A Step Towards Medical Ethics Modeling

Miguel Miranda, José Machado, António Abelha, Gabriel Pontes, José Neves

HAL Id: hal-01054849
https://inria.hal.science/hal-01054849

Submitted on 8 Aug 2014

A Step Towards Medical Ethics Modeling

Miguel Miranda[1], José Machado[1],
António Abelha[1], Gabriel Pontes[2] and José Neves[1]

[1] CCTC, Departamento de Informática, Universidade do Minho
miranda, jmac, abelha, jneves@di.uminho.pt
[2] Centro Hospitalar do Alto Ave, Guimarães, Portugal
gabriel.pontes@chaa.min-saude.pt

Abstract: Modeling of ethical reasoning has been a matter of discussion and research among distinct scientific fields, however no definite model has demonstrated undeniable global superiority over the others. However, the context of application of moral reasoning can require one methodology over the other. In areas such as medicine where quality of life and the life itself of a patient may be at stake, the ability to make the reasoning process understandable to staff and to change is of a paramount importance. In this paper we present some of the modeling lines of ethical reasoning applied to medicine, and defend that continuous logic programming presents potential for the development of trustworthy morally aware decision support systems. It is also presented a model of moral decision in two situations that emerge recurrently at the Intensive Care Units, a service where the moral complexity of regular decisions is a motivation for the analyze and development of moral decision support methodologies.

Keywords: Clinical Ethics, Ethical Modeling, Logic Programming, Artificial Intelligence

1   Introduction

Over viewing the evolution of technology and information systems thematic, a trend of growing pro-activeness and limited intelligence is pushing the role of virtual entities, on a step-by-step basis, higher and higher. Many activities are nowadays performed by automated entities, while supervised by humans. Although most of these virtual entities are still rather limited in learning, adaptation and autonomy, displaying solely reactance to predicted or programmed events, several threads of Artificial Intelligence research methodologies for imbedding further intelligence.

The notion of virtual entity is here used to differentiate entities with higher levels of autonomy, learning, prediction and decision from a mainly reactive and controlled machine. Moreover, considering developments in the area of informatics and Artificial Intelligence (AI) in particular, it must be considered that many of these entities can exist within a single physical machine or even that a single one can be distributed within limitless machines. Therefore, the notion of a virtual entity in this case is similar to the concept of an agent in the area of Multi-Agent Systems.

As virtual entities become more complex and hold critical functions, a justified doubt and concern regarding the impact of actions performed by these entities arises. From the numerous scenarios where they can interact with their surrounding environment, some carry moral consequences and describe ethically intricate actions from a human point of view. From the need to prevent immoral decisions and ensure confidence regarding these virtual entities, further understanding of the capacity of moral agency, moral modeling and the complexity human moral ethics.

Modeling machine ethics can result in further understanding of human ethics itself, either by defining rules and exceptions, or by knowledge extraction, case classification and patterns search over existing cases and outcomes using different algorithms. One can in fact consider that from the numerous methodologies that exist for the study of moral capacity, for each of them different subsequent potential outcomes can be found. While modeling ethics based on defined moral principles can help defining ethical principles and validate the resulting decision process, using learning algorithms and knowledge extraction over existing moral cases and outcomes can deepen the understanding of the underlying moral rules and patterns that may go unnoticed, but define moral decisions. In other words, theses processes aiming to analyze the essence of morality can be used not only to study their simulation/emulation, but also to deepen and evaluate the moral standards and dilemmas in ethically complex systems. The results from these systems are not limited outcome decisions before an ethical complex problem. Using a perspective of decision support or decision optimization, from a knowledge-base (either by previous studied cases or expert input), bearing in mind a specific scenario, similar cases can be aggregated for human user consideration, rules/principles involved in the decision can be induced with a certain degree of certainty, or conditions can be abduced.

There exists no definite solution for modeling ethical virtual entities, and presently several approaches are being presented and some compared against one another. Studying the present study and investigation in the area, different methodologies for modeling moral capabilities using artificial intelligence techniques can be segmented according to their main characteristics [1]. One of the most definite and important disparity in methodologies is the usage of explicit reasoning versus black-box reasoning. In explicit reasoning, the processes underneath a moral decision are clearly defined as principles, rules, exceptions, or other structure defined for one particular modeling. When analyzing AI techniques derivatives of symbolic, sub-symbolic or statistical approaches, there exist some that are able to represent their "line of though", allowing a transparent view of the moral decision process [2].

One of these techniques is logic programming, in which horn clauses contain the formalisms that mold the reasoning within an existing logical predicate. Current research indicates that non-monotic logic, due to its ability to implement defeasible inference, enabling moral principles to add and still diminish the set of conclusions determined by the knowledge base, is an interesting and promising technique to model moral reasoning [3] [4] [5]. By this mean, principles of benevolence and non-malificience can exist in accordance with other principles that are against their value or state an exception for superseding context principles. Regardless of the use of deductive, inductive or abductive logic, the rules used or attained are explicitly defined. However, the usage of each of these techniques of logic programming varies on the objective and context of application.

On the other hand, while using black-box reasoning, the reasoning behind the moral decision itself cannot be perceived in a clear manner. In other words, within the process of a black-box technique, facing a set of inputs, only a set of outputs can be obtained, not the process or reason behind it. That is the case of neural networks, regardless of the methodologies used to attempt to understand the reasoning behind them, the fact remains that no certainty of the processing underneath the trained neural network exists [2]. Although interesting results can be achieved using neural networks trained on existing moral cases and consequently implementing case based reasoning, the understanding of the moral principles within these black boxes is unknown [6]. Different techniques can be used to reverse-engineer neural network's inner structure and imbedded rules, however, the result is not exactly the rules used but rather an induced or a probabilistic set of them [7]. In the end of this reverse-engineer process, it is attained an induced set of rules of a systems that already uses induction or probabilistic methods to train its processing, revealing a certainty of doubt over the extracted rules.

Another divergence in ethical modeling is the learning process of rules or reasoning methodologies in ethical dilemmas. When considering a specific area such as medicine, most of the existing knowledge essential to model moral reasoning is contained in deontological principles or case studies [8]. In either of these cases the core of this knowledge is based on individuals or panels of experts. In light of these sources, the moral decision model can be developed from existing principles, from learnt principles or from hybridization of both sources. While one can consider existing deontological principles as existing principles, learnt principles are those extracted from existing cases. These machine-learning behaviors applied to ethics are a rather complex theme as principle learning may result in immoral principles and depending of the methodology used it may not be possible to clearly understand the underlying principles (e.g. black-box machine learning). Inductive logic programming has also expressed in existing research potential to induce principles and their relations from experts reasoning.[9]

When modeling moral behavior in virtual entities, researchers must always bare in mind the environment that molds its principles. For research purposes selection of an area and a purpose is of the essence in order to evaluate results and contextualize the used approach. With this in mind, the disparity between ideal and real environments in the medical arena creates a complex set of scenarios, which are pressing and interesting to analyze from an ethical point of view. Therefore, this article will address moral reasoning in medicine, and apply it in clinical context.

2   Medical Ethics Modeling – Analysis and Applications

Clinical ethics is an arena of public interest, where themes such as end of life, abortion and refusal or futility of treatment, among others, are constantly discussed as specific dilemmas occur or opinions and believes change. Although the deontic principles of a physician remain centered in the best practice towards the patient, legislation and court decisions mold the parameters of how physicians should behave in specific cases, which bare moral consequences. In fact, the context in which a

morally complex case presents itself may uphold different results. One european study analyzing the frequency and types of withholding and withdraw of life-sustaining therapies within the Intensive Care Units (ICU) of European countries, indicated that different countries and cultures deal in diverse ways with ethical dilemmas arising from these therapies [10]. One can go a step further, and consider the hypothesis that the physicians training and context can as well affect the moral decision making process. In fact, these decisions of withholding and withdraw of therapy, similarly to many other in clinical ethics, are far from an hypothetical situations, they happen frequently in the healthcare arena and allow no time for extensive legal or ethical consulting by the physician responsible for this decision. The moral demanding of clinical staff is overwhelming and can become even more complex and dubious in contexts of intensive and emergency care. Intensivists are constantly presented with new moral dilemmas, which demand for a quick and asserted answer [11]. Medical staff must therefore, be taught and trained to deal with these situations. Studies analyzing moral dilemmas and ethics modeling methodologies can be of help in this mater, to enhance the existing guidelines and understanding of moral decisions.

In the area of medicine, both practice and research activities have been actively overviewed and ultimately limited by existing legislation and court jurisprudence. This legislative effort is deeply connected to the existing moral principles and ethical concerns [11]. However, the existing legal directives can ease a decision concerning a morally complex situation and ethical confrontation, without fear for civil consequences. Some limitations occur on situations, in which decisions that sound ethically sound are limited by law, nevertheless professional conduct codes generally defined the proper conduct within the limits of the law [8].

For centuries, the clinical ethics with roots on Hippocrates principles defined as its main deontological fact, the obligation of the physician to give to the patient all treatments medicine knowledge considered the best fit. Nowadays, the decision is centered on the patients will, moreover, with the development of medical technologies, through their breakthrough and short-comes, physicians have also to take in consideration consequences of physical, mental and financial order [12]. This change of paradigm and the subjacent increment of ethical and civil load to the decisions of clinical staff, is an environment where synergies of medical ethics and AI, in order to understand how moral processing should be designed and how tutoring and decision support systems can be developed and implemented.

One interpretation of the process of learning and practicing clinical ethics is based on a set of corner-stones rules (i.e. moral principles), completed by the interpretation of existing fact in light of the existing numerous case studies. One can therefore consider that the moral behavior of physicians is a complex intertwined system of both rule-based and case based reasoning. Case studies can represent to some extents either rules or specific conditions which classify exceptions. This notion of exception is one of a logical programming point of view, where a context of known and unknown values of an universe can result in an exception to an existing predicate. Case studies can concur with the existing moral principles, alter their relationship, or define a context in which the existing principles were disregarded. When one analysis an ethical case study in medicine, the surrounding context that materializes the moral action defines an example of a decision with moral consequences, where the

boundaries of right or wrong are complex to ascertain. The analysis of such cases is complex, however one should always bear in mind that the existing moral rules and principles of medicine are the barebones of clinical ethics and should not be superseded unless valid exceptions are deemed correct.

From the distinct environments within the medical arena, intensive care medicine embodies an environment where moral decisions are usual and complex. In this specific context, decisions must be taken within short time spans while also regard limited resources and patients in critical conditions [11]. This context enables interest in using moral decision modeling in clinical cases appertaining to the ICU.

3   Modeling Clinical Ethics

With respect to the computational paradigm, it was considered Logic Programming in the form of a Continuous Logic Programming (CLP), once the truth values are defined in the range 0…1, with two kinds of negation, classical negation, ¬, and default negation, not. Intuitively, following the close world assumption, not p is true whenever there is no reason to believe p, whereas ¬p requires a proof of the negated literal. A continuous logic program (program, for short) is a finite collection of rules and integrity constraints, standing for all their ground instances, and is given in the form:

$$p \leftarrow p_1 \wedge \ldots \wedge p_n \wedge not\ q_1 \wedge \ldots \wedge not\ q_m; \text{ and}$$
$$?\ p_1 \wedge \ldots \wedge p_n \wedge not\ q_1 \wedge \ldots \wedge not\ q_m\ (n,m \geq 0)$$

where ? is a domain atom denoting falsity, the p_i, q_j, and p are classical ground literals, i.e. either positive atoms or atoms preceded by the classical negation sign ¬. Every program is associated with a set of abducibles. Abducibles may be seen as hypotheses that provide possible solutions or explanations of given queries, being given here in the form of exceptions to the extensions of the predicates that make the program.

Therefore, being Γ a program in CLP and g(X) a question where X contains variables $X_1 \wedge \ldots \wedge X_n\ (n \geq 0)$, one gets as an answer:

> The answer of Γ to g(X) is true iff
> $\quad$ g(X) → demo(Γ, g(X), true)
> The answer of Γ to g(X) is false iff
> $\quad$ ¬g(X) → demo(Γ, g(X), false)
> The answer of Γ to g(X) is unknown iff
> $\quad not\ \neg g(X) \wedge not\ g(X) \rightarrow demo(\Gamma, g(X), unknown)$

where unknown stands for a truth value in the interval]0…1[. Being Γ a Program, it is possible to define the Minimal Answer Set of Γ (MAS(Γ)):

$\quad$ $\Gamma \vdash s$ iff $s \in$ MAS(Γ)

where $\Gamma \vdash s$ denotes that s is a logical consequence or conclusion for Γ.

Being now AS_i and AS_j two different answer sets of Γ, being E_{ASi} and E_{ASj}, respectively, the extensions of predicates p in AS_i and AS_j, it is defined that AS_i is morally preferable to AS_j ($AS_i \prec AS_j$) where $\prec$ denotes the morally preferable relation, denoting that for each predicate p_1 there exists a predicate p_2 such that $p_1 \prec p_2$ and $E_{ASi} \setminus E_{ASj}$ is not empty ($\setminus$ denotes the difference set operator).

In our approach, the morally preferable relation is based on evolution and it is built on a quantification process of the quality-of-information that stems from a continuous logic program. Indeed, let p_i ($i \in \{1,\dots,m\}$) denotes the predicates whose extensions make a continuous logic program that models the universe of discourse, in terms of the extensions of predicates and let a_j ($j \in \{1,\dots,n\}$) stands for the attributes for those predicates. Let $x_j \in [min_j, max_j]$ be a value for attribute a_j. To each predicate it is also associated a scoring function $V^i j [min_j, max_j] \rightarrow 0\dots1$, that gives the score of predicate p_i assigned to a value of attribute a_j in the range of its acceptable values, i.e. its domain (for sake of simplicity, scores are kept in the continuous interval $[0,\dots,1]$. The quality-of-information with respect to a generic predicate it is therefore given by

$$Q_i = \frac{1}{Card}$$

, where Card denotes the cardinality of the exception set for the predicate pi, if the exception set is not disjoint. If the exception set is disjoint, the quality of information is given by

$$Q_i = \frac{1}{C_1^{card} + \dots + C_{card}^{card}}$$

where C_k^{card} is a k-combination subset, with card elements. The relative importance that a predicate assigns to each of its attributes under observation, $w^i j$, stands for the relevance of a_j for the predicate p_i and it is given by

$$V^i(x) = \sum w^i j V^i j(x)$$

, for all p_i. On the other hand, the predicate scoring function, when associated to a value $x = (x_1,\dots,x_n)$ in a multi-dimensional space defined by the attribute domains, is given in the form

$$V^i(x) = \sum w^i j V^i j(x)$$.

Therefore, it is now possible to measure the quality-of-information that stems from a continuous logic program, by posting Qi values into a multi- dimensional space, whose axes denote the program predicates with a numbering ranging from 0 (at the center) to 1. The area delimited by the arcs gives a measure of the quality-of-information carried out by each problem solution that may be under consideration, therefore defining the process of quantification of the morally preferable relation, as it is stated above in formal terms.

4   Model Behavior

Case 1

Mr. PD is a man with 81 years, a long background of cardiopathy and diabetes is admitted in an ICU with fever, hypertension and dyspnea.   The thorax radiography is compatible with Acute Respiratory Distress Syndrome (ARDS) and the arterial

partial oxygen tension (PaO2) is of 50 mmHg. This condition is often fatal, usually requiring mechanical ventilation and although the short-time mortality in these cases has been decreasing, the probability of mortality is considerably high and moreover this procedure results in a low quality-adjusted survival in the first year after ARDS [8, 13]. At the noon service meeting, while analyzing the current cases, the assistant physician asks the interns whether in light of the survival rates, treatment costs and probable low quality of life, should the ICU resources be used with this 81 years old men.

Case 2
During this meeting Mrs. GB, a woman with 36 years interned at the same hospital due to a car accident and diagnosed with sepsis, Acute Lung Injury (ALI) and Glasgow coma scale of 3, shows breathing complications and needs to be admitted an ICU. The level of its PaO2 and the severity of the ALI indicated a pressing need for mechanical ventilation and intensive care. However the number of beds in the ICU is limited and for this matter Mr. PD would have to be changed to another service. Due to the fragile state of Mr. PD this procedure is problematical, but considering his clinical status, complications and age with Mrs. GB, the greater probability of her to full recover with better quality of life tends to tip the balance from a critical point of view. In light of this context, how should the assistant physician act?

The continuous logic program for the extension of the predicate survival-rate:

```
{
¬survival-rate(X, Y) ← not survival-rate(X, Y ) and
    not exception(survival-rate(X, Y)),
exception(survival-rate(X, Y)) ← survival-rate(X,
unknown-survival-rate),
survival-rate(X, Y) ← ards(X) and pao2(X, low) and
evaluate(X,Y),
exception(survival-rate(gb, 0.5)),
?((exception(survival-rate(X,Y)) or
exception(survival-rate(X,Z))) and ¬(exception(survival-
rate(X,Y)) and exception(survival-rate(X, Z)))
/This invariant states that the exceptions to the
predicate survival-rate follow an exclusive or/
}agsurvival-rate
```

The continuous logic program for the extension of predicate survival-quality:

```
{
¬survival-quality(X, Y) ← not survival-quality(X, Y )
and not exception(survival-quality(X, Y)),
exception(survival-quality(X, Y)) ← survival-rate(X,
unknown-survival-quality),
survival-quality(gb, 0.8),
exception(survival-quality(pd, 0.1)),
?((exception(survival-quality(X,Y)) or
```

```
exception(survival-quality(X,Z))) and
¬(exception(survival-quality(X,Y)) and
exception(survival-quality(X, Z)))
  }agsurvival-quality
```

The continous logic program for the extension of predicate cost:
```
{
¬cost(X, Y) ← not cost(X, Y ) and
      not exception(cost(X, Y)),
exception(cost(X, Y)) ← cost(X, unknown-cost),
cost(gb, unknown-cost),
cost(pd, unknown-cost),
?((exception(cost(X,Y)) or exception(cost(X,Z))) and
¬(exception(cost(X,Y)) and exception(cost(X, Z)))
  }agcost
```

In this specific case we assume that the costs are unknown, so they will be considered as null values for the calculi.

5   Discussion

The extensions of the predicates that make the universe of discourse have to generated and considered in the CLP program construction, in order to have a basis for decision making. This is a bi-directional process because beyond the organizational, functional, technical and scientific requisites, one may have also to attend the ethical and the legal ones, as well as data quality, information security, access control and privacy. This generation is made from the nosocomial Electronic Health Records (EHR). EHR is a core application which covers horizontally the health care unit and makes possible a transverse analysis of medical records along the several services, units or treated pathologies, bringing to healthcare units new computational models, technologies and tools, based on data warehouses, agents, multi-agent systems and ambient intelligence. An EHR is an assembly of standardized documents, ordered and concise, directed to the register of actions and medical procedures; a set of information compiled by physicians and others health professionals; a register of integral facts, containing all the information regarding patient health data; and a follow up of the risk values and clinical profile. The main goal is to replace hard documents by electronic ones, increasing data processing and reducing time and costs. The patient assistance will be more effective, faster and quality will be improved.
Whatever form of an information society related to healthcare we can imagine, it will be based on three basic components, namely raw medical data, reconstructed medical data and derived medical data. Indeed, clinical research and practice involve a process to collect data to systematize knowledge about patients, their health status and the motives of the health care admittance. At the same time, data has to be registered in a structured and organized way, making effective automation and supporting using Information Technologies. For example, from an information repository, one may have collected patient data, which are registered in an efficient, consistent, clear and

structured way to improve disease knowledge and therapy; the medical processes for registering data are complemented with the information interchange between the different physicians that work around the patient; and the clinical data recording are guaranteed in the EHR application and procedural context. Interoperability will allow for sharing information among several information systems.

The process to collect data comes from Problem Oriented Medical Record (POMR) method. This is a format for clinical recording consisting of a problem list; a database including the patient history with physical examination and clinical findings; diagnostic, therapeutic and educational plans; and a daily SOAP (Subjective, Objective, Assessment and Plan) progress note. The problem list serves as an index for the reader, each problem being followed through until resolution. This system widely influences note keeping by recognizing the four different phases of the decision making process: data collection; formulation of problems; and devising a management plan; and reviewing the situation and revising the plan if necessary.

6 Conclusion

Different methodologies for problem solving based on the AI paradigm have been proposed to model ethical reasoning, however we consider that continuous logic programming expresses characteristics that overcome the main shortcomings of other techniques such as black-box techniques. One of the main advantages of using CLP concern the context of ethical modeling itself, as most of the trustworthy knowledge is based on deontological principles and is oriented towards experts consideration. The principle and exception modeling demonstrated presents a modeling clearly understandable by experts, traceable through proof trees and which processing is clearly identifiable, predictable and updatable.

The ultimate goal of using CLP is not to simulate moral reasoning itself, but rather enable decision support architectures, which take into account moral context. That is the reason why the possibility to justify moral decision and doubt on real-time to clinical staff is of the essence. Using such modeling principles, this staff could recur to moral decision support on real time and understand the line of reasoning implicit in the decision advised by the system.

Although a long path has to be walked before such moral aware decision support systems are implemented, this study of moral modeling and representation is of the essence to set the basilar structure in which morality can be imbedded in future systems.

7 References

1.   Tonkens, R., *A Challenge for Machine Ethics*. Minds and Machines, 2009. **19**(3): p. 421-438.
2.   Nugent, C. and P. Cunningham, *A Case-Based Explanation System for Black-Box Systems*. Artificial Intelligence Review, 2005. **24**(2): p. 163-178.

3.   Horty, J.F., *Moral Dilemmas and Nonmonotonic Logic*. Journal of Philosophical Logic, 1994. **23**(1): p. 35-65.

4.   Powers, T.M., *Prospects for a kantian machine*. IEEE Intelligent Systems, 2006. **21**(4): p. 46-51.

5.   Machado, J., et al., *Modeling Medical Ethics through Intelligent Agents*, in *I3E*. 2009. p. 112-122.

6.   Guarini, M., *Particularism and the classification and reclassification of moral cases*. IEEE Intell. Syst., 2006. **21**(4): p. 7.

7.   Floares, A.G., *A reverse engineering algorithm for neural networks, applied to the subthalamopallidal network of basal ganglia*. Neural Networks. **21**(2-3): p. 379-386.

8.   Jonsen, A.R., M. Siegler, and W.J. Winslade, *Clinical Ethics*. 4th Edition ed. 1997: McGraw-Hill.

9.   Anderson, M., S.L. Anderson, and C. Armen, *An approach to computing ethics*. Ieee Intelligent Systems, 2006. **21**(4): p. 56-63.

10.   Sprung, C.L., et al., *End-of-Life Practices in European Intensive Care Units: The Ethicus Study*. Journal Of the American Medical Association, 2003. **290**(6): p. 790-797.

11.   Danbury, C.M. and C.S. Waldmann, *Ethics and law in the intensive care unit*. Best Practice & Research Clinical Anaesthesiology, 2006. **20**(4): p. 589-603.

12.   Serrão, D. and R. Nunes, *Ética em Cuidados de Saúde*. 1998: Porto Editora.

13.   ANGUS, D.C., et al., *Quality-adjusted Survival in the First Year after the Acute Respiratory Distress Syndrome*. Am. J. Respir. Crit. Care Med., 2001. **163**(6): p. 1389-1394.