RESEARCH SUPPORT
SCHOOL OF BUSINESS ADMINISTRATION

AUGUST 1995

TESTING SPEECH ACT THEORY AND ITS APPLICABILITY TO EDI & OTHER COMPUTER-PROCESSABLE MESSAGES

WORKING PAPER #9503-17

SCOTT A. MOORE
UNIVERSITY OF MICHIGAN

Testing Speech Act Theory and its Applicability to EDI & Other Computer-Processable Messages

Scott A. Moore
Computer Information Systems Department
School of Business Administration
University of Michigan
701 Tappan
Ann Arbor, MI 48109–1234

samoore@umich.edu

(313) 763–4385

Abstract

This paper discusses a small empirical study that investigates the relationship between electronic commerce and a linguistic theory called speech act theory (SAT). The study reveals that electronic data interchange messages and inter-application communication messages have the structure predicted by SAT. This should encourage information systems (IS) researchers to continue investigating SAT, IS practitioners to consider basing message structures on a SAT framework, and speech act theorists who support this still-controversial theory.

1 Introduction

In its most compelling form, information systems (IS) research makes contributions in three areas: the field of IS, the practice of IS, and the theory of a reference discipline. The results of this paper teach us how to do electronic commerce better (practice of IS) and the direction future research should take (field of IS). They also give us important insights into a linguistic theory called speech act theory (theory of a reference discipline). These lessons all stem from an empirical study that attempts to determine a relationship between speech act theory (SAT) and electronic commerce.

Any theory that can contribute to the development of more capable systems for electronic commerce could have tremendous impact. Since electronic commerce is based on communication among systems and people, theories of communication might have much to say about it. However, systems such as electronic data interchange (EDI) have developed with little regard for SAT, theories of communication, or even linguistics. This paper reports on the remarkable fact that the message structure implicit in the investigated electronic messaging systems is consistent with SAT. This has consequences for each of the areas mentioned above:

- Field of IS: Researchers have investigated communication and electronic commerce systems based on SAT. They should continue to do so.
- Practice of IS: Developers have created systems that send messages that

implicitly perform speech acts. They should consider making the speech acts explicit so the computer can reason more easily about them.

- Theory of a reference discipline: Speech act (SA) theorists have argued on theoretical grounds that SAT is correct. They should be heartened that this study supported this still-controversial theory.

I conducted a small empirical study of two EDI standards and an inter-application communication (IAC) standard. The purpose of this study was to determine the correspondence between the standards and SAT. SAT is a linguistic theory that describes how language is used. As such, it makes certain predictions about the structure of all utterances. Since electronic messages are a type of utterance, SAT predicts that electronic messages must have a certain structure. This prediction is tested by determining if the electronic messages have that structure. The findings of this empirical investigation supported SAT—messages can be mapped to SAT.

2 Review of Speech Act Theory

Work on SAT began, roughly, with the publication of Austin's *How to Do Things with Words* [4], the text of his William James Lectures at Harvard University in 1955. These lectures specified two very important, though quite general, ideas. The first is a refutation of the then commonly accepted idea that language's only function is to say things that are true or false. Austin felt this was not enough. He believed statements also accomplish something. When people say something, they are not merely saying something but also *doing* something.[1]

Austin also proposed that every (for our purposes) utterance is the speaker's expression of a attitude[2] toward some possibly complex proposition. For example, if the speaker says "It will rain," then typically the speaker is predicting it will rain. The proposition is *it will rain* and the attitude is that of a *prediction.* If the speaker says "Will it rain?" then typically the speaker is asking whether it will rain. In this case, the proposition is the same—*it will rain*—and the attitude expressed is that of a *question.* Thus, speakers can express different attitudes toward the same proposition. SA theorists call these attitudes *illocutionary forces.* Summarizing this idea: every SA has the structure *F(P)*, where *F*, the illocutionary force, is applied to *P*, the propositional content. This is called the *F(P) framework.*

This is a strong claim. SA theorists propose that the outermost operator of every utterance (everything we could possibly say) is not Boolean, not temporal, not even defeasible—it is an illocutionary force. If this is true, then a communication system might benefit from representing the structure explicitly so that the system could reason about it. SA theorists also propose categorization schemes for the illocutionary forces. These schemes differ widely. The classification scheme can be considered a scheme for organizing an object-oriented message hierarchy. Different versions of SAT [6, 18] are more or less suited to be an organizing scheme for such a hierarchy because of the opportunities for inheritance. A less useful hierarchy would be one level deep with no inheritance. A more useful hierarchy would be deeper, allowing messages to inherit properties from other messages. A small number of illocutionary forces that can be categorized into a rela-

[1] And, hence, the title of his lectures.

[2] In the SAT literature, this is referred to as the *propositional* attitude.

tively deep tree would be ideal for a communication system based on SAT.

I work with the version of SAT proposed by Bach & Harnish [5]. The reasons for this choice are that it is representative of other proposals and no one has demonstrated the superiority of another theory. Also, they define a system of inferential communication that is a useful basis for electronic communication [12, 15]. Roughly, under an inferential communication theory, which is also supported by some cognitive scientists and philosophers [17, 22], the recipient must infer what the speaker means and then take the message as a basis of inference for how to act. This is in contrast to Searle's position that communication is a decoding system [18, 19, 20]. Under a decoding theory, the recipient knows precisely what the speaker intends the recipient to do once the message is decoded.

Bach & Harnish propose two major categories of illocutionary forces and six main sub-categories with further subcategorization (see Table 1). The two major categories designate how the hearer should process the utterance.
"Communicative illocutionary acts succeed by means of recognition of intention, whereas conventional ones succeed by satisfying a convention." [5, p. 110] As for the six main sub-categories:

> "*{C}onstantives* express the speaker's belief and his intention or desire that the hearer have or form a like belief. *Directives* express the speaker's attitude toward some prospective action by the hearer and his intention that his utterance, or the attitude it expresses, be taken as a reason for the hearer's action. *Commissives* express the speaker's intention and belief that his utterance obligates him to do something (perhaps under certain conditions). And *acknowledgments* express feelings regarding the hearer or, in cases where the utterance is clearly perfunctory or formal, the speaker's intention that his utterance satisfy a social expectation to express certain feelings and his belief that it does." [5, p. 41]
>
> "*Effectives* effect changes in institutional states of affairs... *Verdictives* are judgments that by convention have official, binding import in the context of the institution in which they occur." [5, p. 110–11]

Table 1 shows that acknowledgments, effectives, and verdictives do not have separate categories of forces below them; the forces can only be distinguished at the verb level.

Each of the illocutionary forces is more carefully defined to distinguish it from other forces of the same type. For example, predictives are defined as follows:

> *Predictives*: (forecast, predict, prophesy) In uttering *e*, *S* predicts that *P* if *S* expresses:
> i. the belief that it will be the case that *P*, and
> ii. the intention that *H* believe that it will be the case that *P*. [5, p. 42]

The definition contains examples of verbs that typically have this illocutionary force. The other forces are defined in a similar fashion.

Communicative				Conventional	
Constantives	Directives	Commissives	Acknowledgments	Effectives	Verdictives
Assertives Predictives Retrodictives Ascriptives Informatives Confirmatives Concessives Retractives Assentives Dissentives Disputatives Responsives Suggestives Suppositives	Requestives Questions Requirements Prohibitives Permissives Advisories	Promises Offers	Verbs: Apologize Condole Congratulate Greet Thank Bid Accept Reject	Verbs: Appoint Nominate Suspend Demote Enlist Apply Resign Abdicate Arrest Indict Etc.	Verbs: Acquit Certify Disqualify Clear Rule Adjudicate Etc.

Table 1: Illocutionary forces defined by Bach & Harnish [5].

SA theorists contend that all utterances can be described within the *F(P)* framework. Bach & Harnish propose the twenty-plus illocutionary forces in six categories listed above. Others propose different numbers of forces in different categories. Bach & Harnish contend that all utterances—verbal, electronic, or otherwise—use one of these illocutionary forces. It is this contention that is investigated in this paper.

3 An Empirical Study

The best way to determine if all utterances can be understood within the *F(P)* framework is to translate all utterances into the framework. Since the number of utterances is unbounded, this is not possible. It is also not clear that a random sampling of utterances is feasible. A more feasible test, and the type used in this paper, is to translate some appropriate sample of utterances. This involves finding existing sets of utterances that reflect the diverse activities performed electronically. The challenge before the scientist is two-fold:

- Find a diverse set of utterances so the results can be regarded with some confidence, and
- Work with an application domain close enough to the actual application area so the results are considered relevant to electronic commerce.

Satisfying these requirements should increase the study's external validity [7]. In this case two EDI standards and an IAC standard were chosen. I chose these because they are existing, rich, diverse commercial standards developed independently of SAT. Each domain serves a different purpose, and the messages within each set differ. Presumably each message set's creators defined it so that a complete range of activities could be performed electronically.

This test should indicate if the prospects for SAT being useful are good or not. SA theorists predict all utterances should map to the *F(P)* framework. In this small empirical study I test this prediction by attempting to classify each message of each standard into one of (or a combination of) the illocutionary forces defined by Bach & Harnish. The

applicability of SAT will be measured by several questions:

1) Can the messages be mapped onto the illocutionary forces?

   SAT will clearly have failed if even one message cannot be mapped onto the *F(P)* framework. SAT predicts all utterances fit this framework so even one failure will contradict SAT. An assumption is that only well-defined messages will be mapped. Text only messages are not included in the sample since what the text says is not part of the standard. It is impossible to categorize these messages not because SAT fails but because the standard provides nothing to categorize.

   Successfully mapping a message does not mean each message will map onto only one illocutionary force. SAT does not predict this about normal language nor is this feature present in normal language. Suppose a person says "Please pass the salt and pepper." This naturally maps onto two separate acts—a request to pass the salt and a request to pass the pepper. It would be surprising if each message mapped onto one illocutionary force.

2) Can the messages be easily mapped onto the illocutionary forces?

   It is one thing to map the messages onto the *F(P)* framework and another thing to do so easily. Mapping is said to be "easy" if 1) each message maps directly to an illocutionary force without having to stretch its definition, and 2) each message is mapped to one, and only rarely two or three, illocutionary forces.

3) Are the depth and breadth of the illocutionary force categorization well-covered?

   The hierarchy of illocutionary forces described by SAT can be thought of as a tree. Mapping the messages onto SAT will be much more useful if both the tree's depth (categories and each level of sub-categorization) and breadth (illocutionary forces) are used, or covered. Property inheritance is more useful to application developers if many categories and sub-categories are used—i.e., the depth is utilized. Using a higher percentage of the forces—i.e., utilizing the breadth—indicates the force is contributing to more of the message's meaning.

   To understand the importance of this point, consider the extreme example in which all messages map into one illocutionary force. This would indicate the illocutionary force does not contribute much to understanding the message. Even if it does contribute, other factors clearly outweigh its importance. The cost of adding this information to the message's representation would probably outweigh its usefulness.

4) Does each standard significantly overlap the other standards?

   The applicability of SAT will be enhanced if each standard uses much of the tree. If the tree were segmented by the standards, then the generality of SAT would have to be questioned. It would cause one to hypothesize that as more message standards are analyzed, more illocutionary forces, sub-categories, and categories will be needed. It would also raise interesting questions about SAT, such as: Is it always the case that certain illocutionary forces are used together? Is this to the exclusion of

other forces? Does this say anything about the complexity of the communication?

This study focuses on these questions. To make it clear what is being investigated, these questions are reworded as null hypotheses:

1) Most messages can be mapped into the *F(P)* framework. The messages that cannot be mapped onto the *F(P)* framework will not be well-defined.

2) Most messages can be mapped easily into the *F(P)* framework.

3-a) The simplistic communication involved in some standards will result in little of the tree being covered by that standard.

3-b) Even the most complex standard will not cover very much of the tree.

4) The tree will not be segmented by different standards.

4. Message Standards

This section describes those application domains and message standards used in the empirical study. These include two EDI standards and one IAC standard. For brevity's sake I focus on the standard for financial transactions.

4.1 EDI – X.12 Standard

The X.12 standard for EDI [1] defines the data and control structures for common business documents such as purchase orders, invoices, and requests for purchase. Companies send these messages to their trading partners. Some messages require a return message; others require the company to send goods; others are purely informational messages.

4.2 Financial Transactions

"S.W.I.F.T. is a world-wide organization working in partnership with its customers to provide them with communication and financial data processing services of the highest quality, security and integrity." [21, p. 5] S.W.I.F.T. developed a standard for sending messages about financial securities, such as trading of securities, settlement of trades, and securities lending and borrowing. Companies send these messages to their trading partners and to other institutions involved in financial transactions. Similar to the X.12 messages, these messages may or may not require responses.

4.3 Inter-Application Communication

The Apple Event (AE) Registry [2] defines standard inter-application communication messages (known as Apple Events on the Apple Macintosh). Applications use these messages to send information to other programs and to get them to perform tasks. Just as with the EDI standards, this standard does not exhaust all possible messages.

When an AE expects information in return, it leaves an electronic "return envelope" with the receiver to put its information into. This envelope can contain an answer to a question or information about a problem encountered by the receiver. Thus, each message that can generate a response implicitly defines another message.

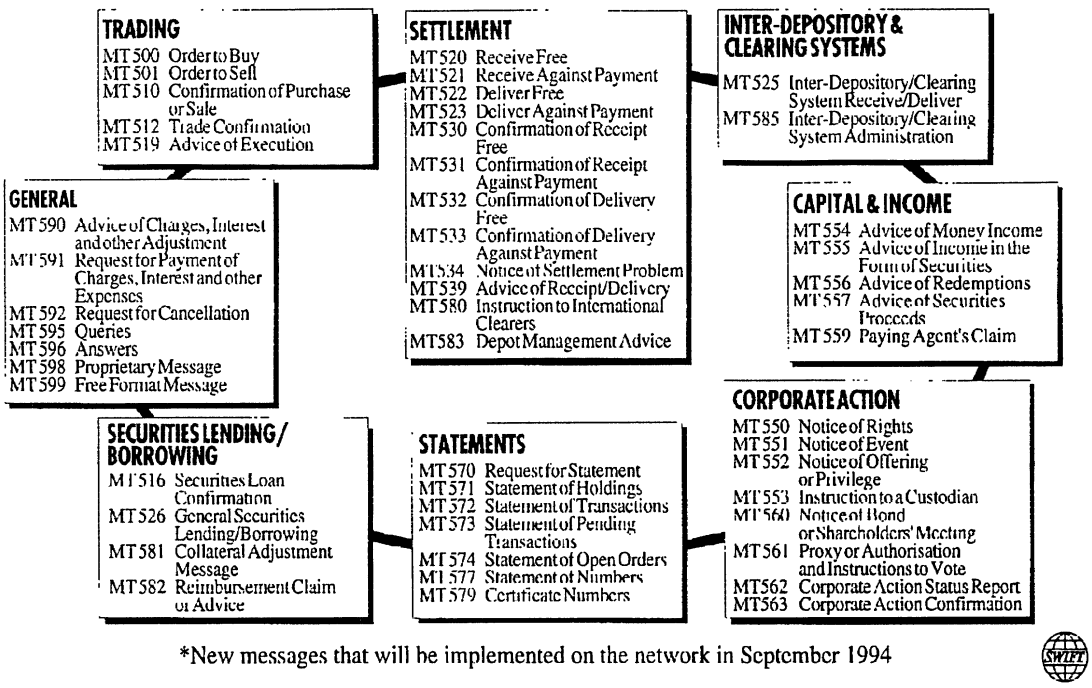

Figure 1: S.W.I.F.T.'s Securities Messages.

5 Results

I mapped each message in each standard to the illocutionary forces defined by Bach & Harnish. For each standard Table 2 displays the number of messages that can possibly perform each illocutionary force. Some messages in the S.W.I.F.T. and X.12 standards can have one of several forces when sent (the "exclusive possibilities" line). Others can have several forces each time they are sent depending on how they are constructed (the "inclusive possibilities" line).

In Table 3 is the categorization of each S.W.I.F.T. message. (Lack of space prevents presentation of similar detail for the other two standards.) Some messages (e.g., MT526) can be used in different ways at different times. Consider the following:

> **MT 526 General Securities Lending/Borrowing**
> This message is sent from one institution to another to list, request, notify, or confirm information relating to Securities Lending/Borrowing. [21, p. 107]

This message is categorized as an informative (to list or notify), a requestive (to request), and a confirmative (to confirm). It holds any one of these forces to the exclusion of the others. Some messages (e.g., MT512) can be used in different ways at the same time. Consider the following:

> **MT 512 Securities Trade Confirmation**
> This message is exchanged between the counterparties to a securities trade to confirm the trade between these parties. It may also include settlement details where necessary. [21, p. 25]

Illocutionary Force	S.W.I.F.T.	X.12	Apple Events
# of messages	50	11	37
Assertives	9	1	
Predictives		1	
Retrodictives			
Ascriptives			
Informatives	20	4	3
Confirmatives	10	2	
Concessives			
Retractives	1		
Assentives			
Dissentives			
Disputatives			
Responsives	8	1	
Suggestives			
Suppositives			
Requestives	7	3	32
Questions			2
Requirements	10		
Prohibitives			
Permissives		1	
Advisories			
Promises		2	
Offers		3	
Acknowledgments			
Effectives	1	1	
Verdictives			
Undefined	2	1	
Inclusive possibilities	4	6	
Exclusive possibilities	9		

Table 2: Number of messages mapped to each illocutionary force.

When this message is sent it can have both a confirmative force and an informative force since it can both confirm and inform in the same message. The rest of the messages are classified in a similar manner.

And now the findings from the empirical study as they relate to each hypothesis.

Hypothesis 1

Most messages can be mapped into the F(P) framework. The messages that cannot be mapped onto the F(P) framework will not be well-defined.

All the well-defined messages were mapped to the *F(P)* framework. The S.W.I.F.T. standard contains two messages (598, 599) that could not be mapped. Message 598 is defined as a "proprietary message" with no further explanation. Message 599 is defined as a "free format" message. Its contents are not interpretable by a machine. Both of these messages are not well-defined and were not included in this test.

Hypothesis 2

Most messages can be mapped easily into the F(P) framework.

The mapping of the messages was easy, particularly for the S.W.I.F.T. and AE standards. To qualify as "easy", the mapping had to be done without stretching the definition of the force and each message had to be mapped to a few forces. Mappings described above for the S.W.I.F.T. messages are representative of the ease with which the mapping was done.

Table 2 contains information about the second requirement. The "Inclusive possibilities" line lists the number of messages that can perform several illocutionary forces at one time within one message. For these messages it is generally clear that two different types of actions are performed by the message (recall the MT 512 example). Also, multiple mappings are not always an indication of complexity (recall the salt and pepper example). In this case, since so few inclusive messages were mapped, I considered this to be easy.

Illocutionary Force	Categorization of Messages
# of messages	
Assertives	559, 571/xor, 572/xor, 573/xor, 574/xor, 577/xor, 579/xor, 581/xor, 582/xor
Predictives	
Retrodictives	
Ascriptives	
Informatives	510/or, 512/or, 519, 525/or, 526/xor, 534, 539, 550, 551, 552, 554, 555, 556, 557, 560/or, 562, 581/xor, 582/xor, 583, 590
Confirmatives	510/or, 512/or, 516, 525/or, 526/xor, 530, 531, 532, 533, 563
Concessives	
Retractives	525/or
Assentives	
Dissentives	
Disputatives	
Responsives	525/or, 571/xor, 572/xor, 573/xor, 574/xor, 577/xor, 579/xor, 596
Suggestives	
Suppositives	
Requestives	525/or, 526/xor, 560/or, 570, 591, 592, 595
Questions	
Requirements	500, 501, 520, 521, 522, 523, 525/6, 553, 561, 580
Prohibitives	
Permissives	
Advisories	
Promises	
Offers	
Acknowledgments	
Effectives	585
Verdictives	
Undefined	598, 599

Table 3: Mapping of each S.W.I.F.T. message. (M/xor: message can be defined to have this force exclusively. M/or: message can perform this action in conjunction with other forces.)

Hypothesis 3a

The simplistic communication involved in some standards will result in little of the tree being covered by that standard.

The study did not find support for this hypothesis. The simplistic communication model of the AE standard leads one to think there is support for this hypothesis. This standard only expresses three illocutionary forces directly. However, two other places within AEs can have illocutionary force. First, the response message can express an illocutionary force (tabulation shown in the first column of Table 4).

Second, the message can express iterated illocutionary forces. Each message fits within the *F(P)* framework. As defined in section 2, *P* is the message's propositional

Illocutionary Force	Inside Response	Inside Request
Assertives	34	1
Predictives		
Retrodictives		
Ascriptives		
Informatives	9	7
Confirmatives		
Concessives		
Retractives		1
Assentives		
Dissentives		
Disputatives		
Responsives		
Suggestives		
Suppositives		
Requestives		
Questions		
Requirements		
Prohibitives		
Permissives		1
Advisories		
Promises		
Offers		
Acknowledgments		
Effectives		
Verdictives		
Simple content (i.e., do something)		24

Table 4: Mapping information for Apple Events.

content. Sometimes this *P* can have a more complex form such as $F_1(P)$ where F_1 is any illocutionary force. For example, you can request that Fred inform Barney that Gary went outside. *F* is request, F_1 is inform, and *P* is that Gary went outside. The second column of Table 4 counts the F_1 for all the AEs of the form $F(F_1(P))$ where *F* is a *requestive* force. The line labeled "Simple content" includes those messages whose content is simply P rather than a more complex $F_1(P)$. Two of the messages (kAEClone and kAECreateElement) both request to inform and request to do something (i.e., in "Simple content" line). This is why 34 messages are in the second column and 32 requestives are shown for Apple Events in Table 2. As the information in Table 4 demonstrates, the seemingly simplistic structure of some standards can hide more diverse message types than is immediately obvious.

Hypothesis 3b

Even the most complex standard will not cover very much of the tree.

This hypothesis was upheld by the study. The X.12 standard covered the largest proportion of the tree and it only covered ten forces and four of six categories. There are at least two possible reasons for this:

1) Lack of expressive power of current message structures limits what they can express, and
2) People do not expect computers to express certain types of messages. For example, a condolence in any type of automated system would be perceived as inappropriate.

Hypothesis 4

The tree will not be segmented by different standards.

The study supported this hypothesis. There was significant overlap between standards.

6 Discussion

6.1 Implications

The results of the empirical study presented in this paper paint an interesting picture. The message structure implicit in three separate electronic communication standards all map

onto the *F(P)* framework proposed by SAT. These standards were not defined with SAT in mind nor was SAT defined with electronic communication in mind. It is hard to think of a reason that the mapping should have been successful except for the possibility that this framework (or something like it) is correct. These three standards were a convenient sample but there is no reason to think that other standards would present a significantly different result. This is strong evidence in favor of SAT.

Not only is SAT supported but the researchers who have proposed that SAT be used as the basis for electronic communication systems are also supported [3, 11, 13, 14, 23, 25, 24]. This study does not indicate the ultimate correctness or utility of SAT. It does indicate that it is feasible to construct electronic messages within the SAT framework.

Previous researchers have demonstrated the benefits of explicitly representing the illocutionary force in electronic messages (e.g., 11, 16, 23, 24). These benefits include better message handling, better message retrieval, and the ability to automate more complex tasks. Previous researchers have also claimed that iterated forces occur naturally within messages [11, 9, 10]. This study found a significant example of iterated operators in the Apple Events messages. If this is a general finding, then it would be a compelling reason to use a formal language for communication that can explicitly (and naturally) represent this information [11, 9, 10].

The current study found no reason to refute the claim that electronic messages have one or more illocutionary forces. I concur with previous researchers who have proposed that the illocutionary force be explicitly represented in electronic messages. Further, I offer this study as evidence that it can be done.

An implication which might be drawn is that forces which had no messages mapped to them are somehow ill-defined, ill-conceived, or somehow faulty. This should not be concluded for at least two reasons. First, the sample is too small to conclude no messages fit into these categories. Second, these are simple electronic messaging systems whose expressive power and purpose are limited. Many normal language expressions would fit into these categories.

The distribution across forces does have practical implications. If a person were to implement a communication system based on SAT, he should determine how he should handle those forces that are most heavily represented.

There are limits to the conclusions that can be drawn from this study. This study does not support the contention that SAT *should* be the basis for electronic communication systems. Just because you can map from the messages to the framework does not mean that this should be done. It is possible that the framework is wrong or that the explicit representation of the framework is not useful. Previous research has indicated that representing the framework is useful for message processing and message retrieval [e.g., 11, 14, 23, 24]. On the other hand little has been concluded about the validity or superiority of any one illocutionary force hierarchy.

6.2 Future research

Though this study indicates the answers to some questions, many more questions remain about the utility of applying SAT to electronic communication systems. Clearly, the *F(P)* framework is not sufficient for a system to understand an electronic message. There are not just 26 (or any finite X) number of messages. A request to paint the house is different from a request to buy two hundred gallons of paint. Both would be represented

as requestives in the *F(P)* framework. A message's content and context—i.e., that information contained in the *P*—must be represented to allow systems to process the message. Defining a general system for representing this information would contribute to the utility of a SAT-based message system.

Another research area is hinted at in the discussion of Apple Events under Hypothesis 3 above. The Bach & Harnish categorization forces all questions into two illocutionary forces: requestives and questions. Questions require a yes or no response. Requestives are all other types of questions. As was shown in Table 4, systems can request to inform, retract, permit, or do. The Bach & Harnish hierarchy draws no distinction between these types of requestives but does separate a yes or no question from the requestive. This seems to be somewhat arbitrary. Researchers need to determine what types of questions they want to ask. Great disagreement in the philosophy literature exists as to what types of questions can be asked (e.g., see 8). Researchers should then determine if an addition to the hierarchy is needed to handle these new types or if it is correct the way it stands.

In addition to this simple question about the Bach & Harnish hierarchy, there is the question of whether or not there is a better hierarchy. One good place to start investigating this question is to map these message standards (and others) to this hierarchy and alternatives. This process can reveal whether the mapping can be done and can also reveal weaknesses or strengths of the hierarchy (as we saw in this study).

A more fundamental question is whether or not SAT is correct. It may be the case that people do not communicate in the manner described by SA theorists. This line of research will not prove that SAT is correct but it could provide some support for the contention that it is correct or incorrect. If a robust, expressive, and powerful communication system can be built based on SAT, then supporters of SAT would have strong evidence that it is correct. On the other hand, if no such system can be built, then supporters of SAT would have to explain the failure. Currently, however, SAT represents the best work of linguists and philosophers of language describing how people communicate. The study described in this paper represents one effort that takes this finding seriously. Much effort remains before we are finished.

Bibliography

[1] American National Standards Institute, I., *American National Standards for Electronic Business Data Interchange*, New York, NY, 1986.

[2] *Apple Event Registry: Standard Suites*, Apple Computer, Inc., Cupertino, CA, Winter, 1992.

[3] Aurämaki, E., Lehtinen, E., and Lyytinen, K., "A Speech-Act-Based Office Modeling Approach," *ACM Transactions on Office Information Systems*, vol. 6, no. 2, 126–152, April 1988.

[4] Austin, J.L., *How To Do Things With Words*, 2nd. Harvard University Press, 1975.

[5] Bach, K. and Harnish, R.M., *Linguistic Communication and Speech Acts*. MIT Press, 1979.

[6] Ballmer, T. and Brennenstuhl, W., *Speech act classification*. Springer-Verlag, 1981.

[7] Cook, T.D. and Campbell, D.T., *Quasi-Experimentation: Design & Analysis Issues for Field Settings*. Houghton Mifflin Company, 1979.

[8] Harrah, D., "The Logic of Questions," in *Handbook of Philosophical Logic*, Gabbay, D. and Guenthner, F., Eds. D. Reidel Publishing Co., 1984, pp. 715–764.

[9] Kimbrough, S.O. and Lee, R.M., "On Illocutionary Logic as a Telecommunications Language," in *Proceedings of the Seventh International Conference on Information Systems,* Leslie Maggi, , San Diego, CA, December 15–17 1986, pp. 15–26.

[10] Kimbrough, S.O. and Thornburg, M., "On Semantically-Accessible Messaging in an Office Environment," in *Proceedings of the Twenty-Second Hawaii International Conference on System Sciences,* University of Hawaii, IEEE Computer Press, 1989.

[11] Kimbrough, S.O. and Moore, S.A., "Message Management Systems: Concepts, Motivations, and Strategic Effects," *Journal of Management Information Systems,* vol. 9, no. 2, 29-52, Fall 1992.

[12] Kimbrough, S.O., Moore, S.A., and Thornburg, M., "On Messaging with Semantic Access in an Office Environment," December 1993, Working papers series, University of Pennsylvania, Wharton School, Decision Sciences Department.

[13] Lehtinen, E. and Lyytinen, K., "Action Based Model of Information System," *Information Systems,* vol. 11, no. 4, 299–317, 1986.

[14] Medina-Mora, R., Winograd, T., Flores, R., and Flores, F., "The Action Workflow Approach to Workflow Management Technology," in *Proceedings of the Conference on Computer-Supported Cooperative Work,* Turner, J. and Kraut, R., ACM SIGCHI & SIGOIS, ACM Press, 1992, pp. 281–288.

[15] Moore, S.A., "Saying and Doing: Uses of Formal Languages in the Conduct of Business," Ph.D. thesis, University of Pennsylvania, Philadelphia, PA, December 1993.

[16] Moore, S.A. and Kimbrough, S.O., "Message Management Systems at Work: A System for Office Communications," 1994, forthcoming in Journal of Organizational Computing.

[17] Nolan, R., *Cognitive Practices: Human Language and Human Knowledge.* Cambridge, MA: Blackwell Publishers, 1994.

[18] Searle, J., *Speech Acts: An Essay in the Philosophy of Language.* Cambridge University Press, 1969.

[19] Searle, J.R., *Expression and Meaning.* Cambridge, England: Cambridge University Press, 1979.

[20] Searle, J. and Vanderveken, D., *Foundations of Illocutionary Logic.* Cambridge University Press, 1985.

[21] *S.W.I.F.T. Securities Market Binder,* Society for Worldwide Interbank Financial Telecommunication, New York, NY, September, 1993.

[22] Sperber, D. and Wilson, D., *Relevance: Communication and Cognition.* Harvard University Press, 1988.

[23] Winograd, T. and Flores, C.F., *Understanding Computers and Cognition.* Norwood, NJ: Ablex, 1986.

[24] Woo, C.C., "SACT: A Tool for Automating Semi-Structured Organizational Communication," in *Proceedings of Conference on Office Information Systems,* Lochovsky, F. and Allen, R.B., ACM SIGOIS & IEEECS TC-OA, ACM Press, April 1990, pp. 89–98.

[25] Woo, C.C. and Chang, M.K., "An Approach to Facilitate the Automation of Semistructured and Recurring Negotiations in Organizations," *Journal of Organizational Computing,* vol. 2, no. 1, 47--76, 1992.