

Integrating Natural Language Understanding and Plan Reasoning in the TRAINS-93 Conversation System *

D. R. Traum[†], J. F. Allen[†], G. Ferguson[†], P. A. Heeman[†], C. H. Hwang[†],
T. Kato[‡], N. Martin[†], M. Poesio[§] and L. K. Schubert[†]

Abstract

This paper describes the TRAINS-93 Conversation System, an implemented system that acts as an intelligent planning assistant and converses with the user in natural language. The architecture of the system is described and particular attention is paid to the interactions between the language understanding and plan reasoning components. We examine how these two tasks constrain and inform each other in an integrated nl-based system.

1 The TRAINS Project

The TRAINS project is a long-term research project to develop an intelligent planning assistant that is conversationally proficient in natural language [Allen and Schubert, 1991]. The TRAINS system helps a user construct and monitor plans about a railroad freight system. The user is responsible for assigning cargo to trains and scheduling shipments, scheduling various simple manufacturing tasks, and for revising the original plans when unexpected situations arise during plan execution. Figure 1 shows a typical initial scenario. The system aids the user in all aspects of this task by interacting in natural language. In particular, the system typically will perform the following tasks:

- Evaluating courses of action, calculating expected completion times, detecting conflicts, and so on;
- Filling in details of the proposed plan that do not require the user's attention;
- Suggesting ways to solve particular subproblems as they arise;
- Presenting and describing the current state of the world and how the proposed plan may affect it;

*This material is based upon work supported by ONR/DARPA under grant number N00014-92-J-1512, the US Air Force under Rome Laboratory research contract number F30602-91-C-0010, and ONR under research grant number N00014-90-J-1811

[†]{traum, james, ferguson, heeman, hwang, martin, schubert}@cs.rochester.edu University of Rochester; Rochester, NY, 14627-0226.

[‡]kato@nttnly.ntt.jp NTT Network Information Systems Labs; 1-2356 Take, yokosuka-shi, Kanagawa 238-03 JAPAN.

[§]poesio@cogsci.ed.ac.uk University of Edinburgh, Centre for Cognitive Science, Edinburgh, EH8 9LW, Scotland, UK.

- Dispatching the plan to the different agents in the world, including the train engineers and factory users;
- Interpreting reports back from the engineers and factory users in order to monitor the progress of the plan and to anticipate problems before they arise; and
- Coordinating the correction and modification of plans with the user.

While we aim to produce a functioning system, the system itself is not really the goal of the effort. Rather, the domain is a tool for forcing our research to address the problems that arise in building a complete dialogue system. For instance, the dialogue module must be able to handle a wide range of everyday discourse phenomena rather than handling a few selected problems of theoretical interest. While this approach focuses our research directions, the solutions that we seek are general solutions to the phenomena rather than specific solutions that happen to work in the TRAINS domain. The TRAINS project currently has several main foci of research:

- Parsing and semantically interpreting utterances as they arise in spoken language, including sentence fragments, repairs and corrections;
- Accounting for the discourse behavior present in natural dialogue;
- Representing the reasoning and control of the discourse agent, including reasoning about plans in the TRAINS domain as well as the plans that drive the system's behavior in the dialogue itself; and
- Providing the knowledge representation and reasoning tools that are needed for planning and scheduling in realistic size domains, including reasoning about time, events, actions and plans.

Four prototypes of the TRAINS system have been developed, one per year starting in 1990. We feel that most of the issues to be discussed at this symposium were raised at one point or the other during our work of developing and refining the TRAINS system. In the final analysis, we can say the system "understands" NL input from the user if it executes the plan as the user expects.¹ More locally, the system can

¹Actually, this condition is too strong, as a user might fail to convey expectations properly, so a more realistic objective might be to execute a plan according to expectations of an *observer* of the conversation.