RecoMedic: Recommending Medical Literature through Argumentation

Andrei Mocanu
University of Craiova, Romania

Xiuyi Fan, Francesca Toni, Matthew Williams
Imperial College London, United Kingdom

Jiarong Chen
Southern Medical University, China

***Abstract*—In medical practice, choosing the correct treatment is a key problem [1]. In this work, we present an online medical recommendation system, RecoMedic, that selects most relevant medical literature for patients with brain metastases. RecoMedic maintains a medical literature repository in which users can add new articles, query existing articles, compare articles and search articles guided by patient information. RecoMedic uses argumentation to accomplish the article selection. Thus, upon identifying relevant articles, RecoMedic also explains its selection. RecoMedic can be deployed using single-agent as well as multi-agent implementations. The developed system has been experimented with by senior medical PhD students from Southern Medical University in China.**

***Keywords*-medical recommendation, argumentation**

I. Introduction

In medical practice, choosing the correct treatment is a key problem. Modern practice has emphasized the role of using explicit evidence to make this decision, and a cornerstone of this evidence is generated from randomized controlled trials (RCT). These compare two (or more) treatments in a cohort of defined patients who are randomly allocated to each treatment arm, thus minimizing bias. However, some areas of medicine may generate many trials (there are over 500 new RCTs/ year on breast cancer in the UK alone), which makes it difficult to identify the optimal treatment for each patient. In such cases, we would like a tool to help us identify the ideal study, which would match each patient characteristic most closely.

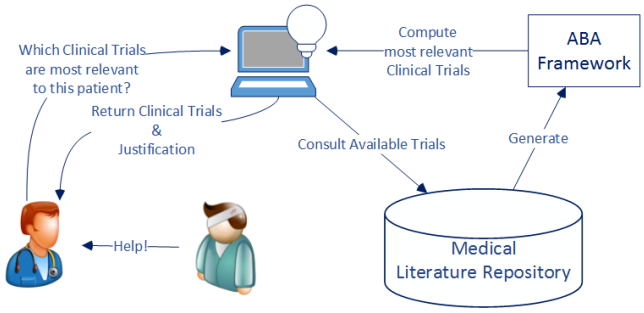

Figure 1. Recommendations & Justifications with Argumentation.

Formal argumentation, as a powerful reasoning tool, has been used extensively in AI in the last two decades (e.g. see [2], [3], [4] for an overview). One unique feature of argumentation is that while performing computation as a form of reasoning, argumentation also gives an explanation to the computation. Thus, argumentation can serve as a versatile tool for applications that need both correct computation as well as transparent explanation.

RecoMedic is a recommendation system which uses argumentation to match individual patients to published clinical trials. Figure 1 describes RecoMedic's main use case scenario: a patient visits a doctor for medical advice; after examination, a set of patient characteristics are collected; to determine the most suitable treatment for the patient, the doctor consults RecoMedic with this patient's characteristics specifically. RecoMedic queries its internal medical literature repository to identify the most relevant medical literature for this patient. Since this query is executed using argumentation (Assumption-based Argumentation (ABA) [5] in particular), RecoMedic not only returns the most relevant literature but also an explanation for this recommendation.

Decision making is a process of selecting good *decisions* amongst several alternatives based on the *goals* met by decisions. RecoMedic views medical literature recommendation as a decision making problem in that the medical literature are alternative decisions and patient characteristics are goals. Thus, RecoMedic uses techniques developed in argumentation based decision making. Decision making with ABA has already been studied in [6], [7], [8], [9], but this is the first time it has been incorporated into a platform directly usable by end-users.

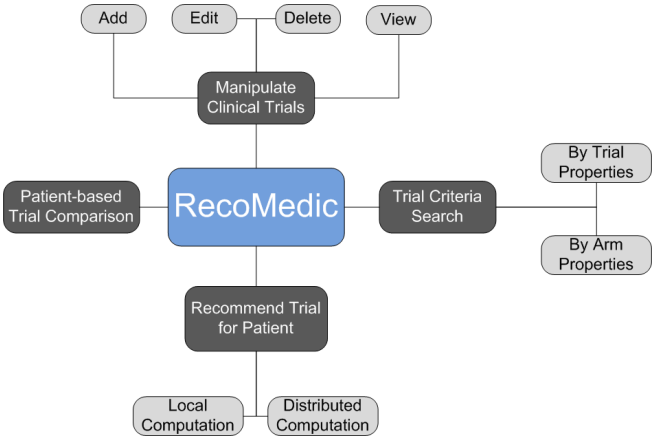

Figure 2. RecoMedic Features

From a functionality point of view, as illustrated in Fig-

ure 2, RecoMedic has four main features: (i) manipulating clinical trials (medical literature), (ii) searching literature in the repository based on certain criteria, (iii) recommending the most relevant literature for given patients, and (iv) comparing medical literature relevance for patients. We describe each of these features in later sections. The rest of this paper is organized as follows. Section II reviews the argumentation-based decision making theory we used throughout this work. Sections III - VI present the four main features of our system, outlined above. Section VII presents the system evaluation conducted with medical literature on brain metastases. Section VIII discusses related works. Section IX concludes.

II. Background

This work relies upon Decision Frameworks and Assumption-based Argumentation (ABA).

Decision frameworks [7] are tuples $\langle \mathtt{D}, \mathtt{A}, \mathtt{G}, \mathtt{T_{DA}}, \mathtt{T_{GA}}, \mathtt{P} \rangle$:
a (finite) set of decisions $\mathtt{D} = \{d_1, \ldots, d_n\}, n > 0$;
a (finite) set of attributes $\mathtt{A} = \{a_1, \ldots, a_m\}, m > 0$;
a (finite) set of goals $\mathtt{G} = \{g_1, \ldots, g_l\}, l > 0$;
a partial order over goals, P, representing the preference ranking of goals;
two tables: $\mathtt{T_{DA}}$, (size $n \times m$), and $\mathtt{T_{GA}}$, (size $l \times m$)[1]:

- for all $d_i \in \mathtt{D}$, $a_j \in \mathtt{A}$, $\mathtt{T_{DA}}[\mathtt{d_i}, \mathtt{a_j}]$ is either:
  1, representing that d_i *has* a_j, or
  0, representing that d_i does not have a_j, or
  u, representing unknown;
- for all $g_k \in \mathtt{G}, a_j \in \mathtt{A}$, $\mathtt{T_{GA}}[\mathtt{g_k}, \mathtt{a_j}]$ is either
  1, representing that g_k is *satisfied* by a_j, or
  0, representing that g_k is not satisfied by a_j, or
  u, representing unknown.

Given a decision framework $DF = \langle \mathtt{D}, \mathtt{A}, \mathtt{G}, \mathtt{T_{DA}}, \mathtt{T_{GA}} \rangle$, a decision $d_i \in \mathtt{D}$ *meets* a goal $g_j \in \mathtt{G}$, wrt DF, iff there exists an attribute $a_k \in \mathtt{A}$, such that $\mathtt{T_{DA}}[\mathtt{d_i}, \mathtt{a_k}] = 1$ and $\mathtt{T_{GA}}[\mathtt{g_j}, \mathtt{a_k}] = 1$. $\gamma(d) = S$, where $d \in \mathtt{D}, S \subseteq \mathtt{G}$, denotes the *set of goals met by* d.

Given a decision framework $edf = \langle \mathtt{D}, \mathtt{A}, \mathtt{G}, \mathtt{T_{DA}}, \mathtt{T_{GA}}, \mathtt{P} \rangle$ the *most preferred decisions* are the decisions meeting the more preferred goals that no other decisions meet, formally defined as follows. For every $d \in \mathtt{D}$, d is most preferred iff the following holds for all $d' \in \mathtt{D} \setminus \{d\}$:

- for all $g \in \mathtt{G}$, if $g \notin \gamma(d)$ and $g \in \gamma(d')$, then there exists $g' \in \mathtt{G}$, such that:
  - $g' > g$ in P, $g' \in \gamma(d)$, and $g' \notin \gamma(d')$.

Assumption-based Argumentation (ABA) *frameworks* [5] are tuples $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$ where

- $\langle \mathcal{L}, \mathcal{R} \rangle$ is a deductive system, with $\mathcal{L}$ the *language* and $\mathcal{R}$ a set of *rules* of the form $\beta_0 \leftarrow \beta_1, \ldots, \beta_m (m \geq 0, \beta_i \in \mathcal{L})$;
- $\mathcal{A} \subseteq \mathcal{L}$ is a (non-empty) set, referred to as *assumptions*;
- $\mathcal{C}$ is a total mapping from $\mathcal{A}$ into $2^{\mathcal{L}} - \{\{\}\}$, where each $\beta \in \mathcal{C}(\alpha)$ is a *contrary* of α, for $\alpha \in \mathcal{A}$.

Given a rule ρ of the form $\beta_0 \leftarrow \beta_1, \ldots, \beta_m$, β_0 is referred to as the *head* (denoted $Head(\rho) = \beta_0$) and $\beta_1, \ldots, \beta_m$ as the *body* (denoted $Body(\rho) = \{\beta_1, \ldots, \beta_m\}$). We focus on *flat* ABA frameworks, with no assumption is the head of a rule.

In ABA, *arguments* are deductions of claims using rules and supported by sets of assumptions, and *attacks* are directed at the assumptions in the support of arguments. Informally, following [5]:

- *an argument for (the claim)* $\beta \in \mathcal{L}$ *supported by* $A \subseteq \mathcal{A}$ ($A \vdash \beta$ in short) is a (finite) tree with nodes labeled by sentences in $\mathcal{L}$ or by τ[2], the root labeled by β, leaves either τ or assumptions in A, and non-leaves β' with, as children, the elements of the body of some rule with head β';
- an argument $A_1 \vdash \beta_1$ *attacks* an argument $A_2 \vdash \beta_2$ iff β_1 is a contrary of one of the assumptions in A_2.

Attacks between (sets of) arguments in ABA correspond to attacks between sets of assumptions, where *a set of assumptions* A *attacks a set of assumptions* A' iff an argument supported by a subset of A attacks an argument supported by a subset of A'.

With argument and attack defined for a given $\mathcal{F} = \langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$, standard argumentation semantics can be applied in ABA [5], e.g.: *a set of assumptions is admissible* (in $\mathcal{F}$) iff it does not attack itself and it attacks all $A \subseteq \mathcal{A}$ that attack it; *an argument* $A \vdash \beta$ *is admissible (in* $\mathcal{F}$*) supported by* $A' \subseteq \mathcal{A}$ iff $A \subseteq A'$ and A' is admissible (in $\mathcal{F}$); *a sentence is admissible* (in $\mathcal{F}$) iff it is the claim of an argument that is admissible supported (in $\mathcal{F}$) by some $A \subseteq \mathcal{A}$.

As shown in [6], [7], [9], ABA can be used to model decision making problems and compute "good" decisions. The ABA framework for computing the *most preferred decisions* in a decision framework is defined as *AF* = $\langle \mathcal{L}, \mathcal{R}, \mathcal{A}, \mathcal{C} \rangle$ for which:

• $\mathcal{R}$ is such that:
for all $d_k \in \mathtt{D}$, $isD(d_k) \leftarrow$;
for all $g_j \in \mathtt{G}$, $isG(g_j) \leftarrow$;
for all $a_i \in \mathtt{A}$, $isA(a_i) \leftarrow$;
for all $g_t, g_r \in \mathtt{G}$, if $g_t > g_r \in \mathtt{P}$ then $prefer(g_t, g_r) \leftarrow$;
for $k = 1, .., n; j = 1, .., m$ if $\mathtt{T_{DA}}[\mathtt{k}, \mathtt{i}] = 1$ then
$hasAttr(d_k, a_i) \leftarrow$;
for $j = 1, ..m; i = 1, .., l$ if $\mathtt{T_{GA}}[\mathtt{j}, \mathtt{i}] = 1$ then
$satBy(g_j, a_i) \leftarrow$;
$met(X, Y) \leftarrow hasAttr(X, Z), satBy(Y, Z), isD(X), isG(Y), isA(Z)$;
$notSel(X) \leftarrow met(X1, Y), notMet(X, Y), notMetBetter(X, X1, Y)$;

[1] $\mathtt{T}[x, y]$ denotes the cell in row labelled x and column labelled y in T.

[2] $\tau \notin \mathcal{L}$ represents "true" and stands for the empty body of rules.

$metBetter(X, X1, Y) \leftarrow met(X, Y1), notMet(X1, Y1), prefer(Y1, Y)$;

• $\mathcal{A}$ is such that:
for all $d_k \in \mathtt{D}$, $sel(d_k)$;
for all $d_k \in \mathtt{D}$ and $g_j \in \mathtt{G}$, $notMet(d_k, g_j)$;
for all $d_k, d_r \in \mathtt{D}, d_k \neq d_r$ and $g_j \in \mathtt{G}$, $notMetBetter(d_k, d_r, g_j)$;

• $\mathcal{C}$ is such that:
$\mathcal{C}(sel(X)) = \{notSel(X)\}$;
$\mathcal{C}(notMet(X, Y)) = \{met(X, Y)\}$;
$\mathcal{C}(notMetBetter(X, X1, Y)) = \{metBetter(X, X1, Y)\}$.

Theorem 1 in [7] sanctions that the aforementioned ABA framework is a sound and complete argumentative computational counterpart for decision making in a way that a decision d is most-preferred iff the argument $\{sel(d)\} \vdash sel(d)$ is admissible.

III. Manipulating Clinical Trials

An essential feature of the system is allowing the users to add new clinical trials and thus improve the selection range for each patient-based medical literature recommendation. The input fields are a collection of representative data and metadata for clinical trials and are classified in two categories: *Clinical Trial Design* and *Patient Characteristics*. The *Clinical Trial Design* category contains administrative fields such as the `Trial ID`, `PMID`, number of arms, clinical trial phase, recruitment area, and trial start/end year, but it also states medical information such as the eligible age/number of metastases/performance status for patients and any excluded histology or extra-cranial disease they may have. The *Patient Characteristics* category includes information for one or more clinical trial arms such as the `Arm ID`, number of patients, performance status range and specific factors such as the percentage of patients with non-small cell lung cancer or the percentage of patients with stable extra-cranial disease.

In order to also accommodate faster (keyboard-only) input we give users the option to type in associated text fields, but the values entered are carefully limited and validated prior to submission to the server (e.g. the possible values for the eligible number of metastases are integers from 0 to 4 and the plus sign meaning more than 4 metastases). Another way by which we improve the input speed is by allowing the users to modify several arm values at once, since many remain unchanged or there are only small differences from one arm to the next. Moreover, in medical literature we may find cases where one or several of the form fields are missing or not applicable, so the fields can also be disabled from the web form. Since it is possible for the `TrialID` and the `PMID` to be missing as well, we added identifiers which are automatically generated and are guaranteed to not be null.

Our users also have the option to edit or delete the clinical trial information they have previously entered into the system. The clinical trials in the system are visible by all users but are only editable by the person who submitted them. Furthermore, the system provides a way for its users to browse literature by supplying the abstract and a direct link to the paper, but it also features an embedded PDF viewer for reading the full text version without leaving the platform.

IV. Performing Criteria Searches

Our collaborators pointed out that while the clinical trials are easily accessible, finding specific information regarding them is not as trivial. Suppose a user wants to retrieve clinical trials with two arms for which the subjects are over sixty years old and at least fifty percent of them have a high Karnofsky Performance Status (KPS) score. Finding the appropriate literature concerning these trials in the old fashioned way would involve a substantial effort on the user's part and would mean browsing through numerous scientific papers. However, since we were already collecting data concerning clinical trials, we can use our database to answer specific queries such as the one we have described.

The system presents the results retrieved for each query in tabular form directly in the web-based UI. The results can be exported into CSV format for offline viewing or manipulation using popular tools such as Microsoft Excel.

V. Computing and Explaining Decisions

The central functionality of RecoMedic is represented by making and justifying patient-based medical literature recommendations, providing transparency through natural language explanations. For our purposes, we decided to use *most preferred decisions* introduced in [7] (see Section II), so our system can reason about patient characteristics according to user-specified preferences on their characteristics.

The possible decisions in our framework ($d_k \in \mathtt{D}$) are represented by the medical literature. The goals ($g_j \in \mathtt{G}$) are the inputs in our system given by the individual patient characteristics that we would like for the literature to reflect. The decisions may have some attributes ($a_i \in \mathtt{A}$) and each goal is satisfied by some attributes according to the tables ($\mathtt{T_{DA}}$ and $\mathtt{T_{GA}}$) that we introduced in Section II. For the current version of our system we focus on four patient characteristics (goals): age, primary disease, number of metastases and performance status score. After consulting with oncologists, we decided to use them as follows: The goal on age can be satisfied by two possible attributes: over eighteen years old (adult) and under eighteen years old (minor) which are extracted directly from the *EligibleAge* parameter of each clinical trial. The goal on primary disease can be satisfied by four possible attributes: breast, lung, testicular and renal cancer which are deduced from the corresponding recruit percentages in the clinical trial arms depending on whether these pass a certain threshold (we chose 60%). The goal on number of metastases can be satisfied by

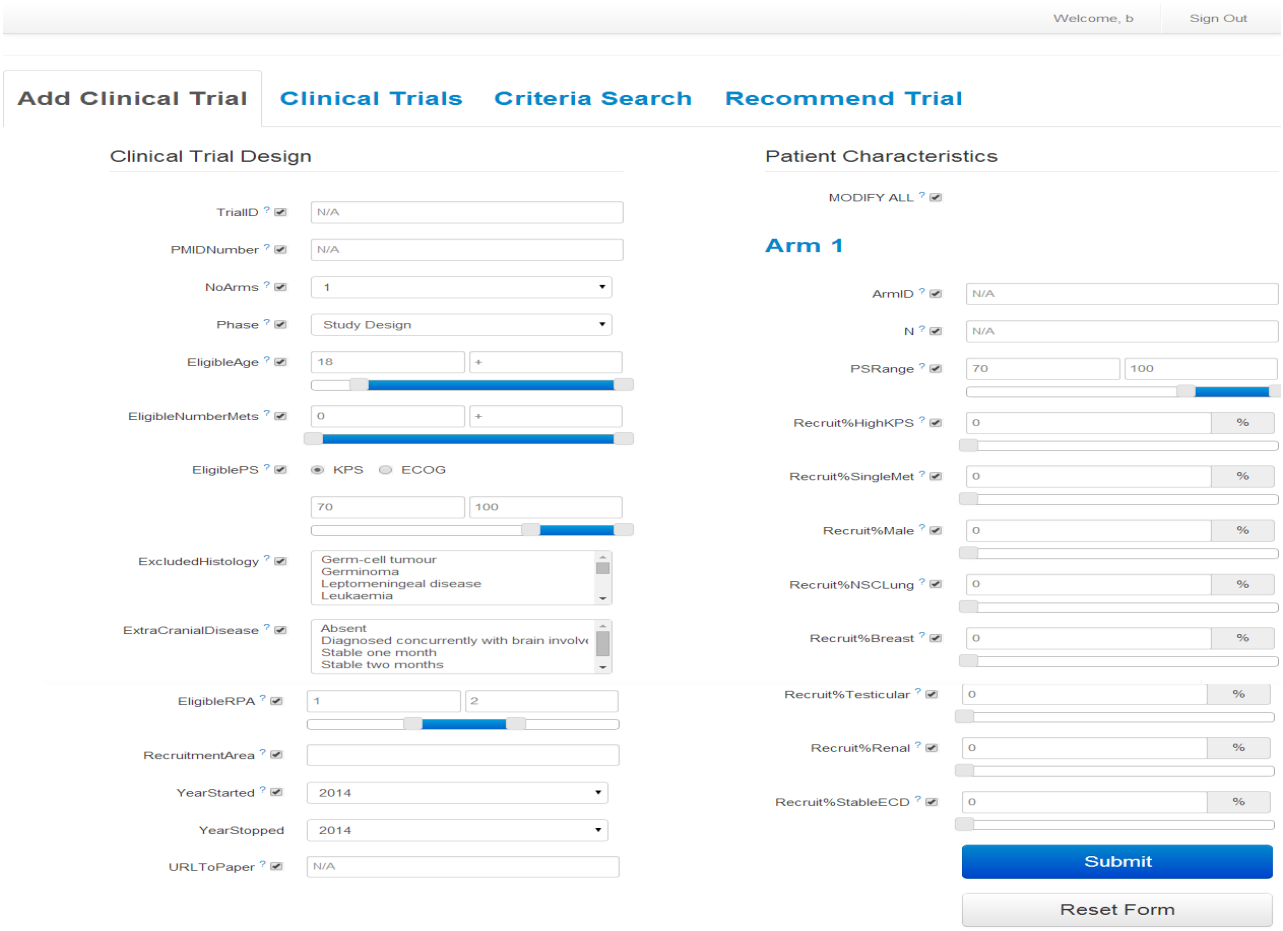

Figure 3. Adding a new Clinical Trial

four possible attributes: no mets, one met, two mets, over two mets which are extracted from the *EligibleNumberMets* parameter of clinical trials. The goal on performance status score can be satisfied by four possible attributes: ECOG PS zero or one, ECOG PS two, ECOG PS three or four extracted from the *EligiblePS* parameter of clinical trials. The performance status scores expressed in KPS scale have been converted to their ECOG equivalent. We have created two user classes for RecoMedic. The first category is represented by healthcare specialists who need not concern themselves with the specifics of ABA. Indeed, we use ABA for recommending the appropriate literature, but we display the recommendation and explanation in natural language as can be seen in Figure 4. The natural language explanation is created by examining the debate graph produced by our computation engine corroborated with facts extracted directly from the generated ABA framework. It presents the patient characteristics according to the user specified ranking and states whether the recommended medical literature matches those certain patient characteristics. The phrasing includes both the user selection and the clinical trial value for that respective attribute, in a manner that is clear and concise.

Figure 4. Recommending a Clinical Trial and Justifying the Decision (Healthcare Specialist View)

Our second category of users, computer scientists, can dive into the intrinsics of ABA by viewing the exact ABA framework that is generated for each query, they can view the debate graph (sanctioning an argument as admissible)

produced by our computation engines (*proxdd* and *grapharg* [10][3]). A sample graph that results from executing a query for this user class is illustrated in Figure 5. Here we selected a 55 year old patient with lung cancer with 3 metastases and performance status (ECOG) 3, with the preferences ranked as follows: primary disease, number of metastases, age and performance status. Our graph shows how the proponent in favor of selecting paper 40 successfully defends against the attacks of the opponent suggesting other papers. For example, the opponent's argument that paper 40 does not address patients with 3 metastases, or that paper 18 addresses a higher preference goal, is refuted by the proponent who, based on the facts in the knowledge base, can state that paper 40 addresses patients with 3 metastases. The bottom three arguments by the opponent in yellow are not attacked because they have been previously defeated by the proponent.

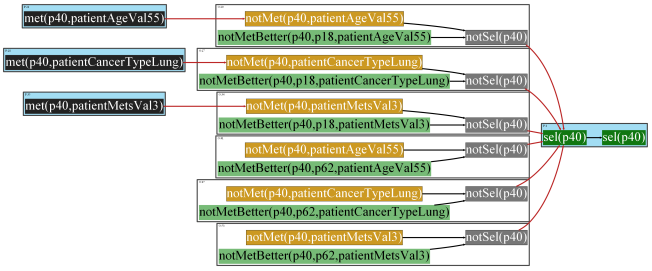

Figure 5. Rendered ABA Debate Graph (Computer Scientist View)

VI. Comparing Decisions

Besides recommending a certain piece of medical literature according to the characteristics of a given patient, we introduce a pairwise comparison feature depicted in Figure 6 that examines two pieces of literature at a time and provides a similarity score based on the selected patient characteristics and the ranked user preferences. Since satisfying the higher ranked goal is more important than satisfying all the goals which are ranked lower, we score the decisions as follows:

$$score_{d_k} = \sum_{j=1}^{n} met(d_k, g_j) \times 2^{n-rank(g_j)}$$

Here, n is the number of ranked preferences and $rank(g_j)$ is an integer between 1 (if g_j is the top preference) and n (if g_j is the last ranked preference). This means that the score of clinical trial d_k (decision) is equal to the weighted sum of the goals that it meets, i.e. where $met(d_k, g_j)$ is equal to 1. If a goal is not met by a decision, then $met(d_k, g_j)$ equals 0 and that term is ignored. We derive a simple similarity score between two decisions using their individual scores. Thus for two decisions d_{k_1} and d_{k_2} the similarity score is:

$$similarity_{d_{k_1}, d_{k_2}} = |score_{d_{k_1}} - score_{d_{k_2}}|$$

[3]Available at http://www.doc.ic.ac.uk/ rac101/proarg/.

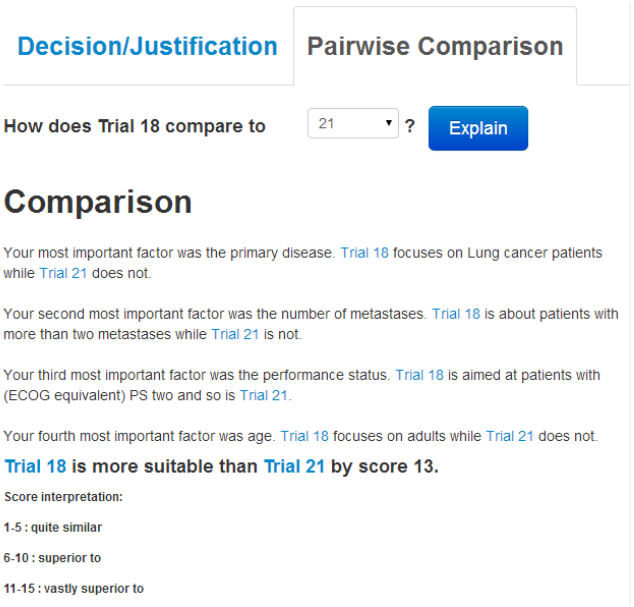

Figure 6. Comparing Clinical Trial Suitability based on Patients

Using the similarity score we can assess how much alike or different are the two given clinical trials. In our case, we are using four patient characteristics and the similarity score can be in the following ranges: **1-5** (the two clinical trials are similar), **6-10** (one clinical trial is superior to the other), **11-15** (one clinical trial is vastly superior to the other).

VII. Evaluation

For the evaluation, we have identified 11 randomized clinical trials on the treatment of brain metastases. The decisions of our model are choices to use a given paper in a diagnosis. Based on the preliminary medical literature data, we have conducted a survey[4] on user experience. A small group of postgraduate medical students from Southern Medical University participated in the evaluation.

The overall impression of the system is positive as 100% of the users believe that RecoMedic serves a genuine need of doctors in searching for clinical trials for patients. 78.75% users state that they easily and fully understand the purpose and functions of RecoMedic. On a scale from 0 to 10, the users rank the UI design of My Clinical Trail, Recommend Trails, Criteria Search and Add Clinical Trial 7.63, 7.44, 7.4 and 7.1, resp. Specifically, none of the users has considered that the overall UI design is complicated.

VIII. Related Work

Our work sits at the intersection between existing work on modeling clinical trials, and logic-based approaches to reasoning with clinical knowledge, and specifically clinical trials. Much of the existing work on modeling clinical trials has been based around defining inclusion criteria to allow for (semi-)automated matching of patients with clinical

[4]Using the online platform available at http://www.sojump.com.

trials [11][5], and has largely been based on ontological approaches. Although this works well for checking whether a patient belongs to an ontological category, it does attempt to model the degree to which different trials which have been conducted match different patients, based on the patients that they actually did enroll (as opposed to the notional inclusion criteria). However, we note that the encoding of trial data is compatible with suggested approaches, such as the "Human Studyome" project[6], although our approach to reasoning goes beyond their aims of knowledge encoding. The BioBIKE project[7] aims to capture and reason with biomedical data, but their focus is on biological and genomic data rather than clinical trials.

Other work that has focused on clinically-orientated reasoning includes many systems that provide medical knowledge representation and reasoning. However, many of these systems require specialist encoding of medical knowledge (e.g. [12], [13]) and hence divorce the clinicians from the underlying clinically-produced knowledge. The work that is closest to ours also works with encoded clinical trial data ([14],[15]), but focuses on reasoning with the trial results in order to decide on the best treatment. The work we have presented here, which focuses on selecting the most appropriate paper, is separate to this, and could be used as an input to their work, providing preferences (in this case based on "paper appropriateness") over conflicting trials.

IX. Conclusions

In this paper, we have presented an online medical literature recommendation system, RecoMedic. Through its easy to use web interface, RecoMedic allows users, primarily doctors, to search through its medical literature repository with patient information. This search is realized with argumentation such that not only the most relevant medical article is identified, but also an explanation to this selection is provided. Standard data repository maintenance features are also supported in RecoMedic, including adding, searching through, and comparing medical papers.

A preliminary user evaluation has confirmed that RecoMedic addresses a genuine need for doctors. The evaluation also confirms that users are able to use RecoMedic with little help and approve its UI design.

In the future, we would like to experiment RecoMedic with more users and extend its features to include not only recommending most relevant articles, but also most suitable treatments. We would also like to explore other search and ranking strategies and explore the use of argumentation in these strategies. We also would like to explore our approach to medical complications other than brain metastases.

Acknowledgements. This work was supported by the strategic grant POSDRU/159/1.5/S/133255, Project ID 133255 (2014), co-financed by the European Social Fund within the Sectorial Operational Program Human Resources Development 2007-2013.

The first author was supported by the aforementioned grant. The other authors were supported by the EPSRC project *Transparent Rational Decisions by Argumentation* : EP/J020915/1.

[5]Available at http://bioportal.bioontology.org/ontologies/OCRE.

[6]Available at http://rctbank.ucsf.edu/.

[7]Available at http://biobike.csbc.vcu.edu/.

References

[1] P. Yildirim, L. Majnaric, O. Ekmekci, and A. Holzinger, "Knowledge discovery of drug data on the example of adverse reaction prediction," *BMC Bioinformatics*, vol. 15, no. Suppl 6, p. S7, 2014.

[2] T. J. M. Bench-Capon and P. E. Dunne, "Argumentation in artificial intelligence," *AIJ*, vol. 171, no. 10-15, pp. 619–641, 2007.

[3] P. Besnard and A. Hunter, *Elements of Argumentation*. The MIT Press, 2008.

[4] I. Rahwan and G. R. Simari, *Argumentation in Artificial Intelligence*. Springer, 2009.

[5] P. M. Dung, R. A. Kowalski, and F. Toni, "Assumption-based argumentation," in *Arg. in AI*. Springer, 2009, pp. 199–218.

[6] X. Fan and F. Toni, "Decision making with assumption-based argumentation," in *Proc. TAFA*, 2013.

[7] X. Fan, R. Craven, R. Singer, F. Toni, and M. Williams, "Assumption-based argumentation for decision-making with preferences: A medical case study," in *Proc. CLIMA*, 2013.

[8] X. Fan, F. Toni, A. Mocanu, and M. Williams, "Multi-agent decision making with assumption-based argumentation," in *Proc. AAMAS*, 2014.

[9] P. A. Matt, F. Toni, and J. Vaccari, "Dominant decisions by argumentation agents," in *Proc. ArgMAS*. Springer, 2009.

[10] R. Craven, F. Toni, and M. Williams, "Graph-based dispute derivations in assumption-based argumentation," in *Proc. TAFA*, 2013.

[11] S. W. Tu, S. Carini, A. Rector, P. Maccallum, I. Toujilov, S. Harris, and I. Sim, "Ocre: an ontology of clinical research," in *11th International Protege Conference*, 2009.

[12] B. G. Buchanan and E. H. Shortliffe, *Rule Based Expert Systems: The Mycin Experiments of the Stanford Heuristic Programming Project*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1984.

[13] M. A. Grando, D. Glasspool, and J. Fox, "A formal approach to the analysis of clinical computer-interpretable guideline modeling languages," *AIM*, vol. 54, no. 1, pp. 1–13, Jan. 2012.

[14] A. Hunter and M. Williams, "Aggregating evidence about the positive and negative effects of treatments," *AIM*, vol. 56, no. 3, pp. 173–190, Nov. 2012.

[15] N. Gorogiannis, A. Hunter, and M. Williams, "An argument-based approach to reasoning with clinical knowledge," *IJAR*, vol. 51, no. 1, pp. 1 – 22, 2009.