THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Planning to Fail, Not Failing to Plan: Risk-taking and Recovery in Task-oriented Dialogue

Citation for published version:

Carletta, J 1992, Planning to Fail, Not Failing to Plan: Risk-taking and Recovery in Task-oriented Dialogue. in *Proceedings of the 14th Conference on Computational Linguistics - Volume 3*. COLING '92, Association for Computational Linguistics, Stroudsburg, PA, USA, pp. 896-900, 14th International Conference on Computational Linguistics (COLING) 1992, Nantes, France, 23/08/92.
<https://doi.org/10.3115/992383.992396>

Digital Object Identifier (DOI):

[10.3115/992383.992396](https://doi.org/10.3115/992383.992396)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Proceedings of the 14th Conference on Computational Linguistics - Volume 3

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

PLANNING TO FAIL, NOT FAILING TO PLAN: RISK-TAKING AND RECOVERY IN TASK-ORIENTED DIALOGUE

JEAN CARLETTA*
University of Edinburgh
Department of Artificial Intelligence
jcc@aisb.ed.ac.uk

Abstract

We hypothesise that agents who engage in task-oriented dialogue usually try to complete the task with the least effort which will produce a satisfactory solution. Our analysis of a corpus of map navigation task dialogues shows that there are a number of different aspects of dialogue for which agents can choose either to expend extra effort when they produce their initial utterances, or to take the risk that they will have to recover from a failure in the dialogue. Some of these decisions and the strategies which agents use to recover from failures due to high risk choices are simulated in the JAM system. The human agents of the corpus purposely risk failure because this is generally the most efficient behaviour. Incorporating the same behaviour in the JAM system produces dialogue with more "natural" structure than that of traditional dialogue systems.

Introduction

There are a great number of different dialogue styles which people use even in very restricted task-oriented domains. Agents can choose different levels of specificity for referring expressions, ways of organising descriptions, amounts of feedback, complexities of explanation, and so on. This work first identifies a number of aspects of task-oriented dialogue along which agents can make choices and identifies these choices in terms of how much effort the agent must expend in order to generate utterances in line with them. In general, expending more effort in building an explanation means that the explainee is more likely to understand it as is; thus we can classify some choices as being "higher risk" than those which take more effort to generate but which are more likely to succeed on the first attempt. Then it identifies a number of recovery strategies which agents use when risky behaviour has led to a failure in the dialogue. The choices which agents make show a trade-off of when effort is expended in the dialogue; agents can either expend effort early in order to head off later difficulty, or take the risk of having to expend more effort in an attempt at recovery. For instance, consider the domain, first described in [5], in which two participants who are separated by a partition have slightly different versions of a simple map with approximately fifteen gross features on it. The maps may have different features or have some of the features in different locations. In addition, one agent has a route drawn on the map. The task is for the second agent to

duplicate the route. The HCRC Dialogue Database [3] contains 128 such dialogues; in this work we examined eight plus a set of dialogues from the pilot study used in Shadbolt's work [17]. Agents who wish to avoid plan failure may structure their explanations carefully and elicit feedback often, behaving similarly to agent A in Shadbolt's example 6.16:

A: have you got wee palm trees aye?
B: uhu
A: right go just + a wee bit along to them have you got a swamp?
B: er
A: right well just go + have you got a waterfall?

On the other hand, agents who are willing to rely on interruptions from the partner and recovery from failure might behave more like agent A in Shadbolt's example 6.11:

A: and then + go up about and over the bridge
B: I've not got a bridge I've got a lion's den and a wood
A: have you got a river?

Either of these approaches is likely to bring the agents to successful completion of the task. However, it is also possible to include too little information in the dialogue, as in the following case, Shadbolt's example 6.21:

A: right + you're going to have to cross the river
B: how?
A: dinnae ken + any way you want...

It is equally possible to give too much information, as in Shadbolt's example 6.27:

B: ah right + erm + oh yes + er + I have a crashed plane marked here + can I + check this + my crashed plane is ABOVE + it's in the BASE of the quadrant + top right hand imaginary quadrant of the + erm + picture + yes er + that SOUNDS too high for me +

A: er

In this case, B provides so much information that A is unable to process it, and they eventually abandon this section of the dialogue. This work looks at the differences between the approaches which human agents use to complete the map task and simulates them using the JAM system. Understanding and comparing the different human approaches to task-oriented dialogue can help us to create more robust computer dialogue agents.

Communicative Posture

Our work extends Shadbolt's analysis of the map task data [17]. He identifies a number of "communicative posture parameters" or aspects of the dialogue for

*This research was supported by a postgraduate studentship from the Marshall Aid Commemoration Commission and supervised by Chris Mellish. The author's current address is HCRC, 2 Buccleuch Place, University of Edinburgh, Edinburgh EH8 9LW, Scotland.

which an agent may make the choice of how to proceed, and classifies the possible settings in terms of risk: for the most part, high risk settings leave the partner to infer information and risk the possibility of plan failure, while low risk settings are more likely to work as planned. He then argues that human agents decide upon their communicative postures according to the Principle of Parsimony, which is "a behavioural principle which instructs processors to do no more processing than is necessary to achieve a goal." (pg. 342) Agents choose the settings for each individual parameter which they believe will prove most efficient. Shadbolt identifies seven different communicative posture parameters. Our own analysis extends his by

- clearly separating out aspects of being a hearer from those of being a listener and
- by making the behaviour of the parameters more independent of each other and subsequently dividing them into sets depending on which part of an agent's planning they affect. We divide utterance planning into different stages similar to Appelt's [4] for this part of the analysis.

The following revised set provides a more solid foundation on which to build the implementation found in the JAM system:

Task Planning Parameters

These parameters affect which task plan an agent chooses. In the map domain, task plans determine the choice of descriptions for sections of the route and for the location of objects.

Ontology: The choice of concepts to use when building an explanation. High risk agents construct simple and short descriptions, providing as little information as they think the partner will allow, while low risk agents provide precise, detailed explanations even if that involves using fairly complex background concepts and introducing new concepts into the dialogue.

Ontological Resolution: The choice of concepts to ask about when hearing an explanation. High risk agents accept the level of detail which is offered to them, while low risk ones ask how concepts are related if they think that the relationship may be an important piece of background for the explanation.

Partner Modelling: Whether or not to heed a model of the partner while building an explanation. High risk agents do not, while low risk agents do, tailoring the explanation for the partner. It takes more effort in the first instance to build an explanation which is tailored to the partner, but the explanation is more likely to succeed without revisions.

Ontology and partner modelling are implemented in the JAM system by means of an evaluation scheme for possible task plans which rates descriptions differently depending on whether these parameters are set to low or high risk. Low risk ontology prefers descriptions which refer to many map objects over simpler ones; if there are several descriptions of equal complexity, low risk partner modelling prefers descriptions which do not refer to map objects that may be unknown to the partner. Ontological resolution is not implemented in the JAM system because

JAM agents are not capable of the spatial reasoning required to determine what other map objects are relevant to a given description.

Discourse Planning Parameters

These parameters affect the structure of the discourse, given the information from the task plan which must be conveyed.

Difference: Whether or not agents assume that their models of the domain are the same unless proven otherwise. High risk agents make this assumption, while low risk agents do not, making them precede new concepts in the dialogue with subdialogues which establish certain knowledge of the partner's knowledge such as direct questions about the status of the concepts. A low risk difference setting makes the dialogue longer and hence requires more effort, but also provides a greater strength of evidence about the partner's beliefs [7] than does relying on the partner's feedback to the explanation itself. This parameter is implemented in the JAM system by means of optional prerequisites on discourse plans which introduce new concepts; low risk agents expand the prerequisites, while high risk agents do not.

Coherence: Whether or not the agents organise their discourse coherently. High risk agents produce utterances in whatever order they think of them, whereas low risk agents try to order them in some way which will make the discourse easier for the partner. This parameter is not implemented in the JAM system because map task participants do not often organise the discourse except as if they were physically following the route. In less well structured domains, it could be implemented using, for instance, RST [11] or focus trees [12].

Utterance Realisation Parameters

These parameters affect the way in which each utterance in the given discourse structure is realised or understood.

Context Articulation: Whether or not the agents signal awkward context shifts. Here context is loosely defined as the goal which is supported by the current part of the dialogue; in the map task, contexts can either be goals of sharing knowledge about a section of the route or the location of an object. High risk agents do not signal awkward context shifts, while low risk agents use meta-comments, changes in diction, or some other means to mark the new context. A limited version of the low risk setting is implemented in JAM which introduces a meta-comment into the dialogue whenever a context shift occurs.

Context Resolution: Whether or not agents ask for clarification of awkward context shifts. Low risk agents ask the partner what the current context is or make their assumptions clear when they are unsure, whereas high risk agents simply choose the most likely context. This parameter is not implemented in the JAM system because JAM agents use a language which does not allow for ambiguity of context.

Focus Articulation: Whether or not agents signal awkward focus shifts. Here, focus is defined specifically for the map task in terms of distance on the

map and semantic relationships among map features. Low risk agents use meta-comments or modifiers on referring expressions to signal awkward focus changes, and high risk agents do not. Focus articulation is not implemented in the JAM system because JAM agents are not capable of the spatial or semantic reasoning required to calculate focus; given these abilities, low risk agents could use some theory of how focus usually moves (such as that of Grosz and Sidner [9]) to determine whether or not signaling a particular shift is necessary.

Focus Resolution: Whether or not agents ask for clarification of awkward focus shifts. Low risk agents ask the partner what the current focus is or mark their assumptions in some other way, whereas high risk agents simply choose the most likely focus. Low risk focus resolution could be implemented by having low risk agents ask for clarification whenever a focus shift does not conform to some theory of focus, with high risk agents "guessing" the current focus.

Specification: Whether or not agents construct referring expressions carefully. Low risk agents generate referring expressions which are roughly minimally unique, whereas high risk agents generate whatever expression comes to mind, even if that expression is under- or over-specific. This parameter could be implemented in the JAM system using, for instance, work by Dale [8] and Reiter [16].

Description Resolution: Whether or not agents decode referring expressions carefully. Low risk agents ask for clarification of ambiguous referring expressions, while high risk agents simply choose the mostly likely referent. This parameter could have an implementation similar to that of the specification parameter, but from the point of view of the addressee.

Meta-Planning Parameter

This parameter affects an agent's choice of how to continue from the current situation in a dialogue.

Plan Commitment: Whether or not agents decide to replan easily. Low risk agents tend to stick to the current plan unless there is sufficient proof that the new plan is better, whereas high risk agents often replan when they encounter failures even without carefully checking the viability of the new plan. Frequent changes in plans are likely to confuse the partner and lead to difficulty in the dialogue, especially if the agent's context articulation setting is also high risk. This parameter is implemented in the JAM system by means of a "replanning threshold" which is added to the estimated cost of a replan and which makes replanning seem less efficient to low risk agents than to high risk ones.

Of course, the choice is not between extremes, but among points on a spectrum which generally reflects the amount of effort to be expended. Shadbolt adapts the Principle of Parsimony to state that agents make the choices which they believe will lead to the lowest effort solution for the entire task. In each case, high risk agents may lose the efficiency advantage which they gained by using less effort initially, if their plans fail and they have to expend more effort to recover from the failures. Recovery strategies are more often needed by high risk agents than by low risk ones.

Recovery Strategies

Our analysis has uncovered the following recovery strategies. Some strategies are only first steps towards finding a solution for the failure, and one, goal adoption, is also useful in other circumstances. We use the same basic definitions for repair and replanning as in Moore's work [13].

Goal Adoption: The agent may infer the partner's goals from some part of the dialogue he or she has initiated and adopt them as his or her own.

Ceding the Turn: The agent may simply not take any action and hope that his or her inaction will force the partner into initiating the recovery.

Elaboration: If an explanation has not been given in enough detail, the explainer may fill in the gaps.

Omission: If an explanation has been given in too much detail, the agents may agree to discard some of the information. This is especially useful in the map task if some description of the route or of the location of an object on the map turns out to hold for one version of the map but not the other.

Repetition: Under any circumstances, an agent may simply repeat whatever action has already failed in the hopes that it will work the subsequent time.

Ignoring the Problem: An agent may ignore a problem and hope that it will disappear.

Repair: If a plan has failed, then checking each of the prerequisites of the plan in turn to see if they are satisfied may lead to a diagnosis. In the map task, plan prerequisites have to do with knowledge about objects on the map. A plan will fail if an agent presupposes that the partner has knowledge which he or she does not have. Since the knowledge transferred in the map domain is so simple, it is sufficient in a repair to re-execute any failed prerequisites, even if the plan has already been completely executed.

Replanning: If a plan has failed, then an agent may attempt an entirely different plan with the same effect. In the map task, this involves using a different description for the information under consideration or trying a different approach altogether.

There are many past systems which have incorporated some form of recovery from plan failure (e.g., [2], [19], [14]). However, very little work has been done on incorporating more than one recovery strategy into the same system. Moore's [13] work allows the use of repair, reinstatement, and replanning, but uses a strict ordering on these strategies to determine which one to try next. Moore's system first attempts any possible repairs, then any reinstatements, and then, only as a last resort, replanning. Neither Moore's ordering nor any other can account for the variety of behaviours which is present in the human map task corpus. In addition, Moore's system only considers replanning when there has been a plan failure, whereas human agents sometimes switch plans when they flesh out enough of the details and discover that the plans which they have adopted are less efficient than they had expected. The solution to these shortcomings is to invoke the Principle of Parsimony and to allow agents at every choice point to decide what to do next based on an estimates

Figure 1: The Structure of a JAM Agent's Planner

of how much effort the remainder of the task will require given each of the possible next actions. This approach is adopted in the JAM system.

The JAM System

The JAM system allows agents to converse about the map task by alternating control between them, following Power [15]. Agents converse in an artificial language which is based on Houghton's [10] interaction frames; these frames specify the forms of two and three move dialogue games for informing, asking wh- and yes-no- questions, opening a subdialogue with a particular topic, and closing a recovery, and also gives plausible belief updates associated with each. An English-like gloss is provided for the benefit of the human observer only by means of very simple templates. Unlike the humans, JAM agents have their communicative postures set before the beginning of a dialogue and can not vary them during its course. Each agent uses five communicative posture parameters (ontology, partner modelling, context articulation, difference, and plan commitment) and three recovery strategies (goal adoption, repair, and replanning). Space will not permit a description of how the parameters are implemented; for more details see [6]. The recovery strategies are implemented within a layered message-passing architecture shown in figure 1 and adapted from MOLGEN [18].

At each layer of the system, operators use the next layer down in order to decide whether or not they are applicable and how to apply themselves. The bottom layer of the system contains plan operators which axiomatise the domain knowledge and the meanings of the dialogue games. The strategy layer contains operators expressing all of the different actions which an agent can take next in the current situation: agents can decide to communicate, infer and adopt one of the partner's goals, plan, replan, or repair. The mode layer contains operators which control which strategy is chosen. The mode operator of most concern to us is the *comprehensive* mode, which always communicates whenever it has something to say, goal adopts whenever it can recognise one of the partner's intentions (using

a simplified version of Allen's plan recogniser [1]), and chooses whether to plan, replan, or repair based on an estimate of the effort needed to complete the entire task if each of the options is taken *next* (for details, see [6]). There are also other modes which eliminate some of the recovery strategies in order to make experimentation with the strategies easier, and one mode which reconstructs as far as possible the choices which Moore's system [13] makes. Finally, the interpreter chooses one of the mode level operators for the duration of the dialogue. If a theory of how agents include and exclude consideration of different recovery strategies were available, it could be implemented in the interpreter and more layers could be added to the system as needed.

Examples

The JAM system generates the following dialogue extract between two agents who both have totally low risk communicative posture settings:

```

mary: i want to talk about the first section.
john: ok.
mary: do you have the palm beach?
john: yes.
mary: do you have the swamp?
john: no.
mary: i want to talk about the swamp.
john: ok.
mary: do you have the waterfall?
john: yes.
mary: the swamp is between the waterfall and the
    palm beach.
john: ok.
mary: the first section of the route goes between
    the palm beach and the swamp.
john: ok.
  
```

In this extract, a low risk context articulation setting leads mary to open the dialogue about each distinct topic (the first section of the route and the swamp). Low risk ontology causes her to choose her to choose fairly complex descriptions of the first section of the route and the location of the swamp (at least for JAM agents), and low risk difference makes her ask john whether or not he has all of the prerequisite knowledge before she gives any descriptions. Finally, low risk plan commitment causes her not to to abandon her description of the first section of the route when she discovers that john's map does not have the swamp. If we run the same dialogue, but replace mary with an agent who has high risk plan commitment and context articulation settings, then the structure of the dialogue changes completely:

```

janet: do you have the palm beach?
john: yes.
janet: do you have the swamp?
john: no.
janet: the first section of the route goes to the left
    of the palm beach.
john: ok.
  
```

In this case, although janet begins using the same plan as mary had, as she explores this plan she discovers that it will not work and decides that it will be less costly simply to abandon it for a description which does not mention the swamp. In another example, two

agents with high risk settings for context articulation, difference, and plan commitment also successfully complete this section of the route, but only by having a plan failure occur and invoking the repair recovery strategy:

carol: the first section of the route goes between
the palm beach and the swamp.
tom: i do not understand.
tom: where is the swamp?
carol: the swamp is between the waterfall and the
palm beach.
tom: ok.
tom: i am done talking about the first section.
carol: ok.

In this dialogue, tom infers that carol has the goal for tom to know how to get around the first section of the route from her initial statement, adopts it, and then diagnoses the failure of prerequisite knowledge about the swamp in order to repair her plan. He chooses this course of action because he believes it will take less effort than the other possibilities (in this case, simply telling carol he does not understand her statement without diagnosing the problem and allowing her to provide a different description of the first section). This failure came about as a result of carol's high risk difference setting, since a low risk difference setting would have made carol ask tom ahead about his knowledge, as mary did for john. Using agents with different communicative postures and interpreters allows the JAM system to simulate many different behaviours which can be recognised in the human corpus.

Conclusions

We demonstrate a number of aspects of dialogue for which agents must choose between expending effort when they create their initial utterances and taking the risk of plan failure, and go on to describe a number of strategies which high risk agents use to recover from failure. A surprising outcome of the human examples is that it is often most parsimonious to risk failure. Agents quickly reach the limits of their resource bounds when they try to avoid possible confusions in the dialogue, and dialogue is such a flexible medium that recovery is relatively inexpensive. In other words, although their behaviour may make them seem to fail to plan, human agents really plan to fail because it is more efficient to do so in the long run. Computer agents who are to interact with human ones should take this into account when they react to their partners' contributions, and it might even be desirable for them to adopt this approach themselves. In addition to the analysis, we simulate some of the choices which human agents make using conversations between two computer agents in the JAM system. These agents, given particular communicative posture choices, try to minimise the total effort that will be expended in the dialogue by always considering all possible actions and taking whichever one they believe will lead to the least cost completion of the dialogue. We leave to further work extensions which would allow the agents to decide not to deliberate about what to do completely, just taking the first action which they "think" of, and which would allow the agents to vary their communicative postures during the course of a dialogue.

References

- [1] J. Allen. Recognising intentions from natural language utterances. In M. Brady and R. C. Berwick, editors, *Computational Models of Discourse*, pages 107-166. MIT Press, 1983.
- [2] J. A. Ambrose-Ingerson. Relationships between planning and execution. *AISB Quarterly*, (57), 1986.
- [3] A. H. Anderson, M. Bader, E. G. Bard, E. Boyle, G. Doherty, S. Garrod, S. Isard, J. Kowtko, J. McAllister, J. Miller, C. Sotillo, H. Thompson, and R. Weinert. The hrc map task corpus. *Language and Speech*, 1992 (forthcoming).
- [4] D. Appelt. *Planning English Sentences*. Cambridge U. Press, 1985.
- [5] G. Brown, A. Anderson, R. C. Shillcock, and G. Yule. *Teaching Talk*. Cambridge University Press, 1984.
- [6] J. C. Carletta. *Risk-taking and Recovery in Task-Oriented Dialogue*. PhD thesis, Edinburgh University Department of Artificial Intelligence, 1992 (forthcoming).
- [7] H. Clark and E. Schaefer. Contributing to discourse. *Cognitive Science*, 13, 1989.
- [8] R. Dale. *Generating referring expressions in a domain of objects and processes*. PhD thesis, Edinburgh University, 1988.
- [9] B. Grosz and C. Sidner. The structures of discourse structure. Technical Report 6097, BBN, 1985.
- [10] G. Houghton. *The Production of Language in Dialogue: A Computational Model*. PhD thesis, University of Sussex, April 1986.
- [11] W. C. Mann and S. A. Thompson. Rhetorical structure theory: A theory of text organization. Reprint Series 190, ISI, 1987.
- [12] K. McCoy and J. Cheng. Focus of attention: constraining what can be said next. In *Proceedings of the 4th International Workshop on Natural Language Generation*, 1988.
- [13] J. D. Moore. A reactive approach to explanation in expert and advice-giving systems. Special Report 251, University of Southern California/Information Sciences Institute, 1990.
- [14] D. Peachey and G. McCalla. Using planning techniques in intelligent tutoring systems. *International Journal of Man-Machine Studies*, 24, 1986.
- [15] R. J. D. Power. *A Computer Model of Conversation*. PhD thesis, University of Edinburgh, 1974.
- [16] E. Reiter. Generating descriptions that exploit a user's domain knowledge. In R. Dale, C. Mellish, and M. Zock, editors, *Current Research in Natural Language Generation*. Academic Press, 1990.
- [17] N. R. Shadbolt. *Constituting reference in natural language: the problem of referential opacity*. PhD thesis, Edinburgh University, 1984.
- [18] M. J. Stefik. Planning and meta-planning. *Artificial Intelligence*, 16, 1981.
- [19] D. E. Wilkins. *Practical Planning*. Morgan Kaufmann, 1988.