Dialectical Argumentation in Causal Domains

Nancy L. Green[1]

Abstract. We present an interactive computational model of argumentation for quasi-causal domains such as genomic medicine. The model is intended to help a lay audience to evaluate an expert's arguments for a claim. The model regenerates the expert's arguments using evidence, domain knowledge, and argument schemes presumed to be available to the expert. In addition, the model enables the user to explore the space of possible attacks on and defenses of the generated arguments.

1 INTRODUCTION

We present an interactive computational model of argumentation for quasi-causal domains such as genomic medicine. The model is intended to help a lay audience to evaluate an expert's arguments for a claim. The model regenerates the arguments using evidence, domain knowledge, and argument schemes presumed to be available to the expert. In addition, the model enables the user to explore the space of possible attacks on and defenses of the generated arguments. In the model, domain knowledge and evidence are represented in a qualitative probabilistic knowledge base (KB). Arguments are generated by instantiation of domain-independent normative argument schemes with beliefs from the KB. The argument schemes reflect the structure of argument proposed by Toulmin [15] and the notion of critical question [19]. User interaction is regulated by a dialogue game based on the structure of the argument schemes.

This interactive model builds upon our ongoing research on natural language argument generation in GenIE, a system for generating first drafts of genetic counseling letters [9]. A letter drafted by GenIE may contain arguments for each of the conclusions reached by the patient's medical providers. GenIE's model distinguishes biomedical domain knowledge, formal argument schemes, and a discourse grammar encoding genre-specific knowledge. The work reported here extends GenIE's approach with the addition of critical questions to argument schemes and use of a dialogue game in place of a discourse grammar. These extensions will enable users to critically examine arguments in letters drafted by GenIE, e.g., to request counterarguments not included in the letter.

For example, suppose that a patient has been told that she has a genetic condition called osteogenesis imperfecta (OI), and that she probably has mosaicism for OI (a new mutation in some of her cells) rather than the autosomal dominantly inherited form of OI (a different mutation, occurring in all of her cells). Assuming that the KB has been updated with the beliefs of the healthcare providers about the patient's case, the patient would be able to interactively explore the arguments relating to her diagnosis as follows. S represents the system's contribution and U represents the user's request. Each agent's dialogue contributions are given in English for ease of illustration although, as shown in Section 5, the implemented interactive prototype currently uses an artificial language for communication between user and system.

1S: *Diagnosis is mosaicism for OI.*
2U: Arguments for mosaicism?
3S: *Patient's skin biopsy showed abnormal collagen.*
4U: So what?
5S: *The mutation associated with mosaicism for OI causes the collagen protein produced by the affected cells to have an abnormal form.*
6U: Another argument?
7.1S: *Patient has mild symptoms of OI such as flexible joints,*
7.2S: *although autosomal dominant OI could also account for her symptoms.*
8U: Arguments against autosomal dominant OI?
9S: *Patient's test result for the autosomal dominant form of OI was negative*
10U: Another argument?
11S: *Patient's family has no history of OI.*
12U: Arguments against mosaicism for OI?
13.1S: *Patient does not have severe OI symptoms,*
13.2S: *although that could be explained if only a small percentage of her cells have the mutation associated with mosaicism for OI.*

To give an overview of the model presented here, the components are a qualitative probabilistic knowledge base (the KB), an argument generator, and a dialogue game, TDGCQ. TDGCQ regulates interaction with a user and sends the argument generator the user's requests for arguments for specified claims. The argument generator sends requests for information satisfying constraints of an argument scheme to the KB. After an argument is returned to the user, the user may request arguments or counterarguments for any of its elements (claim, data and warrant), or qualifications that may weaken it (i.e., corresponding to 'yes' answers to any of its critical questions). In this paper, we describe components of the model from the "lowest" level up. We first describe the KB, mainly in terms of its features relevant to argumentation. Next we present several argument schemes. Then we present the TDGCQ and an example showing interaction in a prototype implementation. Lastly, we describe related research.

[1]University of North Carolina Greensboro, USA., email: nlgreen at uncg.edu

2 QUALITATIVE PROBABILISTIC KNOWLEDGE BASE

The qualitative probabilistic knowledge base (KB) represents an expert's domain knowledge. The KB includes three kinds of knowledge: qualitative probabilistic/causal domain models, quantitative and qualitative probability statements justifying constraints in the domain models, and qualitative probabilistic beliefs about specific individuals. A prototype has been implemented in the domain of clinical genetics. For example, domain models describe potential causal paths from genetic conditions to clinical findings, and inheritance of genetic conditions in families. Epidemiological statistics (such as the frequency of a certain symptom S among individuals with a certain genetic mutation G) may be used to justify a causal link (e.g. between condition G and symptom S) in the domain model. Also, probability statements based on Mendelian theory, which predicts the probability of inheriting a genotype given the genotype of each parent, can be used to justify constraints among genotypes of parents and offspring in the domain model. Beliefs regarding particular patients may include findings (e.g., the patient has chronic respiratory infections), a diagnosis (e.g., she probably has cystic fibrosis, an autosomal recessively inherited genetic disease resulting from mutation of both alleles of the CFTR gene), and beliefs about the source of the patient's genetic condition (e.g., it is likely that she inherited one mutated copy of the gene from each parent). In such a domain, an expert's conclusions are defeasible (i.e., new findings may require belief revision or findings may have alternative explanations) and probabilistic (e.g., although it is likely that the patient has cystic fibrosis, the diagnosis cannot be given with certainty).

The domain models and beliefs about specific individuals are implemented using qualitative probabilistic networks (QPN) [5, 20]. Like a Bayesian network (BN), a QPN is a directed acyclic graph whose nodes represent random variables. However, in a QPN qualitative constraints replace the conditional probability tables of a BN. A qualitative model provides a level of abstraction that is convenient for specifying argument schemes. Note that the domain model is not designed for domain reasoning, e.g., to function as a clinical expert system; and argument generation is not used to update beliefs in the KB. The above statistics and Mendelian probabilities are stored in the KB but not within the QPN formalism. The causal domain models are manually constructed following an abstract, simplified causal model of genetic disease used by domain experts in communication with lay clients. This abstract model was determined by identifying a set of categories (e.g. genotype, test result, symptom) with good inter-rater reliability that can be used to describe the biomedical content of letters in a corpus of patient letters written by genetic counselors [8]. The causal domain models instantiate this abstract model for specific genetic diseases. At runtime, before interaction with a user, the relevant QPN is updated with the expert's beliefs about a specific patient's case: findings, diagnosis, source of the patient's genetic condition (i.e., inferences about genotypes of certain family members), etc.

Qualitative probabilistic constraints for each QPN are defined in terms of the relations of qualitative influence *S*, additive synergy *Y*, and product synergy *X* [5, 20]. A variable *A* is said to have a positive qualitative influence on a variable *B*, written $S^+(A,B)$, if a higher value of *A* makes a higher value of *B* more likely. An analogous definition is given for negative influence, i.e., $S^-(A,B)$, if a higher value of *A* makes a lower value of *B* more likely. The nodes in the causal domain model represent discrete variables. Each arc is described by a positive or negative influence relation. Since not all the variables are Boolean our notation encodes threshold values of the domains explicitly. For example, positive influence is encoded as $S^+(<A,a>,<B,b>)$, i.e., when *A* reaches *a* (or higher), it is more likely that *B* will reach *b* (or higher).

Synergy relations *X* and *Y* are used to express a relationship between a set of variables and another variable. Negative product synergy, $X^-(\{A,B\},<C,c>)$, is used to model the relation between a mutually exclusive set {*A, B*} of potential causes of an event <*C,c*> and the event, e.g., between a set of candidate diagnoses and presence of a certain symptom. Another example of its use in this domain is to model autosomal dominant inheritance; having one mutated allele of a gene fitting this inheritance pattern normally is sufficient to cause a health problem. Thus, if a child has an autosomal dominantly inherited disorder then (since we inherit one copy of each pair of genes from each parent) the child must have inherited the mutation from one or the other of his parents. Zero product synergy, $X^0(\{A,B\},<C,c>)$, is used to model the relation between a set of events {*A, B*} which are jointly necessary to cause the event <*C,c*>; e.g., in autosomal recessive inheritance of a disorder, the child must have inherited one copy of the mutation from each parent. Positive additive synergy, $Y^+(\{A,B\},<C,c>)$, is used to model a situation where *A* enables *B* to result in *C*. Negative additive synergy, $Y^-(\{A,B\},\neg<C,c>\}$, is used to model a situation in which *A* inhibits *B* from resulting in <*C,c*>.

3 ARGUMENT SCHEMES

To derive argument schemes, we analyzed normative arguments from a corpus of genetic counseling patient letters in terms of Toulmin's characterization of argument structure [15]. According to that, the premises of an argument can be classified functionally as data or warrant. The warrant is a principle that licenses the claim given the data; the warrant itself can be justified by other data, called the backing. When deriving argument schemes, we interpreted the text of a patient letter in reference to a presumed conceptual domain model, i.e., a model instantiating the abstract causal model of genetic disease we identified in previous work [8]. Qualitative constraints of the causal domain model (defined in terms of influence and synergy relations described in the previous section) may serve as warrants; the (qualitative belief in the) state of a causal domain model variable may serve as data or claim; and probability statements (in this domain, they would come from the medical research literature or Mendelian theory) may serve as backing for warrants. In some cases we used our interpretation of the writer's intended message, constrained by the domain model, to reconstruct implicit elements, i.e., data, claim, or warrant. (The prevalence of missing elements of arguments in this corpus is described in [10].) The resulting argument schemes are non-domain-specific in the sense that they describe mappings from formal properties of a domain model to Toulmin's functional elements, but do not refer to genetics. Some of the principal argument schemes are shown in Tables 1-3.

To paraphrase part of the scheme labeled E2C, shown in Table 1, an argument for the claim that (it is likely that) the state of *A* is (at least) *a* consists of the warrant that there is a positive influence $S^+(<A,a>,<B,b>)$ and the data that (it is likely that) the state of *B* is (at least) *b*. For example, a variant of this scheme could be used to generate the following argument: *(Data) P has*

frequent respiratory infections. (Warrant) Having cystic fibrosis typically results in having abnormally thick mucous, which typically results in frequent respiratory infections. (Claim) Therefore, P (probably) has cystic fibrosis (two mutated alleles of the CFTR gene). To generate this argument presupposes that the KB contains a variable A representing P's hypothesized CFTR genotype (with two mutated alleles), a variable B representing P's observed symptom of frequent respiratory infections, and a variable C representing P's observed abnormally thick mucous; and a positive influence relation (S^+) from A to C and another positive influence relation from C to B.

To paraphrase part of another scheme, NE2C (shown in Table 2), an argument for the claim that (it is likely that) the state of *A* is less than *a* consists of the warrant that there is a positive influence $S^+(<A,a>,<B,b>)$ and the data that (it is likely that) the state of *B* is less than *b*. A variant of this scheme could be used to generate the following argument: (Data) P's sweat test showed a normal NaCl level. (Warrant) If a sweat test is performed on someone who has cystic fibrosis, the test usually shows an abnormally high NaCl level. (Claim) Therefore, P (probably) does not have cystic fibrosis (P has fewer than two mutated alleles of CFTR). To generate this argument presupposes that the KB contains a variable A representing P's hypothesized CFTR genotype (with fewer than two mutated alleles), a variable P representing P's observed sweat test result (normal), and a variable C representing the action of performing the sweat test on P; and an enabling relation (Y^+) where C enables A to positively influence B.

According to Walton [19], argumentation schemes are specifications of argument patterns that are normative in certain contexts of use; presumptive argumentation schemes are schemes whose conclusions are defeasible. An argument can be challenged by an opponent asking a critical question associated with the argument scheme; the answer to the critical question may defeat the argument. AI researchers have proposed formal approaches to reasoning with argument schemes and the modeling of counterarguments [e.g., 14, 16]. Verheij [16] notes that some of Walton's examples of critical questions amount to questioning the premise or conclusion of an argument scheme; and since any argument can be challenged in this way, it is not necessary to include those types of challenges in the list of critical questions of an argument scheme. Further, he notes that Walton's remaining examples of critical questions represent exceptions to the argument scheme or conditions for its use. The conditions of use correspond to a certain kind of premise (i.e., the warrant in a Toulmin-style analysis), while the exceptions correspond to Toulmin's examples of rebuttal; and only the latter actually need to be specified as critical questions.

For example, an instance of the E2C scheme for a claim A, that the patient's cells have two mutated alleles of CFTR, could be attacked by providing an argument for the claim ¬A, that the patient's cells have less than two (i.e. one or no) mutated alleles of CFTR; or by contesting the data supporting A (e.g., that the patient has frequent respiratory infections); or by contesting the warrant (e.g., disputing the existence of a positive influence relation between that genotype and this symptom); or by showing that an exception holds (e.g., that mutation of some other genotype could have caused this symptom).

To cover this last type of rebuttal, critical questions are included in some of our argument schemes. For example in the E2C scheme shown in Table 1, the critical question (CQ1) is formulated as: *is there a C participating in the qualitative relation $X^-(\{C,A\},<B,b>)$ such that $C \geq c$*? In other words, could there be an alternative explanation C to A as the exclusive reason for $B \geq b$? In the NE2C scheme shown in Table 2, the first critical question (CQ1) is: *is there a C participating in the qualitative relation $Y^+(\{C,A\}, <B,b>)$ such that $C < c$* ? This is intended to describe a situation where, despite $A \geq a$, an enabling condition C does not hold, thereby limiting A's effect on B. The second critical question (CQ2) describes a situation in which one of the necessary conditions is absent. The third critical question (CQ3) is: *is there a C participating in the qualitative relation $Y^-(\{C,A\}, \neg<B,b>)$ such that $C \geq c$*? This describes the situation where, despite $A \geq a$, a treatment C inhibits the causal influence of A on B.

Table 1. "Effect to Cause" argument scheme.

Scheme	**E2C**
Claim	$A \geq a$
Data	$B \geq b$
Warrant	$S^+(<A,a>, <B,b>)$
CQ1	$\exists$ C $X^-(\{C,A\},<B,b>)$: $C \geq c$

Rough gloss: If A (a hypothesized event) occurred then B (another event) is likely to have occurred (in some cases as a result). B has been observed. Therefore A has probably occurred, unless there is an alternative explanation C for B.

Table 2. "No Effect to Cause" argument scheme.

Scheme	**NE2C**
Claim	$A < a$
Data	$B < b$
Warrant	$S^+(<A,a>, <B,b>)$
CQ1	$\exists C$ $Y^+(\{A,C\},<B,b>)$: $C < c$
CQ2	$\exists C$ $X^0(\{A,C\},<B,b>)$: $C < c$
CQ3	$\exists C$ $Y^-(\{A,C\}, \neg<B,b>)$: $C \geq c$

Rough gloss: If A (a hypothesized event) occurred then B (another event) is likely to have occurred (in some cases as a result). But B has not been observed. Therefore A has probably not occurred, unless there is an alternative explanation C1, C2, or C3 for the non-occurrence of B, where C1 is that an enabling condition does not hold, C2 is that a necessary condition does not hold, and C3 is that a mitigating condition does hold.

Table 3. "Increased Risk" argument scheme.

Scheme	**IncRisk**
Claim	$B \geq b$
Data	$A \geq a$
Warrant	$S^+(<A,a>, <B,b>)$
CQ1	$\exists$ C $Y^-(\{A, C\}, \neg<B,b>)$: $C \geq c$
CQ2	$\exists C$ $Y^+(\{A,C\},<B,b>)$: $C < c$

Rough gloss: If A (an event or risk factor) occurred (or is present) then B (another event) might occur (or might have occurred). A has occurred (or is present). Therefore in this case B might occur (or might have occurred), unless CQ1 or CQ2 holds. CQ1 and CQ2 describe a situation where, despite the occurrence/presence of A, B might not occur (have occurred) because A is inhibited by C, or because an enabling condition C does not hold, respectively.

The E2C and NE2C schemes are used for arguments from effects to possible causes. The IncRisk scheme, shown in Table 3, is used for arguments in the opposite direction. The data is about an event that may play a causal role in or signal risk of another event. (For example, obesity is a risk factor for heart disease.) The first critical question (CQ1) describes the situation where, despite $A \geq a$, a mitigating condition $C \geq c$ holds, thereby limiting A's potential to increase risk of $B \geq b$. The second critical question (CQ2) describes a situation in which the risk is mitigated by the absence of the enabling condition.

Now we shall illustrate how these argument schemes (and some variants of them) are used in generating the sample dialogue shown in section 1. An E2C argument for the claim of 1S consists of data given in 3S and warrant given in 5S. Another E2C argument for the same claim consists of data only (7.1S); also a qualification is given (7.2S) based on the critical question (CQ1) of the E2C scheme. (Note that the warrant was not requested by the user for this argument, nor for any of the subsequent arguments in this dialogue.) In 9S the system provides the data of an NE2C argument for the claim of 8U (that the patient does not have autosomal dominant OI). In 11S, the system provides the data of another argument for the same claim, using an argument scheme (UnIncRisk, not shown here) similar to IncRisk but where the data is that the event signaling risk has not occurred. Finally, 13.1S provides the data of an NE2C argument for the claim of 12U (that the patient does not have mosaicism for OI); also the argument is qualified in 13.2S, which gives an affirmative answer to a critical question, CQ2, of the NE2C scheme.

In addition to the schemes shown in Tables 1-3, other schemes have been defined, some of which are illustrated in the previous paragraph. One set of schemes, including the three given in Tables 1-3, describes arguments from cause to effect, or vice versa, witnessed in the corpus. These schemes reflect different possible qualitative influence and/or synergy relations among variables in a KB. (Thus, they extend the causal schemes informally described in [19].) Another set (beyond the scope of this paper) describes how probability statements are used as backing for warrants. A significant contribution of our previous work [9] was the formalization of argument schemes in a way that identified Toulmin's data/warrant distinction with the distinction in a qualitative probabilistic KB between the states of variables and the formal (influence and synergy) relations among variables, respectively. This paper extends that approach with the addition of critical questions to argument schemes. As far as we know, this is the first formal representation of critical questions of causal argument schemes in terms of influence and synergy relations. As we shall discuss in the next section, this extension plays a key role in supporting dialectical argumentation in our model.

4 TOULMIN DIALOGUE GAME WITH CRITICAL QUESTIONS

This section presents a dialogue game (shown in Figure 1) for generating all arguments for and against a claim. State transitions represent legal moves. We call our specification Toulmin Dialogue Game with Critical Questions (TDGCQ) since the moves enable an agent to request an argument supporting or attacking each of the Toulmin-style components of an argument (claim, data, or warrant) as in the Toulmin Dialogue Game [1]. However, unlike that game, TDGCQ also generates attacking moves through critical questions. Another difference is that the moves of TDGCQ are not intended to model the effect on the participants' commitments, i.e., to update the KB. Figure 1 is intended to show how components of an argument can be supported or challenged, and how the supporting arguments or attacks can themselves be supported or challenged in turn. To simplify the presentation of the game, we use a recursive transition network. For example the arc labeled *Request* A_i = *argument(…)* returns a complete argument with components labeled C, D, W, and E. (This replaces sequences of transitions in [1] such as *Why* → *Supply Data* → *So* → *Supply Warrant*.) Also, instead of describing dialogue moves for different roles (Prop, Opp, and Referee in [1]), our network describes the allowable moves of the user since the system has no choice but to try to satisfy the user's request; states are labeled Pro and Con to indicate whether the subsequent move is pro (supporting) or con (attacking), respectively. To simplify the presentation, it is assumed that an argument contains no more than one item each of data, warrant, and critical question and the graph does not show error or terminal states. Finally, note that since there may be zero, one, or more arguments for a claim C, a loop is provided in the graph to allow an agent to iterate over all arguments A_i for C.

For example, if a pro-argument A for claim C includes data D and warrant W, then one also can request pro-arguments for D or W, resulting in a recursive call to TDGCQ in which the claim is D or W, respectively. An argument for a warrant W could be made if the KB contained a probability statement that provided backing for W. Con-arguments for C include pro-arguments for ¬C, as well as arguments that attack D or W, represented by the arcs specifying recursive calls to TDGCQ to find arguments for ¬D or ¬W, respectively. The negation symbol is used in the figure to represent incompatible beliefs; e.g., if P is a statement that G = 2, where G is a variable with domain {0,1,2}, then ¬P represents G < 2. An argument for ¬W, where W is a warrant, could be made if the KB contained a QPN constraint incompatible with W. The exception E is how we refer to the instantiated predicate in a critical question of an argument scheme. As described in Section 3, critical questions are expressed in first-order logic with restricted quantification; when an argument is returned by the argument generator, it may include an instantiation of the predicate (the part following the colon) of one of its critical question. Then, a con-argument for C can be requested by asking for an argument for E, resulting in a recursive call to TDGCQ. An example of interaction through the TDGCQ is given in the next section.

The argument generator is driven by requests from the TDGCQ. (In contrast, in our previous work [9], the argument generator is driven by requests from a discourse grammar.) Given a claim, the argument generator searches for an argument scheme whose claim can be unified with it. For each such scheme found the generator attempts to instantiate the data and warrant with information from the KB consistent with the instantiated claim. After an argument is found, the argument generator attempts to instantiate critical questions as follows. The restriction of each critical question, i.e., the part preceding the colon, refers to a possible qualitative constraint in the domain model. For example, the restriction of CQ1 of E2C is $\exists C\ X^-(\{C,A\},<B,b>)$. Thus, the argument generator will search for a variable C related to B by an X^- relation in the KB. If the restriction can be instantiated consistent with the instantiated claim, data, and warrant of the argument, then the instantiated critical question is returned, without attempting to "answer" the critical question at this point. In subsequent interaction with the TDGCQ, a user may ask for an argument for the instantiated critical question. To continue the above example, the user may ask whether the instantiation of the predicate of CQ1 of E2C, C > c, holds in the KB.

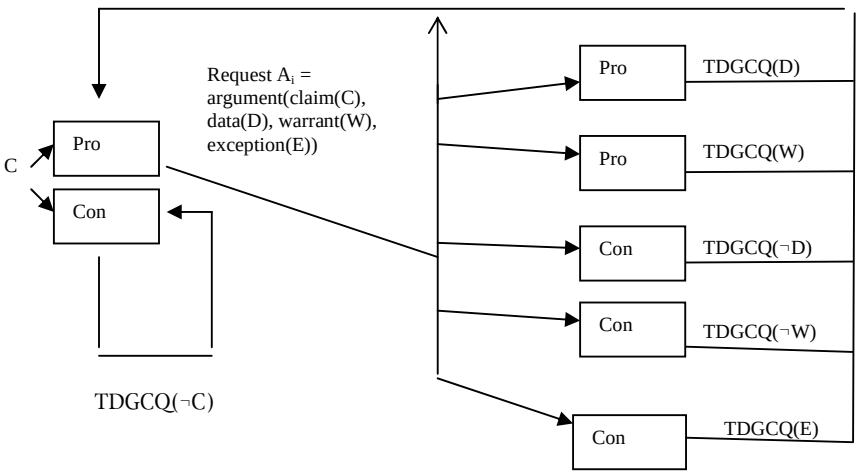

Figure 1. TDGCQ

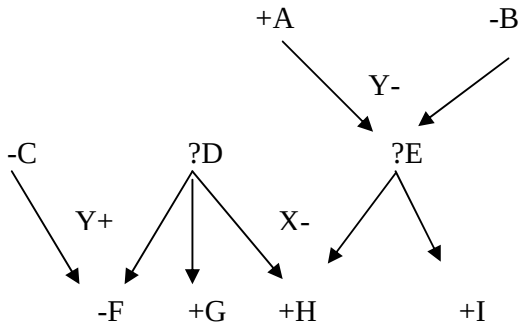

Figure 2. Example of QPN

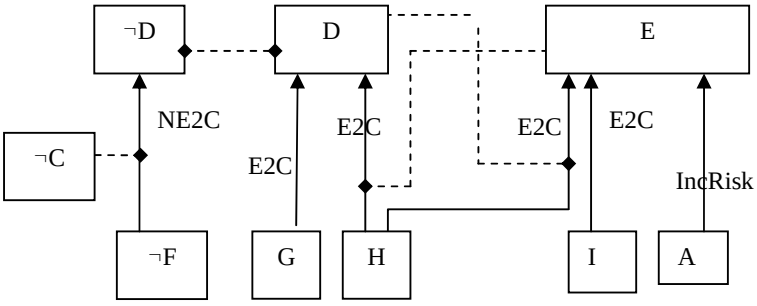

Figure 3. Argument diagram

5 EXAMPLE

In this section we illustrate our approach using a constructed example designed to illustrate many paths through the TDGCQ. The QPN used in the example is shown in Figure 2. To save space, the nodes of the QPN are represented by upper case letters from the beginning of the alphabet; qualitative constraints are not shown in full; and '+' or '-' signifies some degree of positive belief in a value at-or-above-threshold or below-threshold, respectively. For example, nodes D and E could represent alternative possible genetic diagnoses; F, G, H, and I, findings; A, a risk factor for E; B, a factor that mitigates risk of E; and C, an event usually required to enable F.

Figure 3 shows a group of interrelated arguments that can be based on the expert's beliefs represented in the QPN in Figure 2. (Our system is not designed to generate box notation. The figure was manually constructed to illustrate all possible arguments that could be generated from this QPN in answer to a user's moves in the TDGCQ.) The argument diagram notation used in the figure is similar to that used in a number of previous studies [e.g., 17]. Boxes represent claims and data. Warrants are not shown to save space. Arrows, linking data to claims, represent support. In our version of this notation, support arrows are annotated with the name of the argument scheme used, and attacks are indicated using dashed arcs ending in diamond tips. Attacks via critical questions are linked to the support arrow below the name of the argument scheme.

To paraphrase the arguments shown in Figure 3 from left to right, first, an NE2C argument for ¬D is given with data ¬F, but that argument is attacked via the first critical question, CQ1, of the NE2C argument scheme, since the failure of C to enable D could be the reason (rather than ¬D) that ¬F was observed. The argument for ¬D is joined by an attack link to the claim for D, which is supported by two arguments, an E2C argument with data G, and another E2C argument with data H. However, the E2C argument based on H is attacked via the critical question of the E2C scheme since E could be the reason for H. At the far right, three arguments for E are presented. Similar to the E2C argument for D based on H, the E2C argument for E based on H is attacked via the critical question since D could be the reason for H. A second E2C argument for E is given based on I. Lastly, an IncRisk argument is given for E based on A. Since ¬B is observed, that argument cannot be attacked via its critical question.

Figure 4 shows how a user could elicit the arguments shown in Figure 3 via the TDGCQ in our implemented prototype system. Lines have been numbered for reference (S denotes the system and U denotes the user. All lines beginning with the '?' prompt are a user's request.) To gloss the game shown in Figure 4, first, 1U consists of a request for an argument for ~D. In response, the system provides an NE2C argument in 2S. Then in 3U the user requests an argument for ~F, the data of the argument given in 2S. The system replies in 4S that ~F was observed. In 5U, the user requests an argument for ~C, the instantiated critical question of the argument given in 2S. The system answers in 6S that ~C was observed, and so the argument in 2S has been successfully attacked. The system's response to the user request for an argument for D in 7U consists of two arguments (8S and 9S). The user verifies that the data in each of the two arguments can be supported (10U-13S). In 14U, an argument for E, the instantiated critical question of the argument returned in 8S, is requested. The system responds with three arguments for E (15S-17S). Finally, the user requests any arguments against B, a critical question of the argument given in 17S. The system responds in 19S that ~B was observed, and so the challenge via this critical question does not succeed.

1U ? ~D
2S Arg (NE2C): [Claim: ~D], [Data: ~F],
[Warrant: S+(D,F)], [CQ: Y+({D,C},F): ~C]
3U ? ~F
4S Observed
5U ? ~C
6S Observed
7U ? D
8S Arg (E2C): [Claim: D], [Data: H] , [Warrant: S+(D,H)],
[CQ: X-({D,E},H):E]
9S Arg (E2C): [Claim: D], [Data: G], [Warrant: S+(D,G)]
10U ? H
11S Observed
12U ? G
13S Observed
14U ? E
15S Arg (E2C): [Claim: E], [Data: I], [Warrant: S+(E,I)],
16S Arg (E2C): [Claim: E], [Data: H], [Warrant: S+(E,H)],
[CQ: X-({D,E},H):D]
17S Arg (IncRisk): [Claim: E], [Data: A], [Warrant: S+(A,E)],
[CQ: Y-({A, B}, E): B]
18U ?~ B
19S Observed

Figure 4 Example interaction in TDGCQ

6 RELATED WORK

Although domain knowledge is modeled in our system with probabilistic/causal models, the goal of our work differs from research on explanation of reasoning in probabilistic expert systems. The goal of that research is to make the expert system's reasoning plausible to human experts by describing the structure of the causal network and explaining the probability calculations [11]. Another difference is that the probabilistic networks in expert systems are optimized for problem-solving (and may even have been built by automatic methods), while in our work the QPN represents a conceptual model used in expert-lay communication.

Other research has focused on transforming probabilistic networks into rule bases for domain reasoning. Although the goal of that research differs from ours, it is interesting to see its approach to transforming probabilistic domain knowledge to a propositional representation. In [18], evidence nodes of a BN are mapped to atomic propositions and each row of the conditional probability tables of the BN is mapped to a rule retaining the numerical probability information. An argument for a claim is the chain of rules used to derive a claim; the strength of an argument is computed from the joint probability of the argument nodes. In [21], a BN is constructed automatically from medical data. Then the BN is converted to a rule-base. The mapping to the

rule-base uses two translation rules: (1) $p(x \mid y) > p(x)$ maps to the defeasible rule $y \Rightarrow x$, and (2) $p(x \mid y) < p(x)$ maps to the defeasible rule $y \Rightarrow \neg x$. The two translation rules cover the positive and negative influence relations (S^+, S^-) used in our domain models, but this approach does not address other types of qualitative probabilistic constraints addressed in our approach. Also, neither paper relates formal properties of causal/probabilistic models to argument-theoretic notions of warrant and critical question as in our approach.

Some research in natural language generation of persuasive arguments has implemented domain knowledge and the user's presumed domain knowledge (called the "user model") in BN knowledge bases. In NAG [22], the user model is represented in one BN and domain knowledge in another; the user model is employed to predict the persuasiveness (from the audience's perspective) of evidence from the domain model that could be added to an argument. In other research on generating persuasive arguments in natural language there is a focus on addressing the audience's emotions as well as their beliefs [2, 3, 12]. A single BN is used to represent domain knowledge, the user's emotional state, and argument schemes; BN reasoning is used to compute persuasiveness. Although it does not use a Bayesian approach, other research on generating persuasive natural language arguments [7] uses a model of the audience's preferences and values to help plan arguments; persuasive argument strategies from the New Rhetoric [13] are followed. One difference between our work and all of these mainly BN-based approaches to natural language argument generation is that the goal of our work is not persuasion but to enable the audience to see all arguments for and against the experts' claims. Also, although we have considered the role of the audience's affective state in other work, it is not relevant here. Finally, as was the case with the research discussed in the preceding paragraph, none of the BN-based natural language argumentation research relates formal properties of causal/probabilistic models to argument-theoretic notions of warrant and critical questions.

Lastly, there has been other research on argumentation for the domain of genomic medicine. RAGS (Risk Assessment in Genetics) [4] is a rule-based decision-support tool to assist doctors in assessing the qualitative risk that a patient has a particular mutation given data about the patient's family tree. The system's conclusion is explained to the user by listing arguments for and against the assessment that the patient has the mutation. In contrast, risk calculation is outside of the scope of our research, but our model enables a qualitative probabilistic approach to domain reasoning to serve as a knowledge source for argumentation. The REACT system [6] is a decision-support system for planning the medical care of women diagnosed as at risk of developing cancer due to genetic factors. REACT presents arguments for and against treatment options. However, unlike in our approach, the arguments are not generated "on the fly" but are human-authored and stored in a knowledge base.

7 CONCLUSION

We presented a computational model of dialectical argumentation for generating arguments interactively based upon an expert's domain knowledge and evidence represented in a qualitative probabilistic knowledge base. Domain-independent argument schemes map formal properties of the knowledge base to data, warrant, and critical questions of arguments. A dialogue game has been implemented that enables a user to examine support for as well as qualifications of and attacks on the experts' beliefs. In future work, we plan to adapt the natural language generation capabilities developed for the GenIE patient letter drafting system to create an end-user-friendly interface to the interactive model.

ACKNOWLEDGMENTS

This work is supported by the National Science Foundation under CAREER Award No. 0132821. The most recent prototype was implemented by Ms. Kanyamas (Jenny) Navoraphan of North Carolina State University.

REFERENCES

[1] T.J.M. Bench-Capon, T. Geldard, P.H. Leng, A method for the computational modelling of dialectical argument with dialogue games, Artif. Intell. Law 8 (2000) 233-254.

[2] V. Carofiglio, Modelling argumentation with belief networks, in: F. Grasso, C. Reed, G. Carenini, (Eds.), Proc. of the 4th Workshop on Computational Models of Natural Argument (CMNA-04), 2004.

[3] V. Carofiglio, F. de Rosis, Combining logical with emotional reasoning in natural language argumentation, in: C. Conati, E. Hudlicka, C. Lisetti, (Eds.), Assessing and Adapting to User Attitudes and Effect: Why, When and How?, Workshop Proc., 9th International Conference on User Modeling (UM'03), Johnstown, PA, 2003, pp. 9-15.

[4] A.S. Coulson, D.W. Glasspool, J. Fox, J. Emery, Computerized genetic risk assessment and decision support in primary care, Informatics: the Journal of Informatics in Primary Care (May 2000).

[5] M.J. Druzdzel, M. Henrion, Efficient reasoning in qualitative probabilistic networks, in: Proc. of the 11th Nat. Conf. on AI (AAAI-93), 1993, pp. 548-553.

[6] D.W. Glasspool, J. Fox, A. Oettinger, J. Smith-Spark, Argumentation in decision support for medical care planning for patients and clinicians, in: Argumentation for Consumers of Healthcare: Papers from the AAAI Spring Symposium, AAAI Press, Menlo Park, CA, 2006, pp. 58-63.

[7] F. Grasso, A. Cawsey, R. Jones, Dialectical argumentation to solve conflicts in advice giving: a case study in the promotion of healthy nutrition, Int. J. of Human-Computer Studies 53 (2002) 1077-1115.

[8] N. Green, A Bayesian network coding scheme for annotating biomedical information presented to genetic counseling clients, J. Biomedical Informatics 38 (2005) 130-144.

[9] N. Green, Generation of biomedical arguments for lay readers, in: Proc. of the 4th International Natural Language Generation Conference (INLG06), Sydney, Australia, 2006, pp. 114-121.

[10] N. Green, Some pragmatic features of argument presentation, Presented at the 10th International Pragmatics Conference, Gothenburg, Sweden, 8-13 July 2007.

[11] C. Lacave, F.J. Díez, A review of explanation methods for Bayesian networks, The Knowledge Engineering Review 17(2)

(2002) 107-127.

[12] I. Mazzotta, F. de Rosis, Artifices for persuading to improve eating habits, in: Argumentation for Consumers of Healthcare: Papers from the AAAI Spring Symposium, AAAI Press, Menlo Park, CA, 2006, pp. 76-85.

[13] C. Perelman, L. Olbrechts-Tyteca, The New Rhetoric: a Treatise on Argumentation, Notre Dame Press, University of Notre Dame, IN, 1969.

[14] H. Praaken, AI and law, logic and argument schemes, Argumentation 19 (2005) 303-320

[15] S.E. Toulmin, The Uses of Argument, 9th ed., Cambridge University Press, Cambridge, UK, 1998

[16] B. Verheij, Dialectical argumentation with argument schemes: An approach to legal logic, Artif. Intell. Law 11 (2003) 167-195.

[17] B. Verheij, Evaluating arguments based on Toulmin's scheme, Argumentation 19 (2005) 347-371.

[18] G.A.W. Vreeswijk, Argumentation in Bayesian belief networks, in: I. Rahwan et al., (Eds.), ArgMAS2004, LNAI 3366, 2005, pp. 111-129.

[19] D.N. Walton, Argumentation Schemes for Presumptive Reasoning, Lawrence Erlbaum Associates, Mahweh, NJ, 1996.

[20] M.P. Wellman, Fundamental concepts of qualitative probabilistic networks, Artificial Intelligence 44(3) (1990) 257-303.

[21] M. Williams, J. Williamson, Combining argumentation and Bayesian nets for breast cancer prognosis: Unpublished draft of 18 May 2005, http://www.acl.icnet.uk/lab/PUBLICATIONS/ ms413. pdf.

[22] I. Zukerman, R. McConachy, K. Korb, Using argumentation strategies in automated argument generation, in: Proc. of the 1st International Natural Language Generation Conference, 2000, pp. 55-62