Using Arguing Agents to increase the Human Organ Pool for Transplantation

P. Tolchinsky [*] [‡], **U. Cortés**[†] [‡], **J.C. Nieves** [‡], **F. Caballero**[§] and **A. López-Navidad** [§]
Technical University of Catalonia, Barcelona, Spain[‡]
{tolchinsky, ia, jcnieves}@lsi.upc.edu
Hospital de la Santa Creu i Sant Pau, Barcelona, Spain [§]

Abstract

In this paper we describe how a multi agent system (MAS) that supports arguing agents can increase the number of human organs for transplantation. We analyze the current human organ selection process to then propose an alternative which allows discarding less human organs that can safely be transplanted. The proposed human organ discarding process is framed in CARREL, an agent-based organization designed to improve the overall transplant process.

1 Introduction

Human organ transplantation constitute the only effective therapy for many life-threatening diseases. While becoming a commonplace medical event there is a growing disparity between the demand for and the supply of organs for transplantation. Despite this disparity a great percentage of human organs, available for transplantation, are discarded as being considered non-viable for that purpose. Given the importance of this issue, much effort is devoted in finding ways to reduce this gap between demand and supply.

In this paper we present a novel discarding process that relaying on arguing agents can safely allow less human organs to be discarded and thus to augment the organ pool. We describe this new process framed in the CARREL System [Vázquez-Salceda *et al.*, 2003], an agent-based organization designed to improve the overall transplant process of human organs and tissues. CARREL is intended for the tasks involved in managing the vast amount of data to be processed in carrying out:

- recipient selection (e.g. from patient waiting lists and patient records)
- organ/tissue allocation (based on organ and tissue records)
- ensuring adherence to legislation
- following approved protocols
- preparing delivery plans (e.g. using, say, train and airline schedules)

CARREL's behavior could briefly be described as an *agency* that when receiving a tissue request from one hospital tries to allocate the best tissue available from all the known resources (Tissue Banks). And when a donor is detected it tries to allocate each of the transplantable donor's organs to the most appropriated recipient. In this agency different entities (the agents) play different roles that are determined by their goals, rights and duties (see [Vázquez-Salceda *et al.*, 2003]). The relative scarcity of donors has led to the creation of international coalitions of transplant organizations. This new, more geographically distributed, environment makes an even stronger case for the application of distributed software systems, but with the added complication of the necessity to accommodate a complex set of, in some cases conflicting, national and international regulations, legislation and protocols governing the exchange of organs and tissues. It is the last point which underpins our case for the use of so-called electronic organizations whose purpose is to provide over-arching frameworks for interaction in the same way as organizations, or equivalently, collections of social norms, do in the physical world.

In addressing the human organ discarding process two issues must be taken into account i) many criteria for accepting organs do not refer only to the organ and the donor themselves but also to the recipient to which the organ is intended ii) the organ acceptability criteria vary very rapidly and disagreement among the criteria exist. That is, different hospitals and physicians may follow different acceptability criteria. As a consequence, the more potential recipients are considered at the time of evaluating the viability of an organ, the more likely is for that organ to be considered as viable. This motivates the use of agents to effectively distribute the offer of an organ to all the pertinent hospitals, as well as managing the response of each of the contacted hospitals. To handle the possible inconsistency that may raise from contradictory criteria or policy, agents use arguments to represent the acceptability criteria and the relation among these criteria. Thus, the proposed discarding process is managed by agents that exchange arguments.

In §2 we describe CARREL and its context of application.

[*] The first author of this paper is a PhD student

[†] P. Tolchinsky, U. Cortés and J.C. Nieves would like to acknowledge support from the EU Funded project FP6-IST-002307. The views expressed in this paper are not necessarily those of ASPIC consortium.

In §3 we describe the discarding process as it is today and after highlighting some of its critical aspects we present in §4 the novel discarding process framed in CARREL. In §5 we describe the argumentation framework that enables representing the acceptability criteria. To then describe the arguments evaluation process (§6) and how argumentation can help refining the human organ acceptability criteria in §7. In §8 we present our conclusions and future work.

2 The CARREL System

Since 1980 the number of transplant requests has been constantly increasing. As a consequence, the human transplant coordinators are currently facing significant problems in dealing with the volume of work involved in the management of requests, assignation and distribution of tissues and organs. Given the constat progress in transplant-based therapies and the relative success of these therapies, the demand for organs and tissues are expected to raise even more. A review of the coordinator's role and the difficulties faced is presented in [Lopez-Navidad *et al.*, 1997].

Two aspects can be highlighted that make transplantation management a very complex issue: (i) the need to maximize the number of successful transplants due to the scarcity of donors (ii) the complexity of the donor/recipient matching due to the diversity and multiplicity of genetic factors involved in the response to the transplant. The CARREL System is intended to automate many of the tasks that now a days are carried out by human beings. CARREL's design takes the Spanish and Catalan transplant organizations as references, both known to be examples of best practice. Hence, they constitute viable physical institutions on which to base electronic ones. The Spanish organizational model has two levels of action:

Intra-hospital: Where the role of hospital Transplant Coordinator was created to improve the coordination of all the people working at any step of the donor procurement, allocation and transplantation process.

Inter-hospital: Where an intermediary organization and the Organització CATalana de Trasplantaments (OCATT) for Catalonia, Organización Nacional de Transplantes [ONT,] (ONT) for the whole of Spain was created to improve the communication and coordination of all the participating health-care transplant organizations, namely hospitals and tissue banks.

Fig. 1 depicts the inter-hospital level managed by CARREL in which we can identify the entities that interact with the CARREL system. Each TB denotes a tissue bank, each UCTx denotes a transplant coordination unit, the agency that represent a hospital member of CARREL. The ONT and OCATT denote the organ transplantation organizations that own the agent platform and act as observers.

The role of the CARREL Institution can be summarized in terms of following tasks:

T1 to make sure that all the agents which enter into institution behave properly (that is, that they follow the behavioral norms).

T2 to be up to date about all the available pieces in the Tissue Banks, and all the recipients that are registered in the waiting lists.

T3 to check that all hospitals and tissue banks fulfill all the requirements needed to interact with CARREL.

T4 to take care of the fulfillment of the commitments undertaken inside the CARREL system.

T5 to coordinate the piece delivery from one facility to another.

T6 to register all incidents relating to a particular piece.

A hospital becomes a member of the CARREL institution in order to make use of the services provided. In doing so, they accept to respect the norms that rule the interaction inside CARREL. Some of these norms are:

N1 All organ offers and tissue requests should be done through the CARREL institution.

N2 Hospitals must accept the outcomes of the negotiation (assignation) process.

N3 Hospitals receiving an organ or tissue from CARREL must update the institution with any relevant event related to these organs and tissues.

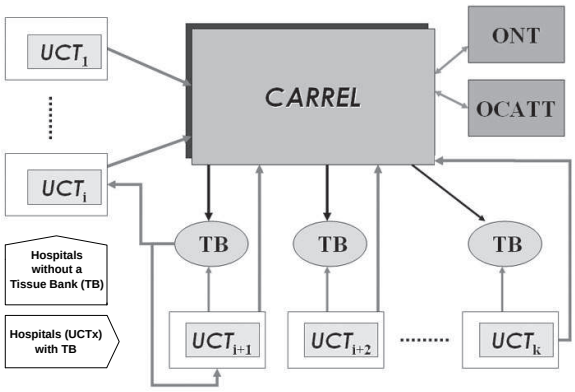

Figure 1: CARREL: An Agent Mediated Organization for Tissue and Organ Allocation.

Each hospital member of CARREL is represented by the Transplant Coordination Unit (UCTx) that manages the intra-hospital level. Each UCTx goal is to successfully culminate any organ and tissue procurement, extraction and implantation process. Each UCTx is modelled as an agency, the roles the different agents play in this agency is presented in [**?**]. We will identify two agents in the UCTx agency and describe their role in the light of the organ discarding process. The TCAx is the Agent representing the Transplant Coordinator of a UCTx. The Transplant Coordinator is responsible of procuring and offering the organs of the potential donors. Based on the potential donor's medical history and tests, the Transplant Coordinator must determine which organs are *viable* and which are *non-viable* for transplantation. The TUAx is the Agent representing the Transplant Unit, they are responsible of the potential recipients, their duty is to extract

the organs from the donors and successfully transplant them to their patients.

3 The Human Organ Discarding Process

From the moment a potential donor is detected until the moment his organs are transplanted there is a filtering process in which the different stakeholders may decline to transplant or to offer for transplantation each of the donor's organs considered transplantable, viz. heart, lungs, liver, pancreas and kidneys (in fact, we only consider solid organs).

The process begins when the Transplant Coordinator (TC) detects a potential donor, in which case, after properly analyzing his characteristics, she informs the OCATT, assuming the TC is located in Catalonia, about the organs she considers viable for transplantation. If the TC is aware of any potential recipient that could match one of the donor's organs, she may consult or even delegate that decision to the professionals in the Transplant Unit (TU) who are responsible of that potential recipient. This should be done before informing the OCATT and it normally happens when the recipient and the donor are located in the same hospital. If the TC considers the organ as viable but its characteristics do not meet the local policy criteria, or no match for the available organ is found among the potential recipients in the waiting lists of Catalonia, the OCATT will offer the organ to the ONT. Otherwise, if a recipient is found, the allocation process takes place and the organ is offered to a Catalan TU that may or may not accept the organ [1]. If refused, the organ will be offered to a different TU until final acceptance or refusal. The TU that accepts the organ has the right to discard it after or during the extraction operation, in which case it is very unlikely to have the organ transplanted. If no TU accepts the organ, it is offered to the ONT.

When an organ is offered to the ONT, a similar process takes place, this time however embracing Spain and not only Catalonia. In case of refusal, the OCATT will offer the organ to transplant organizations in Europe. If these organizations refuse the organ, the OCATT will then offer it in Asia. However, this last step hardly ever occurs. If every organization fails to allocate the organ, the organ will not be retrieved from the donor.

Currently in Catalonia, a world leader in transplantation, between 15 and 20 percent of the livers, 20% of the kidneys, 60% of the hearts, 85% of the lungs and 95% of the pancreas, from donors that were detected, are discarded [OCATT].

Given the gap between organ demand and supply and taking into account how critical a transplant operation is for many people, any attempt to reduce the rate of discards is unquestionable. Even though a drastic improvement in the ratios may not be realistic, any progress will necessarily imply improving the life quality of many people and saving the life of many others.

In what follows we present an alternative algorithm for addressing three important issues: a) to decrease the rate of discards, b)to provide a framework for dealing with a growing number of stakeholders that may follow different policies and medical criteria c) to refine the criteria for accepting human organs for transplantation.

[1] It is worth mentioning that at this stage the offered organ has not yet been extracted. It is after accepting it that a Transplant Unit may extract the organ from the donor

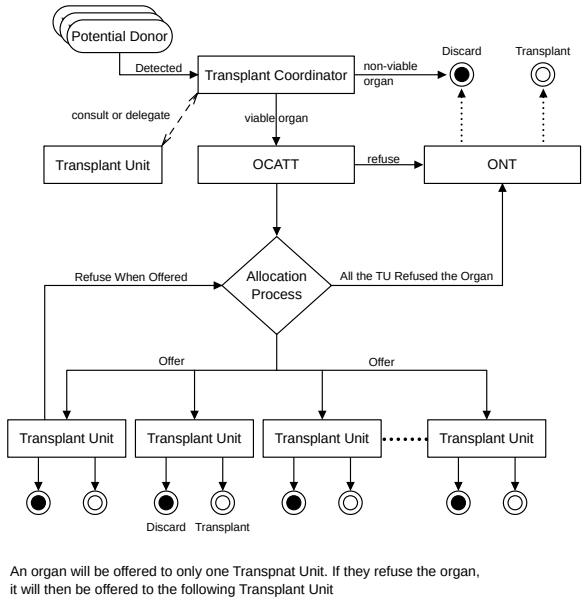

Figure 2: Flow of the Current Organ Discard Process.

4 The new discarding process

The agent-based CARREL platform enables inexpensive communication among multiple parties. Much information exchange that is currently done by time and staff consuming phone calls and faxes could be done much more efficiently in terms of time, money and organization, using a multi-agent based platform. Once this is acknowledged, changes in the conception of the transplant process come alone. This is what we are aiming to do with CARREL.

Our proposal for changing the discard process is to distribute the decision of refusing an organ. Namely, instead of having one TC, or at best a TC with the advice of members of a TU, deciding the viability of an organ, we suggest to enable every TU to be an active part of the decision process in order to warrant minimal waste of scarce resources (organs). The idea is that, without undermining the TC's assessment, even if she considers the organ as non-viable, if there is at least one TU providing valid reasons for transplanting it, this organ should not be discarded without being previously offered to them. An important motivation for this new process is that organs are rarely non-viable or ideal *per se*. The term ideal organ should imply an integral concept that involves donor and recipient characteristics and all the procedure performances between both [López-Navidad and Caballero, 2003].

In the proposed procedure, as soon as the TC detects a potential donor, and after having analyzed his characteristics, the TC will produce a justification for each of the transplantable organs supporting her belief to whether these organs are viable or not. These justifications will arrive to each and every TU member of CARREL that have a potential recipient

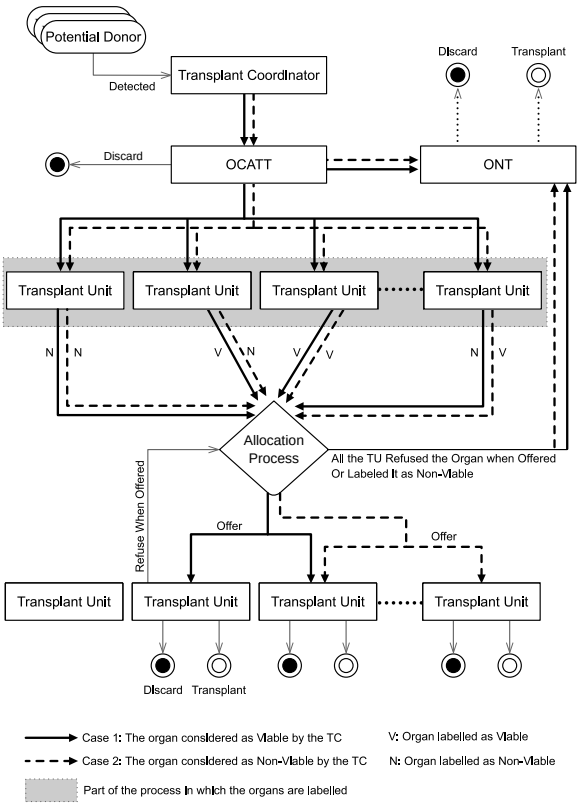

Figure 3: Flow of the Proposed Organ Discarding Process. As we can see, organs initially deemed as non-viable by the Transplant Coordinator (dashed lines) can end up being offered to a Transplant Unit and subsequently being transplanted. Note that an organ is not offered in parallel to the TUs. It is only after a TU refuses the organ that it is offered to the following TU.

that matches one of the organs' characteristics. The TU may decide to transplant an offered organ even if it is against the TC's assessment. If this is the case, and their reasons to transplant the organ are valid, they, or other TUs that also claim the organ as viable, should be able to transplant it.

We now explain how this new process is managed by the CARREL multi-agent system, focusing on how the exchange of justifications among the agents takes place. We also show how and on what basis a justification can be considered as valid. We finally present some ideas on how this proposed process, addressed in the manner we describe in this paper, enables not only to safely and efficiently accept more organs but also to constantly update the human organ acceptability criteria. In [Cortés *et al.*, 2005] other aspects of the process are covered, such the argument generation and reception of the TCs and TUs.

4.1 Discarding an Organ in CARREL

The discarding process in CARREL is carried out by different agents that play different roles. To simplify the description we will only point out the agents whose participation is more relevant for the proposed process, while omitting the agents that play secondary roles. We also omit matters concerning security measures, such as protection of privacy that, although crucial for the applicability of the procedure are expendable to understand the overall process.

In UCTx, the agency that represents a CARREL affiliated hospital; we can identify the Transplant Coordinator Agent (TCAx) and the Transplant Unit Agent (TUAx). Among their tasks, they are responsible of sending and receiving the arguments to and from the TC and TU. The Mediator Agent (MA) that belongs to CARREL is in charge of evaluating the interchanging arguments. The OCATT Agent (OA) will play the transplant organization role in CARREL. For simplicity we will use OA to name the agents that represent any of the transplant organizations, such as OCATT or ONT given that their role in CARREL is essentially the same.

The process will start as it currently does; the TC detects a potential donor and determines which of the transplantable organs are viable and which are not. The TC, located at UCT_i, will provide TCA_i with a justification to why she believes an organ should be considered as viable or not, this is done for each organ. At this stage of the transplant process CARREL considers the organs of a donor as independent. TCA_i will carry the information of an organ to CARREL; this information contains the justification produced by TC as well as the organ's and donor's characteristics, such as the organ type, the organ size, the donor's blood type and the donor's age, etc. Once TCA_i enters CARREL, having passed the security protocols, it enters the transplant organization room where it meets a OA (see fig. 4), representing in this case the OCATT. OA, only on the basis of the organ characteristics will determine whether the organ meets the local policy criteria, in this case the Catalan policy, and if it does, it checks whether the organ characteristics matches any of the potential recipient's needs. If so, the organ is accepted, otherwise, OA will derive TCA_i to an agent representing the following transplant organization, in this case the ONT. This new agent will play the same role but with the difference that the organ discarding policies may vary, and the potential recipients waiting list are different. If all the organizations fail to accept the organ, the organ is discarded. Otherwise OA, sends TCA_i to the evaluation room. If the organ offered by TCA_i cannot be transplanted under no circumstances, for instance if the organ has a malignant tumor, the organ is discarded at the first instance by OA.

In the evaluation room TCA_i meets MA that will send a notification to all the UCTx that have potential recipients waiting for an organ with the same characteristics as the offered organ. Each notified TU, in UCT_j, will send a TUA_j that will provide MA with a justification to their decision indicating why they consider they should accept or refuse that organ. Their justification is built as a response to TCA_i's justification.

If both TCA_i and TUA_j agree, that is, they both consider the organ either viable or non-viable, MA accepts their decision. But if they disagree, MA evaluates TUA_j's arguments, on the basis of TCA_i's arguments, and if it accepts them, it is TUA_j's decision that prevails, otherwise, it is TCA_i's decision which prevails.

This should be happening simultaneously with all the TUs

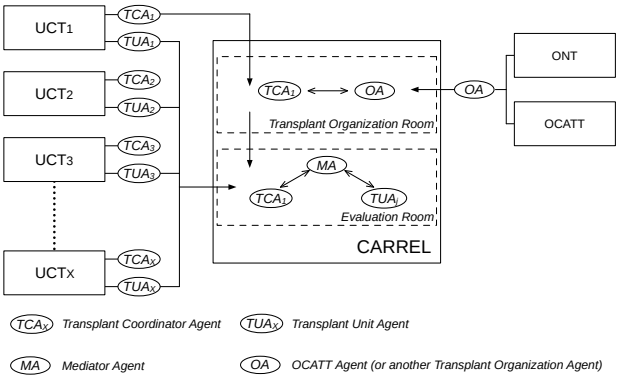

Figure 4: The new Discard Process managed by CARREL.

and for all the transplantable organs of the potential donor. Hence, for each TUA_j, after MA's evaluation, a given organ can be labelled as viable or non-viable depending on the arguments they have provided. In particular, an organ initially offered as non-viable, can be labelled as viable by a TUA_j. In the current discarding process, this organ would have never been offered, preventing many potential recipients from the possibility of benefiting from it. The proposed process not only enables augmenting the human organ pool, but it also has an effect on the allocation process, since TUAs that have an organ labelled as non-viable will not be considered when deciding to which TUA_j to assign this organ.

TUAs to which an organ has been labelled as viable are committed to accept the organ and to successfully transplant it. If a committed TUA_j fails to accept or successfully transplant the organ it will have to justify its decision or action to CARREL. CARREL's policy is to promote the transplantation of as many organs as possible as long as it is safe, i.e. as long as the organs are viable. Thus, any refusal to transplant an organ should be justified. Also, transplant operation must be safe, it is worth noticing that when a transplant operation failure occurs, not only the recipients health is jeopardized, but also the unsuccessfully transplanted organ will most probably be wasted, preventing another potential recipient from benefiting from it. Hence, any failure in the transplant operation must also be justified.

Based on the decisions and actions taken by each TUA_j, as well as on the arguments given to justify them, MA updates a model representing each TUA_j's reputation, in which TUAs with good reputation have usually fulfilled their commitments, thus, have a record of accepting the organs and successfully transplant them, while TUAs with bad reputation have a record of breaking their commitment. Providing valid justifications when breaking commitments helps improving the agents' reputations (while helping to improve the understanding of the domain).

It is worth mentioning that there can be several valid reasons for breaking a commitment. For instance, a TUA_j that initially claimed an organ to be valid, may retract from its claim because the potential recipient to whom the organ was intended had suddenly developed fever[2]. If this were to happen, this TUA_j will be committed to provide CARREL with the appropriate justification.

In the following section we address the representation of the agents' justifications. We start by motivating the use of argumentation for managing the organ acceptability criteria to then describe the argumentation framework that enables agents to handle inconsistency. We illustrate the use of argumentation in some relevant TUA_j-TCA_i interactions.

5 Representing the agents' justifications

In the medical domain, and in particular in the human organ transplantation field, qualified professionals disagree. What may be a sufficient reason for discarding an organ for some qualified professionals may not be for others. Different policies in different hospitals and regions exist. To have a consensus among medical professionals is not feasible and might even be regarded as counterproductive [Kljakovic, 2003]. For that reason we believe that the **K**nowledge **B**ase that represents the human organ **A**cceptability **C**riteria (ACKB) must allow inconsistency, as we want, for instance, MA to be able to derive from the ACKB that two conflicting justifications are acceptable. Namely, we want MA to be able to evaluate the given justifications taking into account each TUAx's acceptability criteria, without undermining CARREL's own acceptability criteria.

Take for example the following scenario: a TCA_i offering a lung of a donor with a smoking history of 40 pack-year, with no chronic obstructive pulmonary diseases and to which no pulmonary fibrosis was found in the chest radiograph (no *COPD*, for short). The TCA_i, based on the facts the donor has no *COPD*, decides that the donor's lung is *viable*. This lung is offered to a TUA_j that decides to reject the organ claiming its *non-viability* based on the donor's smoking history (let us call this the `smoking history` argument). The MA will accept TUA_j's justification given that general guidelines suggest discarding lungs were of donors with a smoking history of more than 20-30 pack-year [OCATT]. Let us suppose this time that the TCA_i, based on the donors smoking history, offers the organ as being *non-viable*. A different $\text{TUA}_{j'}$ may then claim the organ to be *viable* for transplantation as the donor has no *COPD* (let us call this the `no COPD` argument). MA should also accept $\text{TUA}_{j'}$'s justification given that there are qualified physicians that believe that "*any kind of smoking history could be acceptable unless there is a chronic obstructive pulmonary disease or pulmonary fibrosis on the chest radiograph*" [López-Navidad and Caballero, 2003].

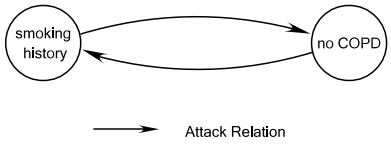

Figure 5: Representing Conflicting Acceptability Criteria

[2]Transplant operation must not take place on patient with fever.

To represent the interaction among the arguments we relay on Dungs' Argumentation Framework (AF) extended with the preference relation between arguments ($>>^{Pref}$)[3]. Dung's AF is defined in [Dung, 1995] as:

Definition 1: An argumentation framework is a pair $AF =< AR, attack >$, where AR is a set of arguments, and $attack$ is a binary relation on AR, i.e. $attack \subseteq AR \times AR$. If $A, B \in AR$ then $A\ attack\ B$ should be read as, argument A attacks argument B

At the moment we have the two arguments `no COPD`, `smoking history` $\in AR$ that attack each other, `no COPD` $attack$ `smoking history` and `smoking history` $attack$ `no COPD`.

Definition 2: A set S of arguments is said to be conflict-free if there are no arguments $A, B \in S$ such that $A\ attack\ B$.

Let us consider also the argument `is viable` claiming the viability of an organ. We can assume that when no contraindications are found, an organ should be considered as viable. If we add this new argument to the existing ones, we have this new relation `smoking history` $attack$ `is viable`. In this case, the set of arguments $S = \{$`is viable`, `no COPD`$\}$ is conflict-free, also the set $S' = \{$`smoking history`$\}$. In this new scenario the argument `no COPD` defends argument `is viable` from `smoking history`. In other words, arguing that the lung *is viable* despite of the donor's *smoking history*, given that *no COPD*, is an acceptable position for a TUA_j

Definition 3: A set S of arguments *attacks* an argument $A \in AR$ if there exist an argument $B \in S$ such that $B\ attack\ A$.

Definition 4: (1) An argument $A \in AR$ is $acceptable$ with respect to a set S of arguments if and only if for each argument $B \in AR$: If $B\ attacks\ A$ then B is $attacked$ by S. (2) A conflict-free set of arguments S is $admissible$ if and only if each argument in S is $acceptable$ with respect to S.

If we extend the original Dung's AF to include the preference relation between arguments (A is preferred to B, $A >>^{Pref} B$) a new relation can be defined. An argument A *defeats* an argument B, if A *attacks* B and it is not the case that $B >>^{Pref} A$. In this extended argumentation framework ($eAF =< AR, attack, >>^{Pref} >$) we redefine the acceptability and admissibility properties:

Definition 5: (1) An argument $A \in AR$ is $acceptable$, in eAF, with respect to a set S of arguments if and only if for each argument $B \in AR$: If $B\ defeat\ A$ then B is $defeated$ by S. (2) A conflict-free set of arguments S is $admissible$, in eAF, if and only if each argument in S is $acceptable$, in eAF, with respect to S.

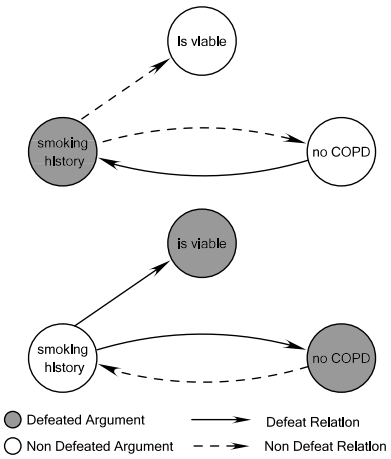

Figure 6: Two possible derivation from the set of arguments $AR_{(TCA_i, TUA_j, O)}$

In this case where no argument is preferred to another, `is viable` is acceptable with respect to S (also w.r.t $\{$`no COPD`$\}$). In this same context, S and S' are admissible sets. But if `no COPD` were to be preferred to `smoking history` then S' would no longer be admissible.

Let us call $AR_{(TCA_i, TUA_j, O)}$ the set of all arguments given by TCA_i and TUA_j when evaluating the viability of the organ O. The initial idea is that MA should accept TUA_j's justification if the set of arguments given by TUA_j, say $\Gamma_{(TUA_j, O)}$, is an admissible set in $AR_{(TCA_i, TUA_j, O)}$. Therefore, if TUA_j claims the organ O to be valid, then `is valid` $\in AR_{(TCA_i, TUA_j, O)}$, whereas if it claims the organ to be non-viable, the argument `is valid` is defeated by $AR_{(TCA_i, TUA_j, O)}$. In our example $AR_{(TCA_i, TUA_j, O)}=$ $\{$`is viable`, `no COPD`, `smoking history`$\}$ both $\Gamma_{(TUA_j, O)}=$ $\{$`is viable`, `no COPD`$\}$ and $\Gamma_{(TUA_j, O)}=$ $\{$`smoking history`$\}$ are admissible sets, thus acceptable positions.

Take now four TUAs to which the lung of the previous example has been offered.

TUA_1: The recipient has Hepatitis C.

TUA_2: The recipient's survival is precarious.

TUA_3: The recipient's survival is precarious and has Hepatitis C.

TUA_4: The recipient has no particularities.

[3]Many other works have already considered the extension of Dung's AF to include preference between arguments. For instance [Amgoud and Cayrol, 1998] and [Amgoud and Cayrol, 2002] from which the $>>^{Pref}$ notation is taken. In fact, in [Cortés *et al.*, 2005] we used Bench-Capon's Value based Argumentation Framework [Bench-Capon, 2003] to represent the human organ acceptability criteria and the preference between them.

Let us also suppose that the offered organ belongs to a donor carrier of the Hepatitis C virus. The TCA_i may offer the organ as *non-viable* giving the arguments `smoking history` and `donor Hepatitis C` [4].

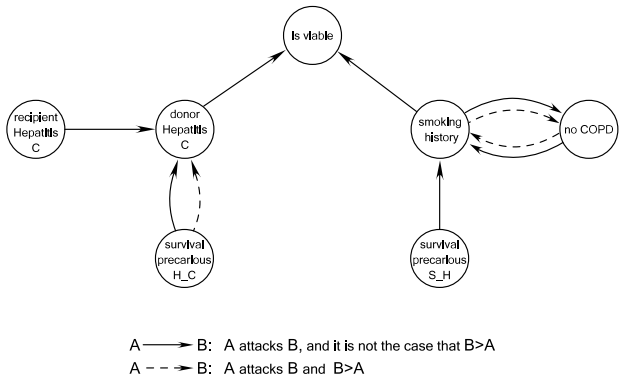

Figure 7: The set of all arguments used in the example. The arrows represent the attacking relation between the arguments allowed by the ACKB. For example, an agent can prefer either `smoking history` to `no COPD` or vice versa (but not the two arguments simultaneously).

In fig. 7 we present all the arguments to be used in this example, indicating the attacking relation and the preference between them allowed by the ACKB. As we can see, a TUA_j may prefer the argument `donor Hepatitis C` to the argument `survival precarious H_C` and vice versa. But the ACKB would not accept any justification based on the preference `smoking history` $>>^{Pref}$ `survival precarious S_H`. That is, the argument `smoking history` is ruled out in the presence of the argument `survival precarious S_H`, from the ACKB perspective.

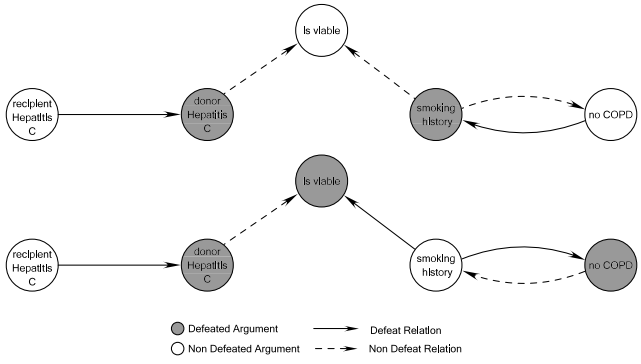

Figure 8: $AR_{(TCA_i, TUA_1, lung)}$, depending on its preference between the two arguments `no COPD` and `smoking history`, TUA_1 can claim the organ to be both viable and non-viable.

TUA_1 may argue that `is viable`, `no COPD` and `recipient Hepatitis C` [5], which conforms an admissible set in $AR_{(TCA_i, TUA_1, lung)}$, hence, it should be accepted (see in fig 8). On the other hand TUA_2 and TUA_3 could both argue that given the precarious conditions of their patients, it is justified to transplant an organ carrier of the Hepatitis C virus [6] (`survival precarious H_C`), similarly, the donor's smoking criterion for discarding a lung is relaxed when the patient survival is precarious (`survival precarious S_H`) [López-Navidad and Caballero, 2003]. Their arguments also conform admissible sets in $AR_{(TCA_i, TUA_2, lung)}$ and $AR_{(TCA_i, TUA_3, lung)}$ respectively, and so, their justifications should be accepted (see fig. 9 and fig. 10). TUA_3 could also argue that $\Gamma_{(TUA_3, lung)} = \{$`is viable`, `no COPD`, `recipient Hepatitis C`$\}$. In fact, TUA_3 cannot justify the non-viability of the organ (see fig. 10).

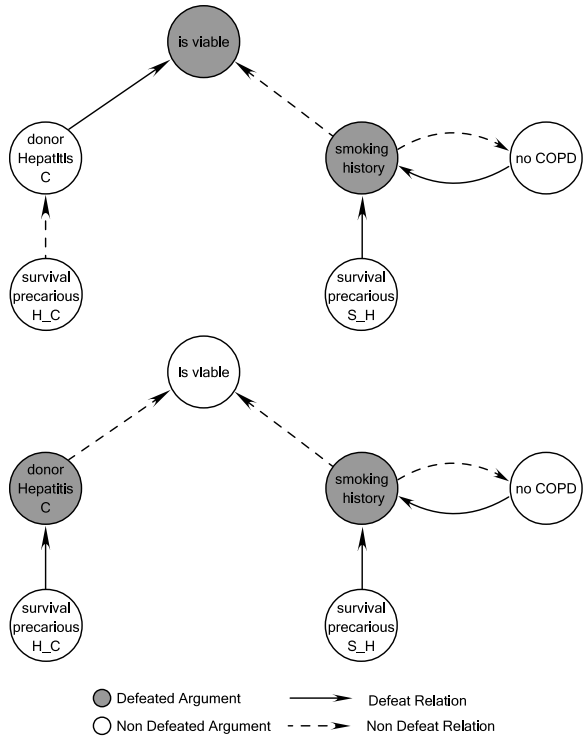

Figure 9: $AR_{(TCA_i, TUA_2, lung)}$, Depending on its preference between the two arguments `survival precarious H_C` and `donor Hepatitis C`, TUA_2 can either claim the organ to be viable or non-viable

On the other hand TUA_4 has no arguments to justify the viability of the lung for its patient (see fig. 11), therefore, the lung will not be offered to it at the allocation process.

Now let us suppose that within the same context the TCA_i

[4] Hepatitis C is an absolute contraindication for being a donor [OCATT,]

[5] If the donor and recipient carry the Hepatitis C virus, the organ can safely be transplanted [López-Navidad and Caballero, 2003]

[6] Organs that carry the Hepatitis C virus can be considered in life-saving transplantation for recipients not infected with this virus [López-Navidad and Caballero, 2003]

offers the organ as *viable*. In this occasion if TUA_1 believes the lung to be *non-viable*, by providing the argument `smoking history` its claim to refuse the organ should be accepted. On the other hand, TUA_2 may argue that `donor Hepatitis C` and its justification for not accepting the organ should also be accepted. Now, TUA_3 does not have valid reasons for refusing the organ, and thus, it will be committed to accept it (which does not means it must).

In the first case, in which the lung was offered as non-viable, TUA_2 argued $\Gamma_{(TUA_2,lung)}=$ {`is viable`, `survival precarious H_C`, `survival precarious H_C`}. These arguments can be accepted because they conform an admissible set in $AR_{(TCA_i,TUA_2,lung)}=$ {`donor Hepatitis C`, `smoking history`, `is viable`, `survival precarious H_C`, `survival precarious S_H`} provided that `donor Hepatitis C` $>>^{Pref}$ `survival precarious H_C`. That is, if TUA_2 claims the lung to be viable it is because it believes the criterion `survival precarious H_C` to be at least not weaker than the `donor Hepatitis C`'s criterion. If in some future organ evaluation TUA_2 makes use of the preference `donor Hepatitis C` $>>^{Pref}$ `survival precarious H_C` in order to support its claim (as it is the case in the example where it claims the non-viability of the lung) TUA_2 would be contradicting its own preferences. In future work we intend to address the agents' commitments with their acceptability criteria preferences.

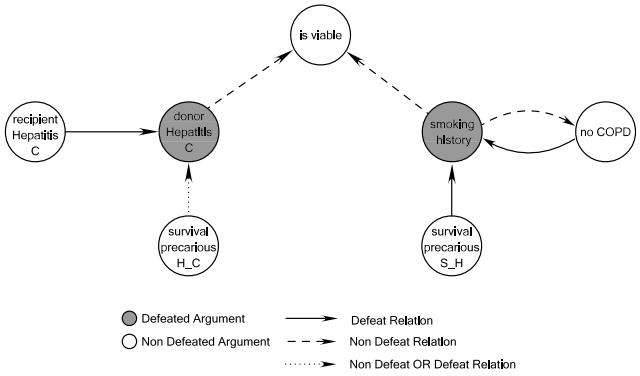

Figure 10: $AR_{(TCA_i,TUA_3,lung)}$, TUA_3 cannot justify the non-viability of the lung for its patient

When a Transplant Coordinator offers an organ she should construct a justification that warrants her conclusion more effectively while providing as much details as possible so that the members of the TU that receive the justification can build their own beliefs based on more solid grounds. The justification should not hide the weaknesses of the decision, *i.e.* not only the arguments that support the claim should be presented in the justification but also the cons, given that the final goal is to help the TU in their judgment. Even if the TC and the TU may disagree they are not in competition; they are both working together for a common goal.

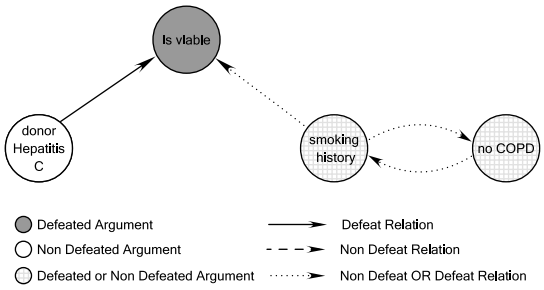

Figure 11: $AR_{(TCA_i,TUA_4,lung)}$, TUA_4 cannot justify the viability of the lung for its patient

In the following section we describe MA's argument evaluation process, this process takes into account two other components in addition to the ACKB. These components enable MA to decide the validity of a justification based on experience. This experience is partly present in previous cases, managed by a Case-Based Reasoning component, and partly present in each of the TUAs reputation model.

6 The arguments evaluation process

MA's argument evaluation process takes into account three components: the ACKB, the Case-Based Reasoning engine (CBRe) and the reputation model of the TUA_j under evaluation (see fig. 12).

As we have seen in the above examples, MA can derive from the ACKB that a justification given by TUA_j is valid if the set of arguments $\Gamma_{(TUA_j,O)}$ is admissible with respect to $AR_{(TCA_i,TUA_j,O)}$, where TCA_i is the agent offering the organ, under evaluation, O. In the previous section we defined $AR_{(TCA_i,TUA_j,O)}$as the union of $\Gamma_{(TUA_j,O)}$ and $\Gamma_{(TCA_i,O)}$. But, in the evaluation process, MA can also add arguments to the set $AR_{(TCA_i,TUA_j,O)}$when pertinent. For example, if the agents' arguments do not refer to a critical matter when deliberating upon the viability of an organ O, as for instance, if the expected travelling time between the recipient's and the donor's locations can damage the organ's integrity and no account of such fact is found in $AR_{(TCA_i,TUA_j,O)}$, MA will add a new argument, or set of arguments, that do reflect this critical matter. Therefore, a justification given by TUA_j is valid if the set of arguments $\Gamma_{(TUA_j,O)}$ is admissible with respect to $AR_{(TCA_i,TUA_j,O)}$, the set of all the arguments given by TCA_i, TUA_j and MA in deliberating upon O's viability.

Continuing the *smoking history* example, if TCA_i were to offer the lung as *viable*, and TUA_3 were to claim the organ as *non-viable* based on the argument `smoking history`, the set $\Gamma_{(TUA_j,O)}=$ {`donor Hepatitis C, smoking history`} would be an admissible set in $AR_{(TCA_i,TUA_j,O)}=$ {`donor Hepatitis C`, `is viable`, `no COPD`, `smoking history`}, but MA, knowing the recipient's potential characteristics, will add the argument `survival precarious S_H` to the set of arguments, and thus, TUA_3's justification, claiming the non-viability of the organ, will be rejected.

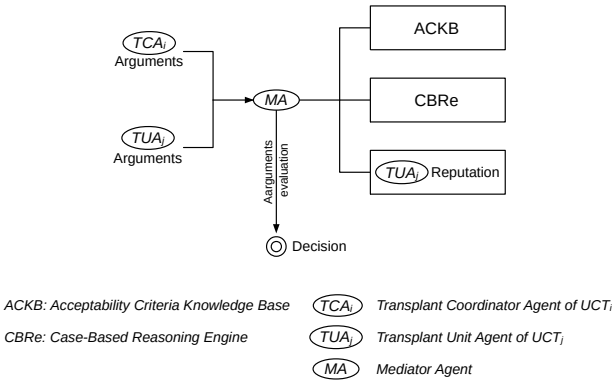

Figure 12: *MA*'s argument evaluation process

The second component that MA takes into account in the argument evaluation precess is the reputation model of the TUA_j under evaluation. Depending on the TUAs' reputation MA may vary the relative weight of the arguments and/or add more or less arguments to the set $AR_{(TCA_i, TUA_j, O)}$.

Take TUA_2, of the *smoking history* example, and suppose it has a long record of refusing organs that it has previously claimed to be viable. Let us also suppose that it has usually justified the refusal claiming *there were logistical inconveniences*. In future evaluations of TUA_2's justifications, MA will bias the evaluation adding the argument `logistical problem` to the set of arguments $AR_{(TCA_i, TUA_2, O)}$. Thus, TUA_2's justifications claiming the viability of an organ will have to be better argued in order to be validated by MA, for instance it will have to explicitly commit that no logistical problems will occur (see fig 13).

The third component, the CBRe, takes part in the arguments evaluation process when there is sufficient evidence to accept or reject a justification based on past cases. In given contexts justifications validated by the ACKB may not be backed up by experience, that is, in practice, these justifications may have a significant number of negative incidences. Similarly, it may happen that arguments not accepted by ACKB may have positive incidences. The CBRe task is to determine, based on previous cases, whether there is sufficient evidence for accepting or rejecting a justification. When the CBRe derives that there is sufficient evidence for determining the validity of a justification, its decision overrides the decision derived from the ACKB.

Continuing the above example we can focus on the two arguments `no COPD` and `smoking history`, which we have considered as equivalent in force, that is, no argument is preferred to the other. As we saw in the example, MA may accept two contradictory justification based upon the same facts. In one case the lung was considered as viable based on the argument `no COPD` whereas on the other case, the `smoking history` supported the claim that the organ was non-viable. We can suppose that after some cases where the argument `no COPD` was used in order to accept a lung, a significant number of these accepted lungs were successfully transplanted. The CBRe may, in some point, determine that there is sufficient evidence for considering `no COPD` as being a better argument than `smoking history`, namely that experience has proven that *any kind of smoking history could be acceptable unless there is a chronic obstructive pulmonary disease or pulmonary fibrosis on the chest radiograph*. Therefore, in future arguments evaluation process, MA may not accept a justification claiming the non-viability of a lung based on the `smoking history` argument on the presence of the argument `no COPD`. Similarly, if CBRe finds the argument `no COPD` as not being a reliable criterion, this may also change MA's decision making.

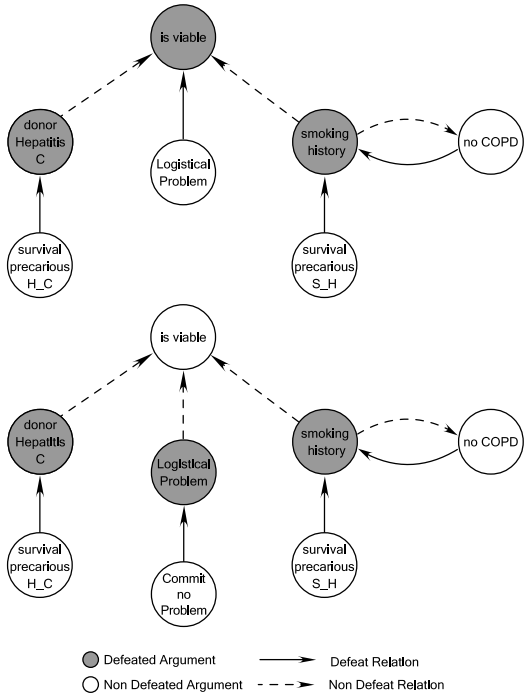

Figure 13: If MA adds the argument `logistical problem` to $AR_{(TCA_i, TUA_2, lung)}$ TUA_2 can commit to not having a logistical problem in the transplant process in order for its justification to be accepted

7 Handling dynamic acceptability criteria

A critical aspect in the human organ acceptability criteria is that they vary very rapidly, usually to refine the acceptability criteria in order to discard fewer organs. We believe that allowing these changes in the acceptability criteria, and even promoting them, is highly important, because preventing the discard of one organ may imply saving a human life. At the same time we must be aware of the risk involved in transplanting non-viable organs.

As long as there is no evidence that one criterion is better than another, two or more conflicting arguments can coexist in the ACKB enabling MA to accept two justifications with contradictory claims. As soon as the CBRe determines that there is sufficient evidence for considering one of the conflicting arguments as weaker or stronger than the others, MA may change its decision when evaluating such arguments, just as we have seen in the previous section's example when

readjusting the two arguments, `smoking history` and `no COPD`, relative strengths.

The CBRe enables to evaluate the validity of the applied acceptability criteria. Allowing the transplant of organs, that if relaying solely on the ACKB, would have been discarded, enables the CBRe to analyze the validity of new criteria, and thus, to eventually accept new arguments. There are two ways in which new arguments given by the agents can be accepted by the MA, and thus added to the CBRe. i) MA may accept justifications of TUAs with good reputation, namely TUAs that have usually accepted the organs and successfully transplant it, even if these justifications are not validated by the ACKB nor the CBRe. ii) If both, the TCA_i and a TUA_j consider an organ as viable it may be transplanted irrespectively of what may be derived from the ACKB or the CBRe. Alse, if the TCA_i considers the organ as non-viable and all the TUAs agreed. These, possibly new, arguments given by the agents claiming the viability (resp. non-viability) of the organs, are added to the CBRe enabling the creation of new acceptability criteria. These new criteria, when validated by the CBRe, can then be used by the agents in their justifications.

Scenario:

Donor's characteristics: Relevant data of the donor, such as his age, gender, blood type, cause of death, viral infections, etc.

Organ characteristics: Specific information of the organ.

Recipient characteristics: Relevant data of the recipient, such as his age, gender, blood type, urgency level, etc.

Logistical characteristics: Location of both donor and recipient, the distance between the two location, expected travel time, etc.

TCAx arguments: The arguments given by the TCAx

TUAx arguments: The arguments given by the TUAx claiming the viability or non-viability of an organ. And if it is the case, the arguments justifying any commitment failure.

TUAx Reputation: A valuation of the TUAx reputation.

Scenario label: A label identifying the case. If there was a disagreement between TCAx and TUAx, and in that case, what kind of disagreement. Whether there are TUAx's arguments justifying their failure to commit.

Solution: MA's decisions

Evaluation: The final result, whether the organ was accepted, successfully transplanted or rejected by the recipient. If the organ is grafted, this feature is continuously updated with the up to date recipient's medical condition.

Figure 14: Case Representation

One of the CBRe's case features is the agents' reputation (see fig. 14), viz. two identical justifications given by two agents with different reputations may produce different outcomes from the CBRe. Given that the MA may accept justifications from agents with good reputation regardless of what ACKB and CBRe may derive, some organs can end up labelled as viable and successfully transplanted even if the arguments claiming their viability were rejected by both ACKB and CBRe. These arguments can be used in future similar *scenarios* (see fig. 14) by agents that also have good reputation, with the difference that this time MA will derive from the CBRe that these arguments are acceptable, because they were used successfully by agents with similar reputation in similar *scenarios*. And as these justifications are used with positive results, the CBRe will accept them for agents with worse reputation. But, as soon as the CBRe accepts a justification (resp. reject) for any agent's reputation while the ACKB rejects it (resp. accept), the CBRe will trigger an update on the ACKB. In forthcoming work we intend to address the CBRe as a proactive component in the acceptability criteria refinement. That is, we want CBRe to generate new arguments. In order to do so, the CBRe has to carefully analyze the cases characteristics (the agents arguments and the donors' and recipients' characteristics) searching for patterns that enable suggesting new criteria.

Returning to the *smoking history* example, let us suppose that the argument `no COPD` is not yet in the ACKB, and let us also suppose that a few lungs were transplanted despite of the donor's smoking history. For instance, to recipients who had a precarious survival. Let us also suppose that some of these transplanted lungs of donors that had *no chronic obstructive pulmonary diseases and to which no pulmonary fibrosis was found in the chest radiograph*. If the transplants of lungs with *no COPD* turn to be as good as the transplants of lungs belonging to donors with *no significant smoking history*, the CBRe may detect this fact and propose `no COPD` as a new criteria. This new criteria can then be used as an argument that would attack the `smoking history` argument for non-viability.

As we explain above, the intention is that all the changes in the human organ acceptability criteria derived from the CBRe will eventually update the ACKB. Once an argument, or structured set of arguments, are added to the ACKB, as new criteria, they are no longer case dependent, and thus, can be used in more general *scenarios* as well as in combination with other arguments stemmed from the ACKB.

8 Conclusion and future work

In this paper we have shown how a MAS can take part in a process as complex and highly sensitive as the human organ transplantation, not only exchanging information efficiently but also reasoning about the exchanged information validity to give support to human makers.

We proposed a novel human organ discarding process that enables to safely augment the human organ pool for transplantation. The agents' argument exchange provides qualitative and efficient communication between the Transplant Coordinators and many Transplant Units distributed in different hospitals that may be located in different countries. As we showed, with the proposed extension to [Vázquez-Salceda *et al.*, 2003], CARREL would not only be able to handle the diverse human organ acceptability criteria that each party may be following, but it would make use of this heterogeneity in order to improve the transplantation process refining and suggesting alternative acceptability criteria. Another highly important aspect of this process is that, if applied, there will be an organized record of all the reasons for accepting and refusing each human organ and the history of each transplanted

organ. This information would be of great value for the research in the human organ transplantation domain as well as being very useful for legal purposes. Further medical motivation for the proposed discarding process is given in [Cortés *et al.*, 2005], in which a first approach to define the discarding process is made, but in which little detail is given about the agents' argument interaction and evaluation.

We are currently working on an argument scheme repository with which we intend to capture a sufficiently wide range of reasoning patterns concerning the decision making over the viability of human organs for transplantation. We believe that, as claimed in [Reed and Walton, 2001], by explicitly handling argument schemes agents can at once broaden the scope of the relevant information, and at the same time, narrow down selection on the basis of the argument schemes detected. We also believe the the application of argument schemes will provide a useful conceptual framework in which to elicit the required argumentation knowledge from physicians with experience in deciding the viability of human organs.

A natural extension to our proposed process is to include the allocation of human organs for transplantation. Hence, in forthcoming work we intend to address the decision making to which Transplant Unit an organ should be offered first. Now a days, in most of the cases[7], this decision takes into account only the hospital in which each patient is registered as a potential recipient and the time he had been waiting for the transplant operation. A more just allocation process should promote the assignation of organs to patient that would make a *better* use of them. Thus, if a human organ is effectively argued to be more suitable and beneficial for a certain patient than to the other potential recipients, these organ should first be offered to the TU responsible of that patient, regardless of the hospital in which this patient is registered and to the time he had been waiting.

References

[Amgoud and Cayrol, 1998] L. Amgoud and C. Cayrol. On the acceptability of arguments in preference-based argumentation. In Gregory F. Cooper and Serafín Moral, editors, *UAI*, pages 1–7. Morgan Kaufmann, 1998.

[Amgoud and Cayrol, 2002] L. Amgoud and C. Cayrol. A reasoning model based on the production of acceptable arguments. *Ann. Math. Artif. Intell.*, 34(1-3):197–215, 2002.

[Bench-Capon, 2003] T. Bench-Capon. Agreeing to differ: modelling persuasive dialogue between parties with different values. *Informal Logic*, 3(22), 2003.

[Cortés *et al.*, 2005] U. Cortés, P. Tolchinsky, J. C. Nieves, A. López-Navidad, and F. Caballero. Arguing the discard of organs for transplantation in CARREL. In *CATAI 2005*, pages 93–105. 2005.

[Dung, 1995] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–358, 1995.

[7] The main exception is urgency-0. Potential recipients that are expected to die in the next 24 hours are given national priority

[Kljakovic, 2003] M. Kljakovic. Clinical disagreement: A silent topic in general practice. *New Zealand Family Physician*, 30(5):358–360, 2003.

[López-Navidad and Caballero, 2003] A. López-Navidad and F. Caballero. Extended criteria for organ acceptance: Strategies for achieving organ safety and for increasing organ pool. *Clin Transplant, Blackwell Munksgaard*, 17:308–324, 2003.

[Loṕez-Navidad *et al.*, 1997] A. Loṕez-Navidad, P. Domingo, and M.A. Viedma. Professional characteristics of the transplant coordinator. In *Transplantation Proceedings*, number 29, pages 1607–1613. Elsevier Science Inc, 1997.

[OCATT,] OCATT. Organització Catalana de Transplantaments (OCATT). http://www10.gencat.net/catsalut/ocatt/en/htm/index.htm.

[ONT,] ONT. Organización Nacional de Transplantes. http://www.msc.es/ont.

[Reed and Walton, 2001] C.A. Reed and D.N. Walton. Applications of argumentation schemes. In H.V. Hansen, C.W. Tindale, J.A. Blair, and R.H. Johnson, editors, *Proceedings of the 4th Conference of the Ontario Society for the Study of Argument*, 2001.

[Vázquez-Salceda *et al.*, 2003] J. Vázquez-Salceda, U. Cortés, J. Padget, A. López-Navidad, and F. Caballero. The organ allocation process: a natural extension of the CARREL Agent-Mediated Electronic Institution. *AiCommunications. The European Journal on Artificial Intelligence*, 3(16):153–165, 2003.