

An Assessment of Dialogue Strategies for a Human Computer Debating System, via Computational Agents

Tangming Yuan and David Moore and Alec Grierson¹

Abstract: This paper reports our research concerning dialogue strategies suitable for adoption by a human computer debating system. In particular, we propose a set of strategic heuristics for a computer to adopt to enable it to function as a dialogue participant. We also consider means of constructing a set of computational agents operationalising the proposed strategy and means of analysing the agent-generated dialogues to facilitate the evaluation of the proposed strategy. It is anticipated that this work will contribute toward the development of human-computer dialogue systems and help to illuminate research issues in the field of strategies in dialectical systems.

1. INTRODUCTION

In formal dialectical systems, the dialogue regulations usually leave some room for choices as to permissible move type and substantive content. It is crucial therefore that a computer system designed to utilise such a system has some means of selecting between the available possibilities. This choice must be based on some suitable strategy. Appropriate strategic knowledge is, then, essential if the computer is to produce high quality dialogue contributions. Further, the importance of strategies in dialectical systems has also been stressed elsewhere (e.g. [1]; [2]; [3]; [5]; [6] [7]; [12]).

A previous paper [17] proposed a set of strategic heuristics for a human computer debating system. Here, we seek to further that investigation in several steps. Firstly, we provide a brief introduction to the set of strategic heuristics. Secondly, we discuss means of evaluating the proposed strategy and construct a set of computational agents able to debate with each other via the proposed strategy. Thirdly, we analyse the dialogue transcripts generated by these agents as they use the set of proposed strategic heuristics. Finally, we summarise and discuss the results of this evaluation.

2. STRATEGIC HEURISTICS

A set of debating heuristics was proposed for our human computer debating system in [17]. The debating system is currently configured as what can be described as a “partially honest” agent, in that the agent is generally expected to speak the truth but it is allowed to insist on its own view for the sake of argument even though it may have more reasons in favour of the user’s view in its knowledge base. The system utilises the dialogue model “DE” [cf. 16]. In a DE dialogue, there are five dialogue situations that the computer might face, defined by the previous move type made by the user: a challenge, a question, a resolution demand, a statement or a withdrawal. The strategic decisions under these dialogue situations are captured at three levels in our current model:

- 1) Retain or change the current focus.
- 2) Build own view or demolish the user’s view.
- 3) Select method to fulfil the objective set at level 1 and 2.

Level (1) and (2) refer to strategies, which apply only when the

computer is facing a statement or withdrawal, while level (3) refers to tactics used to reach the aims fixed at level 1 and 2, and applies in every game situation. The level 3 heuristics for each dialogue situation are given in turn below.

1) A question raised by the user

Questions asked involve questioning an individual statement, e.g. “Is it the case that P?”. In such a situation, the computer is allowed by the DE rules to answer “Yes”, “No” or “no commitment”. Heuristics for the computer when facing a question are proposed as follows.

- (1) If neither P nor $\neg P$ can be found in the knowledge base (KB), then the computer speaks the truth with a “no commitment”.
- (2) If only one of them (P and $\neg P$) can be found in the KB,
 - a. If the computer has previously uttered “no commitment” to the found statement, then it utters “no commitment” to remain consistent.
 - b. Else the computer speaks the truth and utters the found statement.
- (3) If both (P, $\neg P$) are found in the computer’s KB, and assuming that one of them (say $\neg P$) supports the computer’s view and the other (say P) supports the user’s view.
 - a. If the computer has an acceptable support for $\neg P$, then utter $\neg P$.
 - b. If the computer has no acceptable support for $\neg P$, and the computer has not committed to the reasons for P, the computer should say “no commitment”.
 - c. If the computer has no acceptable support for $\neg P$, and the computer has committed to the reasons for P, then the computer should utter P.

2) A Challenge made by the user

There are three DE legal options available in response to a challenge: a resolution demand, a ground, or a withdrawal. The heuristics after a challenge of P are proposed as follows.

- (1) If P is a modus ponens consequence of the user’s commitment, then pose a resolution demand.
- (2) Else if there is only one acceptable ground available in the knowledge base, then state the ground.
- (3) Else if there is more than one acceptable ground available, then state the one that can be further supported.
- (4) Else if all the available acceptable grounds are equally supported, then randomly choose one of the grounds.
- (5) Else if no acceptable ground is available, then withdraw P.

The concept of an “acceptable” ground is specified in [16] and aims at preventing the fallacy of question-begging.

3) A resolution demand made by the user

A resolution demand made by the user concerns an allegation that the computer has committed to an inconsistency in its commitment store. In the most likely event, the computer would face a resolution demand like resolve $\{\neg P, P\}$ in that the computer has committed to both P and $\neg P$. In this situation, the computer is allowed and required to withdraw one of them to keep consistent. It

¹ School of Computing, Leeds Metropolitan University Leeds LS6 3QS United Kingdom Email: {t.yuan; d.moore; a.j.grierson}@leedsmet.ac.uk

is suggested that the computer withdraws the one which is less supported according to its commitment store.

The user might invoke another type of resolution demand (i.e. resolve (Q, $Q \supset P$, why P) or resolve (Q, $Q \supset P$, withdraw P) in the event of the computer's challenging or withdrawing a modus ponens consequence of its commitments. In this situation, the computer is required, by the game DE, to withdraw either Q or $Q \supset P$ or affirm P. Our current prototype, however, is given means of avoiding challenging or withdrawing such a disputed consequence in the first place, and the user is therefore not able to make this type of resolution demand.

4) A "no commitment" made by the user

After a "no commitment", DE places no restrictions on either move type or contents. The computer's options are therefore more open than in the previous situations considered above. The heuristics for when the computer is facing "no commitment P" are proposed as follows.

- 1) If P supports the user's thesis
 - a. If P is a unique support of the user's asserted proposition Q, and Q is not the user's thesis, then challenge Q.
 - b. Else check whether the user retains adherence to the thesis.
- 2) If P supports the computer's thesis
 - a. If the non-committal statement is a modus ponens consequence of the user's commitments, then pose a resolution demand.
 - b. Else switch the current focus.

5) A statement made by the user

After a statement, there is again no restriction on either move types or move contents in DE. In the most likely event, the computer would face a statement which supports the user's view or opposes the computer's view. However, it is possible that the user may unwisely make a statement which supports the computer's view or goes against his own view. The computer may need to deal with these two kinds of statement differently. When the computer is facing a statement (say P) which supports the computer's thesis or mitigates against the user's view, two heuristics are proposed as follows.

- a) If P is a support of the computer's thesis, then use P as the starting point to build the computer's thesis.
- b) Else check whether the user still retains adherence to the thesis.

When the computer is facing a statement (say P) which supports the user's view or mitigates against the computer's view, a set of heuristics is proposed as follows.

- a) If there is an inconsistency (e.g. (P, $\neg P$)) in the user's commitment store, then ask for resolution.
- b) Else if there is a piece of hard evidence in support of $\neg P$, then state the piece of hard evidence.
- c) Else if there is any support of $\neg P$ and the support (say Q) can be further supported, then state $\neg P$, or state Q if $\neg P$ has been uttered, or form a plan of questions aimed at making the user accept $\neg P$.
- d) Else if there is any support of $\neg P$ and the support cannot be further supported, then form a plan of questions aimed at making the user accept $\neg P$.
- e) Else if P is challengeable, then challenge it.

To decide whether a statement is challengeable, the computer needs to consider the nature of that statement (e.g. whether it is a piece of hard evidence) and the relevant DE dialogue rules. If it is not considered hard evidence, and if a challenge would be legally permissible, then it is deemed to be challengeable.

The plan of questions referred to in heuristic (c) and (d) is organised following Walton's [11] scheme of argument from gradualism. The plan can be started by asking a question of a proposition (say A), followed by a series of connected conditionals (say $A \supset B$, $B \supset C$, ..., $C \supset P$) toward the conclusion (say P). Following [6], the computer would hand over the initiative by stating the conclusion P at the end, if the plan is executed successfully, with a view to avoiding a one-sided dialogue.

In practice, the user might reply to the questions in the plan with unwanted answers (i.e. answers unfavourable to the computer's plan). The approach taken here is that the computer tries to remove the obstacles and put the plan back on its track while the initiative is still held. The plan execution process is therefore as follows.

- 1) If a wanted answer is given, then carry on to execute the plan
- 2) If a non-committal answer is given
 - (2.1) If there is an expressed inconsistency in the user's commitment store, then pose the appropriate resolution demand
 - a) If the user affirms the disputed consequence, then continue the plan
 - b) Else abandon this line of questions
 - (2.2) Else abandon this line of questions
- 3) If an unwanted statement (e.g. $\neg P$ rather than P) is given
 - (3.1) If there is an expressed inconsistency in the user's commitment store and the unwanted answer $\neg P$ is an element of the inconsistency, then pose the appropriate resolution demand
 - a) If the unwanted answer is withdrawn, then continue the plan and re-pose the question.
 - b) Else abandon this line of questions
 - (3.2) Else if the unwanted statement is challengeable, then challenge the unwanted statement
 - a) If the unwanted answer is withdrawn, then continue the plan to re-pose the question of P.
 - b) Else abandon this line of questions
 - (3.3) Else abandon this line of questions

This, then, is the set of strategic heuristics currently adopted by our human computer debating system. The issue of whether the proposed strategy can in practice provide adequate services for a computer as a dialogue participant to produce good dialogue contributions cannot be settled on an a priori basis. To assess the appropriateness of a proposed strategy, Maudet and Moore [5] suggest that the strategic heuristics need to be tested, and that a convenient way to do this is via generation of dialogue by the computer itself. There are two possible ways to approach this: one is to enable a human user to debate with a computerised debating system, and the other is to allow two computers to engage in dialogue with each other and study the results. Although both approaches are seen as important to evaluate the appropriateness of the proposed strategy from different perspectives, the approach adopted in this paper is the latter. A pre-requisite of this approach is the construction of suitable computational agents. This will be considered in the next section.

3. COMPUTATIONAL AGENTS

Having decided to use computational agents to conduct this study, it is necessary to consider how to set up such agents. Two computational agents (referred to henceforth as Simon and Chris) to conduct debate with each other have been built by the authors using the Java programming language, reusing the test-bed (DE agents system) developed in [16]. The proposed strategy outlined in the previous section has been incorporated into the agent-based systems. The system architecture is shown in figure 1.

There are six main components of the system. Example output of the *interface unit* is shown in figure 2 and depicts two agents

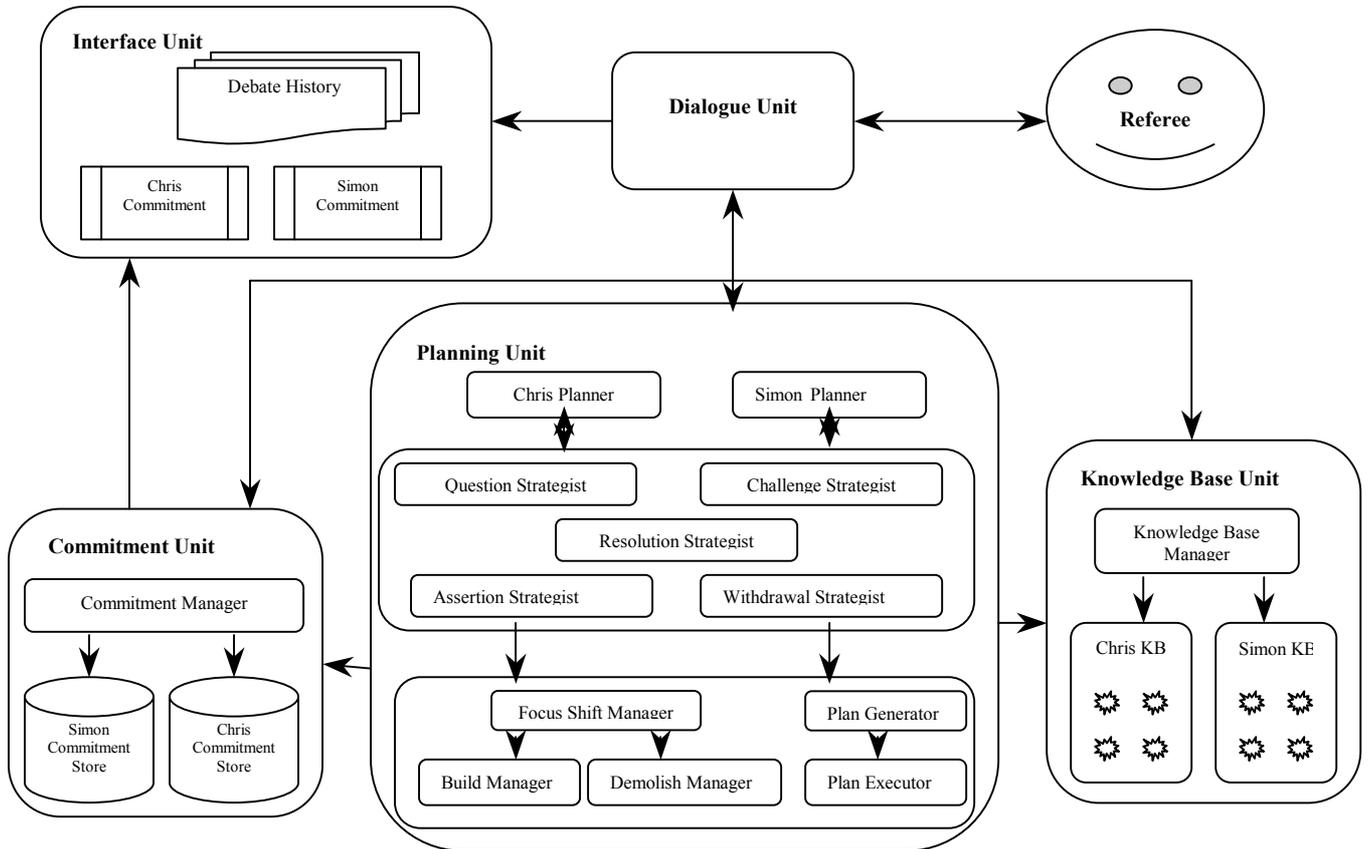


Figure 1 Computational agents system architecture

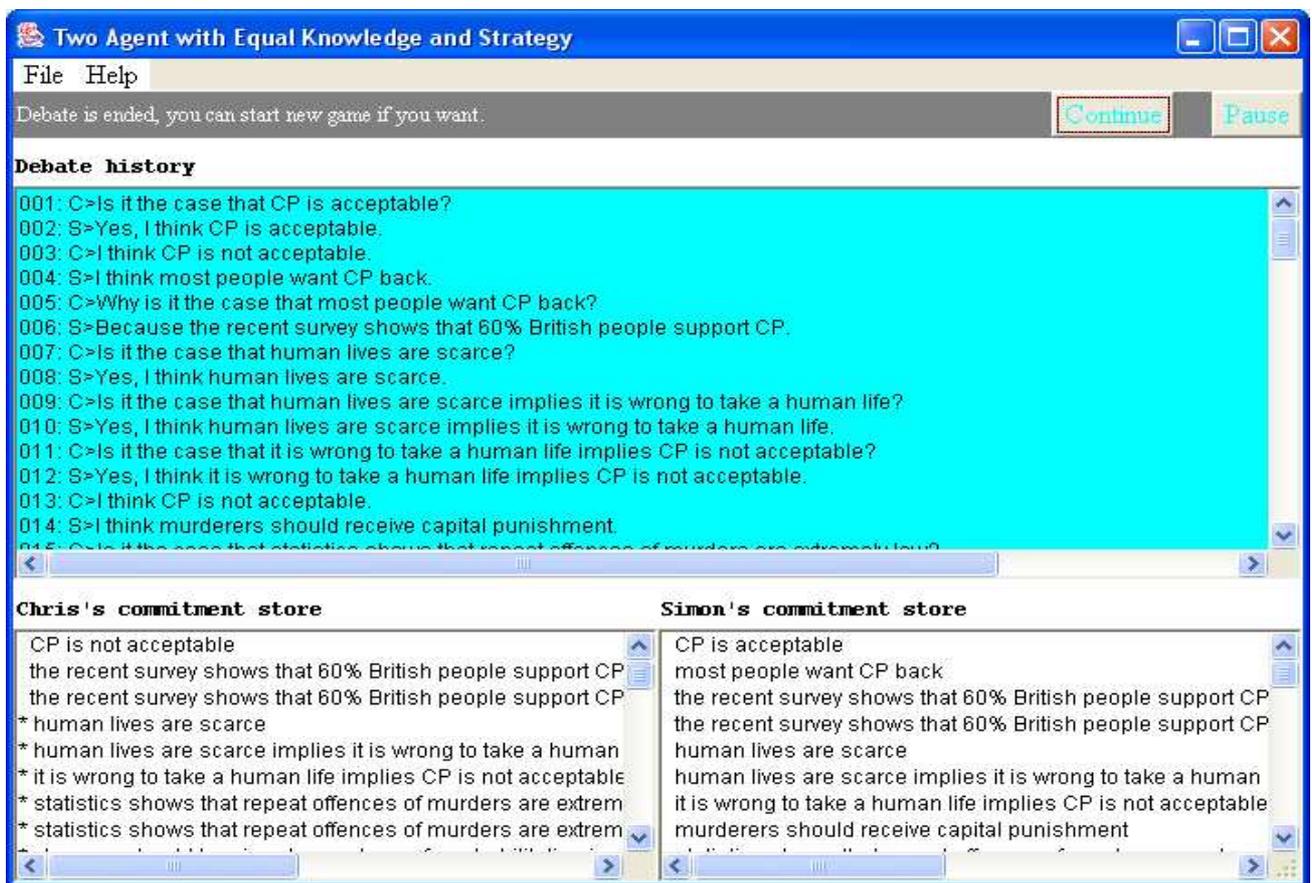


Figure 2 Computational agents system interface

(Chris and Simon) debating capital punishment (CP). This unit provides a dialogue history, which records the debate. The commitment-stores show both agents' commitment contents. In order to control the process of the debate, a "New game" menu item is designed to start the debate, a "Pause" button is available to temporarily stop a debate, and a "Continue" button will carry on the dialogue if necessary. A "Save as" menu item is designed to save the dialogue history and both commitment sets as separate files for subsequent analysis.

The *dialogue unit* (figure 1) can be regarded as the despatch centre of the agent interaction. This unit provides functions to update the dialogue history and manage the turn taking of the agents and referee. It will schedule the corresponding agent to make a move and then the commitment manager to update the commitment stores.

The *referee* is in charge of enforcing the DE rule set and of the termination of the agents' debate; the original DE regime makes no stipulation re winning and losing, but following [6], one agent will lose the debate when it has given up its thesis or explicitly committed to the opponent's thesis. However, it might be possible that one agent runs out of strategy but still adheres to its thesis; under this circumstance, it is suggested that the agent should hand over its turn to its dialogue partner. If both agents run out of strategy and a winner is still not decided, the referee will call off the game (in effect, the dialogue ends in a "stalemate").

The *planning unit* is responsible for generating moves in the light of (i) the knowledge base, (ii) the prevailing state of both commitment stores and (iii) the dialogue rules. Each agent's planner manages assertion, challenge, withdrawal, resolution and question "strategists", each of which is designed to deal with a different dialogue situation following the set of heuristics discussed in section 2. When the agent planner receives calls from the dialogue manager, it will check the current dialogue situation and schedule the corresponding strategist to produce a move. The agent's planner will then pass the move to the dialogue unit to make the agent's contribution.

In addition, there are five components (focus shift manager, build manager, demolish manager, plan generator and plan executor) that are designed to provide special services to the assertion and the withdrawal strategists. The focus shift manager will be called by the assertion or withdrawal strategist to decide whether to change the current focus. The build and demolish managers will be called by the focus shift manager to check whether there are methods available to either build its own positions or attack partner's positions. The plan generator is responsible for generating a set of propositions and forming a line of questions when required by the assertion or withdrawal strategist, the build manager or the demolish manager. The plan executor is responsible for executing a plan. The assertion and withdrawal strategists will constantly look up whether there is a plan under execution, if there is, then they call the plan executor to carry on its execution.

The *commitment unit* is responsible for updating both agents' commitment stores. It contains a commitment manager and two commitment stores, one for each party. The commitment manager will update both agents' commitment stores according to the DE commitment rules. Each commitment store is designed to have two lists of statements, those that have been stated and those that have been merely implicitly accepted. In order to distinguish them from each other, a statement that is only implicitly accepted is marked with an asterisk as shown in figure 2.

The *knowledge base unit* (figure 1) consists of a knowledge base manager and the knowledge bases of the two agents. When the game starts, the dialogue manager will invoke the knowledge base manager to initialise both agents' knowledge bases. The agent knowledge bases contain a set of propositions and consequence relationships between these propositions. These relationships are based on a Toulmin-like structure [8]. The domain knowledge is formalised from Moore's [6] experimental study of DC with human participants.

This, then, is our agent-based system. It was run under the three conditions below:

- (1) One of the agents adopts the strategy and the other uses random argument, and both have the same knowledge base.
- (2) Both agents adopt the same strategy and share the same knowledge base.
- (3) Both agents adopt the same strategy. One of the agents has more knowledge in its knowledge base than the other

It is anticipated that using random argument (1) might reveal certain failures of the heuristics (e.g. unexpected new situations) that might be overlooked by manual use of them. Conditions (2) and (3) may reveal whether relevant issues are well discussed, and condition (3) may also be used to see whether an agent with a smaller knowledge base might turn out to be the loser of the debate since both agents share the same dialogue strategy.

Three dialogue examples (DE4, DE5 and DE6 - full transcripts are located at <http://www.geocities.com/yuantangming/agentdialoguetranscript/StrategicAgentsTranscripts.doc>) have been generated under the condition (1), (2) and (3) respectively for analysis. A prerequisite of such an analysis is a set of evaluation criteria that are independent of the set of debating heuristics. These criteria are proposed next.

4. EVALUATIVE CRITERIA

Five criteria for evaluation are proposed. Criteria (2), (3) and (4) are drawn from [6] and [9]. Criteria (1) and (5) are intuitively seen as important in this context. These criteria are outlined as below:

- (1) *Robustness*. The issue here is whether all dialogue situations are reasonably dealt with. In particular, this concerns whether there are unexpected dialogue situations which have not been considered in the strategy.
- (2) *Equal opportunity*. The issue here is whether each agent has equal opportunity to advocate their point of view. In particular, this criterion concerns whether there are frequent initiative shifts in the process of dialogue, such that the resulting debate transcripts can be described as "mixed initiative" dialogue.
- (3) *Coverage of issues*. The interest here is whether the knowledge in the knowledge base is well revealed and discussed. One of the potential applications of the debate system is to broaden the interaction style of computer based learning systems [15]. In such an educational setting, it might be expected that the system encourages students to look at an issue from different perspectives, and therefore it would be hoped that as many issues as possible are raised.
- (4) *Argument flow*. The issue here concerns whether the dialogues produced are unreasonably disjointed. It is expected that the strategic agents dialogue contributions are clearly related to its dialogue partner's previous utterances, as a result of which, the flow of the developing argument can be deemed acceptable.
- (5) *Defeasibility*. This concerns whether the strategy is making the agent *too* wise to be beaten, and thus leading to difficulties where a human is another agent of the dialogue. A useful debating system, that is, should be able to reasonably lose a dialogue as well as to win, and thus avoid demoralising its human interlocutors (cf. [6]; [9]).

5. ANALYSIS OF RESULTS

This section contains the results of an analysis of the agent-generated dialogue examples DE4, DE5 and DE6. During the analysis, each utterance of the dialogues is considered in turn, via the addition of appropriate annotations in square brackets (see <http://www.geocities.com/yuantangming/agentdialoguetranscript/StrategicAgentsTranscripts.doc> for further details), for example:

001: C>Is it the case that CP is acceptable? ..[C starts the game by

Table 1 Dialogue situations summary

Dialogue situations		Strategic responses	
Move types	Move content turn		
Assertion (74)	Speaker's thesis (18)	DE4-002, DE5-002, DE6-002	Adopt the opposite view
		DE4-016, DE4-024, DE5-034, DE6-006, DE6-021, DE6-031	Build its thesis
		DE5-003, DE5-046, DE6-003	Issue a direct thesis support
		DE6-014, DE6-043	Challenge
		DE5-043, DE6-018, DE6-028, DE6-040	Hand over the turn
	Supporting speaker's view (25)	DE5-004	Issue contradictory evidence
		DE5-020, DE5-047	Issue an objection
		DE4-052, DE5-009, DE5-014, DE6-009	Demolish plan
		DE4-036, DE5-006, DE5-021, DE5-025, DE5-048, DE6-024, DE6-045, DE6-047	Challenge
		DE5-005, DE5-008, DE5-019, DE6-008	Switch the current focus and issue a direct thesis support
		DE4-026, DE5-023, DE5-027	Switch the current focus and build its thesis
		DE4-042, DE4-054, DE5-050	Hand over the turn
	Responder's thesis(1)	DE6-053	End the game
	Supporting the responder's thesis (21)	DE4-012, DE4-038, DE4-040, DE4-048	Use to build its thesis
		DE5-011, DE5-013, DE5-016, DE5-018, DE5-029, DE5-031, DE5-033, DE5-036, DE5-038, DE5-040, DE5-042, DE6-033, DE6-035, DE6-037, DE6-039	Continue its plan execution
		DE4-056, DE4-022	Check partner's thesis adherence
	Handing over turn (5)	DE5-044, DE6-019, DE6-029, DE6-041	Check partner's thesis adherence
		DE5-051	End the game
Unrecognised statements (4)	DE6-004, DE6-012, DE6-016, DE6-026	Check partner's thesis adherence	
Withdrawal (11)	Speaker's thesis (1)	DE4-058	End the game
	Supporting speaker's view (2)	DE6-049	Further challenge
		DE6-051	Check partner's thesis adherence
	Supporting responder's thesis (8)	DE4-014, DE4-018, DE4-028, DE6-011, DE6-023	Drop the plan and build C's thesis.
DE4-030, DE4-050, DE4-044		Hand over the turn	
Question (37)	Game start (2)	DE5-001, DE6-001	Give a positive answer
	Supporting speaker's view (24)	DE4-006, DE4-010, DE4-032, DE5-010, DE5-012, DE5-015, DE5-017, DE5-028, DE5-030, DE5-032, DE5-035, DE5-037, DE5-039, DE5-041, DE6-032, DE6-034, DE6-036, DE6-038	Give a positive answer
		DE4-004, DE5-024, DE6-007	Give a negative answer
		DE4-034, DE6-010, DE6-022	No commitment
	Responder's thesis (9)	DE5-045, DE6-005, DE6-013, DE6-017, DE6-020, DE6-027, DE6-030, DE6-042	Give a positive answer
		DE6-052	Give a negative answer
	Supporting responder's thesis (2)	DE4-008, DE4-020	Give a positive answer
Challenge(11)	DE4-046, DE5-007, DE5-022, DE5-026, DE5-029, DE6-015, DE6-025, DE6-044, DE6-046	Give a suitable ground	
	DE6-048, DE6-050	No commitment	
Resolution demand	No examples generated in the dialogues.		

asking S's view on the issue of Capital Punishment (CP)]
002: S>Yes, I think CP is acceptable.[S chooses to support CP]
003: C>I think CP is not acceptable....[C automatically adopts the opposite view]

This approach to the analysis makes it possible to examine the data under the evaluative criteria discussed in the previous section, and thus to assess whether the proposed strategy can provide adequate services to enable the computer to act as a dialogue participant and produce good dialogue contributions.

Evaluation Criterion 1- Robustness

The interest here is whether all dialogue situations generated by the agents are successfully dealt with by the proposed strategy. In total, the agent systems have generated 153 dialogue situations for the strategic agents to deal with. These dialogue situations are summarised in table 1 and discussed in turn below.

a) Assertion Strategist

74 assertions were generated. They are identified as falling into six categories: speaker's thesis, statements supporting the speaker's thesis, opponent's thesis, statements supporting the opponent's thesis, statements handing over its turn, and unrecognised statements. These are discussed below.

It might be expected that for statements standing on the side of the speaker, the strategy should provide heuristics to attack them. It can be seen from table 1 that the assertion strategist does provide various means of either attacking opponent's view or building its own view, with 7 exceptions of giving up this opportunity. On these exceptional circumstances, the assertion strategist runs out of methods and therefore hands over its turn to the opponent. This can be seen as reasonable since the opponent may have something more to say, but on the other hand, more sophisticated strategy is arguably needed if the strategic agent constantly faces this kind of situation.

For statements standing on the side of the opponent, the assertion strategist is expected to use them rather than to attack them (cf. [12]). It is shown in table 1 that the assertion strategy does provide some means of handling this, e.g. using the strategy to build its thesis or continuing its plan execution, or checking the opponent's thesis adherence. On the occasion of DE6-053, the game ends since the speaker has committed to the opponent's thesis.

It is interesting to see that on 5 occasions the assertion strategist has to decide what to do when its dialogue partner hands over its turn. On 1 of the circumstances, the referee calls off the game since both parties have run out of methods. Concerning the remaining 4 instances, the assertion strategist checks its opponent's thesis adherence. This might be reasonable given the situation that its opponent has run out of moves.

There are four unrecognised statements (from the responder's point of view) generated in DE6, the response from the other agent in each case is to check partner's thesis adherence. This may not be enough in a human computer debate setting if users are allowed to input fresh propositions. There is therefore a need for more sophisticated means to handle this kind of situation.

Generally speaking, the assertion strategist appears to be working well, with the exception of needing more sophisticated means to handle unrecognised statements and the situation of running out of moves.

b) Withdrawal Strategist

11 withdrawals (or no commitment) are present in the transcripts. They are categorised as follows: withdrawal of speaker's thesis, withdrawal of statements supporting the speaker's view and withdrawal of statements supporting the opponent's view. These are discussed in turn below.

On one occasion, the speaker is withdrawing its thesis. The game is therefore ended since the speaker has given up its view.

On 2 occasions, the speaker is withdrawing statements supporting its thesis. The response of the strategic agent is to challenge the statement supported by the withdrawn statement, or assess whether the dialogue partner still adheres with his thesis. The former might be regarded as the strategic agent seeking to demolish the statement having more direct influence to its partner's thesis. The latter may also be seen as reasonable given the situation that its dialogue partner might surrender since it has lost some part of the debate.

There are 8 instances of "no commitment" to statements supporting the opponent's view. On 5 occasions, the withdrawal strategist deals with this by starting another line of argument. However, for the remaining 3 instances, the withdrawal strategist fails to do so. The explanation here is that the withdrawal strategist has run out of methods, and therefore hands over its turn. This needs further consideration if the strategic agent constantly faces this kind of situation.

Given the above analysis, the withdrawal strategist seems to be working satisfactorily with the exception of needing more sophisticated approaches when running out of moves.

c) Challenge Strategist

There are 11 challenges generated. It is shown in table 1 that on 9 occasions the challenge strategist provides a suitable ground following its knowledge structure.

There are however 2 occasions (DE6-048 and DE6-050), on which the challenge strategist gives a non-committal answer. Concerning the former, the strategic agent cannot find a support for the statement in its knowledge base and therefore speaks the truth with a non-committal answer. Regarding the latter, the strategic agent does have a support in its knowledge base for the statement being challenged; however, the support is not an acceptable ground since the partner of the strategic agent has challenged the support and the strategic agent had withdrawn this support from its commitment store during the earlier stage of dialogue. The strategic agent, then, would beg the question were it to answer the challenge with this unacceptable support [cf. 16]. It is therefore reasonable for the challenge strategist to give the non-committal answer rather than to commit the fallacy of question begging.

In sum, then, the challenge strategist seems to be working properly.

d) Question Strategist

In total, 37 questions are generated. They fall into four categories according to the nature of the move content: game start, statements supporting speaker's thesis, responder's thesis, statements supporting responder's thesis. These are discussed in turn below.

Concerning the game start, the responder choosing either view would be considered as reasonable since the other agent always took the opposing view to its partner.

The situation of a statement being questioned that supports the questioner's view might be interpreted as the questioner asking the responder to commit to that statement, according to [13]. The responder might be expected not to commit to it if it has an alternative. It is shown in table 1 that the responder responds 18 times with a "Yes", 3 times with a "No" and 3 times with a "no commitment". Each of the "Yes" responses is accounted for as the only answer found in its knowledge base, and the responder speaking the truth. For the "No" responses, the responder finds the negative answer (which mitigates against the opponent's view) from its knowledge base, and the answer can be further supported, the responder therefore gives the negative answer. For the "no commitment" responses, there are 3 instances (DE4-034, DE6-010 and DE6-022); on the first situation, the responder does have a negative answer (which mitigates against the opponent's view) in its knowledge base, but the responder finds no further support for the negative answer if being challenged, giving a positive response might harm its own positions, the responder therefore issues a no commitment response; on the latter two situations, the responder does not recognise the questioned statement (it cannot be found in its knowledge base), and therefore answers truthfully with a no commitment. The assertion strategist's responses to these questions are therefore reasonable given the above analysis.

It is interesting to see that 2 questions of statements supporting the responder's thesis were generated by the random agent in DE4. As expected, the strategic agent takes advantage of this and gives positive responses.

There are 9 instances where the questioner checks whether the responder still adheres to its thesis. On 8 occasions, the responder insists on its view since it still has acceptable grounds for its view. On 1 occasion (DE6-052), the responder accepts the opposite view. The dialogue fragment for this instance is reproduced below.

032: S>Is it the case that nobody is willing to die?

033: C>Yes, I think nobody is willing to die.

- 034: S>Is it the case that “nobody is willing to die” is a prima facie reason for “CP makes people less likely to commit serious crimes”?
- 035: C>Yes, I think “nobody is willing to die” is a prima facie reason for “CP makes people less likely to commit serious crimes”.
- 036: S>Is it the case that “CP makes people less likely commit serious crimes” is a prima facie reason for “CP is a good deterrent”?
- 037: C>Yes, I think “CP makes people less likely commit serious crimes” is a prima facie reason for “CP is a good deterrent”.
- 038: S>Is it the case that “CP is a good deterrent” is a prima facie reason for “CP is acceptable”?
- 039: C>Yes, I think “CP is a good deterrent” is a prima facie reason for “CP is acceptable”.
- 040: S>I think CP is acceptable.
.....
- 051: C>I don't know why innocent people may get killed.
- 052: S>Is it the case that CP is not acceptable?
- 053: C>No, I think CP is acceptable.

In the above dialogue fragment, the responder (agent C) has no acceptable ground for its thesis since its support has been withdrawn in turn 051. Further, agent C has explicitly committed to the set of propositions and conditionals which implies its dialogue partner's thesis. Agent C therefore makes a concession and accepts the opposite view in turn 053.

The 9 instances of questions involving thesis adherence checking can be seen therefore as being reasonably answered, given the above analysis. In sum, then, the question strategist appears to be working well.

Given the above discussion, it is argued that the various agent-generated dialogue situations can be seen as being reasonably handled by the agent system. Generally speaking, the assertion strategist, question strategist, withdrawal strategist and challenge strategist are working well and can therefore be regarded as satisfying the “robustness” criterion, though more sophisticated heuristics are needed to deal with unrecognised statements and the agent running out of moves.

Evaluation Criterion 2- Equal Opportunity

Of concern here is the issue of initiative. Initiative is relevant because if one dialogue participant is constantly starved of the initiative, he cannot fully or freely advocate his point of view (cf. [6: 229]; [10]).

In DE4, the strategic agent hands over its initiative 9 times to the random agent during the 54-turn dialogue. There are 7 instances of initiative shift during the 52-turn dialogue in DE5, and there are 4 instances of initiative shift during the 54-turn dialogue in DE6. The initiative exchanges can therefore be accounted as frequent. The longest duration of one agent retaining the initiative is from turn DE5-022 to turn DE5-034, in that the agent made two challenges, four questions and one statement to hand over the initiative.

On the whole, both agents have had opportunities to express their point of view, and the strategy therefore satisfies the “equal opportunity” criterion.

Evaluation Criterion 3- Coverage of Issues

Of interest here is whether points implemented in the knowledge base are revealed and discussed. The knowledge base [cf. 17] can be seen as containing three subtopics (deterrent issue, popularity issue and moral issue) which support the view “CP is acceptable” and two subtopics (ethical issue and consequence issue) which support the contrary view.

It can be seen that the issues in the knowledge base are raised during the evolving dialogue. Further, these issues are discussed to the maximum depth possible given the knowledge base in the DE5

and DE6 dialogues (DE4 is not considered here since one of the agents uses random argument). For example, in the following extract from DE5, the popularity issue is discussed from turn 006 to turn 008.

- 006: S>I think most people want CP back.
007: C>Why is it the case that most people want CP back?
008: S>Because the recent survey shows that 60% British people support CP.
009: C>I think it is wrong to take a human life.

However, it has turned out that there is no clear indication of the results of the discussion of each subtopic. On most occasions, the dialogue participants merely move to a new focus without comment (e.g. in the above dialogue fragment, agent C moves to the ethical issue in turn 009 without comment on whether it concedes the popularity point). This Mackenzie [4: 293] sees as a “formalisation of the convention that *silence means assent*”, arguing that “it is observed in many actual dialogues”. This might therefore be acceptable given that the strategic agent can be seen as following this convention.

In sum, we argue that the strategy is acceptable in respect of the “coverage of issues” criterion.

Evaluation Criterion 4- Argument Flow

The analysis here will consider whether the arguments generated by the agents as they follow the strategy are flowing well. If participants' contributions are clearly related to each other or appear logical, then, the flow will be deemed acceptable (cf. [6]).

In order to inspect the protocols for disjointedness, the process of analysis adopted includes a manual insertion of missing premises of incomplete arguments into the machine generated transcripts; this approach is in line with that adopted by [14]. An example of this is the following extract from DE5 (in which presumed missing premises are indicated in italics).

- 003: C>I think CP is not acceptable.
004: S>I think CP is a good deterrent, and “*CP is a good deterrent*” is a prima facie reason for “*CP is acceptable*”.

By doing so, it is easy to see the argument flow in this example - agent S is addressing agent C's utterance.

In DE4, DE5 and DE6, the strategic agent made 133 dialogue contributions. 93 of them involve directly addressing the previous utterance. There are 11 instances of the strategic agent's contributions which involve asking the dialogue partner to continue its line of argument. There are 7 instances of checking whether the dialogue partner still adheres to its thesis given its previous statement. However, there are 22 apparent instances of disjointedness, 10 of which involve building a line of questions toward the negation of the dialogue partner's previous statement, and 12 of which involve a shift of the current focus.

Generally speaking, though, the protocols can be seen as flowing well, - the total of 22 instances of disjointedness represents a relatively small proportion of the 133 strategic agents' dialogue contributions. Further, some of the apparent examples of disjointedness might in fact be reasonable in terms of the underlying strategy of the move maker, e.g. distance strategy (cf. [12]). Consider, for example, the following extract from DE5.

- 014: S>I think it is not always wrong to take a human life.
015: C>Is it the case that human lives are scarce?
016: S>Yes, I think human lives are scarce.
017: C>Is it the case that “human lives are scarce” is a prima facie reason for “it is wrong to take a human life”?
018: S>Yes, I think “human lives are scarce” is a prima facie reason for “it is wrong to take a human life”.
019: C>I think it is wrong to take a human life.
020: S>I think murderers should receive capital punishment.

In the above dialogue fragment, turn 015 appears at first sight to be disjointed from the previous utterance. However, turn 015 can be

seen as agent C starting its distance plan to build the case for “it is wrong to take a human life” (in turn 019) in order to rebut agent S’s view “it is not always wrong to take a human life”. On the other hand, however, turn 20 might be seen as agent S shifting the current focus from the ethical issue to moral issue without explicitly indication.

Overall, then, we argue that the strategy is acceptable in respect of the “argument flow” criterion, with the exception of absence of explicit linking for a focus shift.

Evaluation Criterion 5 Defeasibility

The interest here is whether the agent adopting the strategy is defeasible. It is shown in DE6, the strategic agent C does lose the dialogue and does so in a manner which might be considered reasonable, as opposed to a mere “surrender”. In the process of arriving at this defeat, agent C’s thesis support has been removed by agent S, and agent S provides the prima facie reasons for its thesis in turn DE6-004, DE6-012, DE6-016 and DE6-024, and agent C is explicitly committed to them. The evidence therefore suggests that the strategy is defeasible if the other party of the dialogue makes some good persuasion.

The strategy therefore satisfies the “defeasibility” criterion.

6. DISCUSSION

Before discussing the significance of the evaluations reported above, there are two possible difficulties with the methodology that ought to be discussed. First, it might be argued that only a small number of dialogue transcripts (three in total from three pairs of agents) have been generated for analysis. However, this study is intended not as a statistical enquiry, but rather as an investigation into the detail of the argument generated by the strategy. Further, 165 utterances are generated (DE4: 59; DE5: 52; DE6: 54). Each utterance needs to be considered in depth, and as a result this study does, it is held, provide sufficient data for the purpose of this assessment.

The second difficulty may be that there is a heavy reliance on judgements of quality by the author of the heuristics and the agent-based systems, and that the criteria of quality are themselves intuitively formulated. The judgement issue maybe endemic to the field, and similar criticisms could perhaps be levelled against much of the dialectics literature [6]. Further, computationally generating dialogues from dialectical theories may represent a step forward, and making the various criteria clear and explicit may well localise the issues to relatively narrow concerns at any one time, and this may detract from the judgement element. In addition, these criteria have enabled us to provide a thorough analysis of the data collected, and to leave the results, and the data itself, available for independent inspection.

We argue, then, that the methodology adopted is sound. We believe that the work reported makes a valuable contribution to the fields of dialectical systems and of human-computer dialogue. Concerning the former, we have proposed a set of strategies to be utilised with the dialogue model DE (see section 2 above). Further, because the agent-based system we have built can readily be adapted to function with a different dialogue model and/or a different set of strategies, it potentially provides people working in the field of dialectics with a test bed within which they can experiment with new models and new strategies they develop (cf. [1]; [5]).

The work contributes to human computer dialogue, we argue, in two ways. It indirectly contributes via the contribution to dialectics we have just outlined. Given the usefulness of a dialectical approach to interactive computer systems [5; 15], any development of dialectics per se potentially has a pay-off in terms of human-computer dialogue. Our work also makes a more direct contribution to human-computer dialogue, in that the debate system is a unique system and therefore makes a contribution to the attempt to broaden the human-computer interaction “bandwidth”.

7. CONCLUSIONS AND FURTHER WORK

We have proposed a set of strategic heuristics for a human computer debating system. Means of evaluation of the proposed strategy have been discussed and a set of computational agents constructed as a test-bed to facilitate the evaluation. A qualitative assessment of agent-generated debates has been outlined, and suggests that, generally speaking, the proposed strategy can provide good services enabling the computer to act as a dialogue participant.

Our immediate future work involves user-based studies of the evolving debating system.

8. REFERENCES

- [1] Amgoud, L. and Maudet, N. (2002). Strategic Considerations for Argumentative Agents. *NMR'2002* special session on Argument, Dialogue, and Decision. Toulouse, France.
- [2] Amgoud, L., Parsons, S. (2001). Agent Dialogue with Conflicting Preferences. In *Proceeding of Workshop ATAL01*.
- [3] Bench-Capon, T. J. M. (1998). Specification and Implementation of Toulmin Dialogue Game. *Proceedings of JURIX 98*, GNI, Nijmegen, pp. 5-20.
- [4] Mackenzie, J. (1987). I Guess. *Australia Journal of Philosophy* Vol. 65, No. 3 pp. 290-300.
- [5] Maudet, N. and Moore, D. (2001). Dialogue games as Dialogue Models for Interacting with, and via, Computers. *Informal Logic* Vol. 21, No. 3 pp.219-243.
- [6] Moore, D. (1993). *Dialogue Game Theory for Intelligent Tutoring Systems*. Unpublished Doctoral Dissertation, Leeds Metropolitan University.
- [7] Rahwan, I., McBurney, P. and Sonenberg, E. [2003]: Towards a theory of negotiation strategy (a preliminary report). In: S. Parsons and P. Gmytrasiewicz (Editors): *Game-Theoretic and Decision-Theoretic Agents (GTDT 2003)*, Proceedings of an AAMAS-2003 Workshop, pp. 73-80. Melbourne, Australia.
- [8] Toulmin, S. (1958). *The Uses of Argument*. Cambridge, UK: Cambridge University Press.
- [9] Walton, D. (1984). *Logical Dialogue Games and Fallacies*. University Press of America.
- [10] Walton, D. (1989). *Question-Reply Argumentation*. Greenwood Press.
- [11] Walton, D. (1996). *Argumentation Schemes for Presumptive Reasoning*. Mahwah, N.J., Erlbaum.
- [12] Walton, D. (1998). *The New Dialectics: Conversational Context of Argument*. University of Toronto Press.
- [13] Walton, D. and Krabbe, E. (1995). *Commitment in Dialogue: Basic Concept of Interpersonal Reasoning*. Albany NY: State University of New York Press.
- [14] Walton, D. and Reed, C. (2002). Argumentation Schemes and Defeasible Inference. In *Proceedings of ECAI'2002 Workshop on Computational Models of Natural Argument*, Lyon, France.
- [15] Yuan, T., Moore, D. and Grierson, A. (2002). Educational Human-Computer Debate, a Computational Dialectics Approach. In *Proceedings of ECAI'2002 Workshop on Computational Models of Natural Argument*. pp. 19-22 Lyon, France.
- [16] Yuan, T., Moore, D. and Grierson, A. (2003). Computational Agents as a Test-Bed to Study Philosophical Model “DE”, A Development of Mackenzie’s “DC”. In Press for *Informal Logic*.
- [17] Yuan, T. (2004). *Human-Computer Debate, a Computational Dialectics Approach*. Unpublished Doctoral Dissertation (in preparation), Leeds Metropolitan University.