

A Tool for Capturing Context-Sensitive Judgements in Email Data

Andrew Lampert
CSIRO ICT Centre
Marsfield 2122 Australia
andrew.lampert@csiro.au

Daniel Breese
CSIRO ICT Centre
Acton 2601 Australia
daniel.breese@csiro.au

Cecile Paris
CSIRO ICT Centre
Marsfield 2122 Australia
cecile.paris@csiro.au

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

Abstract *As the complexity and sophistication of document processing tools increases, we can expect to see techniques that go beyond the syntactic and semantic features of documents to consider the more nuanced, context-sensitive aspects of language use that generally fall within the realm of pragmatics.*

The development of such techniques requires data that has been carefully annotated by human judges or annotators in ecologically-valid real-world contexts. To facilitate such in-context annotation, we describe a plug-in for Microsoft Outlook that provides an interface which allows annotators to judge relevant aspects of collections of email messages. Integration into a widely-used application provides a natural environment that closely approximates the real-world environment in which this data is conventionally created and consumed. We describe how such a tool can be used to collect data for the development of a system that automatically detects and classifies requests and commitments in email.

Keywords Natural Language Techniques and Documents, Document Management

1 Introduction

As we build increasingly sophisticated document processing tools, there is a concomitant need to create tools that are sensitive to ever-subtler uses of language. This in turn requires datasets that capture more nuanced judgements about human language use: such concerns go beyond syntax and semantics to what are called *pragmatic* aspects of language use, which explore the meaning of language in its real world context. And to develop these datasets, we need to develop ecologically-valid tools that closely approximate the real-world environment in which the documents under examination are created and used.

In our own research, which is focused on creating tools that assist email users in identifying and managing requests and commitments contained in incoming and outgoing email, we find a need for exactly such

context-sensitive tools. Part of the assistance we aim to provide is a set of statistical classifiers that automatically detect requests and commitments within the text of email messages. Because identifying requests and commitments requires interpretation of the intent that lies behind the language within email messages, it is natural to approach the problem as one of *speech act identification*. In Speech Act Theory [12], speech acts are categories like *assertion* and *request* that capture the intentions underlying surface utterances, providing abstractions across the wide variety of different ways in which instances of those categories might be realised in linguistic form. Our interest is centred on the specific classes of speech acts that represent requests and commitments, where people are placing obligations for action upon themselves or others via actionable content within email messages.

The motivation for our focus on requests and commitments is well explained by observations from ethnographic research into the use of electronic messaging in the workplace [10]:

[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.

This exchange of requests and commitments has been previously identified as a fundamental basis of the way work is delegated and completed within organisations. Winograd and Flores were among the first to recognise and attempt to exploit this with their Coordinator system [15]: their research into organisational communication concluded that "Organisations exist as networks of directives and commissives". It is on this basis that we focus on requests (directive speech acts) and commitments (commissive speech acts).

Building automatic request and commitment classifiers is a complex task, since the function of conveying a request or a commitment does not neatly map to a particular set of language forms, often involving what are

referred to as *indirect speech acts*. While investigating particular surface forms of language is relatively unproblematic, it is widely recognised that “investigating a collection of forms that represent, for example, a particular speech act leads to the problem of establishing which forms constitute that collection” [1]. Gathering manual judgements or annotations from human annotators is the most common way to address this challenge. This data is then used as training data for machine learning algorithms in the statistical classifiers.

To gather human interpretations of the requests and commitments in email data, we have previously undertaken a series of human annotation experiments [8, 6]. These experiments, together with annotation experiments performed by other researchers, have demonstrated that asking human annotators to identify particular speech acts (such as requests and commitments) in an isolated annotation task using a generic dataset such as the Enron email corpus [4] can lead to considerable inter-annotator disagreement [11, 14, 7].

A major reason for such disagreement is the *out-of-context* nature of the annotation task: the annotators do not have access to the wider context around the email messages as they were originally created and consumed. This makes interpreting and classifying the text of the email messages much more difficult and ambiguous. In the absence of the real-world context, annotators bring their own differing assumptions and interpretation biases to the task, none of which may be obviously right or wrong.

One solution to this problem is to capture the intuitions of annotators as they work with real emails in contexts where they are fully aware of their meaning and import. In this paper, we report on a plug-in we have developed for the popular Microsoft Outlook email software, which allows annotation tasks to be performed *in-context*, using a tool that replicates the real-life usage of email and uses an annotators’ own email messages. In Section 2, we discuss the need for in-context annotation in the face of limitations of out-of-context document judgements, particularly for *pragmatic annotation tasks* such as email speech act recognition. We then briefly review the nature of requests and commitments in email in Section 3, to establish the particular set of speech acts we are interested in identifying. Section 4 describes the functionality and features of the Outlook plug-in we have developed, focusing on how it facilitates in-context email annotation, as well as how it allows users to annotate, manage and interact with manually and automatically identified requests and commitments in their email messages. Of course, in-context email annotation of users’ personal email data raises a series of privacy challenges, particularly given our desire to capture user interaction data for evaluation purposes. We address these privacy and logging issues

in Section 5. Finally, we contrast our work with previous email annotation tools in Section 6.

2 The Need for In-Context Judgements

In previous annotation experiments, we have worked with the Enron email corpus, since this is the only large-scale collection of real-world email messages that is publicly available for research use. Despite the appeal of the Enron data, a significant limitation of using such data is that annotation tasks performed on the Enron corpus are *out-of-context*, since human annotators have no access to the real-world context in which the email messages were created or received.

One of the challenges in attempting to build classifiers for automatically identifying requests and commitments (or other speech acts) within the unstructured text of email messages stems from its nature as an annotation task which is fundamentally concerned with *pragmatic* aspects of language. These are generally distinguished from syntactic and semantic aspects: syntax is concerned with the structure of language, semantics with its literal meaning, and pragmatics with language use within a given context [9]. This context includes the linguistic context, the physical context in which the utterance is made, the personal, social and organisational context including the relationship between the email sender and recipient(s), along with other cultural and cognitive contexts, including the shared background knowledge of the participants in the interaction [1]. All these elements figure into the appropriate pragmatic interpretation of an utterance; this dependence on context thus significantly increases the difficulty of classifying requests and commitments. Importantly, pragmatic annotation tasks such as the identification of requests and commitments, by the very nature of these tasks, can never be as reliably performed without access to such contextual information.

Thus, while using the Enron corpus provides real-world email data and allows other researchers to replicate our experiments, it also leaves annotators to guess at the context when interpreting and annotating messages. Additionally, the out-of-context nature of the Enron messages means we are unable to extrinsically evaluate the utility of the requests and commitments we identify, in terms of how useful the results are in helping people manage their real-world email tasks.

In part to address these limitations, we have developed a plug-in for Microsoft Outlook 2007. The plug-in is intended as a general research tool that serves at least three intended uses:

1. As a natural environment for manual in-context and out-of-context *annotation tool* for speech acts (or other phenomena) in email messages;
2. As a *platform for the extrinsic evaluation* of automatic email request and commitment classifiers; and

3. As a useful email task management *application*.

By integrating into Outlook, the plug-in provides an ecologically-valid annotation environment that can be used for performing both out-of-context annotation experiments such as those we've performed using the Enron email data, as well as facilitating in-context experiments that use email messages from annotators' own email mailboxes.

Using the mail of individual annotators creates significant privacy challenges, along with problems of subjectivity and experimental replicability, as discussed in Section 5 below. Importantly, however, in-context annotation of individual collections of email data offers the only way to assess the real-world utility and accuracy of our request and commitment classifiers. Utility can be measured through a combination of implicit and explicit user feedback. More generally, the plug-in offers an example of the type of in-context tools that we believe are required for complex document processing tasks. Specifically, it provides us with a platform for both intrinsic and extrinsic evaluation of our automated request and commitment classifiers, trained on data from human annotation experiments.

3 Requests and Commitments in Email

We are interested in the presence of both requests and commitments at two levels of granularity: the message level, and the utterance or text-span level. At the message level, we are interested in marking each email message with binary classifications that indicate whether or not that message contains any requests or commitments, without specifying the exact text spans that represent the requests and/or commitments. At the utterance or text-span level, we are interested in marking exact spans of text that represent request and commitment speech acts. The latter is a considerably more difficult task, since it is often unclear where the textual boundaries of requests and commitments lie.

As we have discussed in previous work [7], requests and commitments in email are complex. Far from resembling the simple canonical examples that have often been used as definitions for annotation schemes, real world requests and commitments are much more subtle. Interpretations depend on aspects such as conditionality, indirectness, and situational context.

Below, we present a high-level overview of the request and commitment acts in which we are interested. More detailed discussion of the definitions we use can be found in [7, 6].

3.1 Requests in Email

We consider any utterance or series of utterances from an email sender that places an obligation on an email recipient to be a *Request*. Specifically, a request may obligate a 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 [5] can all function as requests. Some linguists have distinguished speech acts that require a physical response from those that require information or a speech act in response (see, for example, [13]). We follow Searle's original approach [12] and do not distinguish between physical and speech act 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.

3.2 Commitments in Email

We consider a *Commitment* to be an offer or promise made by an email sender of 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 those utterances that place an obligation on another person, group of people, or organisation. An example of a third-party commitment is an utterance like *Finn will send you a complete draft by Friday*, where Finn is not included as a participant in the email discussion. Such third-party commitments are common in the enterprise email collections 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.

4 The Plug-in

The plug-in we have developed augments Microsoft Outlook 2007 with functionality that facilitates in-context human annotation of email messages and provides a platform for deploying and evaluating automated document processing tools, in our case, request and commitment classifiers.

To address our research needs, the plug-in includes a range of features for discovering, displaying and interacting with requests and commitments. The methods for display and interaction are the same for both manually identified and automatically identified requests and commitments, which allows us to use the plug-in in several modes: either as a purely manual annotation tool, or as a tool that integrates automatic classifiers (whilst still allowing additional manual annotations and classification corrections). The overall appearance of the interface is as shown in Figure 1. Below we describe the main features of the plug-in.

4.1 Message View Panel

The *Message View Panel* (Figure 2) is displayed at the bottom of the standard Outlook message window when

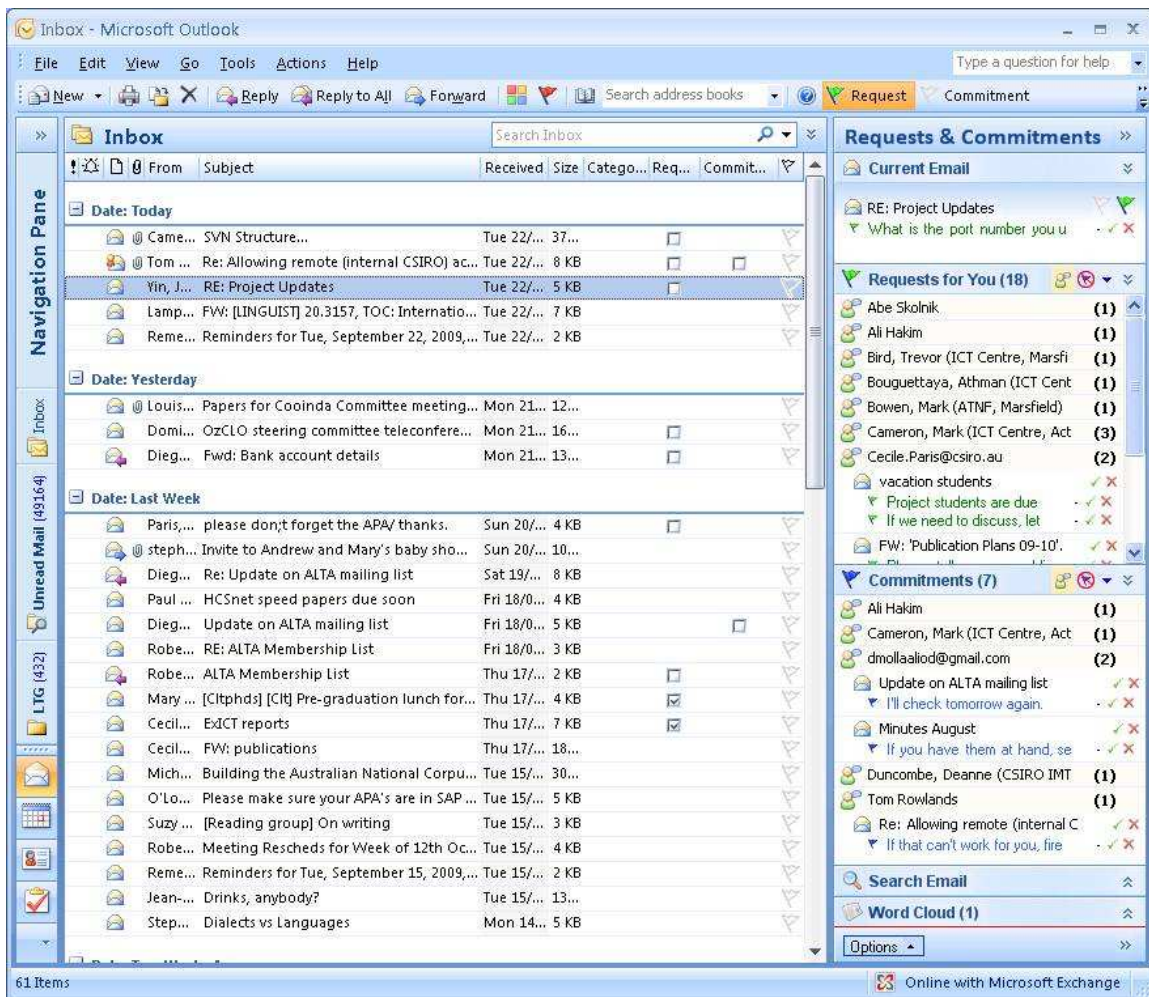


Figure 1: Email Plug-in shown within Microsoft Outlook

an email is opened. The Message View Panel shows details of all the requests and commitments within the open message. Accompanying the requests and commitments are a set of controls, similar to those found in the Task Sidebar, for the user to interact with and manage this set of tasks. For any request or commitment, the user may set a due date, mark it as complete, or delete it (if it has been identified in error). Additionally, as described further in Section 5.2, the user is also able to inspect and modify the current logging policy associated with the message.

To support the Message View Panel, three buttons are added to the standard Outlook ribbon: a Request button, a Commitment button and an Add Task button, as shown in Figure 3. When no message text is selected, clicking on either the Request or Commitment button annotates a message-level request or commitment for the open message, without specifying any particular associated text-span. To annotate a specific text-span as a request or commitment, the user highlights some text (e.g., a phrase or a sentence) and clicks on the request or commitment button. This annotates a fine-grained request or commitment, along with a message-level re-

quest or commitment if one has not already been annotated. The Request and Commitment buttons also function as status indicators, showing whether a request or commitment has been annotated or classified. The Add Task button allows the selected text to be used to create a structured Outlook task that is managed separately from the user's email messages.



Figure 3: Request and Commitment Annotation Buttons in the Outlook Ribbon

4.2 Folder View Indicators

The *Folder View Indicators* (Figure 4) integrate request and commitment information into the standard Outlook folder views. Message-level requests and commitments



Figure 2: Message View Panel

may also be annotated within the folder view using the toolbar buttons shown at the top of Figure 1.




Figure 4: Folder View showing Request and Commitment Fields

When a user views their inbox or folder content, two additional fields are displayed. These indicate the presence of requests and commitments in each message. Because of the nature of the interface, only message-level presence is indicated; no detail of the specific request and commitment text-spans is shown. The field indicators are three-state, indicating that a message contains:

1. No Request/Commitment: a blank Request or Commitment field;
2. One or More Unfulfilled Requests/Commitments: an unchecked tickbox in the appropriate field; or
3. One or More Completed Requests/Commitments: a checked tickbox in the appropriate field.

These fields allows messages to be easily sorted based on their request or commitment content, and allows the user to very quickly obtain an overview of the outstanding tasks in a folder.

4.3 Task Sidebar

The Task Sidebar sits next to the folder view, and contains five information panes, as shown in the right hand panel of Figure 1. The first three panes in the task bar are used to display and interact with the requests and commitments that have been either manually annotated

or automatically classified in email from the user's inbox and other selected folders. The other two panes are the Search Pane and Word Cloud Pane.

4.3.1 Current Email Pane

The *Current Email* pane, shown in Figure 5, contains coloured flags that indicate whether the currently selected message contains any requests and commitments,¹ along with a list of the specific text spans that have been annotated as containing requests and/or commitments. The purpose of this pane is to act as a simple task-focused summary of the message which gives the user a quick, browsable sense of the actionable content in the message. Currently, the pane contains a simple list of extracted requests and commitments; in future work, we aim to explore automatic rewriting of the task content to ensure it makes sense as stand-alone text. This task rewriting functionality will also be useful for migrating tasks to more structured task lists.

Within the Current Email pane, the user can interact with the requests and commitments to mark a request or commitment as complete, delete a request or commitment that has been mistakenly identified, set a due date for a request or commitment, or open the message containing the displayed request or commitment.

4.3.2 Requests Pane

The *Requests* pane (Figure 6) shows the outstanding tasks awaiting action on the part of the mailbox user. Primarily, these include incoming requests from

¹Here, green for requests and blue for commitments.



Figure 5: Current Email pane showing a task-focused summary for an Email Message




Figure 6: Requests pane showing default reverse-chronological ordering of tasks

other email senders. Also included are outgoing commitments made by the mailbox user, since they too require the user's action.

Like most email inboxes, the requests are displayed in reverse chronological order by default. This ordering can be inverted, or the requests can be grouped by the person to whom the task relates (either the sender who made the request, or the recipient to whom the mailbox owner made the commitment). In any of these views, requests can be shown only at the message-level (as in Figure 6), or can include detail of the finer-grained task text spans. Figure 7 shows an example of this finer-granularity display for commitments.

4.3.3 Commitments Pane

The *Commitments* pane shows the outstanding tasks waiting for someone else's action. As with the Requests pane, these include both incoming commitments from other email senders, along with outgoing requests sent by the mailbox user. The aim here is to help a user keep track of actions that others have promised to complete.

Figure 7 shows a chronologically ordered list of outstanding commitments from across the user's incoming and outgoing mail. In contrast to the requests in Figure 6, the commitments are displayed with the detail of finer-grained text-span commitments. Figure 8 shows the same set of outstanding commitments grouped by the person who is to take action.

4.3.4 Search and Word Cloud Panes

The *Search* pane allows the user to search the content of all requests and commitments. This pane is also used for interacting with the content of the *Word Cloud* pane




Figure 7: Commitments pane showing default reverse-chronological ordering of tasks



Figure 8: Commitments pane showing grouping of tasks by the Responsible Person

that shows a visualisation of words in the selected message(s). The word cloud is intended to function as an alternate summary that provides a general overview of the content of one or more selected messages. Clicking on any word in the word cloud either opens the message in which the word occurs (if there is only a single such message), or displays the set of messages containing the word in the Search pane, where individual messages can be opened by clicking on them.

5 Logging and Privacy

In order to understand how users interact with tasks in their email, and to improve the performance of our automated request and commitment classifiers, we have instrumented our plug-in to record user feedback and interaction. We have implemented this logging carefully and deliberately to preserve the user's control and privacy during in-context annotation. Below,

we discuss the details of the instrumentation and privacy-preserving plug-in features.

5.1 Logging and Instrumentation

Apart from providing users with an interface for annotating and interacting with requests and commitments in their email messages, we are also interested in observing and recording how users interact the tasks within their messages.

Given the context-sensitivity of identifying requests and commitments, our intention in automating classification is to begin with a generic classifier that can be retrained and adapted through ongoing feedback from each individual user to tailor the classification to individual data and preferences. This process requires feedback from users about predicted classifications.

One approach to obtaining this feedback would be to employ detailed ethnographic interviews or physical observations and recordings of users in their natural work contexts. Unfortunately, this approach is often impractical on any significant scale due to resource constraints. Instead, we have focused on recording useful feedback we can obtain automatically, based on the explicit and implicit user interactions with our plug-in. Specifically, we log:

- all request and commitment annotations made by the user, either at the message-level or at the text-span level;
- all request and commitment deletion actions, which we interpret as corrections if they delete automatic annotations;
- all request and commitment completion actions;
- email message opening actions; and
- email reply actions.

In order to make use of this feedback data as training data for our automatic classifiers, we also need information about the message content and headers. Where the user permits us to, we also log this information. This allows us to interpret deletions of automatic annotations as false positive classifications and additional manual annotations as false negatives. Task completions act as confirmation of true positive classifications. In all these cases, we need access to the message content to create additional training instances for our classifiers.

In future work we intend to employ active learning methods to proactively prompt users to consider specific messages that would provide maximum information to the automated classifiers as training data.

5.2 Privacy and User Control

Privacy is an extremely important consideration when dealing with email data. Clearly, our desire to gather useful data from plug-in users must be balanced against the need to adequately protect people's privacy and to leave people in complete control of their data.

Currently, we continually log interaction data into a local database on the user's computer. This does not involve the transfer of any data off the user's computer. If the user opts-in to allow us to collect their usage data, we periodically transfer a batch of interaction data to a remote logging server for aggregated analysis.

In order to provide users with control over their data, we offer a range of opt-in and opt-out mechanisms regarding data logging and instrumentation. When the user first installs the plug-in, they are offered the choice to opt-in to allow us to log their interactions and message content. We include an explanation of why this interaction data is helpful for our ongoing research and for improving the performance of the automatic classifiers. Additionally, we inform people that all data we collect is encrypted during transit and kept strictly private and confidential. While most data will be used only for algorithmic improvements (which don't require manual inspection), we note that some data may be viewed by core researchers on the project. We emphasise that no other people will have access to the data.

If the user chooses to opt-in to contribute their interaction and message data, we provide finer-grained opportunities to exclude specific messages. The first method is via an exclusion list that allows a user to exempt messages from particular people or sent to a particular email address. The content of all matching messages is not logged. The user is able to add or remove addresses from the exclusion list at any time. Additional privacy controls are provided in the Message View Panel; these allow the user to see whether the current message will be logged, and to choose to exclude the message from logging.

Finally, prior to transferring any email data, the user is shown a dialogue containing the set of messages whose content is about to be transferred, and is able to inspect and exclude any or all of these messages before any data is transferred remotely. Any data that the user agrees to contribute is encrypted during transit to the remote server, to ensure it is not intercepted in transit. All logged data is stored securely, accessible only by the small group of researchers working directly on the Outlook plug-in project.

Although not yet implemented, we also plan to provide users with a secure web application where at any time they can login and view or delete any or all of their data that they've previously agreed to have logged.

6 Related Work

There has been a lot of work on document annotation tools. Close to our research interests, several tools have been developed for annotating speech acts and other similar phenomena in email and textual conversation. In most cases, annotation has made use of tools that are divorced from the real-world context of use (e.g., [3, 8]). While tools such as the custom web application we developed for our annotators [8] are useful for gathering

out-of-context human annotations, they are difficult to adapt for capturing in-context annotation, since they are not integrated into the user's natural environment. The Outlook plug-in we have developed is intended to help address this gap.

We are certainly not alone in attempting to integrate our research into existing email software: other researchers have looked at building research tools into such environments. However, these efforts have usually focused only on the application of tools, rather than on integrating tools for human annotation of data. Nepomuk [11], for example, includes an Outlook plug-in that integrates email prototype workflow management tools. While this provides a useful platform for deploying the research to users, the plug-in does not support data annotation, and thus does not provide a tool for capturing in-context user judgements.

We are, however, not alone in considering in-context annotation. Recently, Balasubramanian et. al [2] developed a plug-in for Mozilla Thunderbird that is used for in-context evaluation of email leak detection algorithms. To do so, the plug-in logs "all aspects of usage of the extension" [2] and, with the user's express permission, this data is gathered remotely. Even in this case, however, the plug-in is limited to evaluating existing tools, although it could plausibly be extended to gather in-context human judgements or annotations.

7 Conclusion

Complex document processing tasks such as identifying requests and commitments in email text often require interpretation of language use in context. Such tasks require human annotators to have access to a significant amount of contextual information to make accurate judgements, leading to a need for in-context annotation.

To facilitate such in-context annotation for email data, we have developed a plug-in for Microsoft Outlook that allows annotators to use their own email messages as the data for annotation. To address the significant privacy challenges that are created by using such private data, we have integrated a series of software features and research processes that provide the user with methods to control what, if any, data is recorded, while still attempting to elicit necessary data to improve and advance our research goals and to refine our automatic request and commitment classifiers.

While our focus is on requests and commitments in email, we have attempted to design our plug-in such that it could be adapted to accommodate other forms of in-context email annotation. More generally, our plug-in serves as an example of the class of ecologically-valid tools that are needed to approximate the real-world environments in which the documents under examination are created and used, to ensure that annotations and judgements are valid.

References

- [1] Dawn Archer, Jonathan Culpeper and Matthew Davies. *Corpus Linguistics: An International Handbook*, Chapter Pragmatic Annotation, pages 613–642. Mouton de Gruyter, 2008.
- [2] Ramnath Balasubramanian, Vitor R. Carvalho and William W. Cohen. Cut once: Recipient recommendation and leak detection in action. In *Proceedings of EMAIL-08: the AAAI Workshop on Enhanced Messaging*, Chicago, USA, July 2008.
- [3] Edward Ivanovic. Dialogue act tagging for instant messaging chat sessions. In *ACL Student Workshop 2005*, Ann Arbor, Michigan, June 2005. ACL.
- [4] Bryan Klimt and Yiming Yang. Introducing the Enron corpus. In *Proceedings of the Conference on Email and Anti-Spam (CEAS)*, 2004.
- [5] William Labov and David Fanshel. *Therapeutic Discourse: Psychotherapy as Conversation*. Academic Press, New York, 1977.
- [6] Andrew Lampert, Robert Dale and Cécile Paris. The nature of requests and commitments in email messages. In *Proceedings of EMAIL-08: the AAAI Workshop on Enhanced Messaging*, pages 42–47, Chicago, USA, July 13 2008. AAAI Press. Technical Report WS-08-04.
- [7] Andrew Lampert, Robert Dale and Cécile Paris. Requests and commitments in email are more complex than you think: Eight reasons to be cautious. In *Proceedings of Australasian Language Technology Workshop (ALTA2008)*, pages 55–63, Hobart, Australia, 8-10 December 2008.
- [8] Andrew Lampert, Cécile Paris and Robert Dale. 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 2007.
- [9] Charles Morris. *Signs, Language, and Behavior*. Prentice Hall, 1946.
- [10] Denise E. Murray. *Conversation for Action: The Computer Terminal As Medium of Communication*. John Benjamins Publishing Co, 1991.
- [11] Simon Scerri, Siegfried Handschuh and Brian Davis. The path towards semantic email: Summary and outlook. In *Proceedings of EMAIL-08: the AAAI Workshop on Enhanced Messaging*, Chicago, USA, 13 July 2008.
- [12] John R. Searle. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, 1969.
- [13] John Sinclair and Richard Malcolm Coulthard. *Towards and Analysis of Discourse - The English used by Teachers and Pupils*. Oxford University Press, 1975.
- [14] Jan Ulrich, Gabriel Murray and Giuseppe Carenini. A publicly available annotated corpus for supervised email summarization. In *Proceedings of EMAIL-08: the AAAI Workshop on Enhanced Messaging*, 2008.
- [15] Terry Winograd and Fernando Flores. *Understanding Computers and Cognition*. Ablex Publishing Corporation, Norwood, New Jersey, USA, 1st edition, 1986. ISBN: 0-89391-050-3.



Contact Us

Phone: 1300 363 400

+61 3 9545 2176

Email: enquiries@csiro.au

Web: www.csiro.au

Your CSIRO

Australia is founding its future on science and innovation. Its national science agency, CSIRO, is a powerhouse of ideas, technologies and skills for building prosperity, growth, health and sustainability. It serves governments, industries, business and communities across the nation.