The Nature of Requests and Commitments in Email Messages

Andrew Lampert†‡, Robert Dale‡ and Cécile Paris†

†Information Engineering Lab, CSIRO ICT Centre
‡Centre for Language Technology, Macquarie University
Locked Bag 17, North Ryde, NSW 1670 Australia
Centre for Language Technology, Macquarie University NSW 2109
andrew.lampert@csiro.au. rdale@ics.mq.edu.au. cecile.paris@csiro.au

Abstract

It has long been established that many workplace tasks are managed through email communication, and that these tasks involve the exchange of requests and commitments. Users would be better able to manage and monitor tasks in their email if systems could identify the utterances which place responsibility for action on themselves or others. Such systems require a robust understanding of which utterances convey requests and commitments. Previous attempts to classify similar phenomena in email have mostly been at the message level and have lacked detailed and robust category definitions that allow unambiguous classification at the utterance level. To address this gap, this paper presents precise definitions for classifying requests and commitments in email, based on concepts from Speech Act Theory, and informed by the results of two independent manual annotation experiments using data from the Enron email corpus. The specific surface realisation of requests and commitments in email are also considered, with the aim of clarifying how a range of potentially difficult cases should be dealt with. This paper thus contributes a well-grounded definitional basis for the classification of task-oriented speech acts in email.

Introduction

For some time, studies of email usage have highlighted that people routinely use email for managing task requests in the workplace (e.g., (Mackay 1988; Ducheneaut and Bellotti 2001)). With the volume of email ever increasing, users regularly feel overloaded when trying to manage tasks through email (Whittaker and Sidner, 1996; Bellotti et al. 2003). Many struggle to give appropriate attention to the requests and commitments buried in their email that require action or response.

The overall goal of our research is to create tools that assist email users by automatically detecting requests and commitments in incoming and outgoing email. The motivation for this is well explained by observations from ethnographic research into the use of electronic messaging at IBM (Murray 1991):

[Managers] would like to be able to track outstanding promises they have made, promises made to them, requests they've made that have not been met and requests made of them that they have not fulfilled.

More recent studies of task-focused email usage have also identified problems with "keeping track of lots of concurrent actions: One's own to-dos and to-dos one expects from others" using existing email clients (Bellotti et al. 2003). We aim to assist people to identify and monitor outstanding tasks, both for themselves and their correspondents, by focusing attention on requests and commitments that remain unfulfilled. This can be done by augmenting email clients with features such as action-oriented summaries of messages and threads, task-based navigation and visualisations, and dashboards that provide overviews of the state of an email inbox or collection with much greater fidelity than is possible with current tools.

A major challenge in identifying requests and commitments is that they often don't wear their speech act 'on their sleeve'. There is a many-to-many relationship between any request or commitment and its possible realisations. This results in utterances that, from their surface form, appear to be requests or commitments but do not carry such intent, and vice-versa. For example, the utterance *Let me see what I can do* has an imperative form, giving it the appearance of a request; in fact the utterance conveys a commitment to future action by the author. Such utterances create significant challenges for classifying speech acts, particularly in email, which, apart from common ungrammatical usage and misspellings, has been shown to contain a high proportion of indirect speech acts (i.e., speech acts whose surface form does not directly reflect the underlying intent) (Hassell and Christensen 1996). Requests and commitments whose interpretation depends on utterance-external context (e.g., the relationship between the sender and recipient, or previous conversation history) offer further challenges.

Automatic identification of requests and commitments requires a set of definitions that allow human annotators, and ultimately computational algorithms, to objectively identify them. Previous work has attempted to create definitions, but without sufficient detail for our needs. In order to derive a set of robust definitions for requests and

commitments, we performed two independent annotation experiments. Results from these experiments strongly suggest that quite detailed definitions are required to reliably identify requests and commitments at the utterance level. Further analysis of the results led us to the robust definitions for requests and commitments that form the major contribution of this paper.

Related Work

Our work builds on influential ideas proposed by Winograd and Flores (1986) in taking a language-action perspective and attempting to classify speech acts in email. While this differs from the approach of most existing email systems, which routinely treat the content of email messages as simple bags-of-words, there is a growing body of research applying ideas from Speech Act Theory (Austin 1962; Searle 1969) to analyse and enhance email communication. Most of this work focuses on annotating speech acts (more properly, speech-act-inspired units) at the message level, as in (Khosravi and Wilks 1999; Cohen, Carvalho and Mitchell 2004; Leuski 2004; Goldstein and Sabin 2006). Like (Corston-Oliver et al. 2004), however, we observe that a single email message may contain multiple requests and commitments on a range of tasks and topics. We thus focus on a more fine-grained utterance-level classification.

In our attempts to define requests and commitments, we looked to previous work that has classified such phenomena in email. As we outline below, all the existing definitions have significant problems: either being tied to a particular domain, mixing conversation state into the definition of speech acts, or lacking detail that allows unambiguous classification at the utterance level.

One early message-level annotation scheme was defined by Camino et al. (1998). They analysed requests and corresponding answers in the context of exploring the efficacy of structured email messages. They classified requests based on the form of the expected answer (e.g., one of list choice, several of list choice, date, time, free text), and also distinguished requests requiring an email response from those requiring external action. One weakness is in the lack of consideration of conditional requests, other than those that request acknowledgement of sent information being useful or correct. The classification taxonomy also excludes any commitment categories.

Khosravi and Wilks (1999) focus on classifying three classes of requests: Request-Action, Request-Information and Request-Permission. Unfortunately, the cue-phrase-based rules they define for identifying requests are very specific to the computer support domain from which their email corpus was drawn. Their categorisation also mixes aspects of the surface realisation (direct/indirect) with the speech act function (action/information/permission), which as we discuss later, is not desirable.

Cohen, Carvalho and Mitchell (2004) developed separate taxonomies of *Verbs* and *Nouns* to embody their definitions for message-level "Email Act" classification.

The Verb taxonomy defines categories for *Request*, *Commit*, *Propose*, *Amend* and *Refuse* that are all related to requests and commitments. Unfortunately, the definitions mix concepts of conversation state and speech acts in a manner that ties the definition of specific acts to particular conversation states. So, for example, every *Commit* act is defined as being part of the *Conclude* conversation state. This requires the creation of separate, overlapping categories in the taxonomy: *Propose* and *Amend* acts differ only on whether the act occurs at the beginning of a new conversation or during an ongoing conversation, not due to differences in the speech act being performed. This also causes problems in the coverage of the definitions – it is unclear, for example, how a commitment such as *I'll send you the document with further details* that occurs in a conversation-initiating email should be classified. More problematic is the lack of detail in the published definitions, which leaves many cases without a clearly correct classification. Examples include lack of clarity about whether advice or suggestions such as *I think you should include the figures in section 2* should be classified as *Requests*. Similarly, no guidance is given for how to classify conditional requests and commitments, except for a specific class of acts that request response and conditionally commit the sender if the recipient responds (e.g., meeting requests), that are identified as *Propose* acts.

Goldstein and Sabin (2006) also classified speech acts in email using a complex, message-level annotation scheme that includes 23 speech act categories. Like Cohen et al., their taxonomy includes conversation state in the speech act definitions, through the importing of concepts of forward-looking and backward-looking functions from DAMSL (Core and Allen 1997). Many of the speech act categories are distinguished only by their sequence in conversation. For example, requests and commitments that respond to previous acts are assigned different categories depending on whether further response is expected. This distinction is, however, not captured for requests or commitments in an email message that initiates an email conversation. Requests and commitments are also defined as mutually exclusive – a message cannot simultaneously request something from the recipient and commit the message sender. This limitation is problematic for both message and utterance-level annotation. For an utterance such as *Let me know if you'd like a copy of the document* which requests a response and conditionally commits the sender to sending the document, there is no appropriate classification. Our annotation experiments show that such utterances are frequent in business email.

Leuski (2004) offers yet another speech-act-inspired taxonomy used to categorise email messages and to distinguish the roles of different email authors. Leuski's taxonomy focuses on requests, distinguishing four separate request acts based on whether the act seeks information, advice, action or a meeting. No categories are included to capture commitments. A general weakness is the lack of any detail in the category definitions; the only information provided is a single out-of-context example sentence or

phrase for each category. This makes the categories unable to be easily used for manual annotation or computational classification.

The SmartMail system (Corston-Oliver et. al. 2004) attempts to automatically extract and reformulate action items from email messages for the purpose of adding them to a user's to-do list. Unlike the other taxonomies discussed which focus on categorising email at the message level, the SmartMail taxonomy is designed to be applied at the sentence level, and is thus closest to our requirements.

For requests and commitments, there are three categories of interest: *Task*, *Meeting* and *Promise*. No mention is made of the use of the *Promise* category, and no explanation is given for what constitutes a promise. *Meetings* are defined without further detail or discussion as "a proposal to meet". *Tasks* are sentences that "look like an appropriate item to add to an ongoing 'to do' list", with the explicit exclusion of simple factual questions on the basis that the act of responding fulfils any associated obligation (meaning nothing is placed on a task list). Unfortunately, what constitutes a "simple factual question" is not specified, and it is not clear how to distinguish such requests from those that would result in a new action being added to a task-list. For example, an utterance such as *When is the contract expiring?* might be answered immediately without the need to add anything to a task list if the recipient knows the response. Alternatively, it might require the recipient to issue requests to other people or systems; tasks that might reasonably be added to the recipient's task list. In either case, an obligation is placed on the recipient, and unlike Corston-Oliver et al., we believe this should be reflected in the classification of the utterance as a request under either interpretation. Annotators are also restricted to applying a single tag to each sentence, meaning that a sentence cannot embody both a request (*Task*) and a commitment (*Promise*), which as we have illustrated earlier, is an artificial restriction.

In general, none of the existing definitions of requests or commitments provide enough detail to robustly and unambiguously support the classification of requests and commitments at the utterance level. As we outline below, our annotation experiments show that there are many cases that don't resemble the simple canonical examples forming the basis of the surveyed definitions. This may be partly due to the focus on message level classification; while looser definitions may be sufficient to achieve reasonable human agreement that an email contains a request or a commitment, it is more difficult and subjective to determine which of the utterances in the message actually embody the request or commitment. This more challenging problem requires a much more tightly defined set of definitions to guide utterance-level categorisation.

Manual Annotation Experiments

We began our exploration through annotation experiments, to determine levels of human agreement in identifying requests and commitments in email.

First Experiment

Our first experiment used guidelines similar to many of those surveyed. We defined requests as sentences carrying an expectation that the recipient of the email should take action, and commitments as sentences carrying an expectation that the sender is promising future action from themselves or on behalf of another person. Our definitions are more fully described in (Lampert et al., 2007).

For all our annotation experiments, we used the database dump of the Enron email corpus released by Andrew Fiore and Jeff Heer.[1] This corpus is the result of a substantial amount of processing of the contents of the raw Enron corpus, including automated removal of duplicate email messages and normalisation of sender and recipient names. The resulting corpus has just over 250,000 email messages. Like all publicly released versions of the Enron email corpus, no attachments are included. The data annotated in our experiments is all extracted from the message bodies.

Annotators marked 350 sentences from 54 email messages. Inter-annotator agreements between three expert annotators (the authors) are shown in Table 1. The most significant sources of disagreement occurred for conditional commitments and implicit requests.

Table 1:Cohen's Kappa agreements for First Manual Annotation Experiment

	A + B	B + C	A + C	3-Way
Commitments	0.62	0.55	0.45	0.54
Requests	0.79	0.74	0.83	0.78

Second Experiment

Analysis of disagreements from our first experiment led us to include greater detail and guidance for categorising edge cases. In particular our revised guidelines specified that both requests and commitments should be classified as conditional or unconditional, to emphasise the inclusion of conditional requests and commitments. We also asked annotators to mark requests and commitments as either implicit or explicit, hoping to encourage more careful consideration of implicit speech acts. Annotators marked 750 sentences from 350 email messages. Agreements for our second experiment (same three annotators) are summarised in Table 2. Agreement increased for both requests and commitments from our first experiment, particularly for commitments. The improved agreement was due largely to better guidance about how to annotate conditional commitments (and to a lesser degree, requests). Additional edge case analysis of annotator disagreements led to our final definitions, presented in the next section.

Table 2: Cohen's Kappa agreements for Second Annotation Experiment

	A + B	B + C	A + C	3-Way
Commitments	0.85	0.63	0.73	0.74
Requests	0.87	0.74	0.79	0.80

---

[1] http://bailando.sims.berkeley.edu/enron/enron.sql.gz.

Defining Requests and Commitments

The concepts and definitions presented are based on analysis of the results from our two annotation experiments described in the previous section. In particular, analysis of cases of disagreement from these experiments provided insight into the edge cases of classification.

Basic Concepts

The ontological foundation of our taxonomy is the notion of an **action**. Actions are carried out by **agents**; a **request** is, in essence, the placing of an **obligation** by one agent on another agent to carry out the requested action; a **commitment** is the taking on, by some agent, of an obligation to carry out an action. Both requests and commitments may be conditional, in that some state must hold before the request or commitment can be carried out.

In line with the above, requests and commitments are formally defined as tuples of either three or four elements:[2]

Request = ⟨Action, Requestor, Requestee, [Condition]⟩

Committment = ⟨Action, Committor, Committee, [Condition]⟩

We also consider the **realisation** of a request or commitment to be distinct from the request or commitment itself. This allows us to talk of different realisations of the same request or commitment, so that a statement like *I made this request twice already* makes sense. It also allows us to adopt the traditional distinction between direct and indirect speech acts as alternative means of realising requests and commitments. So, we also have:

UtteredRequest = ⟨Direct/Indirect, Utterance, Request⟩

UtteredCommittment = ⟨Direct/Indirect, Utterance, Commitment⟩

Sometimes it is not easy to determine the **locus** of a request or commitment in a text, i.e. the actual piece of an email message that corresponds to the request or commitment as opposed to something else. In the work described here, we restrict ourselves to cases where it is fairly clear that the locus is an individual sentence, but there are cases where this is not so straightforward; see (Lampert et al. in submission) for a detailed discussion of these edge cases. We will generally use the term **utterance** as a size-independent label for the text fragments used to convey requests or commitments.

Requests

For present purposes, we can think of requests as the actionable utterances in an email message, from the point-of-view of a particular email recipient (defined as the owner of the inbox in which we are classifying email messages). A useful, though not exhaustive, diagnostic test for determining whether an utterance corresponds to a request is to ask whether the utterance could result in a task being added to a 'to-do' list. Requests for action, information, permission, confirmation, agreement, evaluation, interpretation, sympathy (Labov and Fanshel 1977) can all, under certain conditions, be considered as actions that might be listed on a to-do list.[3]

In the data we have examined, requests typically place an obligation on the recipient to either schedule or complete an action, or to respond by providing information to the sender. Some linguists have previously distinguished between speech acts that require a physical response from those that require a verbal or information response – e.g., (Sinclair 1975). We follow the approach of Searle's original taxonomy, and make no such distinction in our definition of requests. We thus explicitly include questions requiring an informational response as requests, since they are attempts by the sender to get the recipient to answer, i.e., to perform a speech act. So both of the following constitute requests in our model:

(1) *Please attend the meeting this morning.*

(2) *When does your flight arrive?*

Note that many information requests are often so straightforward or simple to attend to that we would not typically think of them as being candidates for inclusion in a 'to-do' list, perhaps on the grounds that it would take just as much effort to add the item to such a list as it would to respond to it. However, they still require the recipient to act, and so they constitute requests in our approach.

Requests for information are typically surface form questions. Not all questions are requests, of course, **rhetorical questions** being one class of non-request questions.[4]

Another category of question that our annotation guidelines rule out as not being a kind of request are what we refer to as **pleasantries**. These are defined as polite social utterances that, while often resembling requests for information (e.g., *How are things?*), place on the recipient only a very weak or optional obligation to respond or act. One exception to this rule is when an email message contains only sentences that would normally be marked as pleasantries. In such cases, we elevate such sentences to

---

[2] There are other elements we might include in a formal definition of Requests and Commitments; an obvious one is a date by which some action is to take place. To simplify discussion here, we leave the elaboration of these elements for future work.

[3] We might wish to define subcategories of Request corresponding to each of these distinct types. This is compatible with the approach we adopt here, although in our view the consequent complexity it adds to the annotation task means that such sub-categorisation is of questionable value.

[4] Determining whether or not something is a rhetorical question as opposed to a real question may require reasoning about the context; consider a question like *Do we really want to do this?* We will address this issue in future work.

the status of a request, on the basis that ignoring even a weak obligation in this context could be deemed inappropriate.

One interesting consequence of our distinction between an underlying request and its realisation in linguistic form is that it is possible for different requests to be realised in the same way. For example, the sentence *Can you attend this afternoon?* can, in some contexts, be classified as both a request for information ('Are you coming?') and a request for action ('Make sure you attend this afternoon!'). Various factors will influence the classification decision, including details about the organisational relationship between sender and recipient. For example, if the above sentence were uttered by a manager to their direct report, the power relationship could easily dictate that it should be interpreted as a request for action, while this interpretation may be less likely were the email sent from the subordinate to their manager.

Requests *not* to act or inform (also referred to as **prohibitives** by some researchers) are not classified as requests, since they are typically not something that can be actioned, managed or completed in the same way as the task requests in which we are interested.

Commitments

We define commitments as utterances that promise or offer future action by either the sender or an identifiable third-party: in terms of the formal definition provided earlier, although the **Committor** and the **Committee** are often the same person, this need not be the case, as the distinction between the following examples illustrates:

(3) *I'll send you the minutes on Friday.*

(4) *Cathy will send you the minutes on Friday.*

A third-party agent who is committed to act can be either another person, a group of people or some organisation; so an important element of characterising a commitment is the identification of the intended agent of that action; effectively, whose to-do list does the action belong to?

Promises of action by the sender may not explicitly state the sender as the agent of the future action. The deciding factor is whether the recipient would reasonably assume that the sender is taking responsibility for the future action. For example, in one of our annotated emails, the sender notes *More to follow* at the end of their email. Although not stated explicitly, all annotators agreed that this was a commitment to send more information by the sender.

In the case of committing third-parties to future action, there must be an identifiable third-party agent of the action, as in example (4) above. A consequence of this is that we do not consider 'unassigned' actions to be commitments. A typical example would be *Lunch will be served at the meeting*, but ultimately the wider context of the message will determine whether or not there is a specified agent.

An interesting edge case stems from commitments not to act, as in *I won't send any further updates.* Such negative commitments are not considered commitments to act, as there is, by definition, no future action to follow-up on.

Another interesting border case is the collection of socially-motivated sentences that resemble commitments but assign little or no obligation on any party to act. As for requests, we call such sentences pleasantries. The inclusion of such utterances is largely formulaic, and their function is not primarily to commit to any future action, but rather to satisfy the social norms of communication. Variants on *Let me know if you have any questions* are particularly common. We require the context of the entire email message to be considered in order to distinguish between when such a statement is present due to social convention and when it functions as an actual offer or commitment for future action or response.

Conditionality

Both requests and commitments may be conditional, in the sense that the scheduling or execution of the specified action is expected only if some stated condition is satisfied. For both requests and commitments, the condition is usually stated in the utterance containing the request or commitment. It may, however, be stated elsewhere in the email message, and understood to apply to the request or commitment utterance, as in the following example:

(5) *Are you coming to the meeting? If so, send me items for the agenda by 2pm.*

Note also that conditional commitment binds the relevant party to action only when the stated condition is satisfied:

(6) *I'll send a draft if Harry responds before Friday.*

Surface Realisations

As noted, we distinguish underlying requests and commitments from their surface linguistic form. We can then categorise their realisations as either direct or indirect.

Direct requests state the actual request explicitly, either in the form of an imperative command (see example (1) above) or as a question where a response to the literal interpretation of the question would result in a speech act with the required content (see example (2) above). Note that the content may be information, permission, interpretation and so on, as noted in our definition.

Indirect Requests may or may not state the required action, but don't explicitly instruct the recipient to act or respond with the required speech act. Thus, (7) and (8) are direct requests, while (9) is indirect, since a literal response would be a yes/no answer.

(7) *Please send me my curves and trades for Jan 18.*

(8) *What were they?*

(9) *Can you send my curves and trades for Jan 18?*

Similarly, direct commitments state the promise for future action explicitly. Indirect commitments do not state the promise to complete future action. Note that in either case, as for requests, the actual action may or may not be stated

explicitly. So, (10) and (11) represent direct commitments, while (12) is considered indirect.

(10) *I'll send you the document today.*

(11) *I'll do it.*

(12) *Leave it to me.*

Conclusion

We have presented a set of robust definitions for requests and commitments in email communication that we hope will be widely applicable within the email research community. These definitions are based on Speech Act Theory, and informed by the results of two independent experiments in which requests and commitments were manually annotated in approximately 1000 sentences from the Enron email corpus. We cleanly separate the definitions of requests and commitments from aspects of their surface realisation, and thus are able to neatly capture distinctions between alternate renderings of the same underlying request or commitment act. In future work, we aim to develop machine-learning based techniques for the annotation of requests and commitments in line with the characterisations laid out here, further extending the approach described in (Lampert et al. 2007).

References

Austin, J. L. 1962. *How to do things with words*. Harvard University Press.

Bellotti, V., Ducheneaut, N., Howard, M. and Smith, I. 2003. Taking Email To Task: The Design and Evaluation of a Task Management Centred Email Tool. *Computer Human Interaction Conference*, Ft Lauderdale, Florida, USA.

Camino, B. M., Milewski, A. E., Millen, D. R., Smith, T. M. 1998. Replying to email with structured responses. *International Journal of Human-Computer Studies*, Volume 48, Issue 6, pp. 763-776.

Cohen, W. W., Carvalho, V. R. and Mitchell, T. M. 2004. Learning to Classify Email into "Speech Acts". In D. Lin and D. Wu (eds.). *Conference on Empirical Methods in Natural Language Processing*, pp. 309-316. Barcelona, Spain: Association for Computational Linguistics.

Core, M. and Allen, J. 1997. Coding dialogs with the DAMSL annotation scheme. *AAAI Fall Symposium on Communicative Action in Humans and Machines*, pp. 28-35. Cambridge, MA.

Corston-Oliver, S. H., Ringger, E., Gamon, M. and Campbell, R. 2004. Task-focused summarization of email. *ACL-04 Workshop: Text Summarization Branches Out*.

Ducheneaut, N. and Bellotti, V. 2001. E-mail as habitat: an exploration of embedded personal information management. *Interactions* 8(5): 30-38.

Goldstein, J. and Sabin, R. E. 2006. Using Speech Acts to Categorize Email and Identify Email Genres. *Proceedings of the 39th Hawaii International Conference on System Sciences*, pp. 50b.

Hassell, L. and Christensen, M. 1996. Indirect Speech Acts and Their Use in Three Channels of Communication. *Proceedings of the First International Workshop on Communication Modeling - The Language/Action Perspective*, Tilburg, The Netherlands.

Khosravi, H and Wilks, Y. 1999. Routing email automatically by purpose not topic. *Natural Language Engineering*, Volume 5, Issue 03, pp. 237-250

Labov, W. and Fanshel, D. 1977. *Therapeutic Discourse: Psychotherapy as Conversation*, Academic Press.

Lampert, A., Paris, C. and Dale, R. 2007. Can Requests-for-Action and Commitments-to-Act be Reliably Identified in Email Messages? In Proceedings of the 12th Australasian Document Computing Symposium, pp. 48-55. Melbourne, Australia.

Lampert, A., Dale, R. and Paris, C. (in submission). Why Identifying Requests and Commitments in Email Messages is Hard. Submitted to the *Coling 2008 Workshop on human judgements in Computational Linguistics*, Manchester, UK.

Leuski, A. 2004. Email is a stage: discovering people roles from email archives. Proceedings *of Annual ACM Conference on Research and Development in Information Retrieval*, Sheffield, UK.

Mackay, W. E. 1988. More Than Just a Communication System: Diversity in the Use of Electronic Mail. *ACM conference on Computer-supported cooperative work*, pp. 344-353. Portland, Oregon, USA: ACM Press.

Murray, D. E. 1991. *Conversation for Action: The Computer Terminal As Medium of Communication*. John Benjamins Publishing Co.

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

Sinclair, J. and Coulthard, R. M. 1975. *Towards and Analysis of Discourse - The English used by Teachers and Pupils*. Oxford University Press.

Whittaker, S. and Sidner, C. 1996. Email Overload: exploring personal information management of email. *ACM Computer Human Interaction conference*, pp. 276-283. ACM Press.

Winograd, T. and Flores, F. 1986. *Understanding Computers and Cognition*. Norwood, New Jersey, USA: Ablex Publishing Corporation.