

Study on the Principles and Keystones of Artificial Intelligence

Rashmi Singh¹, Singh Divyansh², Bartakke Tanmay³
Asst. Professor¹ and TYBMS^{2,3}

Uttar Bhartiya Sangh's Mahendra Pratap Sharda Prasad Singh College of Commerce & Science, Mumbai, Maharashtra

Abstract: *The field of artificial intelligence study has drawn upon various disciplines such as formal logic, probability theory, decision theory, management science, linguistics, and philosophy. The integration of these fields in AI has required the creation of many improvements and expansions. Computational logic methods are quite potent. When integrated into an agent cycle, computational logic synergistically enhances and surpasses the capabilities of both traditional logic and classical decision theory. A multitude of its techniques can be employed to enhance human intelligence autonomously, without relying solely on artificial intelligence.*

Keywords: Artificial intelligence, Information Technology, ALP, LOT, FOL

I. INTRODUCTION

Computational logic, like other types of logic, comes in a variety of flavours. In this paper, I will concentrate on the computational logic form known as abductive logic programming (ALP).

The ALP agent model, which incorporates ALP into an agent cycle, is an effective model of both descriptive and normative thinking. It includes production systems as a special case as a descriptive model, and as a normative model, it includes classical logic and is compatible with classical decision theory. The ALP agent model's descriptive and normative properties make it a dual process theory that combines intuitive and deliberative thinking. Dual process theories, like most theories, take many forms. put it, intuitive thinking "quickly proposes intuitive answers to judgement problems as they arise", while deliberative thinking "monitors the quality of these proposals, which it may endorse, correct, or override".

In this paper, I will focus on the normative features of the ALP agent model and how they can help us improve our own human thinking and behaviour. I'll concentrate on how it can help us communicate more effectively with others and make better decisions in our daily lives. I will argue that it provides a theoretical foundation for both such English writing style guidelines.

A Brief Introduction to ALP Agents

The ALP agent model is a variation on the BDI model in which agents use their beliefs to satisfy their desires by generating intentions, which are predetermined plans of action. Agents, beliefs, and desires (or goals) are all represented as conditionals in the clausal form of logic in ALP. Beliefs are represented by logic programming clauses, and goals by more general clauses, both with the expressive power of full first-order logic (FOL). The first sentence below, for example, expresses a goal, while the remaining four sentences express beliefs: [5] The ALP agent model is a variation on the BDI model in which agents use their beliefs to satisfy their desires through the generation of intentions, which are predetermined plans of action. In the clausal form of logic in ALP, agents, beliefs, and desires (or goals) are all represented as conditionals. Beliefs are represented by logic programming clauses, and goals by more general clauses, both with full first-order logic expressive power (FOL). For example, the first sentence below expresses a goal, whereas the remaining four sentences express beliefs.

Model-theoretic and Operational Semantics

In the semantics of ALP agents, beliefs describe the world as the agent sees it, while goals describe the world as the agent wishes it to be. Beliefs represent data in deductive databases, while goals represent database queries and integrity

constraints. According to the operational semantics, ALP agents reason forwards from observations and backwards from beliefs to determine whether some instance of a goal's conditions is true, and to derive the corresponding instance of the goal's conclusion as an achievement goal, to make true. Forward reasoning from observations is similar to forward chaining in production systems in that it aims to make the goal true by making its conclusion true whenever its conditions become true. Conditional goals defined in this manner are also known as maintenance goals. Goals are solved by reasoning backwards, looking for a plan of action whose execution solves the goals. Backwards reasoning is a type of goal-reduction strategy, and executable actions are a subset of atomic sub-goals.

Consider the following scenario: I notice a fire. I can then reason with the above-mentioned goal and beliefs, concluding by forward reasoning that there is an emergency and deriving the achievement goal of dealing with it myself, getting help, or escaping. These three options represent the beginning of the search space. The achievement goal by reasoning backward and lowering the target I receive assistance with the subsequent sub-goals. I notify the train's driver and press the alarm button. If the last sub-goal is an atomic action, it can be carried out directly. If the action is successful, both the achievement goal and this instance of the maintenance goal are met.

In model-theoretic semantics, the agent must generate not only actions, but also world assumptions. These assumptions explain why the term abduction is used in ALP. Abduction is the process of developing assumptions to explain observations O . For example, instead of observing fire, I might observe smoke and conclude: there is smoke if there is a fire. The observation is then used to generate the assumption that there is a fire. The forward and backward reasoning then resumes as before.

Observations O and goals G are treated similarly in model-theoretic and operational semantics, with reasoning forwards and backwards to generate actions and other assumptions that make $G \wedge O$ true in them in the world model determined by B . In the preceding example, given $O = \text{there is smoke}$, then there is a fire , pressing the alarm button together with B makes both G and O true. The operational semantics is sound with respect to the model-theoretic semantics. It is also complete with modest assumptions.

Choosing the Best Solution

There may be several alternatives that, when combined with B , make both G and O true. These can have varying values, and the challenge for an intelligent agent is to find the best solution possible given the computational resources available. The expected utility of an action's consequences is used to calculate its value in classical decision theory. The value of an explanation is measured similarly in terms of its probability and explanatory power in philosophy of science. (The more observations that are described, the better.) The same measures can be used in ALP agents to evaluate both candidate actions and candidate explanations. In both cases, candidate assumptions are evaluated by reasoning forwards to generate assumptions' consequences. The task of finding the best is incorporated into the search strategy for reasoning backwards to generate in ALP agents, which uses some form of best-first search, such as A^* or branch and bound. This task is similar to the far simpler problem of conflict resolution in manufacturing systems. Traditional production systems avoid complex decision-theory and abductive reasoning primarily by aggregating higher-level goals, beliefs, and decisions into lower-level heuristics and stimulus-response associations. Lower-level rules and higher-level thinking and decision-making can be combined in ALP agents, as in dual process theories, to get the best of both worlds.

ALP agents, like BDI agents, think while observing and acting, and they do not need to create complete plans before acting. Unlike most BDI agents, who choose and commit to a single plan at a time, ALP agents choose and commit to individual actions. Unlike most BDI agents, ALP agents can pursue multiple alternative plans concurrently to increase their of success. In an emergency, for example, an agent may press the alarm button while also attempting to flee. The search strategy determines whether an ALP agent works on one plan or several alternative plans at the same time. Depth-first search only works on one plan at a time, but other search strategies are frequently preferred.

The ALP agent model can be used to create artificial agents, but it can also be used to describe human thinking and decision-making. However, I will argue in the remainder of this paper that it can also be used as a normative (or prescriptive) model that combines and improves on both traditional logic and classical decision theory.

Clausal Logic as an Agent's LOT

There are three major schools of thought in language philosophy regarding the relationship between language and thought:

The LOT is a private, language-like representation that is distinct from public, natural languages.

The LOT is a type of public language, and our natural language influences the way we think.

Human thought lacks a language-like structure. The ALP agent model is associated with the first school of thought, opposes the second, and is compatible with the third. It is opposed to the second school, in part because the ALP logical model of thinking does not require the existence of natural languages, and in part because Natural language, by AI standards, is too ambiguous and incoherent to be a useful model of human thinking. However, it supports the third school of thought because, as we will see in section 4, it has a connection with t implementation, which hides its linguistic nature.

In artificial intelligence, the notion that some form of logic is an agent's LOT is strongly associated with GOFAI (good old-fashioned AI), which has been partly overshadowed in recent years by connectionist and Bayesian approaches. I will argue that the ALP model of thinking has the potential to bridge the gap between logic, connectionism, and Bayesian approaches. This is due to the fact that ALP's clausal logic is much simpler than standard FOL. has a connectionist implementation that supports Bayesian probability and is related to standard FOL in the same way that the LOT is related to natural language. The first step of the argument is based on relevance, which holds that people understand natural language by extracting the most information for the least amount of processing cost. As a corollary to the theory, the closer a communication is to its intended meaning, the easier it is for a reader (or listener) to extract that meaning. As a result, one way to tell if there is a LOT, look at situations where it can be a matter of life and death that readers understand a communication as intended and with as little effort as possible. We will see that the communication is simple in the case of the London Underground Emergency Notice. because its English sentences are structured either explicitly or implicitly as logical conditionals.

Natural Language and the LOT

The problem of understanding ordinary, everyday natural language dispatches is much more delicate than the problem of understanding dispatches that are designed to be as clear and coherent as possible. This more delicate problem is divided into two corridors. The first step is to determine the communication's intended meaning. To understand the nebulous English judgment " he gave her the book," for illustration, the individualities appertained to buy" he" and" her" must be linked. The alternate step is to represent the intended meaning in a canonical form so that original dispatches are represented also. The following English rulings, for illustration, all have the same meaning the use of a canonical form in a internal representation facilitates latterly logic with the representation. In this case, the common meaning of the colorful rulings could be expressed in either the logical form give(john, mary, book) or in the more precise form. The more precise form can help distinguish between analogous events and books. According to applicability proposition, if you want your dispatches to be easy to understand, you should express them in a form that's analogous to their internal representations.

They should be clear, so that it's easy to prize their meaning, and simple, so that their meaning is close to the canonical form in which they're represented. depending on your meaning. The presence or absence of commas ahead and after the relative clause beginning with the word" which" in written English can indicate the different meanings. They're represented in clausal sense by the difference between conclusions and conditions. exemplifications like these indicate that the distinction and relationship between conditions and conclusions is aabecedarian point of the LOT, and they advance credence to the proposition that commodity.

Standard FOL and Clausal Logic

Various types of logic have been used for knowledge representation in AI, and clausal logic has been proposed as a candidate for the LOT. However, when compared to standard FOL, clausal logic not only stands out due to its simple, conditional form, but it is also just as powerful. It compensates for the lack of explicit existential quantifiers by using Skolemization to give names to individuals who are assumed to exist, such as e1000 and book21 above. In addition, when used in conjunction with the minimal model semantics, it is more powerful than FOL.

Reasoning is also much simpler in clausal logic than in standard FOL, and can be reduced to just forward and backward reasoning for the most part. In addition to the minimal model semantics, clausal logic reasoning includes default reasoning with negation as failure. The relationship between standard FOL and clauseal form may be analogous to the relationship between natural language and the LOT. In both cases, inferences can be divided into two types and carried out in two stages. The first type converts sentences to canonical form, and the second uses the resulting canonical form to reason. The first type of inference rule in FOL (which includes both Skolemization and the replacement of $\text{not}(A \text{ or } B)$ by $\text{not } A$ and $\text{not } B$) can be thought of as converting sentences into clausal form. The second type (including the inference of $P(t)$ from $XP(X)$) is clausal reasoning, and it is built into forward and backward reasoning. As we've seen, there are numerous ways to express the same information in natural language. Similarly, there are an infinite number of arbitrarily complex ways of expressing information in FOL.

In clausal form, there is only one canonical way to express the same information, in this case in the form of two clauses: $\text{feathers}(X) \text{ if } \text{bird}(X) \text{ and } \text{bird}(X)$ (john).

Thus, clausal logic is equivalent to standard FOL, just as the LOT is equivalent to natural language. Clausal logic is a simplified and canonical form of FOL, just as the LOT is a simplified and canonical form of unambiguous sentences in natural language. This analogy strengthens the case for viewing clausal logic as a formalisation of the LOT. [6] In the case of artificial agents in AI, clausal logic has proven to be a practical knowledge representation language, independent of any language an agent may use to communicate with other agents. Clausal logic can also help human agents communicate more effectively by expressing their communications in a form that is closer to the LOT.

Clausal logic can assist people in communicating more coherently by allowing them to connect new and old information. This coherence model takes advantage of the fact that clausal logic lends itself to a connectionist representation, in which information is stored in a goal-belief connection graph.

A connectionist Form of Clausal Logic

The connection graph proof procedure implements clausal logic, similarly to how clausal logic implements FOL by first converting sentences into canonical form, by pre-computing links between conditions and conclusions and labelling links with their unifying substitutions. These links can then be activated later, either forwards or backwards, as required. Links that are frequently activated can be compiled into shortcuts that achieve the same results more directly, similar to heuristic rules and stimulus-response associations.

Although clausal logic is a symbolic representation, the names of the predicate symbols no longer matter once all the links and their unifying substitutions have been computed. All subsequent reasoning can be reduced to the activation of the links and the creation of new clauses, whose new links are inherited from their parent clauses' links. When all of their links have been activated, parent clauses can often be deleted or overwritten.

At any time, any link can be selected for activation. However, most of the time, it makes sense to only activate links when new clauses are added to the graph as a result of new observations, including communications observations. Link activation can be guided by assigning varying strengths to various observations and goals, reflecting their relative importance (or utility).

Furthermore, different weights can be assigned to different links, reflecting statistical information about how frequently their activation has previously contributed to useful outcomes.

The weights on the links can be used to propagate the strength of observations and goals throughout the graph. The resulting proof procedure, which activates links with the highest weighted strength at the time, is similar to [Maes, 1990]. Furthermore, it employs an ALP-style of forward and backward reasoning, as well as a form of best-first search.

The connection graph model of thinking can give the false impression that thinking lacks any linguistic or logical character. However, the distinction between thinking in connection graphs and reasoning in clausal logic is simply the traditional computer science distinction between an optimised, low-level implementation that is close to the hardware and a high-level representation that is close to the problem domain. The mind's connection graph model adds to the argument that thinking occurs in a LOT that is independent of natural language. The LOT may aid in the development of natural language, but its existence is not required.

Representing Uncertainty

Internal links in connection graphs organise the agent's thoughts, while external links ground the agent's thoughts in reality. Observations and the agent's own actions activate the external links. They may also contain references to unobserved world properties. The agent can make assumptions about these properties and attempt to assess their likelihood. The likelihood that an assumption is correct contributes to the likelihood that an agent's actions will result in a specific outcome. You have control over your own actions, but not over the actions of others or the state of the world. At best, you might be able to assess the likelihood that the world is or will be in a particular state.

Better Decision-making

Uncertainty about the state of the world is only one of the complications complicating the decision-making process. Classical decision theory makes simplifying assumptions to reduce this complexity. The most constraining of these is the assumption that all of the alternatives to choose from are provided in advance. For example, if you are looking for a new job, it will assume that all job options are provided and will focus on the problem of determining which of the provided options is most likely to result in the best outcome.

According to other decision analysts, this assumption is not only unrealistic as a descriptive model of human decision making, but it is also ineffective as a normative (or prescriptive) model: To make an informed choice between alternatives, first identify the goals (or problems) that motivate the alternatives. These objectives may derive from explicitly stated maintenance objectives, or they may be hidden implicitly in lower-level heuristic rules or stimulus-response associations.

For example, if you receive a job offer when you are not looking for one, you may be tempted to limit your options to simply accepting or rejecting the offer. However, if you take a step back and consider the larger context of your goals, you may come up with other options, such as using the job offer to negotiate a raise in your current position.

By paying more attention to the goals that motivate the alternatives, decision analysis provides informal strategies for making better choices. The ALP agent model provides a simple framework for formalising such strategies by integrating them with a comprehensive human thinking model. It demonstrates, in particular, how the same expected utility criteria that are used in classical decision theory to choose between alternatives can also be used to guide the search for alternatives in some form of best-first search.

II. CONCLUSION

I've sketched two ways in which the ALP agent model, which is based on many different advances in AI, can be used by ordinary people to improve their own human intelligence. It can help them express themselves more clearly and coherently, as well as make better decisions. I believe that the application of such techniques is a promising area for collaboration between researchers in AI and researchers in more humanistic disciplines in the future.

REFERENCES

- [1]. [Carlson et al., 2008] Kurt A. Carlson, Chris Janiszewski, Ralph L. Keeney, David H. Krantz, Howard C. Kunreuther, Mary Frances Luce, J. Edward Russo, Stijn M.
- [2]. [Hammond et al., 1999] John Hammond, Ralph Keeney and Howard Raiffa. Smart Choices- A practical guide to making better decisions. Harvard Business School Press.
- [3]. [Kahneman, and Frederick, 2002] Daniel Kahneman and Shane Frederick. Representativeness revisited: attribute substitution in intuitive judgment. In Heuristics and Biases- The Psychology of Intuitive Judgement. Cambridge University Press.
- [4]. [Keeney, 1992] Ralph Keeney. Value-focused thinking: a path to creative decision-making. Harvard University Press.
- [5]. [Kowalski, 1975] Robert Kowalski. A proof procedure using connection graphs, JACM, 22(4):572-595.