Chance Discovery Using Dialectical Argumentation

Peter McBurney and Simon Parsons
Agent Applications, Research and Technologies Group
Department of Computer Science
University of Liverpool
Liverpool L69 7ZF United Kingdom
{p.j.mcburney,s.d.parsons}@csc.liv.ac.uk

Abstract

The authors propose the use of a dialectical argumentation formalism for chance discovery in domains where knowledge is distributed across a number of distinct knowledge-bases, as in a system of autonomous software agents. Each agent may have only a partial view of a problem, and may have insufficient knowledge to prove particular hypotheses; our formalism provides a means to aggregate across these partial views in a consistent manner. We identify a novel type of dialogue, which we call a discovery dialogue, and propose a formal model for its conduct. Moreover, we present locutions and rules for the implementation of these dialogues as dialogue-games. In exploring the question of whether this dialogue model may be automated, we consider a genetic algorithm to generate and test hypotheses in the space common to all the agents.

1 Introduction

In 1994-5, the British Government privatized the state-owned national railway monopoly, British Rail. They did this by creating one private company, Railtrack Ltd, to own and operate the physical network of railway track, and then created 25 separate licences, each awarded by competitive tender, to operate train services along these tracks in specific geographic regions. Thus, for example, train services between London and Scotland on the main west-coast line are now provided by Virgin Railways, while those on the main east-cost line are provided by Great North Eastern Railways. In addition, other companies were created to provide specific services to Railtrack and to the train operating companies, for instance, railway communications. The new private companies also outsourced functions which had previously been undertaken within British Rail, such as carriage ownership and track inspection and maintenance. By one estimate there are now more than 100 companies where previously there was just one.

A rail network is a very complex system. What was once a single and unified system, is now fragmented, with disparate responsibilities, distributed knowledge and possibly conflicting interests. No one company in the network now has all the information needed to manage it. The results of this were seen on 17 October 2000, when faulty track caused a derailment at Hatfield, killing four people and injuring 70. Although an inquiry still has to establish ultimate responsibility for the accident, knowledge of the faulty track appears to have been known to the company tasked with network inspection, a sub-contractor to Railtrack [14]. This accident showed the difficulty of managing a rail network as a single entity when responsibilities are divided between many participants, each of whom may have divergent knowledge and interests. Risks and opportunities may not be identified, or may not be acted upon, because the information and intelligence to do so is distributed between multiple agents, and may not be fully shared between them because their economic interests are perceived to diverge. Current trends across the business world encouraging outsourcing, business partnerships and the creation of virtual supply networks mean these problems will be increasingly common.[1]

The problem of interest here is how to identify risks and opportunities (*"Chance Discovery"*) in situations where knowledge is distributed across multiple autonomous agents. We believe that systems of dialectical argumentation, which enable the coherent combination of disparate knowledge types and sources, are applicable to this problem and in this paper we present a formalism for such a system. For simplicity, we assume that the agents involved do not perceive they have divergent economic or other interests, and so are willing to share information fully with each other. We present our formalism in Section 3, after first reviewing the application of argumentation and dialogue games in artificial intelligence in Section 2. In earlier work, we presented a similar dialectical argumentation structure for dialogues over risk in environmental domains, which we which termed a *Risk Agora* [22]. Accordingly, we call the formalism presented in this paper a *Discovery Agora*. Section 4 briefly discusses the use of an evolutionary computational architecture to support automated dialogues in this Agora, and Section 5 presents an example. Section 6 concludes the paper with a discussion of formal properties and future research.

[1] For instance, telecommunications companies must collaborate with their own competitors in order to provide end-to-end services to customers seamlessly. The same pattern is emerging in other industries where competition has been recently introduced, such as the provision of energy and water, and in air transport.

2 Argumentation

In common English usage, an *"argument"* has two meanings: a case for (or against) a particular claim, and a debate between two or more people. Arguments in both senses have been studied by philosophers since at least the time of Aristotle, in a branch of philosophy now called argumentation theory. In this paper, we will use the word *argument* for only the first meaning, and the words *debate* or *dialogue* for the second. Considering arguments as a case for claims, argumentation theory has examined issues such as what constitutes a good or bad argument, under what circumstances is it rational to use non-deductive arguments, and what relationships may exist between distinct arguments. For dialogues and debates, philosophers have explored issues such as how may such debates be organized and structured, what rules are appropriate for different types of interaction, and what impacts arise from variations in these rules. In both areas, philosophers of argumentation have been particularly active since the mid-1950s, perhaps in response to the development of formal, deductive logic in the century before that. A comprehensive review of argumentation theory can be found in [36].

An argument may be compelling without being logically valid, and indeed one may view an argument as a tentative proof. The "proof" is not final because we do not have all the information needed, or because the information we have is uncertain, or because the rules of inference used in the argument are not necessarily truth-preserving. Within Artificial Intelligence, argumentation theory has been applied over the last decade in several ways. The first of these is the modeling of uncertainty, where we can determine our strength of belief in a claim on the basis of the relative strength of the arguments for and against it. Thus, argumentation can provide a qualitative formalism for uncertainty representation, which is an alternative to quantitative measures such as probability theory [19; 25]. Applications of this approach have included systems supporting medical treatment decisions [9; 7] and undertaking risk prediction for new chemicals [6; 34]. Secondly, argumentation systems have been used as proof-systems for non-monotonic reasoning. Because arguments are not definitive, but tentative, they may be over-turned or refuted in the light of new information. Thus, a set of arguments and rules showing their relationships to one another, called an *argumentation system*, can provide a mechanism for deciding whether to accept or reject the conclusions of default rules. These methods are reviewed in [30]. Thirdly, argumentation has been used to model human reasoning when this is non-deductive, for example in legal domains [3]. Much human reasoning uses inference rules which are not deductively-valid (such as reasoning by analogy) and argumentation theorists have found success in developing and applying models for this type of reasoning, so as to better understand it. Elsewhere, we have referred to this use of argumentation as fulfilling an orrery function, on the analogy of mechanical models of the solar system [32]. Fourthly, argumentation has been used to develop systems for debates or dialogues where divergent viewpoints may be represented. This field has become known as Computational Dialectics [31], and systems have been proposed for debates over legal issues [11], urban planning decisions [12], and scientific questions [22]. Moreover, similar dialectical approaches have been proposed for automated dialogues between intelligent software agents [28; 1] and for the automated design of software components [33]. A review of some applications of argumentation in Artificial Intelligence is contained in [5]. In this paper, we will primarily draw upon Computational Dialectics for the modeling of dialogue, and we discuss this next.

In an influential typology, Doug Walton and Erik Krabbe [37] identified six primary types of dialogue, distinguished by their initial situations, the goals of each of their participants, and the goals of the dialogue itself (which may differ from those of its participants). The six dialogue types were: *Information-seeking dialogues*, in which participant seeks the answer to some question from another participant; *Inquiries*, in which all participants collaborate to answer some question to which none has the answer; *Persuasion dialogues*, in which one participant seeks to convince others of the truth of some proposition; *Negotiations*, in which participants attempt to divide up a scarce resource; *Deliberations*, in which participants collaborate to decide what actions to take in some situation; and *Eristic* (strife-ridden) dialogues, in which participants quarrel verbally as a substitute to physical fighting. Most human dialogues may be seen as examples of these six or combinations thereof, although Walton and Krabbe do not claim their typology is comprehensive.

How may specific types of dialogues be modeled? To do this, we draw on the formal dialogue games of philosophy, which were first proposed to better understand fallacious modes of reasoning [13; 21]. These are games between two or more players, where the "moves" made by the players are locutions, i.e. spoken utterances. Recently, such games have been applied in Artificial Intelligence [1; 4], particularly for automated dialogues between autonomous agents, and we have been led to propose a formal model for dialogue-games [24]. In our model, it is assumed that the topics of discussion between the participants are represented in some logical language, whose well-formed formulae are denoted by the lower-case Roman letters, p, q, r, etc. The rules of the dialogue-game can be divided into several distinct types:

Commencement Rules: Rules which define the circumstances under which the dialogue commences.

Locutions: Rules which indicate what utterances are permitted. Typically, legal locutions permit participants to assert propositions, permit others to question or contest prior assertions, and permit those asserting propositions which are subsequently questioned or contested to justify their assertions. Justifications may involve the presentation of a proof of the proposition or an argument for it, and such presentations may also be legal utterances. In multi-agent system applications of dialogue games (e.g. [1; 2]), it is common to impose rationality conditions on utterances, for example allowing agents to assert statements only when they themselves have a prior argument or proof from their own knowledge base.[2] The

[2]Such conditions are similar conceptually to the *feasibility pre-*

dialogue game rules may also permit participants to utter propositions to which they assign differing degrees of commitment, for example: one may merely *propose* a proposition, a speech act which entails less commitment than would an *assertion* of the same proposition.

Combination Rules: Rules which define the dialogical contexts under which particular locutions are permitted or not, or obligatory or not. For instance, it may not be permitted for a participant to assert a proposition p and subsequently the proposition $\neg p$ in the same dialogue, without in the interim having retracted the former assertion. Similarly, assertion of a proposition by a participant may oblige that participant to defend it in defined ways following contestation by other participants.

Commitments: Rules which define the circumstances under which participants express commitment to a proposition. Typically, assertion of a claim p in the debate is defined as indicating to the other participants some level of commitment to, or support for, the claim. In a negotiation dialogue, for example, assertion of an offer may express a willingness to undertake a transaction on the terms contained in the offer. However, depending on the rules of the game, commitment may express merely that the speaker has an argument for p, and this is not necessarily the same as belief in p, nor does it necessarily imply any intention to act. Since Hamblin [13], it is common to track commitments in publicly-accessible stores called Commitment Stores.

Termination Rules: Rules which define the circumstances under which the dialogue ends. These rules may be expressible in terms of the contents of the Commitment Stores of one or more participants.

Instantiating these rules for different types of dialogue has been a recent research focus. For example, Henry Prakken [29] and Walton and Krabbe [37] proposed formal models for persuasion dialogues, Joris Hulstijn [18] models for information-seeking and negotiation dialogues and, in joint work with David Hitchcock [17], we have proposed the first formal model for deliberation dialogues.

3 The Discovery Agora: Formal Structure

With these considerations in mind, we now present our formal structure for the argumentation system for chance discovery, which we have called a *Discovery Agora*. We assume, as above, that the topics of discussion between the participants are represented in some logical language, $\mathcal{L}$, closed under the usual connectives, whose well-formed formulae are denoted by the lower-case Roman letters, p, q, r, etc. Although the participants may believe different sets of axioms (premises) to be true, we assume they have agreed a set of deductive inference rules. We refer to this logical language and the agreed inference rules as the common logic of the Agora. We further assume a denumerable set of autonomous software agents who participate in the dialogue, each of whom is denoted by $\mathcal{P}_i$, indexed by $i \in \mathcal{I}$. We assume in this paper that each agent uses the same logical language and rules of inference, and that they differ only in the information they know to be true (i.e. in their premises). Thus, each participant may know part of the story but not the whole story. We further assume that the agents have no inhibitions about sharing information with each other in the Agora.

In this section we present the formal model for the Discovery Agora. We do this, firstly, in Section 3.1, with an informal discussion of discovery dialogues to motivate our formalism, then, in Section 3.2, a formal model of a Discovery Dialogue in the Agora. Section 3.3 presents the rules for a dialogue game in conformance with the formal model.

3.1 Discovery dialogues

We assume the agents in the Discovery Agora are engaged in dialogue. Which of the Walton and Krabbe [37] dialogue types mentioned above is appropriate to this domain? The closest type would appear to be an Inquiry dialogue, where participants collaborate to ascertain the truth of some question. However, for the domain of Chance Discovery we want to *discover* something not previously known; the question whose truth is to be ascertained may only emerge in the course of the dialogue. This feature is similar to Deliberation dialogues, where the course of action adopted by the participants may also only emerge in the course of the dialogue itself. The other five types of dialogue all begin with some question or issue for discussion. We therefore believe that the dialogue type appropriate to the Chance Discovery domain is not one of the types of Walton and Krabbe. We will call it a *Discovery Dialogue*.

These dialogues differ from Inquiries in another way. In a pure inquiry, the participants would wish to seek the truth, unadulterated by their preferences or emotional responses. This is unlike a Deliberation, where the preferences or emotions of the participants could play an important part in the selection of an optimal course of action. While the participants in a Discovery dialogue are also seeking truth, there may be many possible truths. It would be sensible for the participants to filter the truths they discover by what is interesting, novel or important. Discovering risks, for instance, means identifying potential outcomes with significant and deleterious consequences.

How might a dialogue concerning chance discovery proceed? We could imagine a number of elements to such a dialogue. Firstly, there would be agreement (perhaps implicit) about the purpose of the dialogue; this could be, for instance, to assess the risks inherent in some situation or technology. Next there may be the sharing of relevant information known by each of the participants and the pooling of this knowledge to generate new knowledge. For dialogues seeking to discover consequences or risks, there may also be discussion concerning the possible mechanisms by which such risks or consequences could occur. These mechanisms may be chains of possible scientific causality (as in cellular-level biomedical mechanisms) or metaphorical or analogical modes of reasoning. Legal reasoning concerning the potential motives, opportunity and means of a suspect to commit a crime is another example of such mechanisms. Then, once potential discoveries

condition in the Agent Communications Language of [8, p.48], which specifies conditions under which an agent can be considered *sincere* when transmitting a message.

have been articulated in the dialogue, there may be discussion over their attributes. For instance: Are they equally important? What are their relative consequential losses or benefits? etc. This discussion over attributes may in turn lead to consideration of experiments or data collection activities to verify which of competing hypotheses is more likely correct in explaining causal effects. In human dialogues, of course, such discussions do not occur in a linear fashion, but move back and forth between these various elements as the discussion evolves. Our formal model, to be presented next, will include each of these elements and allow for non-linear dialogues.

3.2 Model of a discovery dialogue

In this section we formalize the discussion just presented. We begin by defining the elements of the discovery dialogue, drawing on the model of an argument proposed by Stephen Toulmin [35].

Purpose: The Purpose of a dialogue is the overall issue or issues which motivated the participants to convene and which governs their dialogue. Examples include the risks or the opportunities of some situation. We assume that a discovery dialogue is initiated by one of the participants with a proposed purpose. However, the other participants may not share the same understanding of the dialogue's purpose, and so this needs to be discussed at the outset.

Data Item: A Data Item is a proposition for which at least one dialogue participant has a proof, using premises in that participant's knowledge base and using the rules of inference of the common logic. Participants who articulate data items will be required to present the arguments for them, if requested in the Agora.

Inference Mechanism: An inference mechanism is a warrant which justifies the drawing of a conclusion from one or more data items. Examples of mechanisms include: the rules of deductive inference of the common logic of the participants; default rules; causal mechanisms; metaphors and analogies; legal precedents, etc.

Consequence: A consequence is a claim arising from the application of an inference mechanism to one or more data items.

Criterion: A criterion is an attribute of a data item or a consequence, which may be used to compare one data item or consequence with another. Examples of criteria include: novelty; importance; costs; benefits; feasibility; etc.

Test: A test is a procedure, generally undertaken outside the Discovery dialogue, to ascertain the truth-value of some unknown variable. Examples include: scientific experiments; data collection exercises; information-seeking dialogues.

Conclusion: A conclusion is a full or partial response to the purpose of the dialogue. For example, conclusion could include significant risks identified in the course of the dialogue or interesting opportunities.

With these elements defined, we now present a formal model of the dialogue itself, which moves through ten stages. Our model is similar in approach to the formal model for deliberation dialogues we developed with David Hitchcock in [17].

Open Dialogue: Opening of the discovery dialogue.

Discuss Purpose: Discussion of the purpose of the dialogue.

Share Knowledge: Presentation of data items relevant to the purpose, drawing only on each participant's individual knowledge base.

Discuss Mechanisms: Discussion of potential rules of inference, causal mechanisms, metaphorical modes of reasoning, legal theories, etc.

Infer Consequences: Identification of the consequences arising from the application of inference mechanisms to the data items presented by the participants.

Discuss Criteria: Discussion of possible criteria for assessment of the consequences presented.

Assess Consequences: Discussion of the data items and consequences against the criteria previously suggested.

Discuss Tests: Discussion of need for undertaking tests of proposed consequences. If such tests are conducted outside the dialogue, the results may be reported back to the dialogue in a **Share Knowledge** stage.

Propose Conclusions: Proposing one or more conclusions for possible acceptance by the participants.

Close Dialogue: Closing of the discovery dialogue.

Agreement is not necessary in these dialogues unless the participants so desire it. If so, the **Propose Conclusions** stage allows a participant to propose a conclusion for possible acceptance, and then allows participants to indicate to the Agora their individual acceptance or otherwise. As mentioned earlier, these stages of a discovery dialogue may be undertaken in any order, subject only to the following constraints:

- The first stage in every Discovery dialogue is **Open Dialogue**.
- The **Open Dialogue** stage may occur only once in any Discovery dialogue. All other stages may occur more than once.
- The only stages which must occur in every Discovery dialogue which terminates normally are **Open Dialogue** and **Close Dialogue**.
- The stage **Discuss Purpose** must precede any other stage, excepting **Open Dialogue** and **Close Dialogue**.
- At least one instance of each of the stages **Share Knowledge** and **Discuss Mechanisms** must precede the first instance of the stage **Infer Consequences**.
- At least one instance of each of the stages **Infer Consequences** and **Discuss Criteria** must precede the first instance of the stage **Assess Consequences**.
- At least one instance of the stage **Infer Consequences** must precede the first instance of the stage **Discuss Tests**.

- At least one instance of the stage **Assess Consequences** must precede the first instance of the stage **Propose Conclusions**.
- The last stage in every Discovery dialogue which terminates normally is **Close Dialogue**.
- Subject only to the constraints expressed in these rules and constraints expressed in the locution-combination rules (articulated below), participants may enter any stage from within any other stage at any time.

Note that the participants may enter the **Close Dialogue** stage more than once in a particular dialogue. This stage (as the locution-combination rules below will indicate) requires participants to indicate that they wish to leave the dialogue. Thus, this stage remains unconcluded, and the dialogue remains open, whilesoever there are at least two participants who wish to continue speaking. It is therefore possible for this stage to be entered multiple times in any one dialogue.

3.3 Dialogue game rules

We now present examples of dialogue-game locutions which, taken together, enable a discovery dialogue to be conducted according to the model just presented. For reasons of space, we do not present all the locutions, nor all the necessary preconditions for, and the consequences of, their utterance.[3] We continue to assume that the subject-matter of dialogues can be represented in a propositional language by lower-case Roman letters. We define *questions* as propositions with one or more free variables, and we represent these by lower-case Roman letters suffixed with a question-mark, e.g *"p?"*. We assume a Commitment Store $CS(P_i)$ exists for each agent P_i. This store contains the various propositions which the agent has publicly accepted, and each store can be viewed by all participants. Entries in the stores are of three sorts: (a) 2-tuples of the form $(type, t)$, where t is a valid instance of type *type*, with *type* $\in$ {*purpose, data item, inference mechanism, consequence, criterion, test, conclusion*}; (b) 3-tuples of the form (c, t, p), where c is a consequence, t is a criterion and p a proposition; and (c) 3-tuples of the form (c_1, c_2, t), where c_1 and c_2 are consequences and t a criterion. The permitted locutions are:

open_dialogue(P_i, p)**:** Participant P_i proposes the opening of a Discovery dialogue to consider the proposed purpose p. A dialogue can only commence with this move.

enter_dialogue(P_j, p)**:** Participant P_j indicates a willingness to join a Discovery dialogue to consider the purpose p. All intending participants other than the mover of **open_dialogue(.)** must announce their participation with this move. Note that neither the **open_dialogue(.)** nor the **enter_dialogue(.)** move implies that the speaker accepts that p is the most appropriate formulation of the purpose, only that he or she is willing to enter into a discussion about it at this time.

propose$(P_i, type, t)$**:** Participant P_i proposes proposition t as a valid instance of type *type*, where *type* $\in$ {*purpose, data item, inference mechanism, consequence, criterion, test, conclusion*}.

[3] These are presented in [26].

assert$(P_i, type, t)$**:** Participant P_i asserts proposition t as a valid instance of type *type*, where *type* $\in$ {*purpose, data item, inference mechanism, consequence, criterion, test, conclusion*}. This is a stronger locution than **propose(.)**, and results in the tuple $(type, t)$ being inserted into $CS(P_i)$, the Commitment Store of P_i. For certain types, utterance of this locution leads to the speaker having a *burden of defence*, i.e. to provide supporting arguments or evidence for the assertion if so requested by another participant.

query$(P_j$,**propose**$(P_i, type, t))$**:** Participant P_j requests participant P_i to provide a justification for his proposal of t as a valid instance of type *type*, where *type* $\in$ {*data item, consequence, test*}, and where $j \neq i$. Similarly, a participant may query an assertion with the command **query**$(P_j$,**assert**$(P_i, type, t))$. In response to either query, P_i must defend his proposal or assertion statement with an utterance of **show_arg(.)**.

show_arg$(P_i, type, t, A)$**:** Participant P_i presents an argument A for proposition p which is type *type* $\in$ {*data item, consequence, test*}. In the case of *data items*, the argument A is a proof from premises in the knowledge base of participant P_i and using deductive inference rules in the common logic of the dialogue. In the case of *consequences*, the argument A comprises one or more sequences of the form (D, I, C), where D is a set of *data items*, I is an inference mechanism and C is a consequence which can be drawn from D using I; all elements of this set must previously been articulated in the Agora by means of appropriate **propose(.)** or **assert(.)** locutions. In the case of *tests*, the argument A also comprises one or more sequences of the form (D, I, C), where D is a set of *data items*, I is an inference mechanism and C is a consequence which can be drawn from D using I, but these need not have been previously presented in the dialogue.

assess(P_i, c, t, p)**:** Participant P_i asserts that when consequence c is assessed on the basis of criterion t, one may conclude proposition p. This locution inserts (c, t, p) into $CS(P_i)$.

compare(P_i, c_1, c_2, t)**:** Participant P_i asserts that consequence c_1 is better or equal to consequence c_2 when they are compared on the basis of criterion t. Each of c_1, c_2 and t must previously been articulated in the Agora by means of the appropriate **propose(.)** or **assert(.)** locutions. This locution inserts (c_1, c_2, t) into $CS(P_i)$.

recommend$(P_i, conclusion, a)$**:** Participant P_i proposes proposition a as a recommended conclusion. This locution inserts $(conclusion, a)$ into $CS(P_i)$.

accept$(P_j, locution)$**:** Participant P_j indicates agreement with the prior locution, *locution*, uttered by another participant. If the prior locution resulted in a change to that speaker's commitment store then the **accept(.)** locution similarly alters $CS(P_j)$.

contest$(P_j, locution)$**:** Participant P_j indicates disagreement with the prior locution, *locution*, uttered by another participant. The **contest(.)** locution is the obverse of **accept(.)** and has no impact on $CS(P_j)$.

retract(P_i, *locution*)**:** Participant P_i indicates retraction of her prior utterance of the locution *locution*. If the prior locution resulted in an insertion into $CS(P_i)$, then the **retract(.)** locution deletes it.

withdraw_dialogue(P_i, p)**:** Participant P_i announces her withdrawal from the Discovery dialogue to consider the governing question p.

For illustration, we give an example of the full articulation of the locution **assert(.)**, showing the required preconditions and post-conditions for this utterance.

Preconditions: **propose**($P_{j_1}, data\ item, D_1$), **propose**($P_{j_2}, data\ item, D_2$), ..., **propose**($P_{j_k}, data\ item, D_k$), **propose**($P_{m_1}, inference\ mechanism, I_1$), **propose**($P_{m_2}, inference\ mechanism, I_2$), ..., **propose**($P_{m_n}, inference\ mechanism, I_n$), although speakers need not be distinct.

Locution: **assert**($P_i, consequence, t$)

Meaning: Participant P_i asserts proposition t as a valid *consequence*.

Response: Another participant P_j ($j \neq i$) may respond with: **query**(P_j,**assert**($P_i, consequence, t$)).
If so, participant P_i must respond with: **show_arg**($P_i, consequence, t, A$),
where $A = (D, I_1, C_1), (C_1, I_2, C_2), \ldots, (C_{n-1}, I_n, C_n)$, and $D = \{D_1, D_2, \ldots, D_k\}$ for some $k \geq 1$.

Commitment Store Update: The 2-tuple $(consequence, t)$ is inserted into $CS(P_i)$.

We now demonstrate that the dialogue game locutions we have defined can be used to undertake a Discovery dialogue in accordance with the formal model we have proposed.

Proposition 1: *Each of the ten stages of the formal model of discovery dialogues presented in section 3.2 can be executed by judicious choice of these dialogue-game locutions.*

Proof. We consider each stage in turn:

1. A dialogue opens with the locution **open_dialogue**(P_i, p) and at least one utterance of **enter_dialogue**(P_j, p), with $j \neq i$.
2. The Discuss Purpose stage consists of utterances of **propose(.)**, **assert(.)**, **accept(.)**, **contest(.)** and **retract(.)**, in each case with the type *purpose*.
3. The Share Knowledge stage consists of utterances of **propose(.)**, **assert(.)**, **accept(.)**, **query(.)**, **show_arg(.)**, **contest(.)** and **retract(.)**, in each case with the type *data item*.
4. The Discuss Mechanism stage consists of utterances of **propose(.)**, **assert(.)**, **accept(.)**, **contest(.)** and **retract(.)**, in each case with the type *inference mechanism*.
5. The Infer Consequences stage consists of utterances of **propose(.)**, **assert(.)**, **accept(.)**, **query(.)**, **show_arg(.)**, **contest(.)** and **retract(.)**, in each case with the type *consequence*.
6. The Discuss Criteria stage consists of utterances of **propose(.)**, **assert(.)**, **accept(.)**, **contest(.)** and **retract(.)**, in each case with the type *criterion*.
7. The Assess Consequences stage consists of utterances of **assess(.)**, **compare(.)**, **agree(.)**, **contest(.)** and **retract(.)**.
8. The Discuss Tests stage consists of utterances of **propose(.)**, **assert(.)**, **accept(.)**, **contest(.)** and **retract(.)**, in each case with the type *test*.
9. The Propose Conclusions stage consists of an execution of **recommend**($P_i, action, a$), possibly followed by utterances of **accept**($P_j, action, a$).
10. Participants may exit a dialogue at any time, by means of the **withdraw_dialogue**(P_i, p) locution. The dialogue itself closes once the second-last participant utters this locution. □

In addition to defining the permitted locutions, we have specified commencement, combination, commitment and termination rules for this game (not presented here for reasons of space). These rules have been specified to accord with principles of dialogue between consenting and reasonable participants proposed by Hitchcock in [16]. Because one of these principles is that the participants should be free to decide the rules of conduct of the dialogue, we have not specified a mechanism for resolution of disagreements. The participants are free to decide this; examples of resolution methods may be voting systems (with majority, plurality or weighted voting schemes, for example) or qualitative argumentation systems, where propositions are classified according to the strength of support they receive in the Agora, as in [20; 25].

4 Generating arguments

We have presented a formal language in which discovery dialogues between autonomous agents may be undertaken. How may we automate such dialogues? If a model of a dialogue can be automated, we say that is has a *generative capability*. Our model for a discovery dialogue permits the participating agents to proceed iteratively, considering partial responses to the governing purposes and refining these by means of assessment and discussion. This iterative view of the dialogue leads us to believe that an evolutionary computational approach may be suitable as a means to automate the dialogue. Evolutionary computational approaches aim to produce a solution to some problem by generating a succession of partial- or near-solutions, which converge to a full solution. For example, a genetic algorithm could be used to generate candidate discoveries which can then be filtered by some agreed means (analogously to survival of the fittest in evolutionary theory) and used to generate new candidates (analogously to reproduction). Drawing on [27], a high-level outline of a genetic algorithm is as follows:

1. We assume we have a means to represent candidate solutions as strings of elements, called *populations of chromosomes*.
2. Computation proceeds in a sequence of iterations, corresponding to successive *generations* of the population.
3. For each generation, we have a means to select survivors from that generation, in a process of selection called *survival of the fittest*. These survivors produce *offspring* who comprise the next generation.
4. We then have a means to produce offspring by combining elements from two or more members of the parent generation (*crossover*).
5. We also have a means to produce offspring by *random mutation* of parent elements.

The Discovery Dialogue formalism presented in Section 3 assumes candidate discoveries are described only in a propositional language. To apply genetic algorithms to this domain, we would need to represent candidates as strings of elements, such that crossover and random mutations made sense. Similarly, the fitness functions, which correspond to the criteria of assessment proposed in the dialogue, would need to be represented as functions on the string elements. One approach to this would be to assume some fixed set of atomic propositions, say n in total, and then have population members be strings of binary sequences, each string corresponding to a well-formed formula involving propositions. The elements of each string would indicate the truth-status of the respective formula in the various logical models for the propositions. Thus, there would be a countably infinite number of strings, with each string consisting of 2^n elements (the number of distinct models of n propositions). A "1" in the i-th position in a particular string would indicate that the formula represented by that string is true in the i-th model. The fitness function, which selects some strings from each generation to survive to the next, could be represented as constraints on the space of models; e.g. only those strings survive that have, say, a "1" in the j-th position.

An evolutionary computational approach, as we have outlined here, could commence with a randomly-selected and finite set of strings, which would then be evolved through multiple generations. In each generation, mutation would introduce additional random strings. In this way, an automated Agora discussion could generate and consider random formulae, thus leading to possible chance discovery. We expect that representations such as these may be more successful in some problem domains than in others. Moreover, these approaches may be most effective in so-called *mixed-initiative systems* [10], i.e. systems which combine both automated processes and human participation. Our work in this area, however, is still too preliminary to report at this time.

5 Example

In this section, we present a simplified example of a Discovery Dialogue, involving prediction of the risk of flooding along the Clarence River region of Northern New South Wales, Australia. The Clarence River is 400 kilometres (km) long and has a water catchment area of 22,700 square km, the largest of any river in south-eastern Australia [15]. Because the catchment area is so wide, and extends over several climate zones, flood rains along one tributary are unlikely to occur simultaneously with rains along another. In addition, the river is tidal 108 km upstream from its mouth at Yamba. In our example, we imagine each participant is an autonomous agent on a different tributary river or at the mouth, with access to rain and river-level information at its geographic site only. The participants are denoted P_G, representing the Guy Fawkes River in the south-west of the catchment, P_T, the Timbara River in the north-west, and P_Y, representing Yamba. We present the dialogue as a sequence of moves, numbered M1, M2,. . ., along with some annotation.

M1: open_dialogue($P_Y, \mathit{Risk\ of\ flooding}$**)**

This move is the first move in the Open stage of the dialogue.

M2: enter_dialogue($P_G, \mathit{Risk\ of\ flooding}$**)**

M3: enter_dialogue($P_T, \mathit{Risk\ of\ flooding}$**)**

With the entry of a second participant, the dialogue may be said to commence. A third participant also enters.

M4: propose($P_G, \mathit{data\ item}, \mathit{Raining\ along\ Guy\ Fawkes}$**)**

M5: propose($P_T, \mathit{data\ item}, \mathit{Raining\ along\ Timbara}$**)**

Two participants report facts known to them locally, thereby commencing a Share Knowledge stage of the dialogue.

M6: propose($P_G, \mathit{inference\ mechanism},$
$\mathit{If\ raining\ along\ two\ tributaries\ then\ risk\ of\ flood\ along\ Clarence}$**)**

M7: assert($P_G, \mathit{consequence}, \mathit{Risk\ of\ flood\ along\ Clarence}$**)**

Participant P_G proposes an inference mechanism, commencing a Discuss Mechanisms stage, and infers a consequence, commencing an Infer Consequences stage.

M8: propose($P_Y, \mathit{test}, \mathit{High\ tide\ expected?}$**)**

Participant P_Y asks if a high tide is expected, thus commencing a Discuss Tests stage.

M9: query(P_T**, propose(**$P_Y, \mathit{test}, \mathit{High\ tide\ expected?}$**))**

Participant P_T asks why this question should be asked.

M10: show_arg($P_Y, \mathit{test}, \mathit{High\ tide\ expected?},$
$\mathit{If\ High\ tide\ and\ flood\ along\ Clarence}$
$\mathit{then\ risk\ of\ major\ flood}$**)**[4]

Participant P_Y explains the reason for testing the tidal level.

⋮

This is not a realistic example of our formalism, as flood prediction can be undertaken more effectively with remote sensing and quantitative estimation techniques. However, the example does illustrate the use of our dialogue formalism in the risk discovery domain. A more realistic application will require considerable effort in coding the domain knowledge.

[4] Note that the syntax of the **show_arg** locution is simplified here.

6 Discussion

This paper has proposed a formal argumentation system for chance discovery in domains where knowledge is distributed across autonomous agents. Our approach has led us to propose a new type of dialogue, which we call a discovery dialogue, for arguments in this domain. We have proposed a formal model for the conduct of discovery dialogues, and presented the locutions and rules for a dialogue game undertaken in accordance with this model. We have also briefly examined the use of evolutionary computational approaches to enable fully automated dialogues to be conducted. These computational approaches will require more investigation before they can be implemented to generate automated dialogues.

In addition to their generative capabilities, there are a number of formal properties one could consider for the system we have presented. Firstly, there are the circumstances under which dialogues terminate. Are they guaranteed to terminate? How robust are terminating dialogues to changes in the rules of the dialogue? Secondly, how quickly do terminating dialogues reach termination? In other words, what is the computational complexity of terminating dialogues? We may seek to determine minimum and maximum bounds on dialogue lengths, along with the expected length. A third set of formal properties relate to comparison of two dialogues in the Discovery Agora. When are two dialogues the same? How similar do two different dialogues have to be in order to be considered equivalent, in some sense. All these questions are the focus of our on-going work.

We are also exploring extensions to the dialogue framework we have presented here. Firstly, we are considering relaxation of the assumption that participating agents are willing to share information. In many real-life applications, as our railway example in the Introduction illustrates, this will not be the case. We expect that interaction between agents in such circumstances could be modeled with persuasion dialogues. Secondly, participants may have different opinions on the inference mechanisms presented in the **Discuss Mechanisms** stage, and the formalism requires some means to discuss these. In earlier work [23], we developed a formalism for agents to argue over rules of inference and a version of this could be used here. Thirdly, we are developing a formalism to enable participants to discuss alternative methods of deciding between different arguments. Such a formalism will enable participants to agree upon a means to resolve differences in opinions on dialogue purposes, inference mechanisms, criteria, etc, as mentioned at the end of Section 3.

Acknowledgments

This work was partly funded by the EU IST Programme through the Sustainable Lifecycles in Information Ecosystems Project (IST-1999-10948) and by the British EPSRC through a PhD studentship. The presentation of this paper at the First International Workshop on Chance Discovery in Matsue, Japan, was made possible by a grant from the Hakuhodo Corporation, of Tokyo. We are very grateful to Mr Makoto Mizuno and Hakuhodo for this support. We also thank David Hitchcock and Doug Walton for discussions on these topics.

References

[1] L. Amgoud, N. Maudet, and S. Parsons. Modelling dialogues using argumentation. In E. Durfee, editor, *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS 2000)*, pages 31–38, Boston, MA, USA, 2000. IEEE Press.

[2] L. Amgoud, S. Parsons, and N. Maudet. Arguments, dialogue, and negotiation. In W. Horn, editor, *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI 2000)*, pages 338–342, Berlin, Germany, 2000. IOS Press.

[3] T. J. M. Bench-Capon, J. Freeman, H. Hohmann, and H. Prakken. Computational models, argumentation theories and legal practice. In C. Reed and T. Norman, editors, *Bonskeid Symposium on Argument and Computation*. 2001 (In press).

[4] T. J. M. Bench-Capon, T. Geldard, and P. H. Leng. A method for the computational modelling of dialectical argument with dialogue games. *Artificial Intelligence and Law*, 8:233–254, 2000.

[5] D. V. Carbogim, D. S. Robertson, and J. R. Lee. Argument-based applications to knowledge engineering. *Knowledge Engineering Review*, 15(2):119–149, 2000.

[6] R.D. Coombes and P. N. Judson. Artificial intelligence systems for predicting toxicity. *Pesticide Outlook*, 7 (4):11–15, 1996.

[7] J. Emery, R. Walton, A. Coulson, D. W. Glasspool, S. Ziebland, and J. Fox. Computer support for recording and interpreting family histories of breast and ovarian cancer in primary care (RAGs): qualitative evaluation with simulated patients. *British Medical Journal*, 319:32–36, 1999.

[8] FIPA. Specification part 2 – agent communication language. Technical report, Foundation for Intelligent Physical Agents, 23 October 1997.

[9] J. Fox, P. Krause, and S. Ambler. Arguments, contradictions and practical reasoning. *Proceedings of the European Conference on Artificial Intelligence (ECAI-92), Vienna, Austria*, 1992.

[10] D. Frohlich and P. Luff. Mixed initiative interaction. In T. Bench-Capon, editor, *Knowledge-Based Systems and Legal Applications*, pages 265–294. Academic Press, London, UK, 1991.

[11] T. F. Gordon. The Pleadings Game: An exercise in computational dialectics. *Artificial Intelligence and Law*, 2:239–292, 1994.

[12] T. F. Gordon and N. Karacapilidis. The Zeno argumentation framework. In *Proceedings of the Sixth International Conference on AI and Law*, pages 10–18. ACM Press, 1997.

[13] C. L. Hamblin. *Fallacies*. Methuen and Co Ltd, London, UK, 1970.

[14] K. Harper. Hatfield: the warnings Railtrack ignored. *The Guardian*, 16 January 2001. (London, U.K.).

[15] Healthy Rivers Commission. Independent Inquiry into the Clarence River System. Final report, Healthy Rivers Commission of New South Wales, Sydney, NSW, Australia, 1999.

[16] D. Hitchcock. Some principles of rational mutual inquiry. In F. van Eemeren et al., editor, *Proceedings of the Second International Conference on Argumentation*, pages 236–243, Amsterdam, The Netherlands, 1991. SICSAT.

[17] D. Hitchcock, P. McBurney, and S. Parsons. A framework for deliberation dialogues. In H. V. Hansen, C. W. Tindale, J. A. Blair, and R. H. Johnson, editors, *Proceedings of the Fourth Biennial Conference of the Ontario Society for the Study of Argumentation (OSSA 2001)*, Windsor, Ontario, Canada, 2001.

[18] J. Hulstijn. *Dialogue Models for Inquiry and Transaction*. PhD thesis, Universiteit Twente, Enschede, The Netherlands, 2000.

[19] P. Krause, S. Ambler, M. Elvang-Gørannson, and J. Fox. A logic of argumentation for reasoning under uncertainty. *Computational Intelligence*, 11 (1):113–131, 1995.

[20] P. Krause, J. Fox, P. Judson, and M. Patel. Qualitative risk assessment fulfils a need. In A. Hunter and S. Parsons, editors, *Applications of Uncertainty Formalisms*, LNAI 1455, pages 138–156. Springer Verlag, Berlin, Germany, 1998.

[21] J. D. MacKenzie. Question-begging in non-cumulative systems. *Journal of Philosophical Logic*, 8:117–133, 1979.

[22] P. McBurney and S. Parsons. Risk agoras: Dialectical argumentation for scientific reasoning. In C. Boutilier and M. Goldszmidt, editors, *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence (UAI-2000)*, pages 371–379, Stanford, CA, USA, 2000. UAI.

[23] P. McBurney and S. Parsons. Tenacious tortoises: a formalism for argument over rules of inference. In G. Vreeswijk, editor, *Workshop on Computational Dialectics, Fourteenth European Conference on Artificial Intelligence (ECAI2000)*, Berlin, Germany, 2000. ECAI.

[24] P. McBurney and S. Parsons. Agent ludens: games for agent dialogues. In S. Parsons and P. Gmytrasiewicz, editors, *Game-Theoretic and Decision-Theoretic Agents: Proceedings of the 2001 AAAI Spring Symposium*, pages 70–77, Menlo Park, CA, USA, 2001. AAAI Press. Technical Report SS-01-03.

[25] P. McBurney and S. Parsons. Representing epistemic uncertainty by means of dialectical argumentation. *Annals of Mathematics and Artificial Intelligence*, (in press), 2001.

[26] P. McBurney and S. Parsons. A dialogue-game specification for discovery dialogues. Technical report, Agent ART Group, Department of Computer Science, University of Liverpool, Liverpool, U.K., 2001 (In preparation).

[27] M. Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA, USA, 1999.

[28] S. Parsons, C. Sierra, and N. R. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.

[29] H. Prakken. On dialogue systems with speech acts, arguments, and counterarguments. In M. Ojeda-Aciego, M. I. P. de Guzman, G. Brewka, and L. M. Pereira, editors, *Proceedings of the Seventh European Workshop on Logic in Artificial Intelligence (JELIA-2000)*, LNAI 1919, Berlin, Germany, 2000. Springer Verlag.

[30] H. Prakken and G. Vreeswijk. Logics for defeasible argumentation. In D. Gabbay, editor, *Handbook of Philosophical Logic*. Kluwer Academic, Dordrecht, The Netherlands, second edition, 2000.

[31] C. Reed and T. Norman, editors. *Bonskeid Symposium on Argument and Computation*. 2001 (in press).

[32] W. Rehg, P. McBurney, and S. Parsons. Computer decision-support systems for public argumentation: Criteria for assessment. In H. V. Hansen, C. W. Tindale, J. A. Blair, and R. H. Johnson, editors, *Proceedings of the Fourth Biennial Conference of the Ontario Society for the Study of Argumentation (OSSA 2001)*, Windsor, Ontario, Canada, 2001.

[33] K. Stathis. A game-based architecture for developing interactive components in computational logic. *Journal of Functional and Logic Programming*, 2000(5), March 2000.

[34] C. A. G. Tonnelier, J. Fox, P. N. Judson, P. J. Krause, N. Pappas, and M. Patel. Representation of chemical structures in knowledge-based systems: the StAR system. *Journal of Chemical Information and Computer Sciences*, 37:117–123, 1997.

[35] S. E. Toulmin. *The Uses of Argument*. Cambridge University Press, Cambridge, UK, 1958.

[36] F. H. van Eemeren, R. Grootendorst, and F. S. Henkemans. *Fundamentals of Argumentation Theory*. Lawrence Erlbaum Associates, Mahwah, NJ, USA, 1996.

[37] D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Series in Logic and Language. State University of New York Press, Albany, NY, USA, 1995.