnagel, 2000).

- Constructs refer to what the system knows or assumes about the situation in which the action takes place. The term is used to emphasise that constructs are artificial, in the sense of being constructions or re-constructions of salient aspects of the situation, and that they are usually temporary. Constructs are similar to the schemata of Neisser (1976) in the sense that they are the basis for selecting actions and interpreting information.

An essential part of control is planning what to do in the short-term, within the system's time horizon. This planning is influenced by the context, by knowledge or experience of dependencies between actions, and by expectations about how the situation is going to develop – in particular about which resources are and will be available to the person. The outcome leads to a sequence of the possible actions, which normally is constructed rather than pre-defined. Frequently occurring patterns therefore reflect the relative constancy (regularity) of the environment rather than the constraints of the performance model.

The control modes correspond to differences in the orderliness or regularity of performance, and each control mode can be associated with a characteristic type of performance. Although the control that a joint system can have over a situation may vary continuously, it is useful to make a distinction between the following four characteristic modes:

- In the scrambled control mode, the choice of next action is basically irrational or random. For humans there is little, if any, reflection or cognition involved but rather a blind trial-and-error type of performance. This is typically the case when situation assessment is deficient or paralysed and there accordingly is little or no correspondence between the situation and the actions. The scrambled control mode includes the extreme situation of zero control.
- In the opportunistic control mode, the salient features of the current context determine the next action. Planning or anticipation are limited, perhaps because the context is not clearly understood or because there is limited time available. Opportunistic control is a heuristic that is applied when the constructs are inadequate, either due to lack of competence, an unusual state of the environment, or detrimental working conditions. The resulting choice of actions is often inefficient, leading to many useless attempts being made.
- The tactical control mode corresponds to situations where performance more or less follows a known procedure or rule. The joint system's time horizon goes beyond the dominant needs of the present, but planning is of limited scope or range and the needs taken into account may sometimes be ad hoc.
- Finally, in the strategic control mode, the joint system has a wider time horizon and can look ahead at higher-level goals. The dominant features of the situation or the interface therefore have less influence on the choice of action. At the strategic level the functional dependencies between task steps and the interaction between multiple goals will also be taken into account in planning.

The scrambled control mode is clearly the least efficient, while the strategic is the most efficient – seen from an overall system perspective. In practice, humans will usually function in what corresponds to an opportunistic or tactical control mode. This represents an equilibrium condition between feedback and feedforward, which corresponds to an efficient use of the available resources. Although the strategic control mode is the optimal one, theoretically speaking, in terms of being able to control a situation, it usually requires so much effort that it cannot be sustained for longer periods of time. In the current approach to modelling, the four modes of control represent regions in a control space and should not be seen as absolute categories.

Basic Cyclical Model

In the following I will consider the use of modelling to describe the control of a vehicle as it moves from point A to point B. The description can be applied to many different ways of moving, ranging from a bicycle to an aeroplane – and even walking. The main example will be that of driving a car, since this is an activity where many people have extensive personal experience. It is also an activity that has become increasingly more difficult over the years, because of the increasing traffic and because of the greater performance capacity of a modern car. Another trend is that more and more technology is put at the disposal of the driver, ostensibly to make driving easier, but sometimes with unanticipated adverse effects. A modern car includes several computers that control not only the performance of the engine, but also how the car performs on the road surface. In addition, other types of information technology are being applied in increasing degree, from infotainment to the movable office.

The basic approach to modelling the control of a vehicle follows the structure of the cyclical model shown in Figure 1. In accordance with the principles of CSE, the objective is to model performance rather than human information processing (thus, cognition according to the pragmatic rather than the axiomatic definition). The focus is on how the joint system acts to achieve its goals while at the same time responding to events in the environment. The cyclical model is based on the principles of the perceptual cycle (Neisser, 1976), but extended from describing perception to describing action and control. The principles of the cyclical model were used as the basis for the Contextual Control Model (COCOM), which describes how the orderliness of performance depends on the level of control, and which provides further details about the selection of actions and the evaluation of events (Hollnagel, 2000).

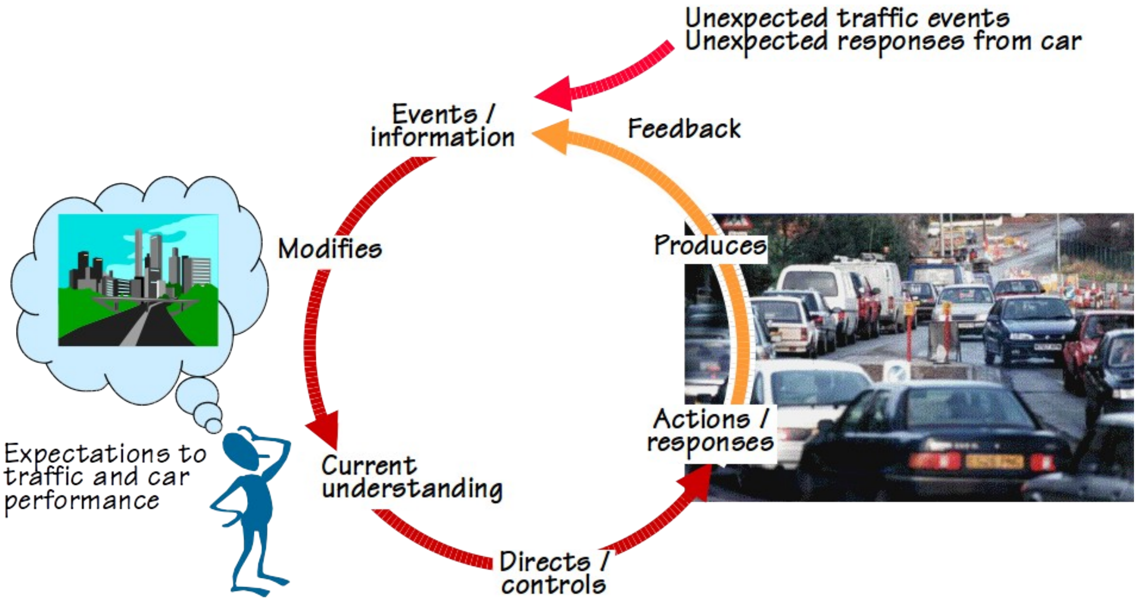

Figure 1: The basic cyclical model of human action.

In contrast to the information processing type of models, the cyclical model describes the performance of the joint cognitive system as a mixture of feedback (error controlled) and feedforward (cause controlled) activities. It provides a way of capturing the dynamic relationships between situation understanding (as constructs), actions (as realised competence), and feedback or information (as events). On a general level the model shows how actions depend on the current understanding (construct), which in turn depends on the feedback and information (events) received by the system, which in its turn depend on the actions that were carried out, thereby closing the circle. The performance of the JCS is reactive in the sense that it is affected by the feedback and the information about the events that take place, and proactive in the sense that actions are determined by the current understanding and therefore also by the expectations of what may happen in the future. On a more detailed level the cyclical model can be used to represent the temporal constraints of dynamic actions, and be extended from individual to co-operative actions (Hollnagel, 1998).

Time And Control

Time refers to the fact that we are dealing with systems and processes that develop and change. This means two things: first, that there is limited time available to evaluate events, to plan what to do, and to do it. Second, that the information that is used needs to be updated and verified regularly because the world is changing. This is one reason why it is unrealistic to describe decision making as a step-by-step process unless the decision steps are minuscule relative to the speed of the process.

Time has generally been treated as a Cinderella in both HCI and HMI (Decortis & De Keyser, 1988). This is also the case for cognitive engineering and cognitive science, despite the obvious importance of time in actual work, i.e., in activities that go on outside the controlled confines of the laboratory. The reason for that is probably the legacy from behaviourism, carried on by human information processing psychology, which focused on how an organism responded to a stimulus or event, rather than on how an organism or system behaved over time.

While it has been known since the days of Donders (1969, org. 1868-69) that mental processes have duration, hence take time, the speed of actions is more important than the speed of mental processes. In other words, the interesting phenomenon is the time it takes to do something, such as recognising a situation or decide about what to do, rather than the time of the component mental processes. One simple reason is that it cannot be assumed that the duration of an action – to the extent that one can talk about this as a meaningful unit at all – can be derived by considering the duration of the elementary or component processes. Even if the internal workings of the mind were sequential, in a step-by-step fashion, the combination or aggregation need not be linear. Human action is furthermore not the execution of a single sequence of steps, but rather a set of concurrent activities that address goals or objectives with different time frames and changing priorities. For example, in order to take decisions a process plant operator needs to be able to reason about temporal information, to reason about changes, to predict the effects of his actions and of the changes he observes, to continuously make reference to what has happened, is happening and might possibly happen, and to coordinate on the temporal axis the actions of several users.

Model Representation Of Time

Effective control requires that the user – and more generally, the controlling system – is able to interpret the events and able to find and choose effective action alternatives. In the cyclical model of human action that is at the heart of CSE (cf. Figure 2) the two arcs called "event evaluation" and "action selection" represent these activities. Associated to the former is the time needed to evaluate events (TE) while associated to the latter is the time needed to select action alternatives (TS). The time needed to accomplish both of these must be seen in relation to the time that is available (TA), as well as the time estimated – and needed – to perform or carry out an action (TP).

In most industrial domains, tasks are force-paced – or process-paced – rather than self-paced. The available time, TA, is limited by the speed of the process and if (TE + TS) exceeds TA this puts severe constraints on the users' possibility to evaluate events and select actions. Some processes such as steel rolling mills, electronic trading, or flying an airplane require rapid or even near instantaneous responses. Other processes such as power generation, land-based transportation or surgery pose less severe demands but still require that actions be taken within a limited time. It would clearly be better if there was ample time, i.e., if (TE + TS) was less than TA, since the user then would have time to refine the current understanding, to plan before acting, hence to be in control of the situation. This can be achieved if the time limitations can be relaxed, for instance by slowing down the process. This is unfortunately only possible in very few cases. A more common approach is to reduce either TE or TS by improving the system and interface design, although this usually has been done in a piecemeal fashion.

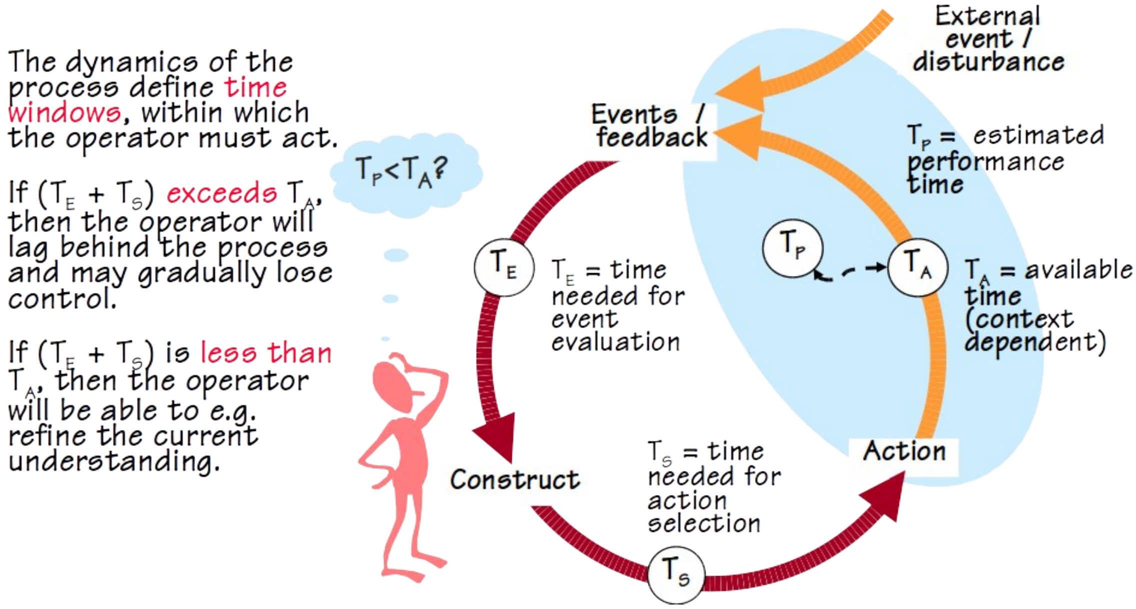

Figure 2: Time and control in the cyclical model.

The predominant position of event evaluation and action selection in the contextual control model provides an easy way of accounting for the coupling between time and control. The two functions can be carried out to various depths depending on the control mode. This corresponds to the basic fact that humans, unlike machines, may be more or less thorough in what they do depending, among other things, on the available time. This provides a single way of expressing of the time-control dependency conceptually as well as a simple way of implementing it computationally.

How To Enhance Control

Using the contextual control model as a basis, it is possible to consider in some detail how control may be facilitated. In the following, this is done by considering each the time parameter by itself.

Time To Evaluate Events

The time needed to evaluate events, (TE), can be reduced primarily by improving the design of information presentation. This has for many years been one of the dominating concerns in the field of human-machine interaction, and a number of high-level design principles have been put forward, including adaptive displays, multimedia interfaces, and ecological interface design. The lofty goal has been to provide the user with “the right information, in the right format, at the right time”, although this in most cases easier to say than to do.

Another common solution is suppression of information, either by filtering information according to logical principles or by using sophisticated graphical information presentation, such as adaptive grouping or task-oriented displays. More sophisticated proposals involve computer supported interpretation and diagnosis, expert systems, and artificial intelligence in various forms. Non-technological solutions to the problem of reducing TE primarily involve education, training, and better use of human resources such as in Crew Resource Management techniques.

Time To Select Actions

The primary means to reduce the time needed to select an action, TS, is by means of procedures and extensive training. Procedures in effect encapsulate prepared decisions, which should relieve users of having to go through the reasoning behind the procedure. Another approach is to improve the human-machine interface, not so much in terms of information presentation but more in terms of the design of controls surfaces and panels, whether hard or soft. Both efforts must obviously be complemented by training, which in itself can improve the ability to estimate the time needed to carry out an action, TP, as mentioned above. More sophisticated solutions involve various forms of decision support systems and computerised predictions. Selecting an action logically involves a prediction of what the outcome of the action will be, and this can in many ways be supported by appropriate system and interface design, as well as by training and experience.

Time To Perform An Action

The time needed to carry out a chosen action, TP, can be reduced by using automation to amplify human performance. The use of automation is in many ways a mixed blessing, since the speed and precision of control actions may be offset by reduced observability and an increase in the number of unexpected events – leading to the dreaded automation surprises and new problems of when to intervene (Dekker & Woods, 1999). Invoking additional resources, such as extra staff, may also help to reduce the performance time although there is an obvious trade-off involved. On the organisational scale rapid deployment forces exemplify this, although such solutions carry a considerable cost.

Available Time

Finally, it may even under some conditions be possible to increase the time available (TA). Although the speed of the process usually is beyond human control or only can be changed within very narrow limits, several other solutions are possible. In many processes safety systems may be activated to buy the users additional time, something, which in the nuclear field is known as the n-minute rule. For organisational processes another solution is to renegotiate deadlines, for instance in software engineering or large construction projects. This requires, however, that the whole development can be slowed down without irreparable damage.

From COCOM To ECOM

Although the basic version of the model is shown with only one loop (which more properly should be a helix), it is acknowledged that actions take place on several levels at the same time. A more neutral expression is that concurrent activities can be described as corresponding to goals at different levels. This notion of multiple levels of activity is different from the distinction between different levels of internal processing (such as skill-based, rule-based, knowledge-based) that is proposed by some information processing models. Most importantly, the different levels of activity refer to different levels of performance rather than to different levels of information processing, hence to a property of the joint system rather than of the user's internalised cognition. The information processing models furthermore assume (1) that the activity types are disjunctive, and (2) that for each type there is only one level of control. The first assumption means that the different types of activity occur only one at a time or interchangeably. The second assumption means that each type of activity is always carried out in the same way, i.e., with the same degree of control.

References

Decortis, F. & De Keyser, V. (1988). Time: The Cinderella of man-machine interaction. IFAC/IFIP/IEA/IFORS Conference on Man-Machine Systems, 14 - 16 June 1988, Oulu, Finland.

Dekker, S. W. A. & Woods, D. D. (1999). To intervene or not to intervene: The dilemma of management by exception. Cognition, Technology & Work, 1(2), 86-96.

Donders, F. C. (1969). On the speed of mental processes. Acta Psychologica, 30, 412-431. (Translated from Over de snelheid van psychische processen. Onderzoekingen gedaan in het Physiologisch Laboratorium der Utrechtsche Hoogeschool, 1868-1869, Tweede reeks, II, 92-120.)

Hollnagel, E. (1998). Context, cognition, and control. In Y. Waern, (Ed.). Co-operation in process management - Cognition and information technology. London: Taylor & Francis.

Hollnagel, E. (2000). Modelling the orderliness of human action. In: N. B. Sarter & R. Amalberti (Eds.), Cognitive engineering in the aviation domain. Mahwah, NJ: Lawrence Erlbaum Associates.

Neisser, U. (1976). Cognition and reality. San Francisco: W. H. Freeman.