2000

Explanations of Information Systems: Can Philosophy Help?

Steven R. Haynes
London School of Economics and Political Science, s.r.haynes@lse.ac.uk

search

Explanations of Information Systems: Can Philosophy Help?

Steven R. Haynes

Departments of Information Systems and Social Psychology

London School of Economics

Houghton Street, London WC2A 2AE

+44 (0) 171 955 6398

s.r.haynes@lse.ac.uk

Abstract- **Accounts of explanation given in the information systems literature, most prevalent in the areas knowledge-based systems and human-computer interaction, generally do not consider the philosophical underpinnings of explanation concepts. The theoretically rich accounts that have emerged from the philosophy of science provide us with a well-developed framework for designing the structure and content of explanations to be provided for information systems in documentation, in help systems, and in embedded explanation facilities when these are provided. The work described in this paper is an attempt to draw some of these philosophical ideas into the realm of information systems by briefly reviewing four of the major models of explanation from the philosophy of science literature: deductive-nomological explanation, functional explanation, rational choice explanation, and pragmatic explanation. Elements are drawn from each of these models and are related to their potential utility in the information systems field. The approach to philosophy of science's contributions to information systems explanation is from the perspective of information systems research and practice, not philosophy, in an attempt to 'pull', not 'push', these ideas into the realm of information systems development and use.**

Introduction

Until recently, little of the information systems research related to explanation has attempted to address the theoretical underpinnings of what it means to explain. Recent efforts in the field of explanations for knowledge-based systems have begun to develop an account of the theoretical foundations of explanation based on an analysis and integration of empirical research in the field of explanations for knowledge-based systems [1], [2]. This paper attempts to provide an additional theoretical anchor for explanation of information systems, one based on a selection from the vast body of work in the philosophy of science, which has addressed problems related to the structure, content, and purpose of explanations. This is not an account of the metaphysics of explanation in the philosophy of information systems science, though this is an interesting topic in its own right, but an attempt to ground the notion of explanation, in the information systems use context, in the relevant philosophical literature.

Explanation as a topic of research in information systems is generally conducted in the sub-field of knowledge-based systems, in particular, expert systems, e.g. [1], [2]. This paper considers explanation in a broader context, as an essential element of all information systems with purpose being to meet the information needs of their respective users. In this sense, explanations are provided by a wide range of communicatory tools including software analysis and design documentation, software manuals, integrated help systems, system training materials, as well as explanation facilities in knowledge-based systems, decision support systems, and intelligent tutoring systems.

Norman [3] describes the most significant problem of human-computer interaction and human-computer communication as the "Gulf of Understanding" that exists between the computer's understanding of the user, and the user's understanding of the computer. To address this problem, a set of questions must be answered to gauge: what information can be used to fill this gulf, where this information comes from, how it is structured, how is it to be communicated, and to identify the major challenges to the success of the resulting explanation product. Though this information is typically provided via system documentation and help facilities, ultimately, attempts to include explanation facilities in information systems involve engineering human-computer interaction. An explanation-providing system embedded within an information system is expected to contribute to user-system fit and thus to facilitate the user achieving their activity goals. Computer-generated explanations are most often applied in two ways: to support the solution of problems within a well-defined domain, and to impart an understanding of a given field to those who do not possess specialized knowledge.

Species of Explanation

The historical and contemporary literature is rich with the work of philosophers who have struggled with providing a complete and precise definition of explanation. Most disciplines include within their corpus attempts to provide a framework for what constitutes an adequate explanation within their field [4] [5]. Research in the field of computer-generated explanation, and information systems in general, has not, for the most part, referenced this work. Although an explication of the full range of philosophical and other theories of explanation is beyond the scope of the paper, this section will examine a cross-section of the most prevalent formal theories to help inform the development of a more well-grounded conceptual framework for the development of systems capable of generating explanations.

search

First, an essential point must be made about the distinction between formal theories of explanation in the philosophy of science and what counts as an adequate explanation in a practical, operationalized context. A continuum may be described between accounts of 'everyday' explanation given in psychology [6] and philosophical accounts of explanation. Whereas the subtleties of everyday explanations between people are so complex that we may never achieve the required level of implicit knowledge recognition in a computer-based system, the philosophy of science is 'merely' concerned with attempts to describe the structure of a best quality explanation [7]. This paper is an attempt to draw from some of these descriptions of best-practice explanation. Among the attributes that are most often related to the notion of explanation, good explanation or complete explanation in the philosophy of science are explanatory power, predictive content, simplicity, and theoretical unification [8]. Friedman [9] enumerates three essential requirements that a formal theory of explanation should address:

1. *It should be general.*
2. *It should be objective, it should not rely on the strength of particular science that is in fashion.*
3. *The theory should relate explanation to understanding.*

The relationship between formal theories that seek to develop definitions of ideal or even just good explanations and what is required in a computer-mediated explanatory exchange is unclear. This paper sets out to review a selection of these theories, identify their essential elements, and to then consider how these elements might relate to the information needs of IS users.

Deductive-Nomological Explanation

Arguably the most important modern theory of explanation is Hempel's Deductive-Nomological (D-N) or Covering Law Model [10]. Almost all contemporary theories of explanation stand in relation to this work [11], [12]. The D-N Model describes explanation as the identification of antecedent conditions combined deductively with the application of general laws. Within this model, the explanation equation is defined as:

L_i	(general laws)
C_i	(antecedent conditions or facts)
______________	(deductively entails)
P	(the phenomenon to be explained)

The D-N Model describes causal explanation, the general laws and antecedent conditions that make up the explanans, that which does the explaining, are considered to cause the explanadum, that which is explained. In the D-N Model, an explanation must fulfil both logical and empirical conditions of adequacy.

Logical conditions:

R1 The explanadum must be a logical consequence of the explanans.

R2 The explanans must contain general laws that are required for the explanandum.

R3 The explanans must be subject to test by experiment or observation.

Empirical conditions:

R4 The sentences making up the explanans must be true, that is, highly confirmed by the available evidence.

Hempel and Oppenheim [10] provide additional metrics for the effectiveness of explanations that fit their model through the notion of systematic explanation. They describe systematic power as contributing to an explanation or prediction of a theory T as the ratio of the amount of information derivable by means of T to the amount of initial information required for that derivation. Their use of the term systematize refers to the process of constructing explanatory and predictory relationships between data. The structure of the covering law model not only defines explanation, but also suggests why we should value theories that fit the model, that is, because of their predictive ability [8].

One of the major problems with the D-N Model is the determination of what counts as a law. Hempel supplies the following definition.

"By a general law, we shall here understand a statement of universal conditional form which is capable of being confirmed or disconfirmed by suitable empirical findings." [5]

In the social sciences, the notion of laws as central to explanation presents a potential problem given the role of intentionality in psychology and human behavior [4]. Hempel allows that law may be too strong a term, for example, in his discussion of historical analysis, and suggests "universal hypothesis" as an acceptable substitute. Other problems with the D-N model include cases where an event fits the model, but is not explained, and cases where the event is explained but does not fit the deductive structure of the model. A classic example of the former is that the length of a shadow does not explain the height of a flagpole, though the length of the flagpole, the position of the sun, and the laws of geometry can be combined into an explanation to fit the D-N model.

A closely related model of explanation is the class of probabilistic explanations, which attempt to address scenarios where causal events stand in a probabilistic relation to their effect or effects. Examples of probabilistic explanations are often given in medicine, where exposure to an infectious agent is given as the cause and explanation of someone being diseased, but not all those exposed to the infectious agent actually get the disease [12]. Two major versions of probabilistic explanation are Hempel's modification of the D-N Model to account for probabilistic relations, the Inductive-Statistical or I-S model [13] and Salmon's Statistical-Relevance, or S-R model [12].

In contrast to the D-N model, which shows that a given event occurred with absolute certainty, the I-S model shows only that given the laws and antecedent conditions, it was highly probable that the event was to occur [14]. A key component of the I-S model of explanation is the requirement for maximal specificity. This requirement holds that an explanation of the I-S form is invalid unless all statistically relevant facts are included in the set of antecedent conditions. For example, an I-S explanation stating that an aircraft of a certain type with a given amount of fuel had a .9 probability of travelling over 500 miles would be invalid if the additional fact that it was flying into a 30 knot headwind was omitted.

A problem that arises in this account and which is addressed by Salmon [12] in his Statistical-Relevance model, described below, is that this requirement for maximal specificity does not take into consideration the relevance of all of the available facts to the explanation that is provided. The Statistical-Relevance (S-R) model of explanation exists on two levels. The first level, statistical relevance, describes the network or matrix of factors and their associated probabilities that make up the explanans. The second level describes the causal force of this array of factors. The S-R model differs from the I-S model in that it does not include Hempel's requirement that the explanandum be highly probable given the explanans. Salmon's model instead states that a given factor is relevant to the explanans if its presence increases the likelihood of the explanandum event occurring. For example, the probability of the average person (A) developing skin cancer (B) is not equal to the probability that a person who spends five hours each day sunbathing (C) will develop skin cancer.

D-N Explanation in IS

The major contribution of Hempel's Deductive-Nomological model is the role of laws and law-like statements in the explanation equation. Assuming the less rigorous definition of such laws, we may admit a range of facts that can contribute to the explanation of information system operations including: hardware and software constraints, published standards, development and user-interface guidelines well-grounded in empirical studies, laws in the form of regulatory frameworks that may bound the development and operation of systems in safety critical domains, and business rules, for example, accounting standards and intra/inter-organizational business process specifications that affect the finished IS product. The professional, socio-cultural, and organizational standards or norms in which individuals operate may also be considered a form of laws or law-like statements. In the context of information systems development and use, these norms may act to guide individuals in their work. Considered in this way, laws take on an importance in the explanation structure in that they provide a reference structure that helps to answer questions about design decisions that impacted the system model.

Technological artifacts such as information systems and their components have been referred to as objects with a dual nature: one physical and the other functional [15]. The physical dimension of information systems may be seen primarily in the constraints that act upon such a system, for example hardware, software, and telecommunications capabilities. A second element of this physical dimension relates to the ergonomics of a system and its fit with human psychological and physical capabilities. Unlike many other engineering disciplines where material constraints play a major role in the form of the finished artifact, information systems, especially software, are bound by few such physical constraints. Logical, intellectual, and social interaction (e.g. team dynamics) constraints play a far more prominent role in the shape of the finished IS product.

Many information systems development practices are driven by the application of standards produced either internally to guide the development project process and deliverables, or externally in cases where regulatory documents play a role in the development process (e.g. safety-critical systems and systems developed to government contract specifications). Consider an IS product developed for the European Community (EC) market and subject to EC standards related to ergonomics. This sentence consists of the antecedent conditions, that the software is being developed for the EC market, and when combined deductively with the law-like statement that is the EC published ergonomic standard, we have an explanation of why certain decisions were made regarding the functionality of the resulting system.

Hempel [5] argues that explanations in the social sciences (specifically, history and sociology) often exclude statements of general laws. They are considered implicit and relate to individual or social psychology that, it is assumed, are familiar to the person to whom the explanation is being given. Though laws may be excluded from the explanation structure, they can be called upon to add credibility or justifying grounds to the assertions of the explanation itself [12]. In this way, laws relate to the design rationale, or argument structure, which can provide a robust description of the reasons why a system operates as it does. Design rationale [16], founded on Toulmin-like argument structures [17], is a form of system design documentation that attempts to model the entire design space including all of the design questions and alternatives considered as well as the criteria applied to their acceptance or rejection. In a given system design space, which is structured according to design rationale formalisms, laws may take the form of a Toulmin warrant backing, where well-grounded facts (laws) are used as an authoritative guide in the application of design criteria.

Functional Explanation

Functional explanations attempt to provide arguments for the existence or persistence of entities (objects, events, or institutions) by reference to the effects, generally the beneficial effects, of those entities. Arguments for the validity of functional explanation are controversial [18] and are especially prevalent in the fields of biology and in the social sciences [11]. The attraction of functional explanation, especially in social science, is based on the assumption that phenomena must have some meaning, that things exist and events occur because of their beneficial consequences [19]. Functional explanations are especially relevant to artifacts, where it is assumed that human design activities and workmanship are applied to serve some purpose.

Dore [20] has examined the concepts of function and cause in an attempt to identify ways in which an analysis of functions can be translated into a statement of causality. In his analysis he differentiates between the cause of something (in his example, social institutions) occurring and something persisting. There is a distinction between the causes of the events that lead to something being created, and the causes of the events that lead to something persisting. This analysis echoes ideas from sociobiology, such as Dawkins' memes [21] and Sperber's cultural cognitive causal chains [22]. Dore argues that it is possible for causal events to be identified as functional. Dore's example involves a discussion of the development of Chinese communes, stating that it is probably possible to find the minutes of the meetings and other documentation of events that led to their creation where the creation was in constant reference to the intended consequences of the communes, their function.

Several arguments against the idea of functional explanation are outlined in [18] and include problems such as a lack of supporting evidence for the mechanisms by which certain features or practices exist. Since these are often not identified, functional explanations are not generally falisifiable. A second problem he identifies is that functional explanations do not show how a given feature came to exist, only how it persists. Against the former, Little [14] argues that in a complete functional explanation, the causal feedback process, the missing mechanism highlighted by Kincaid, must be identified in order for the explanation to be fully coherent. Against the latter, Little urges us to consider a given feature in terms of its current effects, not any future effects that the feature might have. This contributes to the arguments of teleology that are often raised against the validity of functional explanations, a feature's current beneficial effects are what cause it to persist and to disperse, any beneficial future effects are an outcome of this persistence and dispersal.

Functional Explanation in IS

Functional explanations are perhaps most obvious in their relation to the development of technological artifacts such as information systems [14]. They are one of Kroes' [15] two aspects of technological explanations: the physical and the functional. Information systems are (usually) developed to solve a particular problem. Their structure is related to their end goals and the features and processes the system must implement in order to achieve these goals. Kroes describes the design process as that which translates the commercial requirements for an artifact into a description or blueprint for its physical structure. During this process the functional aspects of the artifact are translated into a set of physical characteristics that will achieve the desired functional, and therefore commercial, goals. While Kroes claims that there is little relation between the physical description and the functional description of the artifact at the end of the design process, he argues that a complete design must include a technological explanation which he describes as "an explanation of the function of a technological object in terms of the physical structure of that object" [15].

Kroes excludes from his analysis an important element of a software system design, namely, a logical description of the artifact as distinct from its physical structure and its functional description. In a software system design, physical structure and physical constraints are largely transparent to the designer (except in cases such as operating systems, network control software and other device-dependent systems). Though the designer is likely to employ some physically determined heuristics during the design process, for example, minimize disk storage, CPU processing, and network bandwidth usage, etc., the constraints with the greatest impact on the software design process are logical. These logical constraints are in turn determined by a myriad of factors. The limits of the mind and imagination of the software designer is one such constraint. This constraint is compounded in team design processes where communication of the logical design becomes a constraint [23]. Another important constraint is the ability of the end user to understand and effectively use the artifact that results from the design. Organizational and business factors, in particular cost, also introduce constraints that become important to a comprehensive technological explanation of a designed system.

Rational Choice Explanation

Rational choice theories of explanations are based on work in economics and attempt to describe an entity's behavior (including human behavior) in terms of the perceived benefit of that behavior relative to other possible behaviors. As pointed out by Simon [24], such theories of behavior are simplistic as best given the ambiguities of the beliefs and desires that purportedly drive rational choice decisions. Davidson [25] argues that rationalisations, an actor giving reasons to explain why a particular action was taken, are a type of causal explanation. In particular, "the primary reason for an action is its cause". Central to Davidson's argument is the idea that primary reasons express the intention of the person performing the action. When we give a reason for taking an action, and the reason is true, we expose the belief or attitude that is causing the action. Davidson counters arguments that reasons are not causes of actions by virtue of their conjunction by appeal to the events that led to the formation of the beliefs that underlay the reasons. These beliefs constitute a causal chain, events lead to belief formation, these beliefs, now given as reasons, cause an actor to take particular actions.

In reference to the baseline D-N model of explanation, Davidson allows that causal explanations based on reasons do not necessarily refer to laws (he prefers the weaker claim that causal explanations involve laws rather than entail them) nor are they necessarily purely deductive. Competing beliefs that impact choice of action in a given scenario are selected based on a weighting that is not necessarily predictable or even reliably describable, and the laws that affect a given reason may be so far removed from the belief as to be irrelevant to the explanation.

Rational Choice Explanation in IS

Despite the problematic nature of rational choice theories of explanation, individuals and organizations do, sometimes, make decisions based on their perceptions of the relative value of the potential outcomes. In an information systems development context, a vast array of interconnected and sometimes competing beliefs and desires may interact in the course of a given design decision. Capturing the elements of such an array is probably impossible given the shear number of them that may be present in a group decision context and the fact that many or even most are tacit, inaccessible even to the individual who holds them [26]. Nonetheless, not all decisions are this complex and the primary decision factors can be made explicit within the normal functioning of an information systems development team. In such cases, rational choice explanations do fill a gap in the information systems explanation equation by allowing for the goal maximizing behavior, however simplified by constraints of access and capture, of the individuals as they work in teams.

Pragmatic Explanation

Pragmatic theories frame explanation as "an interest-relative notion" [27]. One of the most vociferous proponent of pragmatic accounts of explanation is Bas van Fraassen [28], [29]. His view is that explanation is best viewed as answering a given why-question from a particular aspect within a particular context. Van Fraassen claims that in answer to the question "why is the light on", both "because I flipped the switch" and "because we are expecting company" are both explanations, depending on the rationale for the original question. To van Fraassen [29],

"Which factors are explanatory is not decided by features of the scientific theory but by concerns brought from outside."

He offers a foil to the hypothetical question in which it is asked whether the length of a shadow explains the height of a pole, most would answer that indeed it does not. However, in van Fraassen's example, the length of the shadow does explain the height of the pole, if the pole is a pointer on a sundial with the desired property that it cast a shadow of a given length. Van Fraassen argues against the idea that explanations can be conceptually analyzed or evaluated based on innate notions such as simplicity, predictive strength, truth, or empirical adequacy that deny the role of context.

Van Fraassen also argues for the centrality of contrast classes and relevance relations in explanation constructions. Contrast classes provide information on why a particular event occurred instead of or in relation to another in its contrast class [8]. In contrast to the similarly named concept employed in Salmon's S-R model, van Fraassen's relevance relation describes events that relate to the event to be explained in terms of the relevance of those events to the purposes of the explanation.

Pragmatic Explanation in IS

One of the most significant problems with explanatory dialogues between humans and machines is the machine's inability to establish the background knowledge possessed by the human user. A key component of human-to-human explanation is that people are able to quickly establish at least a portion of this background knowledge and apply it to the construction of an explanation [30]. Antaki [31] argues that the single most important aspect of an explanation given by an intelligent system is that it be believable. He suggests that in order to be believable, a system providing an explanation must know what the user does not know, know what the user wants to know, know why the user wants to know it, and know what the user will accept as an explanation. Researchers have attempted to address these explanation dialogue and credibility issues through user modeling and attempts to refine the dialogue management capabilities of explanation facilities, e.g. [32] [33]. Though some successes have been achieved in the laboratory, little of this work has been adopted as standard industry practice.

A large proportion of research into computer-generated explanation consists of attempts to meet the challenges of providing interest-relative explanations to end-users with different backgrounds, skills, and knowledge states. These efforts represent an attempt to get to the heart of the human-to-human explanation problem. Though pragmatic theories of explanation highlight the importance of audience and context in any account of explanation, the problems inherent in attempting to approximate human expertise in explanation provision suggest that rather than attempting to mimic human, researchers in computer-generated explanation should consider redirecting their efforts such that the strengths of computers and software systems are best leveraged. As Gregor and Benbasat [1] point out, established theories related to cognitive effort suggest that users of an information system will not use explanation facilities unless they perceive the cost of retrieving and processing these explanations to be outweighed by the benefits. Here, the importance of the 'relevance relation' as described by van Fraassen and Salmon becomes critical. Established models of computer system documentation suggest that users of information systems make parsimonious use of computer system documentation [34]. They seek out information that will assist them in achieving their immediate aims, and quickly give up their search if they do not succeed quickly, preferring instead to explore the systems functionality until they are able to achieve the desired end.

The idea that contrast classes, or what could have been rather than what is, is one of the central points in van Fraassen's account of explanation [28]. Contrast classes are used in explanation to vector a concept in relation to its possible alternatives. Providing an account of why a system feature or process was implemented in a certain way is possibly best expressed in terms of the alternatives and the reasons why they were not selected for implementation. Again, this account bears a close resemblance to the design rationale-based model of system documentation where not only are the end products of design decisions captured, but all of the options considered in the process.

The problem providing a computer system with the ability to recognize context has been highlighted in the literature, perhaps most effectively by Dennett's paper, which described the now famous Frame Problem [35]. This problem can be somewhat mitigated by providing explanations in a format in which the explanation user has control of the information that they extract based on their own knowledge of context rather than the system's. One established way of presenting information in this way is through hypertext, where information is organized into categorized 'trees' that present information in successive levels of detail and allow branching based on immediate information needs. Information structured in this way allows users of information systems to make their own selections based on relevance to current information needs.

Despite the intuitive attraction of providing for the pragmatics of explanation delivery, significant problems exist in achieving this goal. Research into user modeling may yet uncover ways of finessing explanatory dialogue in the systems context [32], [33], but significant barriers to this goal remain. In the user modeling area, the level of complexity involved in classifying users according to some system-use related criteria presents a daunting challenge [36] and some researchers even question the desirability and ethics of capturing the amount of personal information required for the task [37].

search

Discussion

A recent and well-developed account of the nature of explanation in knowledge-based information systems is [1]. In this account, the authors survey the field of explanation in knowledge-based systems and relate explanation facility implementations to theories of cognitive effort, cognitive learning, and Toulmin argument structures. The focus of their work is on the desirability of explanations in intelligent systems, the benefits that accrue to users of these systems when explanations are provided, and the types of explanations that are most useful in the context of human-computer interaction. Explanatory content is one element of explanation type that they identify and they provide a taxonomy of explanation content types which includes:

1. *Trace or line of reasoning content*
2. *Justification or support content*
3. *Control or strategic content*
4. *Terminological content*

This taxonomy was derived from the explanation in knowledge-based systems literature and is a useful framework to apply when considering the nature of explanatory content in the information systems context. However, the thesis of the current paper is that explanation content may be considered from a different, more basic perspective, one derived from the body of work on theory of explanation in the philosophy of science. The motivation for this review is to provide additional theoretical grounding for our efforts at improving user understanding of system design and functionality with the end goal being a contribution to improved human-computer interaction.

The role of causality is considered first given its prominence in the philosophical literature. According to Clancey [38] "causal links" are the most interesting and difficult to explicate components of a system (in this case, an expert system) since they are related to the underlying processes of the application. Causal links in the context of IS use refer not only to the technical causal chains, i.e. inter-modular and inter-application functionality, but to the complexities inherent in the development of designed artifacts within a social and organisational framework. Whether a given system, feature, or process is 'caused' by application of a law, by the need to implement a given functionality, or by goal-maximizing behavior on the part of members of the development team or their user community, all of the explanation models discussed here, with the exception of the pragmatic theory, have a fixed causal basis. In the case of the pragmatic theory, the explanation may be causal but the depth of the given cause is determined only by the goals of the explanation giver in the context of a given explanation request.

Dreyfus and Dreyfus [39] argue that in order to explain how something works, it must first be broken down into the particle level, the component level, and the functional level. This mechanistic, reductionist approach to explanation structure involves showing how a device's components add up to produce the functioning of the whole device. This view of explanation is in contrast to the holistic model. Holistic accounts apply to information systems in cases such as functional or pragmatic explanations where the purpose of the system being explained or the purpose of the explanation itself, overrides the utility of knowing about causal interaction of parts. In the field of AI many expert system researchers have argued that for a system to provide sufficient explanations, it must include knowledge of first principles in its application domain [40]. These first principles constitute the "ultimate explanations" upon which an explanation system must rely and relate closely to the laws and law-like statements required by Hempel.

Based on this explication of a selection of formal theories of explanation from the philosophy of science, a set of principles may be derived to contribute to a conceptual framework for evaluation of explanations and explanation-generation technologies in information systems. This selection consists of the following explanation elements:

The elements below…	Help provide explanations for information systems by…
Laws and Law-like Statements	Relating design decisions to established facts including physical laws, standards, norms, and other reasonably well established universals.
Contrast Classes	Relating design decisions to options that were not adopted, to what could have been rather than what is.
Relevance Relations	Limiting the explanation to elements with direct relevance to the system feature of interest. In an information retrieval context, such relevance relations may be defined through a probabilistic function.
Rational Choices	Showing how the development team and other interested stakeholders affected the form of the system based on their beliefs and desires, which may themselves relate to other explanatory elements.
Functional Purpose	Relating system design decisions to the motivation for the system and its individual features and processes, i.e. the purpose they are meant to serve in the system and domain context. Information systems are designed and developed to support users as they perform tasks in an application domain within an organizational and social context. Explanations of information systems and their features should help users relate design decisions to the underlying motives behind these decisions within the domain and context.
Relative Interests of the Explanation User (Pragmatics)	Help to maintain parsimony in explanation delivery by accounting, as much as possible, for the reasons for the explanation request and interests of the explanation user. Consider the context in which an explanation is provided and map explanatory content to the tasks and goals of the information system user.

Each of these elements contributes something unique and necessary to a complete, best practice explanation of an information system, a feature, or an information

search

transformation process. Explanations of this level of sophistication and detail are necessary given the complexity of the information systems development and use context. They also contribute to the tractability of information necessary for systems that attempt to include explanation-generation capabilities.

Experts who do not possess the ability to draw on a base of increasingly detailed and experiential information in the course of providing an explanation are frequently the subject of credibility problems, especially in safety-critical domains or those that are considered to have a rich, intellectual content [41]. In order to successfully explain a concept at some level of expertise, a clear model of the domain knowledge and a rigorous understanding of the problem solving process is necessary [42]. By drawing on established laws or their empirically well-grounded equivalents in a given field, we help to establish this credibility in the explanations that we provide.

Conclusion

As human-designed artifacts that emerge from a complex socio-technical, organizational, and psychological context, information systems are subject to the web of decision making that characterizes the artificial. In the design of an information system, we hope that this web consists of a series of rational decisions made with an eye to the desired functionality of the resulting artifact and the extent to which it achieves a good fit with the work context and task structure of the target user community. However, the very complexity of this web results in information needs that are rarely met in the context of information systems development and use.

This framework may be further developed, especially within the pragmatics of explanation delivery area, by referencing work in psychology related to theories of everyday explanation [6]. However, before the details of explanation delivery are finessed, we require a rich, theoretically grounded framework to describe the basic elements of explanation and their purpose in the explanation equation. Without such a framework, we are left with little in the way of fundamental, conceptually rigorous criteria for the evaluation of the explications we provide for our information systems products.

References

[1] S. Gregor and I. Benbasat, "Explanations From Intelligent Systems: Theoretical Foundations and Implications for Practice," *MIS Quarterly*, December 1999, in press.

[2] J.S. Dhaliwal and I. Benbasat, "The Use and Effects of Knowledge-based System Explanations: Theoretical Foundations and a Framework for Empirical Evaluation,", *Information Systems Research*, 7, 3, 1996, pp.342-362

[3] D.A. Norman, "Cognitive engineering," in *User Centred System Design: New Perspectives on Human-Computer Interaction*, D.A. Norman. & S.W. Draper, Eds., Hillsdale, NJ : Erlbaum

[4] D. Davidson, "Psychology as Philosophy," in *Readings in the Philosophy of Social Science*, M. Martin and L.C. McIntyre, Eds., Cambridge, MA: The MIT Press, 1994, pp.79-89

[5] C.G. Hempel, "The Function of General Laws in History," in *Readings in the Philosophy of Social Science*, M. Martin and L.C. McIntyre, Eds., Cambridge, MA: The MIT Press, 1994, pp.43-53

[6] C. Antaki, Ed., *Analysing Everyday Explanation: A Casebook of Methods*, London: Sage, 1988

[7] S.W. Draper, "What's going on in everyday explanation?," in *Analysing Everyday Explanation: A Casebook of Methods*, C. Antaki, Ed., London: Sage, 1988, pp.15-31

[8] P. Gasper, "Causation and Explanation," in *The Philosophy of Science*, R. Boyd, R., P. Gasper, and J.D. Trout, Eds., Cambridge, MA: The MIT Press, 1991, pp.289-297

[9] M. Friedman, "Explanation and Scientific Understanding," in *Theories of Explanation*, J. Pitt, Ed., New York: Oxford University Press, 1988, pp.188-198

[10] C.G. Hempel and P. Oppenheim, "Studies in the Logic of Explanation," in *Theories of Explanation*, J. Pitt, Ed., New York: Oxford University Press, 1988, pp.9-50

[11] D-H. Ruben, D-H., "Introduction," in *Explanation*, D-H. Ruben, Ed., Oxford: Oxford University Press, 1993, pp.1-16

[12] W.C. Salmon, "Scientific Explanation and the Causal Structure of the World," in *Explanation*, D-H. Ruben, Ed., Oxford: Oxford University Press, 1993, pp.78-112

[13] C.G. Hempel, "Aspects of Scientific Explanation," in *Aspects of Scientific Explanation and Other Essays*, New York: Free Press, 1965

[14] D. Little, *Varieties of Social Explanation: An Introduction to the Philosophy of Social Science*, Boulder, CA: Westview Press, 1991

[15] P. Kroes, "Technological Explanations: The Relation Between Structure and Function of Technological Objects," *Society for Philosophy & Technology*, 3, 3, 1998

[16] J.M. Moran and T.P. Carroll, T.P., Eds., *Design Rationale: Concepts, Techniques and Use*, Mahwah, NJ: Lawrence Erlbaum, 1996

[17] S. Toulmin, *The Uses of Argument*, Cambridge, UK: Cambridge University Press, 1958

[18] H. Kincaid, "Assessing Functional Explanation in the Social Sciences," in *Readings in the Philosophy of Social Science*, M. Martin and L.C. McIntyre, Eds., Cambridge, MA: The MIT Press, 1994, pp.415-428

[19] J. Elster, "Functional Explanation: In Social Science," in *Readings in the Philosophy of Social Science*, M. Martin and L.C. McIntyre, Eds., Cambridge, MA: The MIT Press, 1994, pp.403-414

search

[20] R.P. Dore, "Function and Cause," in *Readings in the Philosophy of Social Science*, M. Martin and L.C. McIntyre, Eds., Cambridge, MA: The MIT Press, 1994, pp.377-389

[21] R. Dawkins, *The Selfish Gene*, *New Edition*, New York: Oxford University Press, 1989

[22] D. Sperber, *Explaining Culture: A Naturalistic Approach*, Blackwell, 1996

[23] F.P. Brooks, *The Mythical Man-Month: Essays on Software Engineering*, Reading, MA: Addison-Wesley, 1975

[24] H.A. Simon, *The Sciences of the Artificial, 3rd Edition*, Cambridge, MA: The MIT Press, 1996

[25] D. Davidson, "Actions, Reasons and Causes," in *Readings in the Philosophy of Social Science*, M. Martin and L.C. McIntyre, Eds., Cambridge, MA: The MIT Press, 1994, pp.675-686

[26] D.A. Schön, *The reflective practitioner: How professionals think in action*, New York: Basic Books, 1983

[27] H. Putnam, *Meaning and the Moral Sciences*, London: Routledge & Kegan Paul, 1978

[28] B.C. van Fraassen, "The Pragmatic Theory of Explanation," in *Theories of Explanation*, J. Pitt, Ed., New York: Oxford University Press, 1988, pp.135-155

[29] B.C. van Fraassen, "The Pragmatics of Explanation," in *The Philosophy of Science*, R. Boyd, R., P. Gasper, and J.D. Trout, Eds., Cambridge, MA: The MIT Press, 1991, pp.317-327

[30] T. Winograd and F. Flores, *Understanding Computers and Cognition: A New Foundation for Design*, Norwood, NJ : Ablex, 1986

[31] C. Antaki, "Lay explanations of behaviour: how people represent and communicate their ordinary theories," in *Expert Knowledge and Explanation: the knowledge-language interface*, C. Ellis, Ed. Chichester : Ellis Horwood, 1989

[32] A. Cawsey, *Explanation and Interaction: The Computer Generation of Explanatory Dialogues*, The MIT Press: Cambridge, MA, 1992

[33] J.D. Moore, *Participating in Explanatory Dialogues: Interpreting and Responding to Questions In Context,* Cambridge, MA: The MIT Press, 1995

[34] J.M. Carroll, *The Nurnberg Funnel: Designing Minimalist Instruction for Practical Computer Skill*, Cambridge, MA: The MIT Press, 1992

[35] D.C. Dennett, "Cognitive Wheels: The Frame Problem of AI," in *The Philosophy of Artificial Intelligence*, M.A. Boden, Ed., Oxford : Oxford University Press, 1990

[36] R.B. Allen, "User models: theory, method, and practice," *Int'l. J. Man-Machine Studies*, 32, 1990, pp.511-543

[37] D. Browne, D., "Conclusions," in *Adaptive User Interfaces*, D. Browne, P. Totterdell, and M. Norman, Eds., London : Academic Press

[38] W.J. Clancey, "The Epistemology of a Rule-Based Expert System - A Framework for Explanation," *Artificial Intelligence*, 20, 3, May, 1983, pp.215-251

[39] H.L. Dreyfus and S.E. Dreyfus, *Mind Over Machine,* Oxford : Basil Blackwell, 1986

[40] F. Hayes-Roth, D.A. Waterman, and D. Lenat, *Building Expert Systems*, Reading, MA: Addison-Wesley, 1983

[41] B. Chandrasekaran and S. Mittal, "Deep versus compiled knowledge approaches to diagnostic problem-solving," in *Int'l. J. Man-Machine Studies*, 19, 1983, pp.425-436

[42] D.W. Hasling, W.J. Clancey, and G. Rennels, "Strategic explanations for a diagnostic consultation system," *Int'l. J. Man-Machine Studies*, 20, 1984, pp.3-19