

Requests and Commitments in Email are More Complex Than You Think: Eight Reasons to be Cautious

Andrew Lampert
CSIRO ICT Centre
Locked Bag 17
North Ryde 1670
Australia

Andrew.Lampert@csiro.au

Robert Dale
Centre for Language Technology
Macquarie University 2109
Australia
rdale@ics.mq.edu.au

Cécile Paris
CSIRO ICT Centre
Locked Bag 17
North Ryde 1670
Australia

Cecile.Paris@csiro.au

Abstract

Many workplace tasks are managed through email communication, involving the exchange of requests and commitments. Our aim is to build a tool that can automatically identify and manage such requests and commitments. A detailed analysis of real data, however, reveals a range of interesting edge cases that make even human annotation of training data difficult. In this paper, as an important step in the development of annotation guidelines for wider use in the growing email processing community, we identify eight categories of problematic data and propose how they should be handled in annotation and extraction tasks.

1 Introduction

Our aim 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 in the workplace (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.

Other studies have also highlighted that people routinely use email for managing requests and commitments (e.g., (Mackay, 1988; Ducheneaut and Bellotti, 2001)), but struggle to give appropriate attention to requests and commitments that require action

or response because they are buried in their email (Whittaker and Sidner, 1996). 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).

To provide support here, we are working to augment existing email clients with features such as action-oriented summaries of email 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.

In working towards this goal, we have conducted a series of manual annotation experiments, exploring the level of human agreement that is achievable in identifying requests and commitments in email messages. Previous work has often relied on canonical examples as the basis for simple definitions of requests and commitments. Our experiments have found that this level of detail is insufficient to obtain reliable interannotator agreement. Indeed, we have discovered a range of edge cases not well represented by simple examples or definitions. The influence of factors such as politeness lead to a variety of more complex and indirect requests and commitments than other researchers have assumed. Such complexity is found when classifying requests and commitments at both the message-level and the utterance-level.

Our analysis suggests that the range of these edge cases is not an unstructured collection of widely-varying ‘hard cases’; rather, we believe that there

are a number of major distinct categories, and that it is useful to identify these categories.

In this paper, we provide just such a categorisation of these edge cases. For each category we propose how the data should be annotated with regard to the presence or absence of requests and commitments. Note that there are two distinct points here: depending on the intended use of the data, one might make different decisions as to whether a request or commitment is present in a given instance, but the categories remain. Our principle focus here is to identify and define the categories, so that they will be appropriately considered and acknowledged in subsequent work; they cannot be ‘brushed under the carpet’.

In Section 2 we provide an overview of our definitions of requests and commitments. Section 3 provides a brief overview of the email data we are working with and the annotation experiments we have carried out so far. Section 4 identifies and characterises the eight phenomena that we believe need to be explicitly addressed and indicates how we handle these in our annotation scheme. Section 5 makes some concluding remarks and discusses implications of our analysis for automating the identification of requests and commitments in email.

2 Definitions

In defining requests and commitments, we have looked extensively into previous work that has attempted to classify similar phenomena in email and other media. In (Lampert et al., 2008), we described why none of the existing definitions were suitable for our needs. Briefly, many existing definitions deal only with requests and ignore conditionality (which is very common) as a feature (for example, (Camino et al., 1998; Khosravi and Wilks, 1999; Leuski, 2004)). Others define requests and commitments in terms of specific conversation states, requiring the creation of multiple categories for the same speech act in different stages of a conversation. Often not all of the many combinations of speech acts and conversation states are modeled, resulting in uncodable utterances (Cohen et al., 2004; Goldstein and Sabin, 2006).

Previous work on utterance-level classification (Corston-Oliver et al., 2004) relied on short, simple definitions and canonical examples. These lack the

detail and clarity required for unambiguous classification of the complex requests and commitments we find in real-world email. Since our review of related work, Scerri et al. (2008) have noted some similar concerns. Unfortunately, their interannotator agreement for requests and commitments remains low; we believe this could be improved through the careful consideration of the edge cases we outline in this paper.

Conditionality is an important part of our definitions. Conditional requests and commitments require action only if a stated condition is satisfied. Our early annotation experiments, summarised in Section 3 and detailed in (Lampert et al., 2007), show that annotators require guidance about how to classify conditional requests and commitments to achieve even moderate agreement. Others have since replicated this finding (Scerri et al., 2008).

Another issue is the complexity of the relationship between a request or commitment and its possible realisations in surface text. To deal with this, we explicitly distinguish the realisation of a request or commitment from the underlying request or commitment itself. This allows us to adopt the traditional linguistic distinction between direct and indirect speech acts as alternative means of realising requests and commitments. It also allows us to talk of different realisations of the *same* request or commitment, which, as we note in Section 4, is important for dealing with cases where a given request or commitment is stated more than once in a single message.

Below, we give a high-level overview of our definitions for requests and commitments. More detail about the definitions and our treatment of conditionality and the realisation of requests and commitments can be found in (Lampert et al., 2008).

2.1 Requests

We consider a request to be an utterance from an email sender that places an obligation on an email recipient to:

1. Schedule an action, often by adding an entry to a calendar or task list;
2. Perform an action; or
3. Respond with some speech act.

Requests for action, information, permission, confirmation, agreement, evaluation, interpretation, and sympathy (Labov and Fanshel, 1977) can all function as requests. Some linguists have distinguished between speech acts that require a physical response from those that require a verbal or information response (see, for example, (Sinclair and Coulthard, 1975)). We follow Searle’s original approach (Searle, 1969) and do not distinguish between physical and verbal responses. We thus explicitly include questions requiring an informational response as requests, since they represent an attempt by the sender to elicit an action from the recipient, in the form of a speech act.

2.2 Commitments

We consider a commitment to be an offer or promise made by an email sender for future action or response from some specified agent. The agent who is committed is often, but not always, the sender. In contrast to previous work, we include as commitments utterances that place an obligation on another person, group of people, or organisation. Such third-party commitments are common in the email corpus we are working with, and we believe at least some of these to be important commitments to capture. A useful, though not conclusive, diagnostic test for identifying a commitment is whether you might expect to find the future action on the responsible agent’s task list or calendar.

2.3 Considering the End Goal

As a final consideration in our definitions, we explicitly take account of the *purpose* of identifying and extracting requests and commitments, in the sense of what the end application or intended use of the identified material is.

Intended use is a particularly important consideration, since different types of requests and commitments are likely to be considered relevant for different end uses. Within an email client, for example, the requests and commitments that users would like to have highlighted whilst reading an email message (where a user has access to the entire message content and context) are likely to differ from those that should be extracted and aggregated into a separate task list or task dashboard (where the requests and commitments are removed from their original mes-

| Experiment | Msgs | Sentences | 3-way κ | |
|-------------|------|-----------|----------------|--------|
| | | | Request | Commit |
| Sentences 1 | 54 | 310 | 0.78 | 0.54 |
| Sentences 2 | 350 | 750 | 0.80 | 0.74 |
| Messages | 100 | – | 0.84 | 0.78 |

Table 1: Agreement from manual annotation experiments

sage content and context). The set of tasks to include in an action-based summary of a message or thread, displayed in the inbox alongside existing message metadata (such as sender, date, subject and so on) would be different again.

Our goal is *to identify all requests and commitments that would be useful to highlight for a recipient while they are reading an email message*. Using this application scenario as a diagnostic supports our goal of high recall, since this set of tasks is likely to be a superset of the tasks required for other uses.

3 The Data and Annotation Experiments

Our insights in this paper are based on observations from a series of annotation experiments we conducted using data from the Enron email corpus. We employed the database dump of the corpus released by Andrew Fiore and Jeff Heer.¹ This version of the corpus has been processed to remove duplicate email messages and to normalise sender and recipient names, resulting in just over 250,000 email messages without attachments. All data annotated in our experiments is extracted from message bodies.

Table 1 shows an overview of interannotator agreements from our experiments (as Cohen’s κ scores). These agreements refer to binary agreement about whether a specific sentence or message contains a request or commitment. The first two experiments annotated sentences, while the third experiment involved annotations at the message level.

The *Sentences 1* experiment presented in (Lampert et al., 2007) used guidelines similar to many of those found in previous work. The resulting agreement was moderate for requests, and poor for commitments. *Sentences 2* gave more explicit guidance for annotating conditional and indirect requests and commitments, which led to increased agreement, in

¹ Available at <http://bailando.sims.berkeley.edu/enron/enron.sql.gz>

particular for commitments. The *Messages* experiment was conducted to explore the effect of the different unit size (message vs. sentence) on agreement. Agreement for both requests and commitments increased over the sentence-level experiments. We discuss this experiment further in Section 4.1. All experiments were performed over randomly-sampled messages drawn from the Enron corpus with three annotators.

4 Edge Case Phenomena in Email

We report here on systematic challenges discovered through manually annotating requests and commitments in email messages. We organise our discussion around classes of phenomena. We argue that these must be carefully considered when attempting to classify requests and commitments in email, whether by manual human annotation or automatically.

We apply some guiding principles in attempting to resolve cases of ambiguity and controversy. Where appropriate, we prioritise recall, based on the intuition that for our chosen application of highlighting requests and commitments in a message, missing tasks will be more problematic than mistakenly identifying non-tasks.

4.1 Locus Ambiguity

As indicated in Table 1, our experiments show that annotating at the message level achieves higher inter-annotator agreement than at the sentence level. One contributing factor to this difference stems from ambiguity around the actual location of a request or commitment in the text of an email message. We refer to this problem as ambiguity about the *locus* of requests and commitments.

Figure 1 shows an example of an email with ambiguous request locus.² At the message level, our annotators agree that there is a request for Bruce to contact Sean. When asked to find requests at the sentence level, however, the presence of both an indirect and a direct realisation of the same underlying request leads to ambiguity about the locus: all annotators agreed that Sentence 4 is a request, but Sentence 3 was controversial. After discussion, all

²The numbers here have been added to aid in our discussion of the example and are not part of the message as annotated.

```

From:      Sean
To:        Bruce
1. Bruce,
2. Thanks for the chance to meet you
   and the people you work with.
3. I was just wondering where we are
   at now.
4. If you get a chance, either e-mail
   me or give me a call.
5. Thanks again,
6. Sean

```

Figure 1: Ambiguous Locus

annotators agreed that they would mark Sentence 3 as an (indirect) request *if Sentence 4 was not present*. With both sentences present, however, annotators disagreed about whether Sentence 3 is a request or provides background information. In our sentence-level experiments, we fixed the unit of annotation to be a single sentence. Annotators were thus required to mark Sentences 3 and 4 independently; they could not group sentences together as a single request spanning multiple sentences. Annotators had to choose between marking only the direct realisation (Sentence 4) and marking both sentences as requests.

Our guidance for dealing with cases of locus ambiguity is: **each sentence realising a request or commitment, whether directly or indirectly, should be marked**. This is consistent with our guiding philosophy of prioritising recall, as noted earlier. Thus, in Figure 1, both Sentences 3 and 4 are requests.

The email in Figure 2 offers an example of a meeting request with a different type of locus ambiguity. While it seems clear that the email represents a request at the message level, it is unclear which utterance(s) in the email, if any, should be marked as the locus of the request. If we consider the message also to represent a commitment on Stacey’s part to attend, the same problem occurs for identifying the locus of the commitment.

As illustrated in Figures 1 and 2, locus ambiguity is abstracted away when annotating at the message level. This explains the higher interannotator agreement shown in Table 1.

In future work, we plan to explore further the na-

```

From: Stacey White
To: Undisclosed Recipients
Subject: Weekly staff mtg eb3014

CALENDAR ENTRY: APPOINTMENT

Description: Weekly Staff mtg eb3014

Date: 8/21/2000
Time: 2:30 PM - 3:30 PM (Central
Standard Time)

Chairperson: Stacey W White

Detailed Description:

```

Figure 2: Locus ambiguity for a Meeting Request

```

From: Laura Vuittonet
To: Eric Bass, David Baumbach ...
Subject: Prebid Meeting

Today's Prebid Meeting will take
place in EB32c2 at 3pm.

```

Figure 3: A Meeting Announcement Email

ture of locus ambiguity to determine its importance amongst the other phenomena presented in this paper, and to better understand different patterns of locus ambiguity (for example, contiguous indirect-direct realisations and non-contiguous realisations).

4.2 Meetings

Meeting requests and announcements are very common in the workplace email we analysed (both our own inboxes and the Enron corpus). Our current annotation guidelines provide the following principle: **all meeting announcements are requests, in the sense that they are usually implicit requests to attend.**

Following this principle, there are requests in the messages of both Figures 2 and 3. The implicit request in Figure 3 is, however, more readily identified at the sentence level. Structured messages used for exchanging calendar data via email (e.g., iCalendar or vCalendar messages) are also considered to be requests.

A related issue is whether the sender is committing to attend the meeting. This largely depends

on whether the meeting is an activity in which both sender and recipient are involved, and whether the sender is interpreted to be committing to attending or otherwise acting in relation to the meeting. This can only be resolved using the available context.

An alternative approach, employed by other researchers such as Corston-Oliver et al. (2004), is to deal with this ambiguity by creating a separate *meeting* category, distinct from the request category. This, however, introduces the problem of cases that straddle the boundaries between the categories.

4.3 Pleasantries

Utterances like *Let me know if you have any questions* are extremely common in workplace email. Corston-Oliver et al. (2004) made a similar observation, labelling such utterances “formulaic endings”. In some cases, these utterances actually carry the illocutionary force of a request for the recipient to act and/or a commitment from the speaker. In many cases, however, their presence is mostly a matter of conformance with social norms. Where no obligation is assigned, we call these *pleasantries*.

Pleasantries resemble the surface form of requests or commitments, but place very little obligation, or no obligation at all, on the recipient or other identified agent to act or respond. Correspondingly, we have our third principle: **pleasantries are not requests or commitments.**

Like Corston-Oliver et al., we believe the best approach to distinguishing between pleasantries and actual requests and commitments is to consider the context of the entire message. Annotators are instructed to use their judgement to distinguish when utterances such as *Let me know if you have any questions* should be interpreted as mere social convention and when they represent requests for review and comment and/or an offer of future assistance from the sender.

Even with explicit guidance to consider the entire context of a message, there remain cases of unresolvable disagreement between our annotators around the interpretation of specific utterances. Ultimately, the decision is subjective. Automated tools thus need to adapt to the preferences of a specific user and the norms for communication between different email interlocutors.

In our current annotation guidelines, we have cre-

From: Vince Kaminski
To: Ian MacMillan
Subject: Re: Real Options

Ian,

I shall ask our lawyer to prepare a standard non-disclosure agreement.
 ...
Please, don't distribute it: It's an internal document.

By the way, when is your trip?

Vince

Figure 4: Request for Inaction

From: Phillip Love
To: Bruce Mills, Sladana-Anna Kulic, Chuck Ames, Jad Ryder ...
Subject: Weekly staff mtg eb3014

FYI -
 wanted to let you guys know about a few things:
 1. There is a central desk vacation calendar on my desk by the sun machine - it is a little black date book. **Please put any future requests in there so that we can keep track of who needs off when.**

...

Figure 5: Process Instructions

ated a separate category for pleasantries to help quantify the disagreement that arises from this phenomenon; this will support more detailed analysis of the nature of pleasantries.

4.4 Requests for Inaction

Requests for inaction prohibit action or request negated action. They are sometimes called *prohibitives* (Sadock and Zwicky, 1985). An example (with our emphasis added, here and in subsequent examples) is shown in Figure 4.

By definition, requests for inaction do not require action, so one would not expect such utterances to require a new entry in the recipient's task list. As a result, definitions based on the suitability of an action for inclusion in a recipient's task list would ignore requests for inaction. Clearly, however, such requests place an obligation on the recipient, thus our fourth principle: **requests for inaction are considered to be requests.**³

The use of negation illustrates again the complex relationship between the surface text and underlying speech act. Utterances that request action using a negated surface form, such as *Don't forget to send me your comments* (an alternate realisation of the utterance *Please send me your comments*), are requests for action, not inaction.

³Note that our inclusion of requests for inaction as requests differs from our original stance taken in (Lampert et al., 2008).

4.5 Process Instructions

Another class of edge case requests stems from email messages that contain process instructions or hypothetical requests or commitments: instructions of the kind that one might 'file for later use'. An example is shown in Figure 5.

In our discussions over disagreements, one property played a major role in whether an instruction should be marked as a request: the likelihood of the situation leading to the actual execution of the described action. An utterance such as *In the event of a fire, please leave quickly and calmly via the closest emergency exit* is an example of a low-probability-action that our annotators were not likely to mark as a request.

After careful consideration, our current annotations guidelines **leave the decision about specific instances of process instructions to annotator judgement based on the local context**. We expect that analysis of more annotated data will give us a better empirical understanding of the nature of process instructions, and the distinction between instructions that function as requests and are relevant to our proposed email applications, and those that are not.

4.6 Attachment Review Requests

A common use of email is to disseminate documents as attachments. Attachments are commonly, though certainly not always, accompanied by an utterance along the lines of *Please see attached*. In keeping with our goal to include any type of action as a re-

From: Kristian J Lande
To: Chris Stokley
Subject: Re. MEETING W / TIM

Great. Paul asked if you could put together a summary of your accomplishments in an email. Paul said he will shop you around to some different groups here.

...

Figure 6: Reported Request and Commitment

quest, we consider such utterances to be requests to read, archive, or otherwise act on the attached document.

How should we treat attachments where no utterance in the email body directs or requests the recipient to do anything with the attachment? One possible interpretation is that every attachment carries with it an implicit request to read it or otherwise act on its content. The same, however, could be argued for any message that arrives in a user's inbox, resulting in every message being considered to convey an implicit request to be read. Such an interpretation seems unintuitive and inconsistent with our focus on linguistic email content. Thus, our sixth principle is to **only consider as a request attachments accompanied by a textual request (whether implicit or explicit)**.

4.7 Reported Requests and Commitments

Reported requests and commitments are another complication in real-world email. An example email containing both a reported request and a reported commitment is shown in Figure 6. In this case, both are likely to place an obligation: on the recipient (from the request) and on Paul (from the commitment).

Some reported commitments or requests do not actually place an obligation on any agent. An example is shown in Figure 7.

We have found it difficult to find reliable distinctions between reported requests and commitments that should and should not be marked as actual requests and commitments. Consequently, our principle is to **use the local context to determine the function of a reported request or commitment**.

From: Min528
To: Errol McLaughlin Jr.
Subject: Prebid Meeting

Dear Errol,

I'm sorry for the delay. As mentioned in my previous email, your prize was requested to be delivered out late December. However, it seems that there was a confusion on our vendor's side regarding the shipment.

...

Figure 7: Reported Request that is not a Request

4.8 Third-Party Commitments as Requests

Some commitments from the sender place an obligation on one of the recipients (usually one who is cc'ed). According to our definitions, such commitments also function as requests. Consider, for example, the utterance: *My assistant, Shirley Crenshaw, will send you an updated version of my bio* in Figure 8. When Shirley is a recipient of the email, this commitment also functions as a request to Shirley. In contrast, the same email sent without Shirley as a recipient does not function as a request, since it is not communicated to Shirley. It does still function as a commitment for a third party. Our annotation principle is to **consider any commitment that places an obligation on a recipient of the message as both a request and a commitment**.

Commitments functioning as requests are found in both the Enron corpus and in other workplace email we have analysed. They are typically from managers to their personal assistants or direct reports, or sometimes between peers. Because it is not possible in the case of the Enron corpus to know who receives email sent to group aliases or mailing lists, we only consider an individual's email address in determining whether the committed person is a recipient of the email. This issue would not arise, of course, when annotating from the perspective of a specific user's inbox with the aim of identifying obligations for that user.

```
From: Vince Kaminski
To: Vicky Windsor
Cc: Shirley Crenshaw
Subject: Re: Maths course

Vicky,

My affiliation is Enron Corp. My
assistant, Shirley Crenshaw, will
send you an updated version of my
bio.
...
```

Figure 8: Third Party Commitment as a Request

5 Conclusions

In this paper, we have identified eight phenomena in email messages that make the identification of requests and commitments difficult. We believe that any work that attempts to process requests and commitments in email must acknowledge the existence of these phenomena. For each phenomenon, we have provided a definition and indicated how we approach it in our annotation, summarised as follows:

Locus Ambiguity: Each and every sentence which realises a request or commitment, whether directly or indirectly, should be marked, even if this means that two or more realisations of the same request or commitment are marked.

Meetings: All meeting announcements are considered to be requests, in the sense that they are usually implicit requests to attend.

Pleasantries: Pleasantries are not considered to be requests or commitments.

Requests for Inaction: Requests for inaction are considered to be requests.

Process Instructions: Whether or not a specific instance of a process instruction constitutes a request depends on the context and is left to the annotator's judgement.

Attachment Review Requests: Attachments are only considered to constitute requests when they are accompanied by text in the message that can be interpreted, implicitly or explicitly, as a request.

Reported Requests and Commitments: Whether or not a specific instance of a reported request or commitment constitutes an actual request or commitment depends on the context and is left to the annotator's judgement.

Third-Party Commitments as Requests: Any commitment that places an obligation on a recipient of the message is considered to be both a request and a commitment.

While others may want to deal with these categories of phenomena differently, the categories themselves cannot be ignored. As we have tried to demonstrate in this paper, we cannot rely on simple canonical examples of requests and commitments as being representative of the ways in which these speech acts manifest themselves.

The implications of these edge cases for automating the detection of requests and commitments in email are numerous. Many of the phenomena identified are not easily detectable from the surface form of text in an email, often requiring access to contextual information. This context must thus be modelled in the features used when applying statistical machine learning techniques to the problem. We recognise that some categories will still remain ambiguous because they require context that cannot be easily modelled; in such cases, we would expect a tool to resort to a human-in-the-loop mode of operation.

A more significant issue is the evaluation of automatic request and commitment classifiers, given the challenge of reaching reliable human agreement. It is not possible to measure performance against an objective gold standard. Evaluating in the context of real users' mailboxes and task contexts is likely to be more relevant and meaningful than assessing against an abstract gold standard. This, of course, comes with huge challenges around privacy, access to data and experimental repeatability.

A collection of email data, independently labelled by five annotators according to the guidelines described in this paper, is available for download from <http://annotate.thoughtlets.org>.

References

- Victoria Bellotti, Nicolas Ducheneaut, Mark Howard, and Ian Smith. 2003. Taking email to task: The design and evaluation of a task management centred email tool. In *Computer Human Interaction Conference*, CHI, pages 345–352, Ft Lauderdale, Florida, USA, April 5-10.
- Beatrice M. Camino, Allen E. Milewski, David R. Millen, and Thomas M. Smith. 1998. Replying to email with structured responses. *International Journal of Human-Computer Studies*, 48(6):763–776, June.
- William W. Cohen, Vitor R. Carvalho, and Tom M. Mitchell. 2004. Learning to classify email into “speech acts”. In Dekang Lin and Dekai Wu, editors, *Conference on Empirical Methods in Natural Language Processing*, pages 309–316, Barcelona, Spain.
- Simon H. Corston-Oliver, Eric Ringger, Michael Gamon, and Richard Campbell. 2004. Task-focused summarization of email. In *ACL-04 Workshop: Text Summarization Branches Out*, pages 43–50, July.
- Nicolas Ducheneaut and Victoria Bellotti. 2001. Email as habitat: an exploration of embedded personal information management. *Interactions*, 8(5):30–38, September/October.
- Jade Goldstein and Roberta Evans Sabin. 2006. Using speech acts to categorize email and identify email genres. In *Proceedings of the 39th Hawaii International Conference on System Sciences*, page 50b.
- Hamid Khosravi and Yorick Wilks. 1999. Routing email automatically by purpose not topic. *Journal of Natural Language Engineering*, 5:237–250. Cambridge University Press.
- William Labov and David Fanshel. 1977. *Therapeutic Discourse: Psychotherapy as Conversation*. Academic Press, New York.
- Andrew Lampert, Cécile Paris, and Robert Dale. 2007. Can requests-for-action and commitments-to-act be reliably identified in email messages? In *Proceedings of the 12th Australasian Document Computing Symposium*, pages 48–55, Melbourne, Australia, December 10.
- Andrew Lampert, Robert Dale, and Cécile Paris. 2008. The nature of requests and commitments in email messages. In *Proceedings of the AAAI Workshop on Enhanced Messaging*, pages 42–47, Chicago, USA, July 13. AAAI Press. Technical Report WS-08-04.
- Anton Leuski. 2004. Email is a stage: discovering people roles from email archives. In *Proceedings of Annual ACM Conference on Research and Development in Information Retrieval*, Sheffield, UK.
- Wendy E. Mackay. 1988. More than just a communication system: Diversity in the use of electronic mail. In *ACM conference on Computer-supported cooperative work*, pages 344–353, Portland, Oregon, USA. MIT, ACM Press.
- Denise E. Murray. 1991. *Conversation for Action: The Computer Terminal As Medium of Communication*. John Benjamins Publishing Co.
- Jerry M. Sadock and Arnold Zwicky, 1985. *Language Typology and Syntactic Description. Vol.I Clause Structure*, chapter Speech act distinctions in syntax, pages 155–96. Cambridge University Press.
- Simon Scerri, Myriam Mencke, Brian David, and Siegfried Handschuh. 2008. Evaluating the ontology powering smail — a conceptual framework for semantic email. In *Proceedings of the 6th International conference of Language Resources and Evaluation*, Marrakech, Morocco, May 28-30.
- John R. Searle. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press.
- John Sinclair and Richard Malcolm Coulthard. 1975. *Towards and Analysis of Discourse - The English used by Teachers and Pupils*. Oxford University Press.
- Steve Whittaker and Candace Sidner. 1996. Email overload: exploring personal information management of email. In *ACM Computer Human Interaction conference*, pages 276–283. ACM Press.