

Exploring Initiative Strategies Using Computer Simulation

Fan Yang and Peter A. Heeman

Center for Spoken Language Understanding
OGI School of Science & Engineering
Oregon Health & Science University
{fly,heeman}@cslu.ogi.edu

Abstract

We envision that next-generation spoken dialogue systems will be mixed-initiative. However, it is unclear how exactly a mixed-initiative strategy should be designed; under what circumstances should the system take the initiative, and under what circumstances should it let the user do so. The initiative strategies used in human-human conversation are a good starting point, because they are natural for the user to follow. Studying human-human conversation, however, only gives a descriptive account of human strategies. In this paper, we explore the use of computer simulation to better understand human conventions and give an explanatory account. We have two software agents solve a collaborative task using different initiative strategies, the first derived from analysis of human-human dialogues, and two alternatives based on proposals in the literature. Our simulation results show that the former is more efficient than the others. This helps support the explanation that people use an initiative strategy that minimizes collaborative effort.

1. Introduction

The first generation of spoken dialogue systems are mainly system-initiative, in which the system asks questions and the user answers. The usefulness of such systems depends on how structured the task is. As the complexity of the task increases, system-initiative dialogues will not be sufficient, as they will hinder the user's ability to work with the system to accomplish the task.

Next-generation spoken dialogue systems need to be mixed-initiative, in which both the system and the user are allowed to take the conversational lead. The dialogue manager will need to determine what contributions the system should make and when it should make them. For instance, in a collaborative problem-solving domain, the system has to decide whether it should take the lead to propose possible solutions, or it should let the user take the lead. Thus, to build a mixed-initiative dialogue system, we need a strategy that specifies the system's initiative behavior.

We take the stand that computer dialogue systems should be natural for users to follow. Thus instead of creating a strategy based on conjecture, we start with a strategy actually used in human-human dialogues. Once we understand the human conventions, we can implement them in a dialogue system to allow natural human-computer interactivity.

To better apply human strategies in a dialogue system, we not only need to understand what people do in conversation (a descriptive account), but also why people are doing so (an explanatory account). Analyzing human-human dialogues, however, only reveals a descriptive account of human conventions.

It does not reveal the reasons that underlie the human conventions [1]. Computer simulation has been a promising tool for looking inside human conventions [2, 3, 1, 4]. Human languages have evolved over generations. There must be some benefits for people to adopt certain conventions, such as to minimize collaborative effort [5]. Given appropriate models of communication, reasoning, and solution quality, computer simulation can help researchers understand human conventions at a deeper level by justifying the adoption of them.

The long-term goal of our research is to investigate effective strategies of initiative for human-computer dialogues. In our previous work, we examined the Trains corpus [6] and found that initiative is subordinate to dialogue goals [7]. This paper describes our preliminary experiments in using computer simulation to understand the benefits of this conventional strategy. We compare the performance of this initiative strategy against two variations derived from Hagen's claim about mixed-initiative human-computer dialogue [8]. The purpose of this paper is twofold: to explore the use of computer simulation for initiative strategies, and to compare the three strategies.

2. Initiative Strategy in Trains

Initiative is about leading the conversation towards the dialogue goal. Researchers have tried to understand how people manage initiative in conversation. On one extreme, Hagen stated that "taking the initiative equals taking the turn", and thus initiative places no constraints on what conversants can do [8]. In our previous work [7], we examined the Trains dialogues, and found that initiative does not bounce back and forth on a turn by turn basis, as proposed by Hagen. Instead, people have certain conventions in guiding when to take initiative and when to release it. Figure 1 shows a stretch of conversation from the a Trains dialogue, in which two speakers were making a plan to produce and ship orange juice. The user took the initiative in creating the plan while the system just made acknowledgments. After the plan was created, the system showed initiative by evaluating the plan and telling the user that it would take twelve hours.

Such dialogue phenomena happened a lot in the Trains corpus. The speakers collaborated in each joint-goal, in which usually one speaker took the lead in proposing, and the other took the lead in evaluating the proposed plan. According to Grosz and Sidner [9], the joint-goals in a dialogue are the purpose of a discourse segment, which is the segment initiator's intention to create a shared plan [10]. Thus we proposed that initiative is subservient to discourse structure: the speaker who initiates a discourse segment is the one who has initiative during the entirety of the segment; the non-initiator plays a restricted role in achieving the purpose of a discourse segment [7]. We refer to this model as *restricted initiative*. According to this initiative

user:	then how about if I take an engine along with a boxcar over to Corning
system:	mm-hm
user:	add a tanker on
system:	mm-hm
user:	fill the boxcar with oranges and bring it back to Elmira to have it turned into juice
user:	and drop off the boxcar fill the tanker with orange juice and deliver it to Avon
system:	that will take twelve hours
user:	ok

Figure 1: An example dialogue

model, the non-initiator is limited in showing initiative. Initiative is the authority owned by the segment initiator that the non-initiator has to respect. There are several possible hypothesis about why the non-initiator does not equally contribute to the discourse segment purpose:

1. The non-initiator fails to determine the initiator’s sub-goal for the segment till the very end of the segment. In this case, because the non-initiator does not know what the initiator is trying to accomplish, the best the non-initiator can do is to just follow along.
2. Even when the non-initiator knows the initiator’s sub-goal, the non-initiator might still want to wait till the initiator finishes a proposal for the whole plan before commenting on it for efficiency reasons, unless for urgency reasons (c.f. the principles of plan quality and information quality [11]).
3. Other factors, such as actor role, social position, and cultural conversation rules related to politeness might also contribute to the restricted initiative (these rules might exist for efficiency reasons as well).

The first reason is simply because the non-initiator does not know enough to meaningfully contribute, while the second is a deliberative strategy of holding off in order to let the initiator get his entire point across. In this paper, the second hypothesis is examined with computer simulation.

3. Computer Simulation: Domain

To determine whether efficiency can account for the restricted initiative strategy, we run computer simulation in which the user and system are modeled as software agents, and both know what the joint-goal is and are able to think of solutions to solve it. The task and the conversation from the Trains corpus are too complex to model directly. Hence, we use a simpler domain for our simulation, but one that is still complex enough to shed light on the initiative model used in Trains. We use a revised version of the furniture layout task [1, 12].

3.1. Domain Task

The domain task involves two agents, A and B, carrying out a dialogue to arrange furniture in a room. A furniture item has three attributes: type (e.g. sofa), color (e.g. red), and value (1 to 20). Information about all furniture items is mutually known by both agents.

Each agent also has a set of preferences, initially unknown by the other agent, that it wants to satisfy. Six types of preferences are used, as shown below.

- *Item X must be in the room.*
- *Item X must not be in the room.*

- *If item X is in the room, item Y must also be in the room.*
- *If item X is in the room, item Y must not be in the room.*
- *If item X is not in the room, item Y must be in the room.*
- *If an item of type T is in the room, its color must be C.*

The variables X and Y in the preferences can be either specific items, such as a green table, or can just be the type of furniture, as in ‘a table’. Each preference has a value between 1 and 20. The task of the two agents is to pick five furniture items for the room that result in the highest possible score where the score is the sum of the values of the five furniture items, less the value of any violated preference from either agent.

The complexity of the task can be characterized by the following parameters.

I_o : overall number of furniture items
 P_A : number of A’s preferences
 P_B : number of B’s preferences

3.2. Communication Language

To solve the task, we use the following speech act schemas.

- (*propose Speaker Hearer Item*)
- (*inform Speaker Hearer Preference*)
- (*accept Speaker Hearer*)
- (*reject Speaker Hearer*)
- (*acknowledge Speaker Hearer*)
- (*release_turn Speaker Hearer*)

The first five speech acts are from [12], which allow agents to propose items, and to accept or reject them. Agents can also inform each other of their preferences so that they can reach agreement as to what items to include. Agents also must acknowledge proposals and informs. As defined by Whittaker and Stenton, the speech acts of *propose* and *inform* show initiative while the others do not [13].

For turn taking, we include the action “release_turn”, which the agent that currently has the turn can perform to signal that it is relinquishing the turn.

3.3. Agent Reasoning

The furniture arrangement task involves complex reasoning. Assuming that an agent knows all of the preferences, choosing the five best items is an NP-complete problem. Luckily, a greedy strategy can be employed, whose performance is very close to optimal [12]. To get a solution with the greedy algorithm, one finds the single item that results in the highest score and commits to it, and then finds the second item such that it and the first item result in the highest score, etc.

The agents use the greedy strategy in their reasoning. They consider all preferences that they know about: their private preferences as well as the ones that have been communicated by the other agent. We also assume that agents never make mistakes and never forget preferences that have been communicated.

3.4. Problem Solving Assumptions

There are different strategies that agents might adopt to solve the problem in this domain. For example, agents can communicate their preferences with each other, and then create a best plan. Or agents can create a complete plan, and then evaluate it, and then communicate violated preferences, as in the plan mode of [12]. Or agents can decide the items /to move in one by one, as in the item mode of [12]. Or agents might employ some combination of the above three strategies. In this simulation, we focus on the strategy similar to the plan mode, and thus we assume the following:

- Agents will have two joint-goals: to create a plan and to evaluate a plan. Similar to the conversation in Figure 1, agents will first jointly create a complete plan, and then jointly evaluate the plan and communicate preferences. This procedure repeats until both agents agree on a plan.
- Preferences are communicated when violated. Unlike in the artificial domain of the Trains, people often times do not have hard constraints to satisfy. Instead, it is usually a trade-off between difference preferences, which come to people’s mind for comparison when they are violated.
- An agent accepts a plan when the agent has no better alternative or the agent has no private preferences that impact the plan.

4. Initiative Strategies

In Section 3.4, we specify a problem-solving strategy, but we do not specify which agent does which action; in other words, we do not specify the initiative strategy. In this section, we contrast an initiative strategy based on the restricted initiative model with two others, in which agents alternate initiative on a turn-by-turn basis.

4.1. Plan-Based Restricted Initiative (PI)

In the restricted initiative model, one conversant initiates a discourse segment to solve a domain goal, and takes the lead in developing the plan. After the plan is presented, the agents then take on the goal of critiquing it, again, with one of the agents taking the lead. To implement this in the furniture domain with the plan mode, we have each agent first reason about how they would solve the problem given the preferences it knows. Then, one agent, A, takes the initiative and presents its plan one item at a time, with B just acknowledging after each item. Once the plan is fully presented, A releases initiative. B then takes on the goal of evaluating the plan. B first compares A’s plan to the one that B would have proposed. If B does not accept A’s plan, B will take the initiative to inform A of any preferences that are unknown to A and that lower the score of A’s plan but not B’s. Then B rejects A’s plan and releases initiative. A will take the initiative to create another plan. This process repeats until they reach an agreement on a plan. In this strategy, the agent without initiative can take speaking turns, but the turns simply consist of acknowledgments.

4.2. Turn Initiative I (TI1)

As an alternative, we developed a strategy in which agents do not keep initiative over an entire plan. This strategy is consistent with Hagen’s claim that “taking the initiative equals taking the turn”. We present two variations of this. In the first both agents first devise a plan on their own, and then alternate in proposing the five items of a furniture plan, and then alternate in critiquing the resulting plan. In proposing an item, the agent picks the first item from its plan that has not yet been proposed.

4.3. Turn Initiative II (TI2)

The strategy in Section 4.2 has the agents reason about the furniture items as if they will be able to present their plan in its entirety. To overcome this mismatch, we present a second alternative in which agents reason about which furniture item to propose immediately before proposing it, taking into consideration the ones already proposed. Thus A first determines the best item, based on the preferences it knows. Then B determines

the second item, given that the item A proposed is in the plan. Then A determines the third item that is best given the first two. This continues until all five items are chosen, followed by plan evaluation.

5. Computer Simulation Results

To show that the results hold over a range of task difficulties, we examine the three initiative strategies under different simulation parameter settings. In this section, we use $(I_o, P_A - P_B)$ to denote these different settings, for example, (20,40-40) means that each agent initially has 40 private preferences and they choose five out of 20 furniture items to move into the room. For each parameter setting, we create 100 sets of preferences and items and run the simulation experiments for all three initiative strategies.

For evaluating the initiative strategies, we measure solution quality and communication effort. Solution quality is the score of the five items that are agreed upon. Communication effort is the number of speech acts of proposing and informing (showing initiative) in a dialogue.

5.1. Solution Quality

Figure 2 shows the solution quality of the three initiative strategies under different parameters. Figure 2A shows the solution quality as the number of items is varied and the number of preferences of each agent is held at 40. Figure 2B shows the solution quality as the number of preferences is varied and the number of items is held at 20. The solution quality of the three strategies under each setting is very similar and the differences are not statistically significant (paired ttest $p > 0.1$).




Figure 2: Solution Quality

5.2. Communication Effort

Figures 3A and 3B show the communication effort of the three initiative strategies under different parameters. On average, the communication effort of PI is less than that of TI1 and TI2. The difference is statistically significant except under the setting (20,20-20) (paired ttest, $p < 0.01$). In fact, the more complex the task is (higher number of preferences or items), the higher the difference is. For example, for the setting (20,20-20), the difference between PI and TI1 is 1.15 speech acts, which is about 8% of PI; for (20,40-40), the difference is 4.9, 25% of PI; for (20,80-80), the difference is about 10, about 37% of PI.

In terms of speech acts, not only do agents in TI1 and TI2 communicate more preferences (inform) than PI, but they also make more proposals. In the setting of (20,80-80), agents in PI make 14.95 proposals, and in TI1 make 20.10, in TI2 make 20.85. Because each plan takes five proposals, this means that agents in TI1 and TI2 go through more proposal-evaluation cycles.




Figure 3: Communication Effort

Figure 4 tracks how the solution quality of the proposed plans changes as a dialogue progresses for the three initiative strategies with the setting of (20,80-80). We see that for the first plan proposed, the average solution quality of PI (about -417) is lower than the other two (about -407). However, after communicating the preferences, agents working in PI are able to create a better plan (-397 vs -402). PI encourages more useful preferences to be communicated between the agents. Thus we see that although PI has a slow start, it quickly makes up for lost ground, and overtakes the other two strategies.

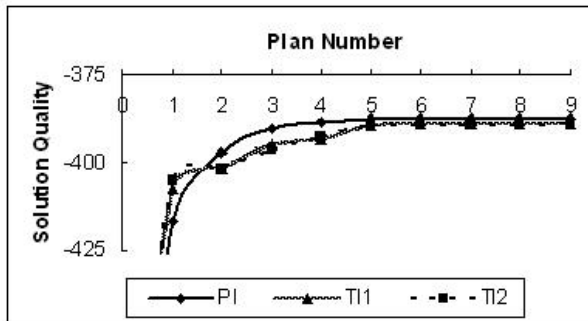


Figure 4: Solution Quality on Each Plan

6. Conclusion

Mixed-initiative does not mean showing initiative freely without any constraints. Instead, we believe that there are certain rules that people follow. Based on our empirical study on the Trains corpus, we proposed that initiative belongs to the speaker who initiates a problem-solving goal. This paper provides further support for that finding, through the use of computer simulation. We set up a domain and implement three plan-based problem-solving strategies, the first consistent with the conventions in the Trains corpus, and two others in which agents alternate initiative on a turn-by-turn basis. We compare the performance of these initiative strategies and find that, under certain domain settings, the restricted initiative strategy produces more efficient dialogues in terms of less communication effort without any loss in solution quality. Additionally, as the complexity of the domain increases, so does the advantage of the restricted initiative strategy.

A good initiative strategy for computer dialogue systems should have two features. First, it should be natural for users to follow. This reduces the training effort on users and encourages users to work with the system. Second, it should be efficient in problem-solving. According to the Paradise evaluation plat-

form, dialogue length is one of the most important factors that lead to improved user satisfaction [14]. The simulation experiment, by showing that the strategy of restricted initiative is also efficient, suggest that this strategy could also work for a computer dialogue system: with the computer playing the role of proposing a plan for the user to evaluate.

Computer simulation opens a new window for examining human language. Our preliminary experiments shed light on using computer simulation to better understand human conventions. In our future work, we plan to extend our current experiments and introduce more sophisticated cognitive models on human reasoning and human memory, to better simulate human behavior. This should allow us to better understand the applicability of the restricted initiative model, and the circumstances in which it does not apply.

7. References

- [1] M. A. Walker, "The effect of resource limits and task complexity on collaborative planning in dialogue," *Artificial Intelligence Journal*, vol. 85, pp. 181–243, 1996.
- [2] R. Power, "The organization of purposeful dialogues," *Linguistics*, vol. 17, pp. 107–152, 1979.
- [3] J. Carletta, "Planning to fail, not failing to plan: Risk-taking and recovery in task-oriented dialogue," in *Proceedings of 14th International Conference on Computational Linguistics*, Nantes France, 1992, pp. 896–900.
- [4] C. I. Guinn, "An analysis of initiative selection in collaborative task-oriented discourse," *User Modeling and User Adapted Interaction*, vol. 8, pp. 255–314, 1998.
- [5] H. H. Clark and D. Wilkes-Gibbs, "Referring as a collaborative process," *Cognitive Science*, vol. 22, pp. 1–39, 1986.
- [6] P. A. Heeman and J. F. Allen, "The Trains spoken dialogue corpus," Linguistics Data Consortium, CD-ROM, 1995.
- [7] F. Yang, P. A. Heeman, and K. Hollingshead, "Towards understanding mixed-initiative in task-oriented dialogues," in *Proceedings of 8th ICSLP*, 2004, pp. 217–220.
- [8] E. Hagen, "An approach to mixed initiative spoken information retrieval dialogue," *User Modeling and User Adapted Interaction*, vol. 9, pp. 167–213, 1999.
- [9] B. J. Grosz and C. L. Sidner, "Attention, intentions, and the structure of discourse," *Computational Linguistics*, vol. 12, no. 3, pp. 175–204, 1986.
- [10] K. E. Lochbaum, "A collaborative planning model of intentional structure," *Computational Linguistics*, vol. 24, no. 4, pp. 525–572, December 1998.
- [11] M. Walker and S. Whittaker, "Mixed initiative in dialogue: An investigation into discourse segmentation," in *Proceedings of 28th ACL*, 1990.
- [12] F. Yang and P. A. Heeman, "Using computer simulation to compare two models of mixed-initiative," in *Proceedings of 8th ICSLP*, 2004, pp. 213–216.
- [13] S. Whittaker and P. Stenton, "Cues and control in expert-client dialogue," in *Proceedings of 28th ACL*, 1988, pp. 123–130.
- [14] M. A. Walker, D. J. Litman, C. A. Kamm, and A. Abella, "PARADISE: A framework for evaluating spoken dialogue agents," in *Proceedings of 35th ACL*, Madrid, Spain, July 1997, pp. 271–280.