Information in the Study of Human Interaction

Keith Devlin* and Duska Rosenberg†

May, 2006

Information as an analytic tool

This chapter describes one way that information — as a conceptual entity — may be used (by an analyst, as a tool) in a study of human interaction. (Actually, most of what we say will apply to interaction in general, for instance human–machine interaction, but our examples will be taken from human interaction.) The "analyst" here may be a professional social scientist (as is the case for our main technical example), or could be an ordinary person trying to make sense of a particular interaction. When applied to such latter cases, our article also provides insight into much of the common talk about "information" that takes place in today's "information society", and in that way our essay can be viewed as an analysis of the rational structure that lies behind (and is implicit in) the modern, information-oriented view of the world.

To give a very simple example, suppose Alice (A) issues the instruction "Sit down" to Bill (B). We may view this as an attempt by A to achieve a particular action by B. A makes this attempt by herself carrying out a particular action, namely uttering certain words. The analysis could proceed by examining why A chooses the particular words she does, why B interprets those words the way he does, and what action B carries out as a result and why. Typically, this might be done by identifying social norms that describe (or prescribe) how people use language to achieve their ends. (An example we shall examine in some depth later in the paper will show just how such an analysis may proceed.)

But there is another way we could analyze the same interaction; namely as being mediated by the transmission of information from A to B. In the alternative, information-based approach, we analyze Alice and Bill's interaction in terms of the issuance of certain information by A, its reception by B, and the consequences of this transmission in terms of the actions of the two participants.

What is gained (or lost, or obscured) by the introduction of the mediating notion of information? Which (if any) approach is better (for what purpose), and why?

An analogy might help to explain the distinction between the two approaches.

*CSLI, Stanford University, devlin@csli.stanford.edu.
†Royal Holloway University of London, D.Rosenberg@rhul.ac.uk

Suppose we want to study a wrestling match between two people. Then we would most naturally analyze the interaction in terms of the forces each exerts on the other. In contrast, if we want to examine a game of tennis between the two individuals, it is more appropriate (and surely more productive) to look at the way the ball is batted from one to the other. Why, in the second case, do we not analyze the game in terms of the forces each player exerts (through the racket) on the ball? After all, the ball is an inert object; the entire play of the game is dictated by the actions of the two players, just as it is in the wrestling match.

The reason we analyze the tennis match in terms of the motion of the ball, is precisely that the ball does indeed *mediate* between the actions of the two players in the tennis match. Mediation of a human–human interaction, even if by an inert object, changes things sufficiently that a framework appropriate for analyzing one form of interaction may be unsuitable for analyzing another. This is why newspaper accounts of tennis games typically include descriptions of the motion of the ball as well as the two players.

In the case of a human–human linguistic interaction, however (such as the "Sit down" example we just gave), we seem to have an entirely free choice between two different forms of analysis. We can adopt one of several traditional (non-information-based) approaches, focusing on the (descriptive or prescriptive) rules and protocols that describe or prescribe how interaction is done, the choice of words each participant makes, and the way each understands the words spoken by the other. This corresponds to the way we analyze the wrestling match, where we look at the various capacities each participant brings to the encounter and the manner in which those capacities result in the physical interaction that ensues. Or we may equally well consider the linguistic interaction as a transmission of information. This would correspond to our analysis of the tennis match, with the information passed from one person to the other at any one stage being the analogue of the tennis ball.

Of course, as with any analogy, it is important to recognize the limitations of the comparison. In the case of a tennis game, the same ball gets passed back and forth between one player and another; in human interaction, considered as mediated by an exchange of information, *different* information is conveyed at each stage.[1] In a typical human *linguistic* interaction (such as a conversation), for instance, there is something physical passed from one participant to another at each stage, namely the individual utterances (tokens); but these are not the information, rather they (can be said to) *carry* the information. Part of any formal account of information exchange has not

[1]Interaction also involves feedback – implicit information – that helps A and B coordinate their actions, which is necessary if they work together on a shared task, or perform any kind of joint action. Conversation can be viewed as a joint action whereby participant establish mutual understanding, or, in the words of Clark [2], "common ground". However, joint action does not always involve using language. Imagine two people carrying a plank. Each of the individual movements is felt through the movement of the plank, which can be said to carry information about the participants' moves. We can also call this "feed-through" (Dix in [9]). In this context the medium plays an important part. If, instead of the plank, the two people carry a mattress, the medium does not transmit information about their movements in the same way. The characteristics of information in interaction and joint action therefore depend significantly on the medium or the way the interaction is mediated.

only to include a definition of information, but also provide a mechanism for how tokens can in fact carry information. (The theory we make use of in our account, situation theory, does just that.) The purpose of the analogy is to distinguish between the conceptualization or analysis of a wrestling match as an *unmediated interaction* and a tennis match as being *mediated by a neutral object*, namely the ball.

Is one approach better than the other, and if so how? The answer is that each offers advantages the other does not. For some purposes, a descriptive analysis is better, on other occasions the information-based approach is more suitable. In some cases, carrying out *both* forms of analysis may result in greater understanding.

This distinction between the two analytic approaches is not unlike the one that arises in several different guises in physics, between "action at a distance" and the transmission of a particle. For example, do we think of gravity in terms of geometric distortion of space–time or as the transmission of gravitons? Again, is light a wave (a perturbation in the fabric of space–time) or a particle (a photon)?[2]

The distinction is not merely one of theoretical interest to the analyst; it gets at a fundamental feature of the way we conceive of and live in our current world. Today, much of our everyday thinking, writing, and talk about human activities is couched in terms of information. Yet, this way of talking about the world is relatively recent. The change was brought about largely by the development of various communication technologies — printing, the newspapers, postal services, dictionaries and reference books, radio, telephone, television, photocopiers, the Internet — that, by mediating human–human interaction, made possible (indeed encouraged) an information-based (tennis game or particle) way of thinking about communication.[3] We say a little more about the development of the modern, popular concept of information in just a moment.

Information

As indicated by some of the other articles in this collection, the word "information" has several different meanings, including a fundamental entity (closely related to entropy) that exists in the universe, a measure of order in the universe, a number of (different) mathematical concepts, and the less precise but more common, everyday (and, of particular relevance to this article, socially constructed) notion implicit in terms such

[2]Our mathematical treatment of information, described later, takes this analogy a step further by regarding information as made up of discrete items called "infons." Indeed, the invention of that word, by Devlin, was motivated entirely by that analogy.

[3]The use of a concept of information as a mediator is not restricted to communication. In human action and interaction (especially in computer-mediated communication), we are not talking only about transmission of information. We define the concept of mediation further, to include sharing of information as well. This refers in particular to the use of information to coordinate action, express communicative intent, and ultimately create trust, identity with a group or a community, and shared culture, all of which are essential features of social life. Sharing information is different from exchange and utilization; in particular, sharing is more profoundly social than transactional. When we exchange information, nothing changes unless the exchange causes some kind of change in the cognition of individuals involved. When we share information, then the information that is shared changes, because the act of sharing gives rise to new and different information.

as "information desk," "departure/arrivals information," and "Can you give me some information about renting bicycles in Amsterdam?"

In this article, we take as our concept of information the socially constructed, everyday notion mentioned last in the above list. In the more technical part of the paper, we shall make that everyday notion a little more precise by way of a mathematical definition, and use that additional precision to examine in some detail the way that information may be used to analyze human interaction (and, more generally, human action).

We make no attempt to provide a comprehensive overview of the topic. The field is far too broad for any short survey such as this to come even close to completeness. Rather we shall outline the main themes and *illustrate* the way information can play a role in an analysis of a social phenomenon.

It will be helpful to begin with a few brief (and hence simplified) remarks about the origins of the notion of information we shall focus on.

Prior to the nineteenth century, the word "information" (which first appeared in the English language in the fourteenth century) was used to refer to a knowledgable or informed individual. For example, the term "man of information" would translate into modern English as "man of learning", "well educated man", or "well informed man".

During the nineteenth century, the generally accepted conception of information shifted from something possessed by an individual (if indeed it was conceived as some*thing* that could be "possessed") to one of a public commodity — something (and in this case definitely a *thing*) that could be shared. The cause of this shift in meaning can be traced to the growth of communication technologies, in particular the publication of mass market newspapers in the early eighteenth century and onwards. With the appearance of newspapers, and also dictionaries, encyclopedias, and general reference books and the introduction of postal services, the telegraph, and later the telephone, it was possible to identify (or conceive of) a "substance" or "commodity" that could be called "information".

That substance was largely autonomous, having an existence outside the individual human mind. It came in identifiable chunks. For instance, newspapers impose the same physical structure (a block of text within a very small range of size) on every topic reported on, be it politics, war, sport, theater, science, or whatever. Moreover, the organizations that produce newspapers, reference books, and the like provide an institutional "stamp of approval" on the information they impart, giving it the air of being neutral, free of bias and personal perspective or interpretation — "the truth."

The nineteenth century concept of information was thus an itemized one that was largely identified with its representation. It became possible to talk in terms of "amount of information." Information was also true; otherwise it would be called *mis*information.

With the rise of itemized, autonomous information, it was no longer appropriate to use the term "information" to describe personal facts. For instance, only in very special circumstances would a person today say "Alice provided me with the information that

she enjoyed last night's movie." Rather one might say "Alice told me she enjoyed the movie," this fact neither being public property nor having an "institutional stamp of authority" that would grant it the status of information.

With the nineteenthc century shift in meaning, information also came to be viewed not as the *result* of a person being informed, but its *cause*.[4]

The modern everyday conception of information is different again. Whereas the nineteenth century notion was closely tied to the "containers" of information that gave rise to the notion — the books, encyclopedias, newspapers, etc. — the concept of information that arose around the middle of the twentieth century *transcends its representation*. Moroever, whereas nineteenth century information was, by definition, true, the same cannot be said for today's concept.

The modern notion of information did not fully develop until the 1970s, although the beginnings of the shift can be seen as far back as the 1940s. Like its predecessor, this new notion also developed as a result of changes in communication technologies — in this case the development of the digital computer and the growth of the many associated electrical and electronic "information and communication" media that are now part of our everyday lives.

Today, most of us think of information as a commodity that is largely independent of how it is embodied. It can be bought, sold, stolen, exchanged, shared, stored, sent along wires and through the ether, and so forth. It can also be *processed*, using *information technologies*, both concepts that would have sounded alien (and probably nonsensical) to anyone living in the nineteenth century, and even the first half of the twentieth.

The separation of information from its various representations is what made it possible for contemporary technology guru Ted Nelson to make his oft-repeated observation "Paper is just an object that information has been sprayed onto in the past."

The way present-day society conceives of information today is well captured by the following passage from *Business Week* (special issue on "The Information Revolution") in 1994:

> We can glean it [information] from the pages of a book or the morning newspaper and from the glowing phosphors of a video screen. Scientists find it stored in our genes and in the lush complexity of the rain forest. And it's always in the air where people come together, whether to work, play, or just gab.

It is the use of today's concept of (disembodied) information as a means to understand (and, when done more formally, analyze) human interaction that is the subject of this paper.

[4]Nunberg [7], a good reference for much of the present discussion, observes that a similar shift in meaning occurred when the terms *mystery* and *horror* began to be used to describe literary genres.

How does information arise?

A fundamental question to be answered at the start is, how is it possible for something in the world, say a book or a magnetic disk, to store, or represent, information? This question immediately generalizes. For, although we generally think of information as being stored (by way of representations) in things such as books and computer databases, any physical object may store information. In fact, during the course of a normal day, we acquire information from a variety of physical objects, and from the environment.

For example, if we see dark clouds in the sky, we may take an umbrella as we leave for work, the state of the sky having provided us with the information that it might rain. On Halloween night in North America, a light on in the porch provides the information that it is acceptable for children to approach the house and ask for candy; no light indicates that the householders do not want to be disturbed. In rural parts of North America, setting the flag on the mailbox in the upright position indicates to the mail carrier that there is outgoing mail to pick up.

How can an object or a collection of objects encode or represent information? How can part of the environment encode or represent information? For instance, how does smoke provide information that there is fire, and how do dark clouds provide information that it is likely to rain? Part of the explanation is that this is the way the world is: there is a systematic regularity between the existence of smoke and the existence of fire, and a systematic regularity between dark clouds in the sky and rain. Human beings and other creatures that are able to recognize those systematic regularities can use them in order to extract information. The person who sees dark clouds can take an umbrella to work, the animal that sees smoke on the horizon can take flight.

Notice that we are definitely talking about information in these examples, not what the information is about. For example, people or animals that see smoke do not necessarily see fire, but they nevertheless acquire the information that there is a fire. And the sight of dark clouds can provide the information that rain is on the way long before the first drop falls.

In general then, one way information can arise is by virtue of systematic regularities in the world. People (and certain animals) learn to recognize those regularities, either consciously or subconsciously, possibly as a result of repeated exposure to them. They may then utilize those regularities in order to obtain information from aspects of their environment.

What about the acquisition of information from books, newspapers, radio, etc., or from being spoken to by fellow humans? This too depends on systematic regularities. In this case, however, those regularities are not natural in origin like dark clouds and rain, or smoke and fire. Rather they depend on regularities created by people, the regularities of human language.

In order to acquire information from the words and sentences of English, you have to understand English — you need to know the meanings of the English words and you need a working knowledge of the rules of English grammar. In addition, in the case of

written English, you need to know how to read — you need to know the conventions whereby certain sequences of symbols denote certain words. Those conventions of word meaning, grammar, and symbol representation are just that: conventions. Different countries have different conventions: different rules of grammar, different words for the same thing, different alphabets, even different directions of reading — left to right, right to left, top to bottom, or bottom to top.

At an even more local level, there are the conventional information encoding devices that communities establish on an ad hoc basis. For example, a school may designate a bell ring as providing the information that the class should end, or a factory may use a whistle to signal that the shift is over.

The fact is, anything can be used to store information. All it takes to store information by means of some object — or more generally a configuration of objects — is a convention that such a configuration represents that information. In the case of information stored by people, the conventions range from ones adopted by an entire nation (such as languages) to those adopted by a single person (such as a knotted handkerchief). For a non-human example, DNA encodes the information required to create a lifeform (in an appropriate environment).

People also have the ability to obtain information from a configuration of objects in a particular context. An example is a hotel key rack. The original purpose of the key rack is to store keys. However, because it is commonly understood that each room in a hotel has a key, the number of keys on the key rack gives information about the size of the hotel. Because the traditional key racks were also used to store passports, messages, bills, a glance at the key rack can result in obtaining information about guests who are in their rooms, who have just checked in, who are about to leave. In this respect, an object such as a key rack can be said to carry information because of the way it is used by a community of people who share experience of hotels — hotel employees, guests, visitors and others.

For a more modern example, a management consultancy today employs ever increasing number of mobile workers. Since the consultants travel a lot, information about their whereabouts is quite important. If, for example, A's mobile phone is on the charger rack, most of his colleagues will assume he is in the office. Otherwise, the mobile phone would not be there. The phone charger carries that information for people who understand work practices in the organisation and can make reasonably accurate assumptions about the meaning of their colleagues' actions. Information in this context is often related to knowledge and understanding — the phone charger is what is often called a "common artefact that functions as a focus of interaction. It can only fulfil this function, however, if there is shared understanding of how it is used.

This is more than convention — the result of some kind of mutual agreement by a group of people that "table" will refer to an object with a flat surface and 1, 3 or 4 legs. Information carrying capacity of common artefacts is more dynamic, as it arises from action and interaction whose significance is understood by a given community.

To make any progress in understanding information in a precise, scientific way, we

need, first, to provide a precise, representation-free[5] definition of information, and, second, to examine the regularities, conventions, etc. whereby things in the world represent information. This is what two Stanford University researchers, Jon Barwise and John Perry, set out to do in the late 1970s and early 1980s. The mathematical framework they developed to do this they named Situation Theory, initially described in their book *Situations and Attitudes* [1], with a more developed version of the theory subsequently presented by Devlin in [4]. We shall provide a brief summary of part of situation theory in due course.[6]

One question that arises naturally in a study such as ours is whether information really exists. Perhaps talk of information is just that: so much twentieth and twenty-first century talk. This is a fascinating question, and one that we will touch on again at the end of the chapter. For the purposes of our discussion, however, we may sidestep the issue, and remain completely agnostic as to whether information has any kind of real existence. To do this, we can adopt what we shall call the *information stance*. This refers to the information-based way of thinking about (and analyzing) human action that we shall outline. When we adopt the information stance, we agree to talk *as if* information really exists and we approach human action and interaction in terms of the creation, acquisition, storage, transmission, exchange, sharing, and utilization of information. In adopting such an approach, we are taking our lead from the philosopher Daniel Dennett [3], who sidestepped many thorny questions about intentionality by viewing it as a stance ("the intentional stance") that may be adopted for various purposes.

Situation theory

In situation theory, recognition is made of the partiality of information due to the finite, *situated* nature of the agent (human, animal, or machine) with limited cognitive resources. Any agent must employ necessarily limited information extracted from the environment in order to reason and communicate effectively.

The theory takes its name from the mathematical device introduced in order to take account of that partiality. A *situation* can be thought of as a limited part of reality. Such parts may have spatio-temporal extent, or they may be more abstract, such as fictional worlds, contexts of utterance, problem domains, mathematical structures, databases, or Unix directories. The distinction between situations and individuals is that situations have a *structure* that plays a significant role in the theory whereas individuals do not. Examples of situations of particular relevance to the subject matter of this paper will arise as our development proceeds.

The basic ontology of situation theory consists of entities that a finite, cognitive agent individuates and/or discriminates as it makes its way in the world: spatial loca-

[5]Of course, our theoretical framework will have to have its own representations. The theory we will use adopts the standard application-domain-neutral representation used in science, namely mathematics.

[6]However, since situation theory is not the focus of this paper, our description will be very partial; we introduce just those situation-theoretic concepts and tools we require for our present purposes.

tions, temporal locations, individuals, finitary relations, situations, types, and a number of other, higher-order entities.

The objects (known as *uniformities*) in this ontology include the following:

- *individuals* — objects such as tables, chairs, tetrahedra, people, hands, fingers, etc. that the agent either individuates or at least discriminates (by its behavior) as single, essentially unitary items; usually denoted in situation theory by $a, b, c, \ldots$
- *relations* — uniformities individuated or discriminated by the agent that hold of, or link together specific numbers of, certain other uniformities; denoted by $P, Q, R, \ldots$
- spatial *locations*, denoted by $l, l', l'', l_0, l_1, l_2$, etc. These are not necessarily like the points of mathematical spaces (though they may be so), but can have spatial extension.
- *temporal locations*, denoted by $t, t', t_0, \ldots$. As with spatial locations, temporal locations may be either points in time or regions of time.
- *situations* — structured parts of the world (concrete or abstract) discriminated by (or perhaps individuated by) the agent; denoted by $s, s', s'', s_0, \ldots$
- *types* — higher order uniformities discriminated (and possibly individuated) by the agent; denoted by $S, T, U, V, \ldots$
- *parameters* — indeterminates that range over objects of the various types; denoted by $\dot{a}$, $\dot{s}$, $\dot{t}$, $\dot{l}$, etc.

The intuition behind this ontology is that in a study of the activity (both physical and cognitive) of a particular agent or species of agent, we notice that there are certain regularities or *uniformities* that the agent either individuates or else discriminates in its behavior.[7]

For instance, people individuate certain parts of reality as *objects* ('individuals' in our theory), and their behavior can vary in a systematic way according to spatial location, time, and the nature of the immediate environment ('situation types' in our theory).

We note that the ontology of situation theory allows for the fact that different people may discriminate differently. For instance, Russians discriminate as two different colors what Americans classify as merely different shades of blue.

Information is always taken to be information *about* some situation, and is taken to be in the form of discrete items known as *infons*. These are of the form

$$\ll R, a_1, \ldots, a_n, 1 \gg \ , \ \ll R, a_1, \ldots, a_n, 0 \gg$$

[7] This is true not only of individuals but also of groups, teams, communities. If A and B are engaged in a dialogue or a conversation, or indeed any other form of joint action, they recognize uniformities as individuals in a similar ways. Socially, they negotiate the precise meanings of these, so that they can agree the exact shape of the uniformities that apply in the situation they are in.

where R is an n-place relation and $a_1, \ldots, a_n$ are objects appropriate for R (often including spatial and/or temporal locations). These may be thought of as the informational item that objects $a_1, \ldots, a_n$ do, respectively, do not, stand in the relation R.

Infons are items of information. They are not things that in themselves are true or false. Rather a particular item of information may be true or false *about a certain part of the world* (a situation).[8]

Given a situation, s, and an infon σ, we write

$$s \models \sigma$$

to indicate that the infon σ is made factual by the situation s, or, to put it another way, that σ is an item of information that is true of s. The official name for this relation is that s *supports* σ.

It should be noted that this approach treats information as a *commodity*. Moreover a commodity that does not have to be true. Indeed, for every positive infon there is a dual negative infon that can be thought of as the opposite informational item, and both of these cannot be true (in the same situation).

Over the years, several people have misunderstood the role of infons in situation theory, and more generally have misunderstood the purpose of the situation-theoretic ontology, so it is worth making a few remarks here. A fundamental assumption underlying the situation-theoretic approach to information is that information is not intrinsic to any signal or to any object or configuration of objects in the world; rather information arises from ineractions of agents with their environment (including interactions with other agents). The individuals, relations, types, etc. of the situation-theoretic ontology are (third-party) theorist's inventions. For an agent to carry out purposeful, rational activities, however, and even more so for two or more agents to communicate effectively, there must be a substantial agreement first between the way an agent carves up the world from one moment to another, and second between the uniformities of two communicating agents. For instance, if Alice says to Bob, "My car is dirty," and if this communicative act is successful, then the words Alice utters must mean effectively the same to both individuals. In order for a successful information flow to take place, it is not necessary that Alice and Bob share exactly the same concept of "car" or of "dirty," whatever it might mean (if anything) to have or to share an exact concept. Rather, what is required is that their two concepts of "car" and of "dirty" overlap sufficiently. The objects in the ontology of situation theory are intended to be theorist's idealized representatives — prototypes — of the common part of the extensions of individual agent's ontologies. In consequence, the infons are theoretical constructs that enable

[8] One of the advantages of the framework and notation provided by situation theory is that it allows us to express partial information about complex relations. For example, the relation *eat* presupposes agent, object, instrument, place, time, but much of this information can remain implicit, as in "I'm eating." This makes it possible to choose which aspect of the structure to emphasize in a given instance of interaction. And this choice of emphasis also carries information in its own right, since it is recognised and interpreted as attitude or intent.

the theorist to analyze information flow. (In terms of our tennis-ball analogue of communication, the tennis ball — the infon — is a figment of the analyst's imagination, but one that facilitates a useful and meaningful analysis of a communicative act.)

Moving on now, situation theory provides various mechanisms for defining types. The two most basic methods are type-abstraction procedures for the construction of two kinds of types: situation-types and object-types.

Situation-types. Given a *SIT*-parameter, $\dot{s}$, and a compound infon σ, there is a corresponding *situation-type*

$$[\dot{s} \mid \dot{s} \models \sigma],$$

the *type* of situation in which σ obtains.

This process of obtaining a type from a parameter, $\dot{s}$, and a compound infon, σ, is known as *(situation-) type abstraction*.

For example,

$$[SIT_1 \mid SIT_1 \models \langle\langle \text{running}, \dot{p}, LOC_1, TIM_1, 1 \rangle\rangle]$$

Object-types. These include the basic types *TIM, LOC, IND, REL*^{n}, *SIT, INF, TYP, PAR,* and *POL*, as well as the more fine-grained uniformities described below.

Object-types are determined over some initial situation.

Let s be a given situation. If $\dot{x}$ is a parameter and σ is some compound infon (in general involving $\dot{x}$), then there is a type

$$[\dot{x} \mid s \models \sigma],$$

the *type* of all those objects x to which $\dot{x}$ may be anchored in the situation s, for which the conditions imposed by σ obtain.

This process of obtaining a type $[\dot{x} \mid s \models \sigma]$ from a parameter, $\dot{x}$, a situation, s, and a compound infon, σ, is called *(object-) type abstraction*.

The situation s is known as the *grounding* situation for the type. In many instances, the grounding situation, s, is the world or the environment we live in (generally denoted by w).

For example, the *type* of all people could be denoted by

$$[IND_1 \mid w \models \langle\langle \text{person}, IND_1, \dot{l}_w, \dot{t}_{now}, 1 \rangle\rangle]$$

Again, if s denotes Jon's environment (over a suitable time span), then

$$[\dot{e} \mid s \models \langle\langle \text{sees, Jon}, \dot{e}, LOC_1, TIM_1, 1 \rangle\rangle]$$

denotes the type of all those situations Jon sees (within s).

This is a case of an object-type that is a type of situation.

This example is not the same as a *situation-type*. Situation-types classify situations according to their internal structure, whereas in the type

$$[\dot{e} \mid s \models \langle\langle \text{sees, Jon, } \dot{e}, LOC_1, TIM_1, 1 \rangle\rangle]$$

the situation is typed from the outside.

Types and the type abstraction procedures provide a mechanism for capturing the fundamental process whereby a cognitive agent classifies the world. Applying the distinction between situation types and object types to interaction phenomena, we may say that we all recognise that the relationship between situation-type *fire* and the situation-type *smoke* obtains only if both are in the same place at the same time. This is then a part of the shared knowledge among members of the same group or community that is often assumed and therefore rarely articulated. Situation theory offers a mechanism for articulating these assumptions by means of defined constraints. *Constraints* provide the situation theoretic mechanism that captures the way that agents make inferences and act in a rational fashion. Constraints are linkages between situation types. They may be natural laws, conventions, logical (i.e., analytic) rules, linguistic rules, empirical, law-like correspondences, etc.

For example, humans and other agents are familiar with the constraint:

Smoke means fire.

If S is the type of situations where there is smoke present, and S' is the type of situations where there is a fire, then an agent (e.g. a person) can pick up the information that there is a fire by observing that there is smoke (a type S situation) and being aware of, or *attuned to*, the constraint that links the two types of situation.

This constraint is denoted by

$$S \Rightarrow S'$$

(This is read as "S *involves* S'.")

Another example is provided by the constraint

FIRE *means fire.*

This constraint is written

$$S'' \Rightarrow S'$$

It links situations (of type S'') where someone yells the word FIRE to situations (of type S') where there is a fire.

Awareness of the constraint

FIRE *means fire*

involves knowing the meaning of the word FIRE and being familiar with the rules that govern the use of language.

The three types that occur in the above examples may be defined as follows:

$$\begin{aligned} S &= [\dot{s} \mid \dot{s} \models \langle\langle \text{smokey}, \dot{t}, 1 \rangle\rangle] \\ S' &= [\dot{s} \mid \dot{s} \models \langle\langle \text{firey}, \dot{t}, 1 \rangle\rangle] \\ S'' &= [\dot{u} \mid \dot{u} \models \langle\langle \text{speaking}, \dot{a}, \dot{t}, 1 \rangle\rangle \wedge \langle\langle \text{utters}, \dot{a}, \textit{fire}, \dot{t}, 1 \rangle\rangle] \end{aligned}$$

Notice that constraints link types, not situations. However, any particular instance where a constraint is utilized to make an inference or to govern/influence behavior will involve specific situations (of the relevant types). Constraints function by capturing various regularities across actual situations.

A constraint

$$C = [S \Rightarrow S']$$

allows an agent to make a logical inference, and hence facilitates information flow, as follows. First the agent must be able to discriminate the two types S and S'. Second, the agent must be aware of, or behaviorally attuned to, the constraint. Then, when the agent finds itself in a situation s of type S, it knows that there must be a situation s' of type S'. We may depict this diagrammatically as follows:

$$\begin{array}{ccc} S & \stackrel{C}{\Longrightarrow} & S' \\ s : S \uparrow & & \uparrow s' : S' \\ s & \stackrel{\exists}{\longrightarrow} & s' \end{array}$$

For example, suppose $S \Rightarrow S'$ represents the constraint *smoke means fire*. Agent $\mathcal{A}$ sees a situation s of type S. The constraint then enables $\mathcal{A}$ to conclude correctly that there must in fact be a fire, that is, there must be a situation s' of type S'. (For this example, the constraint $S \Rightarrow S'$ is most likely reflexive, in that the situation s' will be the same as the encountered situation s.)

A particularly important feature of this analysis is that it separates clearly the two very different kinds of entity that are crucial to the creation and transmission of information: one the one hand the abstract types and the constraints that link them, and on the other hand the actual situations in the world that the agent either encounters or whose existence it infers.

For further details of situation theory, the reader should consult [4].

An example of human interaction

In his seminal article [10], published in 1972, the sociologist Harvey Sacks sought to illustrate the role played by social knowledge in our everyday use of language. He took the following two sentences from the beginning of a child's story

The baby cried. The mommy picked it up.

and examined the way these two sentences are normally understood, paying particular attention to the role played by social knowledge in our interpretation of the story.[9]

As Sacks observes, virtually every competent speaker of English understand this story the same way. In particular, we all hear it as referring to a very small human (though the word 'baby' has other meanings in everyday speech) and to that baby's mommy (though there is no genitive in the second sentence, and it is certainly consistent for the mommy to be some other child's mother). Moreover it is the baby that the mother picks up (though the 'it' in the second sentence could refer to some object other than the baby).

To continue, we are also likely to regard the second sentence as describing an action (the mommy picking up the baby) that follows, and is caused by, the action described by the first sentence (the baby crying), though there is no general rule to the effect the sentence order corresponds to temporal order or causality of events (though it often does so).

Moreover, we may form this interpretation without knowing what baby or what mommy is being talked of.

Why do we almost certainly, and without seeming to give the matter any thought, choose this particular interpretation? Sacks asks.

Having made all of his observations, Sacks explains [10, p.332]:

> My reason for having gone through the observations I have so far made was to give you some sense, right off, of the fine power of a culture. It does not, so to speak, merely fill brains in roughly the same way, it fills them so that they are alike in fine detail. The sentences we are considering are after all rather minor, and yet all of you, or many of you, hear just what I said you heard, and many of us are quite unacquainted with each other. I am, then, dealing with something real and something finely powerful.

It is worth pausing at this point to emphasize our purpose in working through Sacks' example in some detail, as we shall do momentarily. After all, as Sacks himself notes, "the sentences we are considering are ... rather minor." Yet, from the point of view of understanding the complexities of human interaction, the example embodies many of the key issues that arise. As Sacks himself observes, almost all of us understand the two sentences the same way. We do so despite the fact the practically none of that understanding is within the sentences themselves; it depends on our experience — what Sacks calls the 'fine power of a culture'.

One way to analyze the way the sentences are (normally) understood is to explicate the social relationships that are not overtly expressed. Sacks himself studied the semantic strategies people use in communication. He showed how they may draw upon their knowledge of the social systems in order to arrive at shared interpretations of the actions they observe (or imagine, as in the case of the example of the child's story).

[9] We first discussed Sacks' example in our research monograph [5]. Much of the technical material in this article is taken from that monograph.

His main concern was to explain how shared social norms make such actions intelligible and interpretable (cf. [6, p.327]).

An alternative approach — which is the one we shall adopt here — is to identify the informational and cognitive *structures* that lead to the understanding, in particular the relational structures where relations that apply in a given situation represent the regularities the agent discriminates. The underlying structural form is indicated by the diagram on page 13. Our analysis has two main components. To identify which types S and S' are used and identify which constraints C connect those types. Paralleling Sacks' analysis in our framework, we formulate *rules* that explicate how his "fine power of a culture" leads to the choice of types used to describe or understand the event or action. We use the type structure (i.e., the information-supporting structure) to explicate how that same "fine power of a culture" guides the interpretation in a structural way.

Because the example, even though it may seem mundane, encompasses all of the main elements of human interaction, either form of analysis will result in insights and methods that have wide applicability.

The importance of such studies goes beyond the internal goals of social science. For, the better our understanding of human action and interaction, the better we will be able to design information and communication technologies. For this particular application, structural analyses are particularly well suited, of course. Descriptive analyses were created to enhance understandimg, not to design technologies. To bring that understanding closer to desgn, we need to be able to use a different framework, which is what we explore here.

In our analysis of the example, we shall concentrate on both speaker and listener, as we seek to describe the mechanisms they invoke to achieve successful communication. One of the advantages that is gained by including the information flow as part of our study is that we are able to pull apart the speaker and listener actions, and track the manner in which the speaker invokes mechanisms that enable the listener to correctly interpret the utterance.[10]

We should note that our analysis assumes that the speaker's perspective has been determined. That is, we shall not, at this stage, ask ourselves *why* the speaker chooses the particular form of words she does, an issue closely related to the question why we see things in a certain way, but rather shall use the framework of situation theory to track the way the speaker and listener cooperate in order for the communicative act to be successful.

By carrying out our analysis in terms of information flow inspired by the framework of situation theory, we will be able to achieve a level of granularity that is conceptually (and intellectually) closer (compared with standard descriptive analyses) to the concept of information that is the concern of those working with Information Technology, Information and Communication Technology, etc.

[10] Just as a description of a tennis game can be given in terms of the individual actions of the two players in a way that is simply not possible for a wrestling match.

It is important to observe that, although Sacks' example concerns a linguistic event, his analysis (and the information-mediated alternative account we subsequently present here) is not a linguist's analysis — neither he nor we are doing a syntactic or semantic analysis. (In particular, we are not doing situation semantics, the application of situation theory that motivated its original development of situation theory by Barwise and Perry.) Our focus is on the *interaction* between the speaker and the listener and between the speaker and what he or she hears. We seek to highlight how information-mediated analysis can lead to the development (or uncovering) of information structure (more precisely the information-supporting structure).[11]

An information-based analysis of the Sacks example

In order to carry out our analysis, we need to introduce some situation-theoretic structures to represent the way that information flows from the speaker to the listener.

Reference to babies and mommies is captured in our framework by means of the types:

$$\begin{aligned} \text{'baby'} &= T_{baby} = [\dot{p} \mid w \models\!\ll \text{baby}, \dot{p}, t_{now}, 1 \gg], \\ \text{'mommy'} &= T_{mother} = [\dot{p} \mid w \models\!\ll \text{mother}, \dot{p}, t_{now}, 1 \gg], \end{aligned}$$

where $\dot{p}$ is a parameter for a person. (In these type definitions, the situation w is "the world", by which we mean any situation big enough to include everything under discussion. It is purely a convenience to think of this situation as the world, thereby providing a fixed context for the type definitions.)

We observe (as did Sacks in his original analysis) that both babies and mommies have different aspects. For instance, a baby can be thought of as a young person or as a member of a family, and a mommy can be viewed in relation to a child or to a father. These aspects, which affect the choice of words speakers make and the way listeners interpret them, are captured in our framework by the hierarchical structure on types (types of types, types of types of types, etc.).

Let:

$$\begin{aligned} T_{family} &= [\dot{e} \mid w \models\!\ll \text{family}, \dot{e}, t_{now}, 1 \gg], \\ T_{stage\text{-}of\text{-}life} &= [\dot{e} \mid w \models\!\ll \text{stage-of-life}, \dot{e}, t_{now}, 1 \gg], \end{aligned}$$

where $\dot{e}$ is a parameter for a type.

The activity of crying is closely bound to babies in the stage-of-life type, so when the listener hears the sentence "The baby cried" he will understand it in such a way that

(1) $$T_{baby} : T_{stage\text{-}of\text{-}life}.$$

That is to say, this item of information will be available to the listener as he processes the incoming utterance, and will influence the way the input is interpreted.

[11] This information structure plays a role in our analysis somewhat parallel to, though very different from, the social structure of Sack's analysis.

Since the reader may be familiar to uses of "types" in other disciplines (such as computer science), where they are generally rigid in nature, we should stress that in situation theory, any type will typically be a member of an entire stucture of types, and the applicability of a particular type may well depend upon two or more levels in the of-type hierarchy. For instance, the applicability of the type T_{baby} will be different when it is considered in the light of being in the type $T_{stage\text{-}of\text{-}life}$ as opposed to being in the type T_{family}. In the former case, individuals in the type T_{baby} will be typically and naturally associated with the activity of crying (type T_{crying}); in the latter case they will be typically and naturally associated with having a mother (2-type $T_{mother\text{-}of}$). (In situation-theoretic terms, these associations will be captured by constraints that link types. Those constraints are in general not universals, rather they may depend on, say, individual or cultural factors.) This particular distinction will play a significant role in the analysis that follows.

One immediate question concerns the use of the definite noun phrases 'the baby' and 'the mommy'. Use of the definite article generally entails uniqueness of the referent. In the case of the phrase 'the baby', where, as in the Sacks example, no baby has previously been introduced, one would normally expect this to be part of a more complex descriptive phrase, such as 'the baby of the duchess's maid', or 'the baby on last night's midnight movie'. So just what is it that enables the speaker to open an explanation with the sentence 'The baby cried'? It could be argued that an implicit suggestion for an answer lies in his later discussion of proper openings for 'stories', but this is a part of his article we do not consider here.

For a situation-theoretic analysis, there is no problem here. The situation theorist assumes that all communicative acts involve a *described situation*, that part of the world the act is *about*. Exactly how this described situation is determined varies very much from case to case. For example, the speaker may have witnessed, read about, or even imagined the event she describes. In the Sacks example, the speaker *imagines* a situation in which a baby cried and its mother picked it up. Let s denote that situation.[12]

The situation s will be such that it involves one and only one baby, otherwise the use of the phrase 'the baby' would not be appropriate. In starting a communicative act with the sentence 'The baby cried', the speaker is informing the listener that she is commencing a description of a situation, s, in which there is exactly one baby, call it b. (Whether or not b is a real individual in the world, or some fictional entity, depends on s. This does not affect the way our analysis proceeds, nor indeed the way people understand the utterance.)

The principal item of information about the described situation that is conveyed by the utterance of the first sentence 'The baby cried' is

[12]It does not affect the mechanics of our analysis whether you think of situations as objects in the speaker and listener's realm — possibly as things they are aware of — or purely as theorist's objects in an abstract ontology adopted to study interaction. All we need to know is that these situations are definite objects available to the theorist as part of a framework for looking at the world. In the case where situations are regarded purely as theorist's abstractions, s will *correspond* to some feature of the interaction—you can think of s as providing us with a *name* for that feature.

$$s \models \ll \text{cries}, b, t_0, 1 \gg$$

where t_0 is the time, prior to the time of utterance, at which the crying took place. In words, in the situation s, the baby b was crying at the time t_0.

Notice that, in the absence of any additional information, the only means available to the listener to identify b is as the referent for the utterance of the phrase 'the baby'. The utterance of this phrase tells the listener two pertinent things about s and b:

(2) $$b : T_{baby} \quad \text{(i.e. } b \text{ is of type } T_{baby})$$

where T_{baby} is the type of all babies, and

(3) b is the unique individual of this type in s.

Now let's consider what additional information is conveyed by the utterance of second sentence, 'The mommy picked it up.' Mention of both babies and mommies invokes the family type, T_{family} . This has the following structural components that are relevant to our analysis:

$M(x)$	the property of x being a mother
$B(x)$	the property of x being a baby
$M(x, y)$	the relation of x being the mother of y
T_{mother}	the type of being a mother
T_{baby}	the type of being a baby
$T_{mother\text{-}of}$	the 2-type that relates mothers to their offspring

In the type T_{family}, the type $T_{mother\text{-}of}$ acts as a fundamental one, with the types T_{mother} and T_{baby} being linked to, and potentially derivative on, that type. More precisely, the following structural constraints[13] are salient in the category T_{family}:

$$\begin{aligned} T_{mother} &\Rightarrow \exists \dot{y} T_{mother\text{-}of} \\ T_{baby} &\Rightarrow \exists \dot{x} T_{mother\text{-}of} \end{aligned}$$

where

$$T_{mother} = [\dot{x}, \dot{y} \mid w \models \ll \text{mother-of}, \dot{x}, \dot{y}, t_{now}, 1 \gg].$$

What do these mean? Well, $T_{mother\text{-}of}$ is a 2-type, the type of all pairs of individuals x, y such that x is the mother of y (at the present time, in the world). The first of the above two constraints says that the type T_{mother} *involves* (or is *linked to*) the type $\exists \dot{y} T_{mother\text{-}of}$. This has the following consequence: in the case where $T_{mother} : T_{family}$ (i.e. T_{mother} is of type T_{family}) and $T_{baby} : T_{family}$, the following implications are salient:

(4) $$p : T_{mother} \rightarrow \exists q \, (p, q : T_{mother\text{-}of})$$

[13] The notion of constraint used here extends that described in [4].

(5) $$q : T_{baby} \;\rightarrow\; \exists p\,(p, q : T_{mother\text{-}of}).$$

These two implications are not constraints. In fact they do not have any formal significance in situation theory. They are purely guides to the reader as to where this is all leading. (4) says that if p is of type T_{mother} (i.e. if p is a mother), then there is an individual q such that the pair p, q is of type $T_{mother\text{-}of}$ (i.e. such that p is the mother of q). The salience of this implication for an agent $\mathcal{A}$ has the consequence that, if $\mathcal{A}$ recognizes that p is a mother then $\mathcal{A}$ will, if possible, look for an individual q of which p is the mother. Analogously for (5).

To continue with our analysis, as in the case of 'the baby', in order for the speaker to make appropriate and informative use of the phrase 'the mommy', the described situation s must contain exactly one individual m who is a mother. In fact we can make a stronger claim: the individual m is the mother of the baby b referred to in the first sentence. For if m were the mother not of b but of some other baby, then the appropriate form of reference would be 'a mother', even in the case were m was the unique mother in s. We can describe the mechanism that produces this interpretation as follows.

Having heard the phrase 'the baby' in the first sentence and 'the mommy' in the second, the following two items of information are salient to the listener:

(6) $$m : T_{mother}$$

(7) $$m \text{ is the unique individual of this type in } s.$$

In addition, we shall show that the following, third item of information is also salient:

(8) $$m \text{ is the mother of } b.$$

Following the utterance of the first sentence, the listener's cognitive state is such that the type T_{baby} is of type $T_{stage\text{-}of\text{-}life}$. This type has categories that include T_{baby}, T_{child}, $T_{adolescent}$, T_{adult}, all of which have equal ontological status within the stage-of-life type, with none being derivative on any other. But as soon as the phrase 'the mommy' is heard, the combination of 'baby' and 'mommy' switches the emphasis from the type $T_{stage\text{-}of\text{-}life}$ to the type T_{family}, making salient the following propositions:

(9) $$T_{baby} : T_{family}.$$

(10) $$T_{mommy} : T_{family}.$$

In the T_{family} category, the various family relationships that bind a family together (and which therefore serve to give this type its status as a type) are more fundamental than the categories they give rise to. In particular, the types T_{baby} and T_{mother} are derivative on the type $T_{mother\text{-}of}$ that relates mothers to their babies.

Now, proposition (9) is the precondition for the salience of implication (5), namely

$$q : T_{baby} \;\rightarrow\; \exists p\,(p, q : T_{mother\text{-}of}).$$

Substituting the particular individual b for the variable q, we get

$$b : T_{baby} \;\rightarrow\; \exists p\,(p, b : T_{mother\text{-}of}).$$

But by (2), we know that

$$b : T_{baby}.$$

Thus we have the salient information

(11) there is an m such that $m, b : T_{mother\text{-}of}$.

The use of the definite article in the phrase 'the mommy' then makes it natural to take this phrase to refer to the unique m that satisfies (11). Thus the listener naturally takes the phrase 'the mommy' to refer to the baby's mother. This interpretation is reinforced by the completion of the second sentence '. . . picked it up', since there is an expectation that a mother picks up and comforts her crying baby. This explains how the fact (8) becomes salient to the listener.

It should be noticed that the switch from the salience of one set of constraints to another was caused by the second level of types in the hierarchy. The constraints we were primarily interested in concerned the types T_{mother} and T_{baby}. These types are part of a complex network of inter-relationships (constraints). Just which constraints in this network are salient to the agent is governed by the way the agent encounters the types, that is to say, by the type(s) of those types—for instance, whether T_{baby} is regarded (or encountered) as of type $T_{stage\text{-}of\text{-}life}$ or of type T_{family}. By moving to a second level of typing (i.e. to types of types), we are able to track the way agents may use one set of constraints rather than another, and switch from one set to another. The first level of types allows us to capture the informational connections between two objects; the second level allows us to capture the agent's preference of a particular informational connection. This level of uncertainty is needed or else there could be no negotiation in interaction.

Our analysis thus explicates the information *structure* that the speaker and listener implicitly make use of in order a communicative act to succeed. In particular, it highlights the crucial roles played not only by constraints (the key players in a situation semantic analysis) but also by the internal and hierarchical type-structures. This latter feature is quite new, and takes the analysis a considerable distance from situation semantics. We believe it is a significant tribute to the care Barwise, Perry, and their colleagues gave to the choice of the ontology for situation theory as a framework to support the study of information and of natural language semantics that it proves to be adequate for a detailed analysis of human interaction such as the one presented here.[14]

An example from industry

The fundamental nature of the issues emmbodied in the Sacks example means that the methods we employed in our anaysis have much wider applicability. For instance,

[14]Both Barwise and Perry expressed on many occasions a desire to extend their work to look at action and interaction, but they never made such a step.

in the late 1980s and early 1990s, we analyzed what had gone wrong when a large manufacturer and supplier of mainframe computer systems had tried to automate part of its own information system, namely the data collected in the standard form (the Problem Report Form, or PRF) filled in when an engineer was called out on a repair job.

The PRF was a simple slot-and-filler document on which could be entered various reference numbers to identify the customer and the installed system, the fault as reported by the customer, the date of the report, the date of the engineer's visit, the repair action he took, and any components he replaced.

The PRF was a valuable document, providing the company with an excellent way to track the performance of both their computer systems and their field engineers, as well as the demand for spare parts. In particular, by analyzing the data supplied by the forms, the company could identify and hopefully rectify the weakest components in the design of their systems.

Because of the highly constrained nature of the PRFs, the highly focused nature of the domain — computer fault reporting and repair — and the fact that the important technical information on the forms was all entered by trained computer engineers, the PRFs formed the basis of a highly efficient source of information for all parts of the company. In the early days, when the PRFs were paper documents, experts faced with reading the forms frequently encountered great difficulty understanding exactly what had gone wrong with the customer's system and what the engineer had done to put it right. The PRF was a shared artefact – the focus of interaction between many departments: customer services, spare parts, diagnostics, etc. Information flowed naturally and any uncertainties were cleared up in conversation. When the PRF was computerised, it became an information record in a database, and the information flow between people was interrupted. If a particular PRF had (what company employees referred to as) "good information" in it, it could be easily interpreted and understood well enough to lead to action. If it contained "bad information", it presented a problem.

When an expert system was introduced, the expectation was that it would introduce intelligence into the interrupted information flow, so that the PRF could continue to function as mediated by the expert system. But this did not happen. Things got worse; the information flow was disrupted. Different people (agents) had different perspectives on the information in the PRF. The database representation of this information did not allow for different perspectives, it only encoded what the database designer specified and in the form that the designer specified. There was therefore no flexibility that could allow individual perspectives to be recognized and negotiated, and for people to establish shared understanding.[15]

Applying extensions of the techniques used to analyze the Sacks example, we were able to carry out a detailed analysis of the way social and cultural knowledge affected the information conveyed by the PRFs. This led to a restructuring of the procedures surrounding the completion and use of the documents, resulting in better information

[15]Following Perry and Israel [8], we can say that a PRF had the information potential that agents could pick up, but the database could not.

flow and improved efficiency in the company.

Furthermore, the additional problem our analysis addressed was to relate the structure of the document to its broader uses in the organisation as a whole. We viewed the PRF as a resource that facilitates (or obstructs, as the case may be) the interaction between different sections of the organisation. In this context, the social significance of the document needs to be understood so that the information flow between different sections may be organised and managed.

An investigation into the uses of the document, as opposed to its structure, brought to light the need to develop a dual perspective — what we called the *document intension* and the *schema of investigation*. The document intension is an "information-structure-skeleton" of the PRF that captures the communicative intent of the various sections of the document, through the use of the constraints that formalize the informational links within the document (essentially its underlying type structure). The schema of investigation traces the information pathways a reader of the document creates in the process of interpretation, schematically presented in Figure 1.

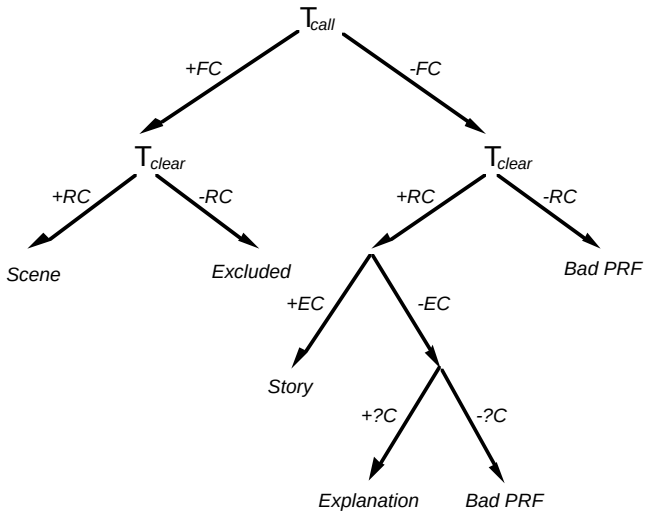

Figure 1: Interpretive grammar

The schema captures formally how the successive application of constraints leads to "perfect" information in the "scene", when everything fits — on the far left of the tree — and also to the "bad PRF" on the far right of the tree. These examples illustrate the strategies that computerized resources capture easily.

However, most of the everyday cases analyzed were not so clear cut. Going from left to right in the tree in Figure 1, if the fault description is clear and the remedial action is not, this would be interpreted as the engineer not knowing his job. Needless to say, no PRF among the hundreds analyzed gave this information explicitly. The most frequent and the most challenging examples were those in the middle of the tree, where the fit had to be established between the fault description, the appropriate remedial action

and the resources used in implementing the remedy. This is where most of human interrpretive effort was focused. Sadly, this is also where computerized tools are still grossly inadequate as they are not responsive to the human uses of the information stored in them.

An empirical study of the uses of the PRF in the organization showed that the information contained in the document was needed to support the decisions of customer services, fault diagnosis, spare parts, change management, and production, as well as intra-organizational partnerships with the suppliers of various parts of the total system delivered to the customer.

We should stress that, even though both examples presented in this paper, the baby and the fault record, concern particular interactions in particular contexts, the analytic methods used to analyze them, described above, are able to capture the underlying regularities, or uniformities, and hence can be generally applied. This is where the information-based, structural approach can offer advantages over purely descriptive analyses. If our interest were solely the understanding of human action and interaction, that advantage might be of little consequence. It can become significant, however, if we are interested in the design of tools and resources that embody this understanding, and in the organisation of work that recognises the importance of social relationships in everyday practice.

For further details on the PRF example, we refer the reader to our monograph [5].

The utility and relevance of the information stance

In our analysis of the Sacks example, we showed how a communicative interaction can be analyzed in terms of information flow, using the framework of situation theory. What makes viewers see a scene the way they do, and why do they choose the precise form of words they use to convey information about that scene? Information may be regarded as (and arguably is, if it is anything at all beyond a manner of speaking) an inert commodity exchanged by the two participants in a linguistic interaction. Hence, adopting the information stance allows us to tease apart the two individuals in a communicative interaction. This allows us to analyze the actions of each with respect to the information transmitted or received, thereby shedding light on some of the intricacies of everyday human–human interaction.

The price that we might pay for this increased level of analytic precision is twofold. First, for some interactions, viewing the interaction as mediated may impede or even skew the analysis. (To go back to an earlier analogy, it would be possible to analyze a wrestling match in terms of some postulated "particles of force" that the two protagonists emit toward one another, but this is unlikely to lead to a better analysis, and in fact will probably obscure the interaction.) Second, it is at least arguable that information simply does not exist — that it is just a way of talking — and that it is more intellectually honest to stick to what is really there.

This last point might be a significant objection if the results of an information-based analysis could be presented only in terms of information. However, when it is

possible to adopt the information stance for the purposes of carrying out an analysis, and then present the conclusions without reference to information, as we could with our examination of the Sacks example, that objection surely melts away.

There remains the question as to whether information really does exist. If the matter were to be settled by popular vote, the answer would surely be a resounding "Yes." Indeed, we suspect it would be almost unanimous. This is clearly a significant observation for the relevance of information-based analyses of social phenomena. If information is universally accepted in society, then it is legitimate to analyze social phenomena in terms of information. The results of that analysis (*presented in terms of information*) may then be legitimately presented as conclusions relevant to social science concerns.

In other words, an analysis of a social phenomenon based on information has validity in and of itself. It need not defend itself by an appeal to the information stance. (Although that remains a valid methodological approach.) Given a socially accepted notion of information that appears reasonably stable, an analysis like our study of the Sacks example we presented above turns out to be more than just a "what if" argument, where the intermediate steps are mere ephemera to be discarded once the final conclusion is reached. Rather, each step in the analysis establishes a genuine truth about the world — a truth about the information that flows from one agent to another. Viewed in this way, such an information-based analysis of human action is both valid and genuinely, qualitatively different from other forms of sociological analysis. That difference can be of significance when it comes to applications. A particular strength of the information-based approach, based on a mathematical framework such as situation theory, is that it allows for a formal treatment that can be informed by insights from sociology while at the same time yielding an analysis *that can be applied to the design of information and communication technologies*. This, in fact, was the reason we developed our analytic technique and carried out our analysis of the Sacks example in the first place. The work desrined here represents a genuine, original, and on this occasion highly successful, application of the modern concept of information to the development of understanding of a certain domain of human activity leading to a successful engineering design.

We end on a more speculative note. Today's concept of information assumes — and encourages us to believe — that information has some form of existence beyond its various physical encodings, each of which is often viewed as a container. Futurist commentator John Perry Barlow, co-founder of the Electronic Frontier Foundation is quoted[16] as having said:

> "So far we have placed all of our intellectual protection on the containers and not on the contents. And one of the side effects of digital technology is that it makes those containers irrelevant. Books, CDs, filmstrips — whatever — don't need to exist anymore in order to get ideas out. So whereas we thought we had been in the wine business, suddenly we realized that all along we've been in the bottling business."

[16] In C. W. Beardsley, Uncorking Fine Wine: Copyright Laws, *Mechanical Engineering*, August 1994.

The suggestion is that today's digital technologies will completely separate the information from its various representations, which are seen as containers. "Information wants to be free" is a popular rallying cry. But perhaps — and this is definitely where our sympathies lie — social scientist Paul Duguid had it right when he observed that such talk is akin to saying we want to remove the banks and still have the river.

For all that (today's conception of) information has a form of existence, we lean toward the view that what that existence really amounts to is a collective acceptance of the information stance. That is to say, it really is just a way of conceiving of and talking about various aspects of our world.

References

[1] Barwise, J. and Perry, J. *Situations and Attitudes*, Bradford Books, MIT Press (1983).

[2] Clark, H. *Using Language*, Cambridge University Press (1996).

[3] Dennett, D. *The Intentional Stance*, MIT Press (1989).

[4] Devlin, K. *Logic and Information*, Cambridge University Press (1991).

[5] Devlin, K. and Rosenberg, D. *Language at Work: Analyzing Communication Breakdown in the Workplace to Inform Systems Design*, Stanford University: CSLI Publications and Cambridge University Press (1996).

[6] Gumpertz, J. & Hymes, D. (eds). *Directions in Sociolinguistics, The Ethnography of Communication*, Holt, Rinehart and Winston Inc. (1972).

[7] Nunberg, G. (ed), *The Future of the Book*, University of California Press (1996).

[8] Israel, D and Perry, J, What is Information?, in *Information, Language and Cognition: Vancouver Studies in Cognitive Science, Vol. I*, University of British Columbia Press (1990).

[9] Rosenberg, D. and Hutchinson, C. (Eds.) *Design Issues in Computer-Supported Cooperative Work*, London: Springer-Verlag (1994).

[10] Sacks, H. On the Analyzability of Stories by Children, in [6], pp.325–345.