Taking Decisions into the Wild: An AI Perspective in the Design of iDMSS

HAMID R. EKBIA

Computer Science Dept.

and

The Redlands Institute

University of Redlands

e mail: hamid_ekbia@redlands.edu

Introduction

Decision Support Systems (DSS) have evolved significantly during the last four decades. However, their capabilities are still very limited. Elgarah et al. (2002), for instance, writing on a project to develop a DSS for urban infrastructure decision making for the city of Houston, report that they "know of no DSS design methodology suitable for use in such a complex, conflict-filled situation as this." This is an alarming observation in the face of decades of research and practice on DSS. This article seeks to suggest remedies to this situation.

The major reason for the current shortcomings of DSS, I argue, has to do with the techno-centric nature of the development of these systems. An overview of the history of DSS reveals that its development has been largely driven by changes and innovations in computer technologies such as data and knowledge bases, expert systems, software agents, and more recently web-based tools. This reveals the dominance of a techno-centric view in DSS development that is also manifested in its relationship to Artificial Intelligence (AI). Both of these areas were originally influenced by the ideas of Herbert Simon — DSS through Simon (1960)'s seminal work in management science and AI through his work, with Allen Newell (1961, 1976), on human problem-solving and on the Physical Symbol Systems' Hypothesis (PSSH). However, the two areas have followed rather different development paths for a good part of their history, as Simon himself observed many years ago (1987).

As it turns out, the AI community, having discovered the limits and flaws of PSSH, has made serious departures from its underlying premises and assumptions, while the DSS community has remained largely committed to the traditional concept of decision making formulated by Simon — or, at least, it has not seriously questioned some of the basic premises and assumptions of his views. Given that the DSS community usefully understands AI as "a reference discipline for DSS research" (Goul *et al. 1992)*, it might be worthwhile to revisit the relationship between the two fields once again. This is the main thrust of this chapter. Unlike previous similar attempts, however, we will go about this in a top-down fashion. That is, rather than starting with AI tools and technologies that might prove useful for DSS, we start with a study of conceptual developments in

AI, examine their implications for DSS, develop a new conceptual framework, and then arrive at a design methodology that would support the framework.

To this end, we trace the development of the two fields from their common origin to the present time. In particular, we draw upon recent developments in AI to propose an embedded, action-oriented, and improvisational approach to the design of intelligent decision-making support system (iDMSS). The methodology that emerges, best characterized as *decision-making in the wild*, is in line with recent developments in DSS and in software development — e.g., active DSS (Shim *et al.* 2002), the *WinWin Negotiation Model* (Boehm *et al.* 1995) — and with some of the major empirical findings in the history of the field itself — e.g., the superiority of prototyping and evolution-based methodologies (Alter 1978, Alavi 1984, Mahmood and Medewitz 1985).

The chapter continues in the next section with a brief historical outline of the development of AI and DSS, focusing mainly on the ideas of Simon. This follows with a discussion of recent developments in AI, particularly in the area of situated cognition. Next, we examine the implications of these developments for the conceptualization and design of iDMSS. This results in a conceptual framework that we then apply to outline a new methodology.

Background: The Common Origins of AI and DSS

The emergence of AI as a reference discipline for DSS has been previously explored (Goul et al. 1992, Eom 1998). Rather than "emergence," however, the relationship between the two areas might be better described as *co-origination.* For not only did the two disciplines emerge at about the same time (late 1950's), they both have a common beginning in the works of Herbert Simon. Simon is deservedly known as a pioneer of AI and as the founder of decision science as we know it today. His book *The new science of management decision*, which brings these two major strands of his thinking and research together, provides a good context for the joint understanding of Simon's views and of the history of these two areas. CAPS title case

In the opening chapter of his book, which deals with the impact of computer technology on the processes of management, Simon discerns three dimensions of disagreement among experts on the degree of this impact: a technological dimension, a philosophical dimension, and a socioeconomic dimension. Roughly speaking, expert opinions vary depending on how small or big they envision the impact. Aligning the first two dimensions, Simon recognizes four possible schools of thought in this respect (see Table 1), and characterizes his own position as "fairly extreme along all dimensions" — namely as a technological radical, an economic conservative, and a philosophic pragmatist (1977: p. 6):

I believe that in our times computers will be able to perform any cognitive task that a person can perform. I believe that computers already can read, think, learn, create. I believe that computers and automation will contribute to a continuing, but not greatly accelerated, rise in productivity, that full employment will be maintained in the face of that rise, and that mankind will not find life of production and consumption in a more automated world greatly different from what it has been in the past. (pp. 6–7).

Almost forty years after these prophecies, it is indeed sobering to see how many of them have not yet materialized (e.g., full employment and disappearance of poverty and deprivation that he predicts later in his book), and how many others are still debated (e.g., rise in productivity, employee empowerment, etc.; cf. Kling 1996). Of particular relevance to the present discussion, however, is Simon's radical view about the capabilities of computers, especially as they relate to organizational and management processes. To understand the source of this radicalism, we need to study Simon's views of human cognition and of organizational decision-making.

Socioeconomic

Philosophical
Technological

	conservative	*radical*
conservative	Computers are limited in power, and business is done as usual	Computers are limited in power, but there will be plentitude of goods and services
radical	Computers equal human in terms of capabilities, but business is done as usual	Computers equal humans in terms of capabilities and will replace humans

Table 1. Simon's classification of views on the impact of computers on management

Cognition as Means-Ends Analysis

As mentioned earlier, Simon's view of human cognition is best represented in his work with Newell on human problem-solving. The main idea behind this hypothesis is that, "human thinking is governed by programs that organize myriads of simple information processes — or symbol manipulating processes if you like — into orderly, complex sequences that are responsive to and adaptive to the task environment and the clues that are extracted from that environment as the sequences unfold" (Simon 1977, p. 68). This mentalistic view of human cognition is based on a number of key assumptions in what is now called *classical AI*. As Agre (1997) points out, mentalism's simple answer to all questions of psychology is: put it in the head. "If agents need to think about the world, put analogs of the world in the head. If agents need to act in situations, put data structures called 'situations' in the heads. If agents need to figure out what might happen, put simulations of the world in the head' (p. 51). In short, the basic method of mentalism is to reproduce the entire world inside the head.

Newell and Simon's idea is an exemplar of the mentalistic view that emphasized problem-solving in the abstract. In their work on the computer program *GPS* (General Problem Solver), Newell and Simon (1972) referred to *means-ends analysis* as the key component of human thought process — that is, the analysis of the difference between what we need and what we have. In their studies of human problem-solving under laboratory conditions, they used "thinking-aloud protocols" in order to tap into their subjects' internal thought processes — a methodology that turned into a pillar of knowledge engineering for years to come. Simon argued that, "Problem solving proceeds by erecting goals, detecting differences between present situation and goal, finding in memory or by search some tools or processes that are relevant to reducing differences of these particular kinds, and applying these tools or processes" (1977, p. 70). He also emphasized that these tools (or heuristics) are "subject-matter free," in the sense that they apply to any problem that can be cast into an appropriate general form.

In short, Simon viewed cognition as a heuristically guided search activity within an abstract mental space. As we shall see next, it was a short step from this view to the idea that a good part of management decision-making can in principle be performed by computers and to a putative methodology that would implement it.

The Automation of Management

> *Within the near future — perhaps in the next generation — we shall have the technical capability of substituting machines for any and all human functions in organizations.* (Simon 1977: 16)

On the basis of the above premises, Simon envisioned the *automation of management* in the spirit of what had previously happened in factory automation. He assessed this development as a "technological revolution of the decision-making process" (1977, p. 31). To motivate this, Simon classified organizational decisions into two major categories: "programmed" (structured) and "unprogrammed" (unstructured). The first category, which refers to the routine and cut-and-dried decisions of heads and middle managers, belongs in the realm of operations research and its formal and mathematical techniques such as linear programming. The second category, which consists of the basic, once-for-all, and unusually consequential decisions of presidents and top managers, falls into the domain of artificial intelligence or "heuristic programming" (as Simon sometimes preferred to call the new discipline of AI). In the same manner that humans reach "unprogrammed decisions by reducing them to a series of programmed decisions," according to Simon, executives can similarly do this by following a number of well-defined, but interwoven, phases.

This conceptualization of organizational decision-making gave rise to Simon's well-known four-phase model of decision-making — intelligence, design, choice, and review — which dominated most design methodologies in DSS for years to come (Gorry and Morton 1971). Simon describes the main focus of each of these phases as follows:

1. *Intelligence*: Survey the economic, technical, political, and social environment to identify new conditions;
2. *Design*: Invent, design, and develop possible courses of action for handling situations
3. *Choice*: Choose among alternative actions already developed to meet an identified problem and already analyzed in terms of their consequences
4. *Review*: Assess the outcomes of past actions as part of a repeating cycle that leads again to new decisions

I shall demonstrate later how *in practice* decision-making processes differ from this conceptualization. But even a superficial examination of current organizational processes would illustrate that how far indeed we are from "substituting machines for any and all human functions in organizations." This gap between Simon's predictions and current reality might be indicative of serious flaws in his original conceptualization of thought processes in humans and of decision-making processes in organizations. The task of discovering those flaws falls onto the shoulders of the AI and DSS communities, respectively.

The Divergent Paths of AI and DSS

Despite all the changes and accomplishments in DSS tools and techniques, the conceptual framework laid out by Simon seems to have largely remained intact, dominating design methodologies in DSS throughout decades. At least, to best of my knowledge, no one in the DSS community has explicitly and systematically questioned the premises and assumptions behind that framework. AI, in the meantime, has gone through serious challenges and changes that greatly depart from the classical notion of cognition as heuristic problem-solving. This has created a conceptual chasm between the two disciplines that researchers have tried to bridge every once in a while (Goul et al. 1992, Eom 1998). Nevertheless, contacts between AI and DSS has largely remained at the level of tools and techniques, partly because of the techno-centric character of development in DSS (in contrast to scientific and philosophical aspirations of AI; cf. Ekbia forthcoming(a)). This, of course, does not imply that AI has followed a steady and smooth development path. In fact, as different authors have argued, the development of AI has been tumultuous, strained, and stifled with false starts (Bloomfield 1987, Crevier 1993, Agre 1997).

The Development of AI: From Classical to Situated View

The development history of AI can be roughly divided into four periods dominated by four major approaches, which I am going to call: i) the classical approach, ii) the knowledge-based approach, iii) the connectionist approach, and iii) the situated approach. Although there are overlaps among these, each one is differentiated by basic features and premises (see Table 2).

Table 2. Major approaches to AI and their key idea

Approach	*Main idea*
Classical	Symbol manipulation
Knowledge-based	Knowledge
Case-based	Reminiscence
Connectionist	Distributed computation
Situated	Embodiment and embeddedness

The Classical Approach

Classical AI viewed cognition as *abstract* (physical embodiment is irrelevant), *individual* (the solitary mind is the essential locus of intelligence), *rational* (reasoning is paradigmatic of intelligence), and *detached* (thinking is separated from perception and action) (Smith 1999). Early AI, as exemplified by PSSH, made long strides by relying on these principles. The systems built upon these were relatively successful in abstract problem solving, but they miserably failed in dealing with more concrete tasks and domains that seemed mundane and trivial for human beings (at least at the beginning) — e.g., recognition of letters of alphabet, translation between languages, navigation in non-idealized terrains, and so on.

The failure of the classical approach to tackle these issues and to deliver its promises resulted in a decline of interest in AI research on the part of funding agencies in early 1980's, leading practitioners to look for practical problems to solve.

The Knowledge-based Approach

As knowledge was conceived to be the key to such endeavor, a new class of artifacts ("expert systems") and a new group of practitioners ("knowledge engineers") appeared on the scene. The task of a knowledge engineer is threefold (Hayes 1990: 201):

> To elicit from an expert — i.e., "a human being whose head contains knowledge" — the knowledge (they intuit) they deploy in dealing with their area of expertise;
> To formalize this knowledge, using a suitable representational language, in a knowledge base;

To compare the performance of this formalization with that of the expert for the purpose of "correcting" machine behavior — i.e., bringing it closer to the expert's introspective account.

The early success of expert systems bestowed upon AI the respect it was longing for, not only in academia, but also in business, where, according to some accounts, billions of dollars were invested in expert systems for manufacturing, financial services, machinery diagnosis, and so on. But this success was very limited because of the fragility of these systems. Despite apparent sophistication and expertise in specialized areas like medicine, expert systems demonstrated a clear lack of understanding of very basic facts that a human being takes for granted.

Attempts to rectify this situation have largely failed to this date, as best exemplified by the Cyc project (Lenat and Guha 1990). Cyc was motivated by the idea that a vast amount of knowledge is the key to intelligence. What we need to achieve human-level intelligence, the creators of Cyc believed, is to provide a machine with enough *commonsense knowledge* for it to be able to continue the process of knowledge acquisition on its own. Therefore, they embarked on the creation of a huge knowledge base that after a certain point was meant to learn on its own and go "beyond the frontiers of human knowledge." This turned out to be an elusive objective due to, among other things, the vast, implicit, and contextual character of a great deal of human knowledge (Smith 1991). Organizations face similar issues in dealing with knowledge (Blackler 1995).

The knowledge-based approach was concomitant with two other views. One is *case-based reasoning,* which emphasizes reminiscence rather than knowledge (Schank 1982). The other is the so-called *planning* view, which considers human activity as fundamentally planned and well thought out. Like the knowledge-based approach, both of these views have faced insurmountable obstacles (Suchman 1987, Hofstadter 1995).

The Connectionist Approach

The connectionist approach (re)emerged in the 1980's and captured the imagination of many researchers, including some in DSS. The main feature of this approach was its opposition to explicit forms of knowledge and its emphasis on brain-like architectures, but it remained committed to most of the principles of classical AI. In particular, it retained the notion of mental representation that was central to classical theories of mind. Therefore, despite the initial fervor, connectionism could not distract AI from some of its fundamental premises (Clark 1997, Clancey 1997). This was left to later developments, especially those associated with the *situated* view of cognition.

The Situated Approach

The situated view reverses many of the assumptions of the previous approaches, especially those of classical AI. That is, it considers intelligence to be *embodied* (physical embodiment is important), *embedded* (the immediate natural and social environment matters), *action-oriented* and largely *improvisational* (Agre 1997, Clancey 1997, Clark 1997, Smith 1999).

These conceptual shifts in AI approaches (which we shall discuss in more detail shortly) are consequential not only for our understanding of human intelligence but also, I argue, for the design of artifacts such as computerized decision support systems. To be able to articulate these consequences, we also need to follow the development of DSS since its co-origination with AI.

The Development of DSS: From Data to Knowledge

As mentioned at the outset, DSS have significantly evolved during the last few decades. In fact, the DSS community (being commendably self-reflexive) has both summed up and anticipated the development of the field at various junctures. Accordingly, the definition of DSS has changed from support technologies in semistructurred domains (Keen and Scott Morton 1978) through interactive data models (Sprague and Carlson 1982) and group decision support systems (DeSanctis and Gallupe 1987) to adaptable and domain-specific representational models (Turban 1992) and social DSS (Turoff et al. 2002). More recently, Shim et al. (2002) have discerned a trend toward the personalization of DSS user interface, the use of Web-based technologies, and ubiquitous computing. These authors have also prescribed the development of active and intelligent systems as a promising path for the future.

A careful examination of DSS development reveals a close parallel between the conceptualizations of DSS and the development of computer technologies and tools (Shim *et al.* 2002). In the era of data processing and management information systems (MIS), for example, the emphasis in DSS was on databases and data models. Later on, with the advent of expert systems and executive information systems (EIS), the scope of DSS extended to group and corporate levels. Then, the growing interest in knowledge brought about the notion of organizational learning and knowledge management. Most recently, the expansion of the World Wide Web and wireless technologies is giving rise to web-based DSS and to new conceptualizations of decision making from multiple perspectives. As Table 3 illustrates, this development has been dominantly bottom-up and technology-driven, with the available technical tools giving rise to new conceptualizations.

Stage	*Approximate Period*	*Dominant Concept of DSS*	*Technologies*
I	1960's–1970's	Data modeling and problem solving	Databases, MIS
II	1980's	Collaborative and Group Decision Support (GSS)	Knowledge bases, expert systems, EIS
III	1990's	Organizational learning and Knowledge Management	OLAP, data warehouse, data mining
IV	2000's	Web-based and active DSS	Internet, client-server tools, software agents

Table 3: The development of DSS in relationship to computer technology

A comparison of Tables 2 and 3 illustrates that AI and DSS have followed rather independent development paths, but they have converged at certain points in terms of concepts, methods, and techniques (see Figure 1). The last such point was coincident with the trend toward knowledge-based DSS. Goul *et al.* captured this point by anticipating and articulating the above trend in their proposition, which emphasized supporting human decision-making "by selectively incorporating machine-based expertise in order to deliver the potential of DSS in the knowledge era" (1992, p. 1268). As we saw, AI has moved beyond the so-called knowledge era to situated interaction, and in order for the two disciplines to maintain their relationship we need to align them once again. The remainder of this paper is an attempt in this direction.

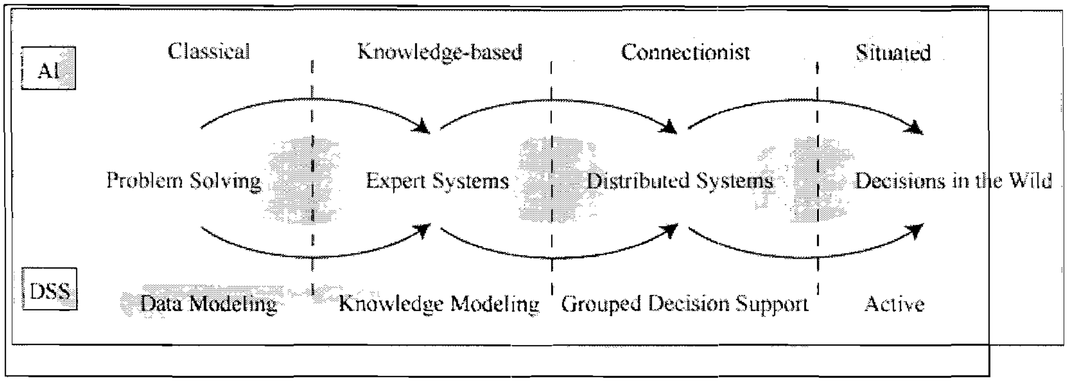

Figure 1. The intermittently convergent paths of AI and DSS development

A Situated Approach to Design

To go about our goal of alignment, we can follow one of the two strategies (see Figure 2):

a) *Bottom-up*: look for the most recent tools and techniques in AI that might be useful for DSS, and then find the methods and concepts that would emerge from them;

b) *Top-down*: bridge the conceptual gap between AI and DSS by providing a new conceptualization of DSS, and then find suitable methods and techniques that would support the new conceptualization.

The bottom-up approach has been the strategy of choice in most previous attempts, with certain advantages and disadvantages accrued to it. We follow the top-down strategy here, hoping to demonstrate its advantages throughout the following discussion. In line with the main theme of the current volume, henceforth our discussion will focus on intelligent decision-making support systems (iDMSS).

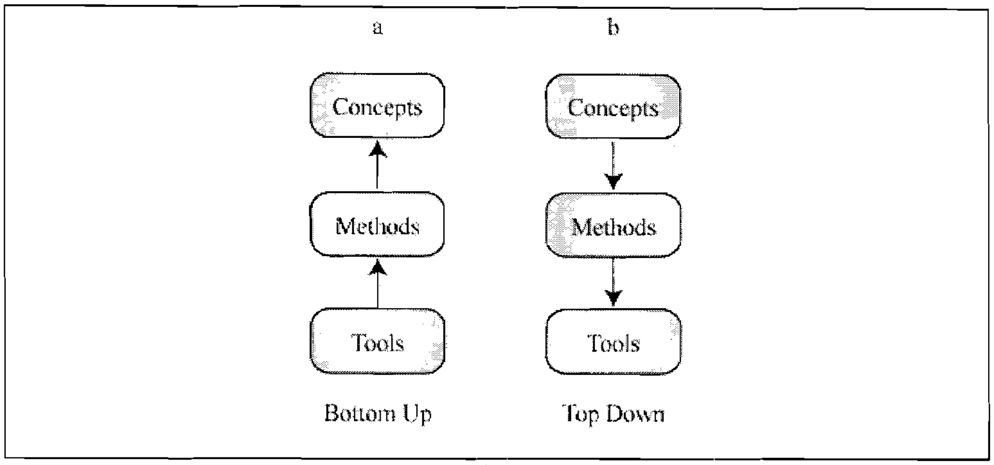

Figure 2: Two alignment strategies

For decision support systems to be intelligent, they should be incorporated in their human environment in as seamless a manner as possible. For this purpose, they should be built on the same principles on which human cognition operates. In other words, the design of these systems should follow, to the extent possible, the model of human cognition. Situated AI, we may recall, is based on four key principles — namely, that cognition is: i) embodied, ii) embedded, iii) action-oriented, and iv) improvisational. The first principle, although very important in the design of AI systems such as robots, is not directly relevant to iDMSS design, so we do not pursue it here. Let us therefore discuss the other three principles in terms of their significance for the design of iDMSS.

Embedded Design

The idea of embedded cognition is that meaning can arise only through an intelligent system's direct interaction with the world, where interaction is broadly understood as perception and action. The classical view marginalized perception and action as secondary issues, on the one hand, and, on the other, considered them as separate processes mediated by a "brain" or central processor. The situated view, to the contrary, takes perception and action as the centerpiece, and pushes central processing and internal representations of the world to the margin. Furthermore, it considers perception and action as tightly intertwined.

This shift in our understanding of perception and action has important implications for the design of iDMSS. A key tenet of the problem-solving paradigm in DSS was its emphasis on mental representations of external situations (Simon's "intelligence" phase). According to this view, people deal with the outside situations by building (more or less) faithful models of them in their "heads." Therefore, all thinking (e.g., decision making) consists mainly of the manipulation of these internal models and symbols. Problems are in our heads, as are solutions to problems. According to the situated view, however, problems are not so much in the head as they are in external situations. In other words, what we often have to deal with are "problematic situations," not problems as mental models of those situations. A problematic situation is one that is "disturbed, ambiguous, confused, full of conflicting tendencies, obscure, etc." (Dewey 1938). This means that problems do not present themselves as given, rather "they must be constructed from the materials of problematic situations which are puzzling, troubling, and uncertain" (Weick 1995: 9).

An embedded design methodology (not to be mistaken with "embedded computing") would therefore start with problematic situations rather than with mental problems. It would seek to create a balance between the local, situated knowledge of participants and stakeholders, on the one hand, and global facts and procedures, on the other (Walsham 2001: 108–130). The failure to maintain this balance might generate unpredictable side effects in an organization or in the whole society. Walsham (2001) reports a case study of a DSS implemented for the corporate lending process of a large UK bank, where the system was intended to provide lenders with an analysis of a borrower's capacity. The system calculated the probability of the loan defaulting on the basis of similar previous cases in its knowledge base. The parameters of these calculations were weighted to reflect data gathered by knowledge engineers through elicitation of "best practices" from selected "good" managers. In discussing the impacts of the DSS use, Walsham reports that the users "were expected to become dispassionate loan workers, …where loan judgments were seen as adhering to global standards rather than local contingencies" (p. 128). In light of this, Walsham finds it difficult to make a definite assessment of the DSS either in terms of "efficiency, effectiveness and profitability" of the bank or in terms of their profound societal effects — for example, the change in the profile of companies (such as small businesses) that are affected by the system. Such intricacies in terms of the interaction of local and global knowledge are involved in almost any substantial DSS, and are therefore taken seriously by an embedded design approach.

Action-Oriented Design

Closely related to the above discussion is the issue of human interaction with the world. The classical view of AI considered this as a linear process where the outside world shapes our thoughts, which we then turn into plans that in turn influence the outside world. This is the essence of Simon's four-part scheme of decision-making. However, thoughts and actions are much more intertwined than this linear picture portrays. People do not face a situation as a given, rather they enact and produce the situations of which they are a part. As Garfinkel describes them, "*in the course of a career of actions*, [people] discover the nature of the situations in which they are acting... [T]he actor's own actions are first order determinants of the sense that situations have, in which, literally speaking, actors *find* themselves" (1967: 115). In other words, the traditional view of decision-making as one of weighing a set of given alternatives (Simon's "choice" phase) might actually be putting the cart before the horse. For example, Garfinkel found out that jurors do not first evaluate the harm, then allocate blame, and finally choose a remedy. Rather, they first decide a remedy and then settle on the "facts" from among alternatives that justified the remedy. In short, they *retrospectively* justify a decision on grounds other than (or beyond) facts.

These observations are particularly important in today's uncertain, rapidly changing, and information-laden environments of decision-making. The magnitude as well as the inherent uncertainty of data available today turns their full assimilation by decision-makers into an unrealistic expectation. Under these circumstances, the whole notion of *planning* (Simon's "design" phase) as a centralized and guided process might be an ungraspable chimera. What people need in most situations is less, not more, information; and what iDMSS needs to provide first and foremost is the facilitation of action not the accumulation of knowledge. This is in line with Shim *et al.* (2002)'s prescription, originally conceived by Keen (1987, p. 121), for active DSS and their emphasis on "screening, sifting, and filtering" of data.

Improvisational Design

This brings us to the third and final aspect of the situated approach to design. In contrasting the planning and situated views of AI, Agre (1997, p. 7) argues that, human activity is "fundamentally improvised; ...People conduct their activity by continually redeciding what to do." Given that life is almost routine too, it seems that *human activity is improvised and routine at once*, with improvisations relying on routines. In other words, rather than repeatedly performing the same actions according to preset plans, we utilize our relatively stable relationships with the environment — e.g., in the way we organize furniture, our paper documents, or our computer desktops — as a backdrop for our moment-to-moment interactions with the world.

Brought into the realm of design, this notion of improvisation shifts our attention away from products to the *processes* of technical work (Agre 1997, p. 15). Central to this view are "a willingness to forego planning in favor of acting, an openness to re-assembly of and departures from routines, and a preference for process rather than structure" (Walsham 2001: 51; cf. Weick 1998). Unlike the rigid norms of classical life-cycle models of software engineering, which were often dictated by bureaucratic imperatives, this approach prescribes a more reflective attitude toward software design.

In Search of a Design Methodology

Having laid out the key conceptual principles of situated design, we now need to explore appropriate methods that would support these concepts. It has long been known that design methods have a direct impact on the success of DSS (Mahmood and Medewitz 1985). Based on the above principles, I propose a five-step method consisting of problem setting, bricolage, coordination, narration, and simulation, a brief discussion of which follows.

Problem Setting

"When we set a problem, we select what we will treat as the "things" of the situation, we set the boundaries of our attention to it, and we impose upon it a coherence which allows us to say what is wrong and in what directions the situation needs to be changed. Problem setting is a process in which, interactively, we name the things to which we will attend and frame the context in which we will attend to them" (Schön 1983: 40).

The first step is problem setting, the purpose of which is to turn a problematic situation into a problem. This is an ongoing, iterative, and reversible process during which the actors try to achieve a common understanding of the situation. As Schön describes above, this mainly consists of naming the objects (or even the situation as a whole) and setting the boundary of attention. The product of this step would be a "laundry list" of main objects and a schematic outline of the major issues, constraints, and concerns.

Bricolage or Assembly

Once the objects are named, the next step would be to assemble them in an improvisatory manner. This consists not only of expert knowledge and combinations of already existing pieces of technology — hardware, software and facilities — but also of "of appropriate work practices, skills, training, communications ..." (Büscher and Morgensen 1997, p. 79). Since part of this cannot be formalized (in the sense of encoding and storing in a knowledge base), various representation schemes should be used to capture and preserve the components — e.g., database tables, knowledge base (KB) rules, frames (schemas), hypertext, images, photos, etc. The participatory character of this phase is also very important.

Coordination

Following the initial assembly of bricolage, a process of filtration and refinement is needed to eliminate or reduce redundancies, mismatches, and conflicts. An explicit attempt should also be made at this stage to transform, to the extent possible, informal representations to formal ones. Ideally, the goal of this stage is to develop an agreed-upon *model*, although in practice this might be difficult to achieve. There is always a residue of local, implicit, and informal knowledge that cannot be formalized and needs to be presented in an informal manner.

Narration

The purpose of narration or story-telling is to give structure and meaning to an otherwise incoherent ensemble of data, objects, models, etc. One of the problems of using DSS effectively is the complexity of models and the challenge that this poses to users in terms of comprehension and accessibility. Narration can help users make sense of the models using tools that could be formal or informal, mathematical or descriptive, textual or visual, and so on.

Simulation

Simulations have recently attracted considerable attention among social and organizational scientists. Simulations are useful not only because they provide the opportunity for controlled experimentation — i.e., playing around with, and observing the impact of, parameters without affecting the real world — they are also useful because they "enable the observation and recording of the background of planning, decision-making, and evaluation processes that are usually hidden" (Dörner 1996: 9–10). In short, simulations can be conceived as "social labs" that provide an effective way of learning different from the known alternative of *learning by doing* (Bousquet *et al.* 1999).

Although simulations can take the form of mathematical modeling, what is usually intended by the term is either simulation *in ludo* (by people) or *in silico* (by computers), or a combination thereof. The first type of simulation normally takes the form of *role-playing games* (RPG), where people are assigned roles (often different from their real-life status) and given the opportunity to make decisions and to observe the short-term and long-term consequences of those decisions. This helps participants arrive at a shared representation of the problem, which then facilitates the process of collective decision-making. Bousquet *et al.* (1999), who report on the application of RPGs in different projects, describe their various uses in training, research, and policy making. Becu et al. (2003) have noticed the usefulness of simulations as "very efficient communication media."

The crucial point is the experimental and, preferably, visual nature of this phase. Whatever the form, the purpose is to give users a chance to experiment with ideas, to understand the consequences, to be able to associate with others, and so on.

Retrospection as Validation

This is probably the most challenging phase, and its purpose is for actors to relive and review the decision-making process by trying to explain it to others. In this manner, retrospection is utilized as a means of validating, legitimating, and making the process transparent. In this fashion, issues of uncertainty will be at least partly handled. Needless to say, this whole process is iterative and incremental, as people might discover in the act of retrospection that critical aspects are missing, conflicts are still outstanding, goals are not achieved, and so on.

In sum, the methodology outlined above supports the conceptual principles laid out earlier. Technically, This method is also strongly supported by emerging technologies such as those suggested by Shim *et al.* (2002) — e.g., web and networking technologies, ubiquitous computing, and enterprise information systems (EIS). This methodology is also in line with previous empirical findings in the history of DSS — e.g., the advantages of prototyping (Alavi 1984) and of evolutionary methods over other alternatives — and finally, with recent developments in software engineering — e.g., Boehm's Theory W (Boehm et al. 1995).

Conclusion: Taking Decisions into the Wild

What is missing in the design of many contemporary IS is an effective link between the planned, automated decision process and all those tacit aspects, such as the because-of motives, or past experience, which give meaning to the development and implementation of a decision. This is why automated procedures tend so often to be underutilized, for they do not match changing circumstances, badly mimic the know-how of even a novice, feel unnatural and clumsy, seem to lack meaning and be out of context, and are full of loopholes which have to be filled by [improvised] human intervention. (Cibora 1996, p. 375; in Walsham 2001)

AI and DSS have both undergone dramatic changes since their common inception about half a century ago. In the course of this development, the two fields have mutually informed each other, but most of their dealings have taken place at the technical level. The DSS community, in particular, has always stayed "on top of" the latest technical developments in AI and other computer technologies. However, as Cibora contends in the above quote, this has not always resulted in the most effective, meaningful, useful, and user-friendly systems. This article is an attempt to remedy this situation by aligning DSS with recent developments in AI. Following a top-down approach, we began at the conceptual level and derived basic design principles that, in turn, led us to methodological (and partly technical) levels. The outcome of this study is a perspective on the design of iDMSS with the following highlights:

By introducing activities into the picture, this perspective emphasizes the *process* of decision making rather than its *product*;

By downplaying mental models, it reduces the cognitive load of deliberation on decision makers

By starting from the external situation, it makes it more likely and probably easier for multiple decision makers to arrive at a common representation of the problem (which is arguably a major step toward consensus-building)

By incorporating improvisation, it increases the likelihood of more reflective and adaptable outcomes;

These are the crucial features of the conceptual and methodological approach presented here. In summary, the big lesson of recent developments in AI for the design of iDMSS is that a good deal of human cognition takes place not in individual, detached, and disembodied brains, but in the contextual interactions among embodied human beings — or, put metaphorically cognition takes place *in the wild* (Hutchins 1995). This is the lesson that the DSS community should take most seriously. *From the boardroom to the floor and from the headquarters to the field* — this should be the motto of future iDMSS designers.

Acknowledgements

The author would like to thank the Redlands Institute and its director Jordan Henk for the continuing support and invaluable input, research analysts Paul Burgess and Frank Davenport for very useful interactions, and Lindsey Devlin for help with the diagrams.

References

Agre, Philip E. *Computation and Human Experience.* Cambridge University Press, 1997.

Alter, Steve. Development patterns for decision suppport systems. MIS Quarterly, September, 33-42, 1978.

Alavi, Maryam. An assessment of the prototyping approach to information systems development. Communications of the ACM, 27(6), 556-563, 1984.

Becu, N., Bousquet, F., Barreteau, O., Perez, P., & Walker, A. A methodology for eliciting and modelling stakeholders' representations with agent-based modelling. In D. Hales, B. Edmonds, E. Norling & J. Rouchier (Eds.), Multi-agent-based simulation iii - 4th international workshop (Vol. Lecture Notes in Artificial Intelligence 2927, 131-148. Melbourne, Australia: Springer, 2003.

Blackler, F. Knowledge, Knowledge Work and Organizations: An Overview and Interpretation, *Organization Studies*, 16(6), 1021-1046, 1995.

Boehm, B.W., P. Bose, E. Horowitz, and M.J. Lee. Software Requirements Negotiation and Renegotiation Aids: A Theory-W Based Spiral Approach. *Proceedings of ICSE-17*, 243–253, 1995.

Adrienko, Natalia, and Gennady Adrienko. "Intelligent Support for Geographic Data Analysis and Decision Making in the Web." *Journal of Geographic Information and Decision Analysis* 5.2 (2001): 115-28.

Bhargava, Hemant K., and Daniel J. Power. "Decision Support Systems and Web Technologies: A Status Report." *America's Conference on Information Systems*. Boston, Massachusetts, 2001. 7.

Bloomfield, B. P. *The Question of Artificial Intelligence*. London, UK: Croom Helm, 1987.

Bousquet, F., Barreteau, O., Le Page, C., Mullon, C., & Weber, J. An environmental modeling approach : the use of multi-agent simulations. In F. B. a. A. Weill (Ed.), Advances in environmental modeling 113-122, Elsevier, 1999.

Büscher, M. and P.H. Morgensen. Mediating Change: Translation and Mediation in the Context of Bricolage. In: T. McMaster et al. (eds.), *Facilitating Technology Transfer Through Partnership: Learning from Practice and Research*. London: Chpman Hall, 1997.

Ciborra, Claudio U. Improvisation and Information Technology in Organizations. In J.I. DeGross, S. Jarvenpaa and A. Srinivasan, eds. *Proceedings of the 17th International Conference on Information Systems*. Baltimore: ACM, 1996.

Clancey, William J. Situated Cognition: On Human Knowledge and Computer Representations. Cambridge, UK: University of Cambridge, 1997.

Clark, A. *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA: MIT Press/Bradford Books., 1997.

Crevier, D. The Tumultuous History of The Search for Artificial Intelligence. New York: Basic Books, 1993.

DeSanctis, G. and Gallupe, B.A. A foundation for the study of group decision support systems. *Management Science*, 1987, 33(5), 589–609.

Dewey, John. *Logic: The Theory of Inquiry*. Carbondale, IL: Southern Illinois University Press, 1938.

Dörner, D. *The Logic of Failure: Recognizing and Avoiding Error in Complex Situations*. Cambridge, MA: Perseus Books, 1996.

Ekbia, Hamid R. *Artificial Dreams: Science and Engineering in Tension.* Forthcoming (a)
_________ Rethinking DSS: Lessons from Environmental Management. Forthcoming (b)

Elgarah, Wafa, James F. Courtney, and John D. Haynes. "A Dialectical Methodology for Decision Support Systems Design." *35th Annual Hawaii International Conference on System Sciences.* Big Island, Hawaii, USA: IEEE Computer Society, 2002. Vol. 256.

Eom, Sean B. "An Overview of Contributions to the Decision Support Systems Area from Artificial Intelligence." *Proceedings of the Americas Conference on Information Systems.* Baltimore, 1998.

Forgionne, Guisseppi A. "Decision Technology Systems: A Vehicle to Consolidate Decision Making Support." *Information Processing and Management* 27.6 (1991): 679-797.

Forgionne, Guisseppi A., Jatinder N.D. Gupta, and M. Mora. "Decision Making Support Systems: Achievements, Challenges and Opportunities." *Decision Making Support Systems: Achievements and Challenges for the New Decade.* Eds. M. Mora, G. Forgionne and J.N.D. Gupta. Hershey, PA: Idea Group, 2002. 392-403.

Gorry, G.A. and M.S. Scott Morton. A framework for management information systems. *Sloan Management Review*, 13(1), 5070, 1971.

Goul, Michael, J. Henderson, and F. Tonge. "The Emergence of Artificial Intelligence as a Reference Discipline for Decision Support Systems Research." *Decision Sciences* 23 (1992): 1263-76.

Gupta, Jatinder N.D., and S.K. Sharma. *Intelligent Enterprises for the 21st Century.* Harrisburg, PA: Idea Group Publishing, 2004.

Hayes, Patrick J. The Naive Physics Manifesto. In M. A. Boden (Ed.), *The Philosophy of Artificial Intelligence* (pp. 171–205). Oxford, UK: Oxford University Press, 1990.

Hofstadter, Douglas R. *Fluid Concepts and Creative Analogies.* New York: Basic Books, 1995.

Keen, P.G.W., and M.S. Scott Morton. *Decision Support Systems: An Organizational Perspective.* Addison-Wesley Series on Decision Support. Reading, Massachussets: Addison-Wesley Publication Company, 1978.

Kling, Rob. *Computerization and Controversy: Value Conflicts and Social Choices.* San Diego, CA: Academic Press, 1996.

Lenat, Douglas. B., & Guha, R. V. *Building Large Knowledge-Based Systems.* Reading, MA: Addison Wesley, 1990.

Mahmood, Mo A. and Jeanette N. Medewitz: Impact of design methods on decision support systemws success: an empirical assessment. Information & Management, 9(3), 137-151, 1985.

Morton, Alec, Fran Ackermann, and Valerie Belton. *Distributed "Group" "Decision" "Support": A Study of Some Key Concepts*. Management Science : Theory, Method and Practice Series. Glasgow: University of Strathclyde, Department of Management Science, 2001.

Newell, Allen and Herbert A. Simon. GPS, A Program that Simulates Human Thought. *Lernende Automaten,* H. Billing (ed.) 109–124, Munich: R. Oldenbourg KG, 1961.

_________ *Human Problem Solving.* Englewood Cliffs, N.J.: Prentice-Hall, 1972.

_________ Computer science as empirical inquiry: symbols and search. *Communications of ACM,* 19(3): 113–126, 1976.

. Pittsburgh, Pennsylvania: Design Research Center, Carnegie-Mellon University, 1982.

Pereira, Guimarães A. , and S. Corral Quintana. "From Technocratic to Participatory Decision Support Systems: Responding to New Governance Initiatives." *Journal of Geographic Information and Decision Analysis* 6.2 (2002): 95-107.

Rinner, Claus. "Web-Based Spatial Decision Support: Status and Research Directions." *Journal of Geographic Information and Decision Analysis* 7.1 (2003): 14-31.

Schank, Roger C. *Dynamic Memory: A Theory of Learning in Computers and People.* Cambridge, UK: Cambridge University Press, 1982.

Schön, D. A. The reflective practitioner : How Professionals Think in Action. New York: Basic Books, 1983.

Shim, J.P, Merrill Warkentin, James F. Courtney, Daniel J. Power, Ramesh Sharda, and Christer Calsson "Past, Present, and Future of Decision Support Technology." *Decision Support Systems* 33.2 (2002): 111-26.

Simon, Herbert Alexander. *The New Science of Management Decision.* Englewood: NJ: Prantice-Hall, 1977.

_________ Two heads are better than one: The collaboration between AI and OR. *Interfaces,* 17(4), 8–15, 1987.

Smith, Brian C. The Owl and the Electric Encyclopedia. *Artificial Intelligence,* 47(1–3), 252–288.

_______ Situatedness. In *MIT Encyclopedia of Cognitive Science* (MITECS). Cambridge, MA: MIT Press, 1999.

Sprague, R.H. and E.D. Carlson(eds.). *Building effective decision support systems.* Englewood Cliffs, NJ: Prentice-Hall, 1982.

Suchman, L. *Plans and Situated Actions.* Cambridge, U.K.: Cambridge University Press, 1987.

Turoff, M., Hiltz, S. R., Cho, H.-K., Li, Z., & Wang, Y. Social Decision Support Systems. Paper presented at the 35th Annual Hawaii International Conference on System Sciences, Big Island, Hawaii, 2002.

Walsham, Geoff. *Making a World of Difference: IT in a Global Context.* West Sussex: UK: John Wiley & Sons, 2001

Weick, Karl E. *Making Sense of the Organization.* Oxford, UK; Malden, MA: Blackwell Publishers, 2001.

Weick, Karl E. *Sensemaking in Organizations.* Thousand Oaks: Sage Publications, 1995.

_______ Improvisation as a Mindset for Organizational Analysis, *Organizational Science*, 9(5), 543–555, 1998