UNIVERSITÄT KOBLENZ · LANDAU

Institut für Wirtschaftsinformatik
Fachbereich Informatik
Universität Koblenz-Landau

ULRICH FRANK

REFLECTIONS ON THE CORE OF THE INFORMATION SYSTEMS DISCIPLINE

November 1998

Arbeitsberichte des Instituts für Wirtschaftsinformatik Nr. 14

UNIVERSITÄT KOBLENZ · LANDAU

Institut für Wirtschaftsinformatik
Fachbereich Informatik
Universität Koblenz-Landau

ULRICH FRANK

REFLECTIONS ON THE CORE OF THE INFORMATION SYSTEMS DISCIPLINE

November 1998

Arbeitsberichte des Instituts für Wirtschaftsinformatik Nr. 14

Anschrift des Verfassers/
Address of the author:

Prof. Dr. Ulrich Frank
Institut für Wirtschaftsinformatik
Universität Koblenz-Landau
Rheinau 1
D-56075 Koblenz

Arbeitsberichte des Instituts für Wirtschaftsinformatik
Herausgegeben von / Edited by:

Prof. Dr. Ulrich Frank
Prof. Dr. J. Felix Hampe

Bezugsquelle / Source of Supply:

Institut für Wirtschaftsinformatik
Universität Koblenz-Landau
Rheinau 1
56075 Koblenz

Tel.: 0261-9119-480
Fax: 0261-9119-487
Email: iwi@uni-koblenz.de
WWW: http://www.uni-koblenz.de/~iwi

iwi Institut für Wirtschaftsinformatik
Fachbereich Informatik
Universität Koblenz-Landau

Abstract

Information systems research is still a young discipline that has not accomplished a coherent profile yet. Without emphasizing the necessity of *one* common profile, this report discusses a number of candidate profiles or cores of the discipline. Firstly, a recent proposal made by Weber will be evaluated. Weber argues that "deep structure phenomena" of information systems constitute the core of the discipline. Slightly different from Weber, this paper is based on the assumption that information systems research aims at concepts which support the design, introduction, use and maintenance of information systems that are suited to increase an organisation's *efficiency*. Therefore essential inhibitors of successful automation will be identified. Against this background, the report describes and analyses three prototypical approaches to accomplish a higher level of automation in organisations: a *reductionist*, an *inductivist*, and a *constructivist* approach. While there is evidence that the constructivist approach is most effective - and therefore a favourite candidate to constitute the core of the discipline, it will be shown that it implies severe epistemological problems. It will be argued that there is hardly any straightforward solution to these problems. However, for the information systems discipline to establish a profile of its own, it has to accept them as a pivotal challenge.

1. Introduction

The information systems discipline is still characterised by a remarkable diversity. Among other things, it includes research topics such as the economic evaluation of information systems (for instance: [RaAa96], [Saa96]), the prerequisites of creativity in information systems organisations ([Cou94]), the development of information models and even the investigation of sexual harassment via e-mail ([SiWa97]). This diversity is accompanied by a number of different research methods. This "so-called free-for-all situation" ([BaLa92], p. 87) has caused various authors to emphasise the need for a coherent profile of the discipline ([BaLa92], [Hir92]). While it is arguable whether there has to be one common profile all researchers of a particular discipline feel committed to, it is certainly a good idea to discuss possible candidates for profiles that allow to identify research communities within a discipline. A delphi study conducted a few years ago [KöHe95] produced a number of suggestions for profiles of the information systems discipline in german speaking countries ("Wirtschaftsinformatik"). The majority of experts (academics as well as practitioners) voted for the analysis of organisational issues as the core of the discipline. Other suggestions included labels like "information science", "innovation science" or "system development". In a recent monograph [Web97], Weber suggests to concentrate on conceptual modelling as the core of the information systems discipline. Although his perspective does not reflect a main stream of research in the discipline, this remarkable outline of a possible core of the discipline deserves our attention. This is for two reasons: Firstly, there is no doubt that information and its representation plays a central role in the discipline ("nomen est omen"). Secondly, Weber's suggestions are well suited to foster a fruitful discussion because they are specific and - at least in part - provocative. Therefore our reflections on the core of the discipline will start with a review of Weber's book.

We regard the core of a discipline being defined by the subject (in our case: information systems) and the knowledge we want to gain about the subject in the long run. Hence, the core of a discipline corresponds to a long term research orientation or strategy - which includes assumptions about appropriate conceptualisations and research methods. Beside the diversity of particular research foci, there is one implicit topic that is hardly questioned as an important, although not necessary, orientation of information systems research: the development of frameworks or theories that are suited to improve efficiency by increasing the level of automation in organisations. Against this background, the approach suggested by Weber is certainly helpful, but - as will be shown in this paper - not sufficient. Instead there is need for research strategies that focus more directly on the application of information systems in order to contribute to more efficient business processes. While we can think of a plethora of approaches to focus on particular aspects of this orientation, there are only a few generic research strategies. In this paper we will analyse three prototypical strategies which are based on specific ideas about the core of the information systems discipline: a reductionist, an inductivist, and a constructivist research strategy. In practice, they may occur together in combined configurations. The results of this investigation suggest that a constructivist approach - which is quite common for the actual evolution of corporate information systems - is most promising. However, at the same time it is accompanied by severe epistemological challenges for the information systems discipline. We will argue that these problems can hardly be solved in a satisfactory way. Facing them, however, recommends to reflect and eventually change the way information systems research is organised.

2. Conceptual Models of Information Systems as the Core of the Information Systems Discipline?

According to Weber, the brief history of information systems research has been characterized by diffuse patterns of work:

> "The problems and phenomena studied by IS researchers became increasingly diverse. Almost any phenomenon that was vaguely associated with computers was fair game. A new liberalism prevailed to the point where at times a casual observer of the discipline might have been forgiven for wondering what relevance the research had for the IS discipline as opposed to, say, the psychology, economics, or organizational behavior disciplines." ([Web97], p. 3)

While this "disciplinary anarchy" ([Web97], p. 10) could be regarded as an expression of intellectual freedom, Weber argues that it contributes to the "parlous state" ([Web97], p. 13) of the IS discipline. Therefore he recommends to search for common orientation for the discipline. That requires to articulate the essence of core of the discipline, "the body fo knowledge that leads others to recognize it and to acknowledge it as being *distinct from* other disciplines ..." ([Web97], p. 27) - we could also say: the *core competence* of the discipline. In order to accomplish a coherent profile, Weber suggests to "surrender the past" ([Web97], p. 29) and start from scratch. Based on the assumption that information systems research deals with the design, implementation, use, and maintenance of information systems, Weber starts his endeavour with an analysis of the central terms "information", "system" and "information system". Being a declared follower of objectivist research methods, his intention is to provide definitions which are as precise/formal as possible. Therefore it comes to not much surprise that the concepts Weber suggests are based on the work of the philosopher Mario Bunge. In his seminal, seven volume "Treatise on Basic Philosophy", Bunge analyses the structure of models and their relationship to the things they represent. Refering to Bunge, Weber assumes that the world is made up of "things" which "really exist in the world" ([Web97], p. 34). We perceive things through their properties. Within a model, properties are represented by attributes. However, an attribute does not have to correspond directly to a property. Instead, it may be an abstraction over a set of properties. A "property in general" is the set of permissible "properties in particular". Other terms are "class of things", "event", "state", "lawful state". The basic terms are described in a number of formal postulates and definitions. For instance:

"Postulate 2.2: Let S be a set of things, and let $\{T_i\}$ be a set of nonempty sets, the elements of which may be equal to or different from S. Then:

(a) a property in general can be represented by a proposition of the form:

$A: S \times T_i \times ... \times T_n \rightarrow$ Propositions including A

(b) a property in particular of a thing $s \in S$ can be represented by a value of an attribute A (s, $t_1, ... , t_n$), where $t_i \in T_i$" ([Web97], p. 36)

This approach leads to the specification of concepts which can be used to create information models. Notice that there is a clear relationship to concepts in conceptual modelling - although Weber does not mention it: a "property in general" corresponds to a type or a class, a "property in particular" to an instance. A "law" can be regarded as a constraint and a "lawful" state is similar to the concept of an "invariant" ([Mey97], p. 364).

The term "information" which is also defined in a formal language is distinguished from "data" with respect to the effect it has on the recipient: An external event is regarded as data, if it does not change the recipients state; otherwise it is regarded as information ([Web97], p. 63). This definition implies that the distinction between data and information is context dependent. Furthermore it is based on the assumption that there is no external event that is not data or information. Weber regards an information system as a "system of represenations": "Information systems represent the history of things in the real world in terms of the ways we have chosen to conceive these things." ([Web97], p. 65) From our perspective, these definitions, while arguable in some detail, are not essential. However, based on his notion of an information system, Weber differentiates three levels of abstraction within information systems: by "surface structure phenomena" he apparently thinks of those artefacts that are used to allow for the interaction with a user: interface widgets, report formats etc. They do no lie in the core of the discipline. "Physical structure phenomena" include platform specific features, like data codes, low level communication protocols etc. They are rejected as a candidate for the core of the discipline, too. Instead, Weber argues that the core is constituted by "deep structure phenomena" ([Web97], pp. 79). They are constituted by those concepts that describe the characteristics of the "real world phenomena that the information system is intended to represent. ... They allow us to discern the meaning embedded within the information system - the essence of the system, and the reason for its existence."

Although being somewhat unusual, this focus seems to make sense at first sight. However, it does not seem to be entirely convincing. Firstly, it rises the question how the representation of deep structure phenomena should look like. While a conceptual model seems to be appropriate, there are obviously information systems that do not come with such an abstraction. Instead, they may be described by program code - machine code at the worst - only. If we think of complex systems, even conceptual models will make it difficult to discern the meaning of a system - not to speak of machine code. In addition to that, the semantics of an information system cannot entirely abstract from other levels of abstraction - namely data types or generic functions provided by the implementation language and/or the operating system. Nevertheless, we agree with Weber that a conceptual level is the most appropriate abstraction for the information systems discipline. We also agree with Weber that a discipline that studies conceptual models needs suitable criteria to judge a conceptual model's quality. We do not, however, follow Weber when he states that the information systems discipline should not take into account the surroundings of an information system, hence the organisation it is embedded in, the business processes it is used by or the corporate strategy it should support:

> "An information system is a system that can be studied in its own right, independently of the way it is developed and deployed in its organizational and social context." ([Web97], p. 84)

At first sight, his argument may make sense: Those aspects are subject of other disciplines. However, these other disciplines usually do not consider the relationship of their subject with information systems. The focus Weber suggests also raises the question how to distinguish between computer science and the information systems discipline: It is characteristic for computer science to abstract from the context an information system is used in and focus on internal, usually formal properties of a system instead. In order to evaluate a model's quality independently from its context, Weber suggests two criteria: *Accuracy* and *Completeness*. Accuracy is accomplished if the meaning of a concept in a model is the same as the meaning of the corresponding concept in "the user's model of the real world system that the information system is supposed to represent." Completeness on the other hand requires that an information system represents all the meaning that is part of the user's conception of the domain of interest. Weber suggests two ways to evaluate these criteria. Firstly, they can be analysed by observing an information system's behaviour. Secondly, they can be evaluated by analysing representations of the information system, such as conceptual models or program code. Such an approach does not seem to be satisfactory in the end. An information system will usually serve different users who in turn may come to deviating judgements. The evaluation of a system would require to somehow aggregate individual judgements. In addition to that, we have to take into account that individual judgements may change over time. While these are problems, the evaluation of artefacts can hardly avoid, the demand for completeness seems to be misleading: For a good reason, the representations that constitute an informations systems are *abstractions* of mental models of a particular domain. In other words: Usually you would not even try to accomplish a complete representation of a mental model.

Focusing on the quality of the representation of mental models alone is not satisfactory. In this case the quality of an information system would depend on the quality of corresponding mental models which in turn would not be subject of further investigation. Information systems, however, are not an end in themselves. Instead they should serve certain objectives. As stated already in the introduction of this report, we assume that the general, abstract goal of information systems is to increase an organisation's efficiency by increasing the level of automation. The three prototypical research strategies are based on this assumption.

3. The Reductionist Approach: Formalizing Individual Knowledge

Lifting the level of automation in organisations suggests a concentration of effort on tasks which cannot so far be performed by machines. Since those tasks are performed - more or less successfully - by human beings, the automation of individual intellectual capabilities would allow to accomplish a higher level of automation in organisations, too. In the context of this paper, we call a research strategy reductionist if it focuses on the automation of individual competence. The classical example of a reductionist approach is the strategy that is inherent in Artificial Intelligence. It aims at reconstructing human intelligence or relevant parts of it with machines. If it is successful, such a strategy would allow for "... substituting machines for any and all human functions in organizations." ([Sim64], p. 594). We cannot decide whether or not it is possible to achieve this ambitious goal. Therefore, it is not useful to discuss this possibility. However, the lack of success in the respective research has dampened the early enthusiasm. As a consequence, a new, less ambitious approach emerged in the eighties: Instead of reconstructing human intelligence, the "expert system" approach focused on the identification and formalisation of *knowledge* which can be interpreted by a machine in order to automate intellectual tasks like diagnosis, planning and decision making in general. Although expert systems failed to fulfil the daring promises made in the eighties, formalisation of knowledge is still a pivotal strategy in fostering automation and supporting efficient decision making. Within the information systems domain, *managerial* decision making is of special importance. In the following section, we will analyse possibilities for and limits of automating managerial decision making through the formalisation of managers' knowledge.

3.1 Formalizing Managerial Decision Making: The Pivotal Challenge

The variety of tasks that are performed by managers is immense. Some of those tasks have been automated already, many others seem not to allow for automation. For the purpose of our investigation, it would be sufficient if there was an essential, common ingredient of managerial competence that, at the same time, resisted to being formalized. Usually, a manager has to act in an environment that is characterized by imponderability, ambiguity and risk. Therefore, it seems reasonable to identify the ingredient we are looking for as the ability to cope with *contingency*. The term contingency denotes aspects of reality - like events or processes - which we cannot be certain about: We do not know for sure when they will occur or what outcome they will produce. Therefore, acting in a contingent environment implies to select from options we cannot take for granted: They are possible, but not necessary, or - as Luhmann ([Luh84], p. 152) puts it: The notion of contingency denotes something given (experience, expectations, thoughts) with respect to its possible "Anderssein" (being different). Rescher ([Res68], pp. 229) uses the term "possible worlds" to characterize this phenomenon. In case of mutual contingent expectations, Luhmann speaks of "doppelte Kontingenz" (double contingency, [Luh84], pp. 148).

If dealing with contingency is a core task of managers, the question remains what it takes to act successfully in a contingent environment. Instead of analysing this question directly, we assume that the basic assumption of the reductionist approach is appropriate: Professional decision making under contingency is usually based on knowledge which can be formalized. Additionally, we simplify our investigation by avoiding the ontological question about the nature of human knowledge. Instead, we focus on formal representations of "knowledge" only. This is feasible because we are primarily interested in the question whether these representations are sufficient for successful automation.

3.1.1 Remarks on "Vague" Knowledge

The representation of knowledge for the purpose of automation requires a formal language. Different from a natural language, a formal language is a system of symbols that is generated by a calculus. It is based on a well defined alphabet of characters and a clear syntax. The semantics of propositions expressed in a formal language is provided by an unambiguous interpretation.[1] In traditional logic, an interpretation defines whether a proposition is true or false. For this reason, it seems that formal languages are not appropriate to describe knowledge about contingent domains: It is an essential feature of those domains that we are not certain about all their relevant aspects. In order to allow for the representation of so called "vague" knowledge, research on expert systems has focused on a formal languages that are not restricted to a binary concept of truth - like probabilistic logic and fuzzy logic.

Nilsson ([Nil86]) suggests a probabilistic logic that allows to assign truth values between zero and one. Nevertheless, he argues that his approach does not hurt the *tertium non datur* of traditional logic: Any proposition can be applied for a set of possible worlds. Within one subset of those worlds the proposition is true, within the other it is false. Hence, the probability for the truth of a proposition is defined as the relative frequency of the worlds in the first subset. In order to allow for logical deduction, Nilsson defines the probabilities resulting from basic logical operations:

Negation	$p(\neg A) = 1 - p(A)$
Conjunction	$p(A \wedge B) = \min(p(A), p(B))$
Adjunction	$p(A \vee B) = \max(p(A), p(B))$
Implication	$p(A \Rightarrow B) = \max(1 - p(A), p(B))$

1. .Notice that there may be more than one interpretations. An interpretation is sometimes called a model.

If the implication is not certain, it can be assigned a probability, too. In this case the probability of the implied proposition has to be within the following interval:

$$p (A \Rightarrow B) + p (A) - 1 \leq p (B) \leq p (A \Rightarrow B)$$

Different from probabilistic logic, fuzzy logic gives up the traditional tertium non datur and replaces predicates by fuzzy sets: An object can be assigned a particular feature in different degrees of intensity. Hence, objects may have different "degrees of membership" within a fuzzy set ([Gog68], p. 332). The proponents of fuzzy logic often use natural language designators such as "*very likely*", "*very unlikely*", "*likely*" etc. ([Zah75], p. 408) to characterize a degree of membership. Nevertheless, in the end the degree of membership is expressed by a number - a "truth value" - between zero and one. Similar to the probabilistic logic suggested by Nilsson, there are rules for the calculation of truth values resulting from logical operations ([Zah75], p. 410). Fuzzy logic certainly deserves more detailed considerations of some of its features - like the fact that tautologies may have a truth value less than one and contradictions may have a truth value larger than zero. However, for our investigation it is sufficient to note that fuzzy logic like probabilistic logic requires to express confidence by numbers.

At first sight, a logic that allows to express uncertainty seems to be well suited to represent vague knowledge about contingent domains. However, there are severe objections against the use of such representations for the automation of intellectual tasks. The degree of certainty (or membership) will usually be assigned by human judgement. Contingency means that we do not know exactly how reliable our assumptions are. Assigning a number nevertheless implies to deal with a paradox in the end: certainty about uncertainty. Moreover, such an approach suffers from the fact that most people are not qualified to deal with probabilities. A plethora of research on judgement under uncertainty has revealed numerous sources of bias (for instance: [TvKa71], [KaSl82], [Hog80]): "... people systematically violate the principles of rational decision making when judging probabilities ..." ([SlFi77], p. 169). This may result in constellations that are hardly acceptable for the purpose of automating decision making: Somebody expresses confidence in a certain proposition and denies another one which is logically equivalent. While it is not too surprising that many people are having difficulties to deal with probabilities, specialized research indicates that highly qualified domain experts do not perform well either ([LiFi82]).

3.2 Challenges and Limits of Automation via Knowledge Representation

Although the problems to be expected from representing vague knowledge dampen the hope for knowledge based systems that allow to automate decision making under contingency, they are not sufficient to compromise the basic assumption of the reductionist approach: The human ability to deal with contingency is based on knowledge which, in principle, can be formalized. A well known objection against this assumption claims that there is knowledge that we simply cannot express - not to speak of formalize it. Wittgenstein uses a metaphor to illustrate this argument: We know how high the Mont Blanc is - and we are able to describe this knowledge in a comprehensive way. This is, however, different with our knowledge about the sound of a clarinet ([Wit71], §63). While this argument - "knowing but not being able to tell" - certainly denotes a severe challenge, it does not exclude the possibility of retrieving knowledge that seems to be tacit.

Judgement in the face of contingency in social systems requires to take into account a plethora of information. While some of this information may have a clear meaning (like revenues, profits etc.), other information requires additional interpretation (like information about the performance of employees or the trustworthiness of customers). Often, it will be necessary to actually *perceive* aspects of reality in order to judge them properly. Interface devices that are able to monitor reality - like cameras or microphones - are certainly not sufficient since they do not detect any semantics. Even software that allows for the analysis of natural language input would not be of much help, since it completely neglects the intention of the speaker and the background of his statements. However, in order to make appropriate decisions in a contingent environment, it is essential to judge a speaker's confidence in his statements and whether or not he may be lying.

These thoughts have a severe impact on the basic assumption of the reductionist approach since they suggest that knowledge alone is not sufficient for coping with reality. Moreover, there is evidence that knowledge and perception cannot be separated: Knowledge is a prerequisite for selective perception, perception in turn is necessary not only to gain but also to apply knowledge. Simon ([Sim83], p. 28) speaks of a "feedback tie", according to Weick ([Wei80], p. 159) "Action, perception, and sense-making exist in a circular, tightly coupled relationship ...".

Today, we know too little about the physiological and intellectual capabilities that constitute perception to allow for their automation. Nevertheless, similar to the basic assumption underlying early research in Artificial Intelligence, we cannot exclude the possibility. There is, however, another aspect of managerial decision making that denotes a definite limit of the reductionist approach. There is evidence that managers sometimes make decisions without substantial knowledge. If there is no knowledge, there is not chance to use it as a basis for automation. While this is not surprising, it is important to note that humans can nevertheless be successful in coping with contingency. Since managerial competence is "hard to judge" ([Per86], p. 11), selecting managers is based on "beliefs" ([Pfe78], p. 24). Therefore, their judgements - even if they are not based on substantial knowledge - may be commonly accepted as an appropriate orientation. Hence, managers can be successful by constructing "'objective' features of their surroundings" ([Wei80], p. 164). Self-fulfilling prophecies are an obvious example. In other cases - shared values, language games - the influence is more subtle. There is no doubt that this symbolic function of decision making under contingency does not allow to be substituted by machines: It is essentially related to social interaction.

4. The Inductive Approach: Searching for Generic Patterns of Information Use

The limits encountered for the reductionist approach do not compromise it in general: There is certainly still individual knowledge that can be formalized. Even if it does not allow to completely automate decision making, it may help with increasing the level of automation in organisations by *supporting* human decision making. However, we should not expect a substantial shift in the level of automation if we focus on the reductionist approach only. While the reductionist approach is focusing on the automation of individual skills, the following approach is less ambitious. It tries to discover common patterns of existing information use. In order to stress this aspect, we call it the inductivist approach although it is not necessarily based on induction as a research method. Note that this is a prototypical strategy that one can hardly find in a pure sense.

4.1 Prospects

Beside the difficulties to automate intellectual capabilities to cope with contingency, one major inhibitor of automation in organisations is the lack of common concepts to represent and process information. There is an immense diversity in information concepts that are used within software and that are relevant for information interchange. Almost any software used in organisations is based on its own concepts of entities like "customer", "address", "invoice", "order", etc. Different from the reductionist approach, formalizing those concepts is not a substantial problem. Nevertheless, the variance of existing concepts hinders automation. On the one hand, it imposes a problem for information exchange (hence: for the *integration* of information systems). If the receiver of a message uses a concept that includes more information than the concept used by the sender, a transformation will either be accompanied by the lack of information or by the risk that is related to the reconstruction of missing information. Even in the case that there are no semantic differences, an automatic transformation may fail because it is too expensive - in the case of n participants with individual concepts, there is need to implement n (n - 1) transformations. On the other hand, the lack of common concepts hinders the *reusability* of software. With increasing levels of software reuse, the cost per copy is decreasing. This would in turn open up new areas of economic automation.

The idea of an inductivist approach is based on the assumption that there are common patterns of information use. Discovering them would allow to establish reference concepts which would decrease transaction costs and foster software reuse.

4.2 Feasibility and Pitfalls

The variance of requirements and of actual information use indicates that it would be a remarkable challenge to find concepts that are suitable for all organisations. This is especially the case for concepts that incorporate a high level of semantics (information content): The higher the information content of a concept the less the probability that it represents a particular facet of reality in an adequate way. In principle, however, it would be possible to find generic concepts, since - at a point in time - it is a finite number of cases that has to be taken into account. In order to cover a high amount of variance, those concepts would have to be very complex and therefore hard to implement and use. Without any doubt, identifying suitable generic concepts would cause tremendous costs. Independent from the challenges an inductive strategy would have to face, there are reasons why its results - if it would ever be successful - were of doubtful use. Actual information use is an expression of the ways tasks and processes are organized. Usually, (business) processes as well as the documents that they use have not been designed - not to speak of optimized - with the capabilities of modern information technology in mind. In order to fully exploit the potential of this technology, it is more promising to rethink traditional ways of performing processes and using information, instead of simply reproducing them. Schank ([Sch85], pp. 23) uses an analogy to illustrate this thought:

> "The first users of cars and computers had to struggle to make these completely new machines operate within the limits of the systems that were designed for an earlier world. ... The computer industry’s perspective suffers from the same lack of creativity and long-term vision. Computers are severely limited by the world views and ideas that have preceded them."

Considering both, the tremendous effort that it would cause and the doubtful outcome, a purely inductivist - we could also say: a descriptive - approach to increase the level of automation in organisation is hardly an attractive option.

5. The Constructivist Approach: Preparing Reality for the Information Age

There is not doubt that suitable common concepts - we could also speak of a *common language* - would help to increase the level of automation in organisations. If we want to avoid the pitfalls imposed by replicating existing information use, we have to design new, artificial concepts that take into account the special advantages and restrictions of information technology. We call this approach constructivist to express that it aims at creating new shapes of reality. This is the case for two intertwined reasons. First and more obvious, the approach includes the analysis and possible redesign of processes performed in organisations. Second, the introduction of new concepts to describe reality will influence the way we perceive and speak about the affected reality.

The brief history of information technology offers numerous examples of constructivist approaches: The definition of character codes as well as many standards that define the representation and semantics of digitized information. However, so far those efforts have been mainly concentrated on concepts that incorporate little semantics. Only recently, there are attempts to standardize high level concepts as well - like so called business objects ([OMG96]). Standardization processes are dominated by standardization organisations and powerful stakeholders like governments and large corporations. In the following section, we want to look at the constructivist approach as a *research* strategy which has a number of epistemological implications that are less relevant for pragmatic approaches.

5.1 Reducing Contingency

Similar to the inductive approach, the constructivist approach aims at fostering integration and reusability of information systems. Different from the inductive approach, actual ways of using and producing information are not taken for granted. Instead, a constructivist strategy is based on the assumption that the effective exploitation of the potential provided by modern information technology recommends to redesign traditional means of communication and cooperation. In other words: It recommends to construct reality by introducing new ways of coordinating cooperative work. Where the inductivist approach assumes that the variance in using information and expressing it through languages is a necessary reflection of the variety of tasks to be taken into account, the constructivist approach relies on the presumption that variance in actual information use and related coordination mechanisms is the result of an - at least partially - arbitrary process. For this reason, reducing variance by introducing new common concepts to handle information would not necessarily cause dysfunctional effects. Moreover, if the processes they are to be used in were thoroughly designed, they would contribute to more efficiency. In any case, those common artefacts would allow for a high level of integration and reusability. There are numerous examples where concepts/artefacts were constructed to exploit the potential of computerized information systems rather than merely mapping existing concepts. For instance, languages and architectures to describe electronic documents, such as ODA, SGML, HTML ([App91], [GoRu90]), and especially OpenDoc ([Ope94]) are based on a notion of document that is significantly different from the traditional notion.

Where the reductionist approach aims at reproducing human capabilities to cope with contingency, the constructivist approach would reduce contingency - in other words: variance, ambiguity, and risk by introducing common infrastructures for organisations. This is very similar to the traditional notion of organizing which is characterized by Weick "as a consensually validated grammar for reducing equivocality ..." ([Wei80], p. 3) In order to foster automation, the focus of the constructivist approach is on creating *formal* concepts to handle information. Concepts required to exchange and manipulate information would be provided on an appropriate level of semantics. Since it can hardly be expected that their is usually only one general concept that fits all organisations, a certain variety would have to be provided. In the end, a common formal business language would allow for the automation of processes that cannot be automated in an economic way today.

5.2 Feasibility

The constructivist approach requires to specify information concepts on a high level of semantics (a common concept of an integer may be necessary but is certainly not sufficient). In order to define the meaning of a concept used within an application domain, we need an appropriate representation of this domain: Who is using a concept for what purpose? What is the relevant context (associated concepts, related processes etc.). The formal or at least: semi-formal specification of these concepts could be accomplished in any of many (semi-) formal languages. The representations should be intuitive for everybody who participates in the organisation of a domain. At the same time, it should provide an appropriate input for developing software. *Conceptual modelling* is regarded as the best option to combine both goals. In recent years, a number of authors ([Zac87], [Kat90], [SoZa92], [Pet93]) have suggested enriched modelling frameworks - often named ”enterprise modelling” (a term which, however, is not used in a unique way). Such approaches usually suggest a number of views on the enterprise and intend to capture the relationships between these views. The development of enterprise models seems to be an appropriate strategy to stress the constructivist approach, since they combine the design of information systems and the related tasks of analysing and (re-) designing a firm's organisation and strategy.

Among other things, a method for enterprise modelling consists of a set of modelling languages that are used to model different views of an organisation - for instance: a static information model (like an object model), information exchange models (like message flow diagrams), models of business processes, etc. Multi Perspective Enterprise Modelling (MEMO), for example, offers a number of graphical modelling languages which are specified and integrated through metamodels which in turn are defined by concepts of a common meta-metamodel ([Fra97], [Fra98]). A method for enterprise modelling could be used to develop generic reference models - for instance of certain processes in banks or in insurance companies together with concepts of the information needed within these processes. Meta concepts like specialization or aggregation could be used to create more special models.

While the vision of generic enterprise models may be appealing, it is certainly a complex endeavour that raises the question whether it is feasible at all. At first sight, the obvious variety suggests that generic reference models can hardly satisfy the peculiarities of many organisations. Also, the high amount of time that it takes to develop reference models may cause them to be outdated as soon as they could be established. Indeed, a few attempts in the past - mainly focusing on reference models for manufacturing enterprises ([Sav90], S. 136, [ESP91]) were not successful from a pragmatic point of view: None of them has ever been implemented in a substantial number of organisations. However, there is evidence that organisations are able - and willing - to adapt their processes and information concepts to reference models. Vendors of so called standard business software - like SAP - emphasize that their software is based on optimized reference processes. For this reason, they recommend that their customers should adapt to those reference processes in order to gain competitive advantage. While it is not clear whether customers follow this recommendation in order to improve their processes or simply to take advantage of affordable standard software, it is a matter of fact that more and more organisations are adapted to reference systems.

5.3 Epistemological Problems

While any of the prototypical approaches we have discussed so far can be identified - either isolated or more likely combined - in actual efforts to increase the level of automation in organisation, we are looking at them as research strategies of the information systems discipline. Beyond the subtle differences between theories of knowledge[1], we propose that the results (typically: hypotheses or theories) a research approach produces should satisfy a number of requirements. Among other things,

- they should provide general statements about classes of entities, not statements about single entities,
- they should provide new knowledge,
- they should provide comprehensible procedures to check them against criteria commonly regarded as relevant. This implies the ability to compare competing research results in a way that allows to select the most appropriate option.

Different from traditional goals of (empirical) research, the results of research related to our three research strategies are - at least in part - not just statements that can be checked by confronting them with reality (no matter what theory of truth one prefers). Instead, we are dealing with artefacts - knowledge bases and inference engines, information models, process models and may be corresponding software. While we cannot simply apply the criteria listed above to artefacts, we still assume that the creation of artefacts can be a research goal. The notion of scientific research is based on the idea of progress – in terms of growing knowledge and improving technologies. Progress, however, implies the existence of criteria that allow to discriminate between competing options – be it explanations of reality or artefacts that are to fit actual or future real world domains. A research discipline that does not seriously care about such criteria risks to sacrifice its identity.

Beyond the specific difficulties we have encountered for the reductionist approach, it seems that it does not satisfy the request for new knowledge since it aims at formalizing existing knowledge. However, in principle the reductionist approach could lead to new explicit statements about reality by identifying and describing knowledge that was only tacit before. In this case, the evaluation of a knowledge base could be very similar to the evaluation of hypotheses or theories in science: Every proposition of the knowledge base would have to be confronted with reality. Although we do not propagate a naive realism, such an evaluation should usually lead to acceptable results. This is different, if a knowledge base contains propositions which are assigned a probability - especially if those "probabilities" are in fact thought to express a certain level of confidence. As we have already argued above, there is no entirely satisfactory way to evaluate probabilities. Evaluating the outcome only - like the answers given by a system

1. .The most important theories of knowledge in Germany are "Kritischer Rationalismus" (for instance: [Pop35], [Alb68]), "Kritische Theorie" (for instance: [Hab81]), and "Konstruktivismus" (LoSc75]) - the first sometimes also referred to as "neo-positivistic", the latter as "hermeneutic". In my opinion, the main difference between these methodologies can be seen in the procedures suggested to justify (or falsify) theories.

- is not satisfactory either. If an answer itself is assigned a probability, its evaluation faces the same problem as the evaluation of uncertain propositions in the knowledge base. While testing how well the answers/recommendations with the highest probability actually work seems to be a pragmatic approach, it is hardly acceptable from a epistemological point of view: A set of "appropriate" answers can still be based on inappropriate propositions/probabilities, leaving the risk of failure with any additional try.

The inductivist approach tries to discover general concepts of information and general patterns of information use. This is very similar to traditional empirical research. The evaluation of results produced by this approach seems to be straightforward: General statements (about concepts and processes) have to be checked against reality by asking whether they are suited to describe concepts and processes in actual organisations in an appropriate way. If we abstract from the problems of checking statements about reality against reality, there is still another problem to be faced with the inductivist approach: In the end, it is to provide (semi-) formal artefacts (like conceptual models, specifications or software). Although those artefacts are based on statements about reality, they require a specific evaluation. In order to design (software) artefacts, we need to build abstractions - or models - of the real world. Those models should provide both, an intuitive representation of the real world domain they are directed at, and an appropriate basis for the implementation of software. Evaluating a conceptual model against these requirements is a delicate task. A model can be regarded as intuitive, it if corresponds to individual patterns of perception and conceptualisation. Those patterns, however, are not only difficult to identify, they also vary from person to person as well as over time. Software engineering requirements are often not coherent as well. Therefore, designing conceptual models usually includes decisions which are in part arbitrary. In recent years there has been growing awareness of those problems. There are a few publications that suggest criteria/measures for evaluating the quality of conceptual models ([MoSh94], [Lin94], [KrLi95]) - none of them being convincing.

Moreover, developing conceptual models of real world domains imposes the challenge to evaluate modelling *languages*, since a modelling language (its semantics, abstract syntax and graphical notation) has a pivotal impact on the quality of models. Although we are able to reflect upon language, for instance by distinguishing between object and meta level language, our ability to speak and understand a language is commonly regarded as a competence that we cannot entirely comprehend ([Lor96], p. 49). Therefore any research that either aims at analysing a language and its use or at inventing new "language games" (i.e. artificial languages and actions built upon them), has to face a subtle challenge: Every researcher is trapped in a network of language, patterns of thought and action he cannot completely transcend - leading to a paradox that can hardly be resolved: Understanding a language is not possible without using a language. At the same time, any language we use for this purpose will bias our perception and judgement – or, as the early Wittgenstein put it: "The limit of my language means the limit of my world." ([Wit81], §5.6).

Such considerations may seem to be of philosophical nature only. However, they characterise precisely one dilemma that evolves from the evaluation of modelling languages. Evaluation implies the knowledge of a modelling language. The more you know about the requirements of modelling, the more likely will you be able to understand existing concepts or to point at missing features. Modelling languages are usually designed by people who have gained an outstanding experience with the use of such languages. Usually they will design a language to fit their preferences. While both, the design and the evaluation of modelling languages require modelling experience, it is exactly this experience that will have a tremendous influence on the outcome of those efforts. Not only that this sort of bias hinders the invention of new modelling paradigms, at the same time – and this is an additional problem – we can assume that many people who are affected by modelling languages have perceptions and cognitive styles different from the language experts.

The constructivist approach inherits all the epistemological problems encountered for the inductivist approach. In addition to them, it causes another tremendous problem. Different from the inductivist approach, it does not aim at representing parts of the real world. Instead, it aims at designing new ways of performing collaborative tasks and creating artefacts that fit those future worlds. The pivotal goal of (re-) designing the world is the reduction of contingency. It is hard to tell whether or not this goal has been accomplished, since reducing a system's contingency is always accompanied by a temporary growth of complexity ([Luh67], p. 109) which in turn may cause additional contingency. For instance: If sources of ambiguity are eliminated from a system, somebody who was used to this sort of ambiguity may find the new precise concepts even more contingent as long as he has not understood their meaning.

The evaluation of artefacts has to face two challenges. First, any construct has to be evaluated against the basic assumption that it is suited to foster automation without causing side effects that would overcompensate for the benefits of automation. While this is difficult enough, it is not sufficient: There should also be a way to compare alternative suggestions to (re-) construct the world. Second, there are (implicit) value judgements: Those constructs will affect they way we handle and perceive information, the organisation of our work - and the language we speak. While ignoring individual interests is hard to accept, identifying them in a reliable way (humans are able to learn) is hard to accomplish.

In this situation, one could leave the evaluation to evolution: Those alternatives are most suitable that survive/dominate in the end. Although this "best practice" approach to evaluation is rather common in the information systems domain, it is no satisfactory option. First, it does not allow for an ex ante evaluation which would be desirable because the realisation of a particular design just for the purpose of testing it will usually be no option. Second, and more important, best practice means to sacrifice scientific standards for criteria which are common (and maybe appropriate) for markets. At the same time, emphasising the need for scientific evaluations of artefacts has to face objections, too. On the one hand, it can easily be mistaken for a positivistic attitude: prescribing those ignorant practitioners what language is best for them. On the other hand, considering the little influence the information systems discipline takes these days on the actual evolution of information systems, there seems to be hardly a chance to propagate artefacts produced by the discipline in practice.

Unfortunately, it seems that there is no convincing solution to these problems. The current situation is not satisfactory for another reason, too: It is almost common practice that artefacts (models, modelling languages etc.) are presented at conferences without being thoroughly discussed and evaluated by members of the corresponding scientific community. In our opinion, there is only one chance to overcome this problem (although we are not very optimistic): We have to put more emphasis on rational discourses that aim at a common evaluation - and comparison - of competing designs. In principle, those discourses have to be open for everybody who will live and work in the worlds to be constructed.

6. Conclusions

The prototypical approaches we have discussed so far confront us with a peculiar situation: The more promising the strategy to accomplish higher levels of automation at affordable cost, the bigger the epistemological challenges we have to face. One may argue that a constructivist approach should not be considered as a subject of scientific research - similar to the argument that engineering cannot be measured by the same criteria as science (for a discussion of this question with respect to Kuhn’s theory of science see [WeWi97]). However, that would hardly solve the problem that there is lack of suitable criteria to evaluate and compare alternative artefacts to be used within information systems. It is no option either to neglect the construction of reference artefacts. Those artefacts are created and propagated anyway - by powerful vendors and standardization organisations. They are an essential subject of the information systems discipline. Therefore, it seems a necessity that this discipline takes care of their evaluation. In fact, we suggest that the constructivist approach together with the establishment of procedures to evaluate the corresponding artefacts constitute the core of the information systems discipline.

In the light of the constructivist approach it is certainly not sufficient to concentrate on the "deep structure phenomena" of information systems only - as suggested by Weber. Instead, exploiting the potential of information systems recommends to (re-) organize reality and to describe it with appropriate, wide spread artificial languages. That requires to stress a cross-disciplinary approach which takes into account aspects like human perception (of language and information artefacts), the psychology of organizing, language as a design instrument (software engineering) and as a medium for social interaction and human thought. If it is to be more than pure lip service, such an approach requires suitable interfaces between the disciplines that are involved. According to our experience, the (semi-) formal languages which are used to design conceptual models provide a well suited interface to computer science. On the other hand, the conceptual models themselves should foster the communication with those disciplines who are in charge of the domain being represented - provided the models offer an intuitive representation. Such a separation of concerns does not jeopardize the information systems discipline’s profile. Instead, the core competence of the discipline would be that of a mediator with an objective - and an add-on - of its own: the design, implementation, maintenance and use of information systems to improve organisational performance. That requires concepts that are based on competencies from other disciplines but that cannot be developed by these disciplines alone. In order to foster communication and cooperation with its neighbour disciplines, information systems research has to develop appropriate interfaces, in other words: concepts or languages.

In addition to the cooperation with other disciplines, the constructivist approach would have another impact on the organisation of information systems research. The attempt to establish generic artefacts that could serve as references for a wide range of domains requires an amount of resources that can hardly be covered by a single institution. That implies the need to bundle research efforts and - at the same time - to improve the chances for a fruitful competition of alternative options. Against the background of the (social) reality of universities, these demands may seem naive. However, they are the only chance to oppose the large amount of redundancy caused by isolated research. Also, they offer the only chance for researchers at universities to maintain their profile against the growing dominance of the information systems industry and its subtle influence on research and teaching.

References

[Alb68] Albert, H.: Traktat über kritische Vernunft. Tübingen 1968

[App91] Appelt, W.: Document Architecture in Open Systems. The ODA Standard. Berlin, Heidelberg u.a: Springer 1991

[BaLa92] Banville, C.; Landry, M.: Can the Field of MIS be Disciplined? In: Galliers, R. (Hg.): Information Systems Research. Issues Methods and Practical Guidelines. London et al.: Blackwell 1992, S. 61-88

[Cou94] Couger, J.D.: Measurement of the Climate for Creativity in I.S. Organizations. In: Nunamaker, J.F.; Sprague, R.H. (Hg.): Proceedings of the 27th Annual Hawaii International Conference on System Sciences. Bd. IV, Los Alamitos, Ca.: IEEE Computer Society Press 1994, S. 351-357

[ESP91] ESPRIT Consortium AMICE (Ed.): Open system architecture: CIM-OSA AD 1.0. Brussels 1991

[Fra97] Frank, U.: Enriching Object-Oriented Methods with Domain Specific Knowledge: Outline of a Method for Enterprise Modelling. Arbeitsberichte des Instituts für Wirtschaftsinformatik. Nr. 4, Koblenz 1997

[Fra98] Frank, U.: The MEMO Object Modelling Language (MEMO-OML). Arbeitsberichte des Instituts für Wirtschaftsinformatik. Nr. 9, Koblenz 1998

[Gog68] Goguen, J.A.: The Logic of Inexact Concepts. In: Synthese, vol. 19, 1968/69, pp. 325-373

[GoRu90] Goldfarb, C.; Rubinsky, Y.: The SGML Handbook. Oxford: Clarendon Press 1990

[Hab81] Habermas, J.: Theorie des kommunikativen Handelns. Vol. 1: Handlungsrationalität und gesellschaftliche Rationalisierung. Frankfurt/M. 1981

[Hir92] Hirschheim, R.: Information Systems Epistemology: An Historical Perspective. In: Galliers, R. (Hg.): Information Systems Research. Issues Methods and Practical Guidelines. London et al.: Blackwell 1992, S. 28-60

[Hog80] Hogart, R.M.: Judgement and Choice: The Psychology of Decision. New York 1980

[KaSl82] Kahneman, D.; Slovic, P.; Tversky, A. (Eds.): Judgement under Uncertainty. Heuristics and Biases. Cambridge 1982

[Kat90] Katz, R.L.: Business/enterprise modelling. In: IBM Systems Journal, Vol. 29, No. 4, 1990, pp. 509-525

[KöHe95] König, W.; Heinzl, A.; Poblotzki, A.: Die zentralen Forschungsgegenstände der Wirtschaftsinformatik in den nächsten zehn Jahren. In: Wirtschaftsinformatik, Heft 6, 1995, pp. 558-569

[KrLi95] Krogstie, J.; Lindland, O.I.; Sindre, G.: Towards a Deeper Understanding of Quality in Requirements Engineering. In: Iivari, K.; Lyytinen, K.; Rossi, M. (Eds.): Proceedings of the 7th Conference on Advanced Information Systems Engineering (CaiSE ’95). Berlin et al.: Springer 1995, pp. 82-95

[LiFi82] Lichtenstein, S.; Fischhoff, B.; Phillips, L.: Calibration of Probabilities: The State of the Art to 1980. In: Kahneman, D.; Slovic, P.; Tversky, A. (Eds.): Judgement under Uncertainty. Heuristics and Biases. Cambridge 1982, pp. 306-334

[Lin94] Lindland, O.I.; Sindre, G.; Sølvberg, A.: Understanding the Quality in Conceptual Modeling. In: IEEE Software, vol. 11, no. 2, 1995, pp. 82-95

[Lor96] Lorenz, K: Sprache. In: Enzyklopädie Philosophie und Wissenschaftstheorie. Hg. von J. Mittelstraß. Bd. 4, Stuttgart, Weimar: Metzler 1996, pp. 49-53

[LoSc75] Lorenzen, P.; Schwemmer, O.: Konstruktive Logik, Ethik und Wissenschaftstheorie. Mannheim, Wien, Zürich 1975

[Luh84] Luhmann, N.: Soziale Systeme. Grundriß einer allgemeinen Theorie. Frankfurt/M.: Suhrkamp 1984

[Luh64] Luhmann, N.: Soziologische Aufklärung. In: Soziale Welt, Vol. 18, 1967, pp. 97-123

[Mey97] Meyer, B.: Object-Oriented Software Construction. 2nd. Ed., Upper Saddle River, N.J.: Prentice-Hall 1997

[MoSh94] Moody, D.L.; Shanks, S.: What Makes a Good Data Model? Evaluating the Quality of Entity Relationship Models. In: Loucopoulos, P. (Eds.): Entity-Relationship Approach - ER'94. Business Modelling and Re-Engineering. 13th International Conference on the Entity-Relationship Approach. Berlin, Heidelberg etc.: Springer 1994, pp. 94-111

[Nil86] Nilsson, N.J.: Probabilistic Logic. In: Artificial Intelligence, vol. 28, 1968, pp. 71-87

[OMG96] Object Management Group: Common Facilities RFP-4: Common Business Objects and Business Object Facility. OMG, TC 13CF/96-01-04, 1996. (http://www.omg.org/library/public-doclist.html)

[Ope94] The OpenDoc Design Team: OpenDoc Technical Summary. Component Integration Laboratories, 1994 (http://www.cilabs.org/)

[Per86] Perrow, C.: Complex Organizations: A Critical Essay. 3rd ed., Glenview, Ill. 1986

[Pet92] Petrie, C.J. (ed.): Proceedings of the First International Conference on Enterprise Integration Modeling Cambridge, Mass.: MIT Press 1992

[Pfe78] Pfeffer, J.: The Ambiguity of Leadership. In: McCall, M.W.; Lombardo, M.M. (Eds.): Leadership: Where else can we go? Durham, NC 1978, pp. 13-34

[Pop35] Popper, K.R.: Logik der Forschung. Wien 1935

[RaAa96] Ragowsky, A.; Ahituv, N.; Neumann, S.: Identifying the value and importance of an information system application. In: Information & Management, No. 2, Vol. 31, 1996, S. 89-102

[Res64] Rescher, N.: Topics in Philosophical Logic. Dordrecht 1964

[Saa96] Saarinen, T.: An expanded instrument for evaluating information system success. In: Information & Management, No. 2, Vol. 31, 1996, S. 103-118

[Sav90] Savage, C.M.: Fifth generation management - integrating enterprises through human networking. Digital Press 1990

[Sch85] Schank, R.C.: The Cognitive Computer. On Language, Learning and Artificial Intelligence. Reading, Mass. 1985

[Sim64] Simon, H.A.: The Corporation: Will It be Managed by Machines. In: Leavitt, H.J.; Pondy, L.R. (Eds.): Readings in Managerial Psychology. Chicago, London 1964, pp. 592-617

[Sim83] Simon, H.A.: Search and Reasoning in Problem Solving. In: Artificial Intelligence, vol. 4, 1983, pp. 181-201

[Sim76] Simons, H.W.: Persuasion. Reading, Mass.: Addison-Wesley 1976

[Sip97] Sipior, J.C.; Ward, B.T.: Sexual Harassment via Employee E-mail: An International Issue or just a United States Quirk? In: Galliers, R. et al. (Hg.): Proceedings of the 5th European Conference on Information Systems. Cork 1997, Vol. I, S. 313 - 320

[SlFi77] Slovic, P.; Fischhoff, B.; Lichtenstein, S.: Behavioral Decision Theory. In: Annual Review of Psychology, vol. 28, pp. 1-39

[SoZa92] Sowa, J.F.; Zachman, J.A.: Extending and formalizing the framework for information systems architecture. In: IBM Systems Journal, Vol. 31, No. 3, 1992, pp. 590-616

[TvKa71] Tversky, A.; Kahneman, D.: The Belief in the "Law of Small Numbers". Psychological Bulletin, vol. 76, 1971, pp. 105-110

[VaDe95] Van der Spek, R.; De Hoog, R.: A framework for a knowledge management methodology. In: Wiig, K. (Ed.): Knowledge Management Methods. Arlington/Texas: Schema Press 1995, pp. 379-393

[WaUn91] Walsh, J.P.; Ungson, G.R.: Organizational Memory. The Academy of Management Review, 16 (1), 1991, pp. 57-91

[Web97] Weber, R.: Ontological Foundations of Information Systems. Melbourne: Coopers&Lybrand 1997

[Wei80] Weick, K.E.: The Social Psychology of Organizing. 2^{nd} ed., Reading, Mass. Et al.: Addison-Wesley 1980

[WeWi97] Wernick, P.; Winder, R.L.: Software Engineering as a Kuhnian Discipline. In: Winder, R.L.; Probert, S.K.; Beeson, I.A. (Eds.): Philosophical Aspects of Information Systems. London: Tayler&Francis 1997, pp. 117-129

[Wii93] Wiig, K.: Knowledge Management. The Central Management Focus for Intelligent-Acting Organizations, vol. 2., Arlington/Texas: Schema Press 1993

[Wit71] Wittgenstein, L.: Philosophische Untersuchungen. Frankfurt/M.: Suhrkamp 1971

[Za87] Zachman, J.A.: A framework for information systems architecture. In: IBM Systems Journal, Vol. 26, No. 3, 1987, pp. 277-293

[Zah75] Zadeh, L.A.: Fuzzy Logic and Approximate Reasoning. In: Synthese, vol. 30, 1975, pp. 407-428