Elaborating MCP (Part I)

John McCarthy
Computer Science Department
Stanford University
Stanford, CA 94305
jmc@cs.stanford.edu
http://www-formal.stanford.edu/jmc/

1997 Aug 11, 4:39 p.m.

Abstract

The *Missionaries and Cannibals* puzzle presents several problems in the logical expression of common sense reasoning. These include generality of axiomatization of common sense facts, nonmonotonic reasoning, Gricean implicatures and elaboration tolerance. We concentrate on the last of these.

We chose this *Drosophila* for research in elaboration tolerance rather than a practical problem for the same reason *Drosophilas* are chosen for many genetic experiments rather than pigs.

1 Introduction

The basic missionaries and cannibals situation (abbreviated MCPS0) is:

> **Statement: Three missionaries and three cannibals come to a river and find a boat that holds two. If the cannibals ever outnumber the missionaries on either bank, the missionaries will be eaten.**

The basic missionaries and cannibals problem (MCP0) adds to MCPS0 the question:

Problem: How shall they cross?

Saul Amarel proposed [Ama71]: Let a state (*mcb*) be given by the numbers of missionaries, cannibals and boats on the initial bank of the river. The initial situation is represented by 331 and the goal situation by 000.

Most AI texts that mention the problem accept this formulation and consider tree searches that give the solution:

$$331 \rightarrow 310 \rightarrow 321 \rightarrow 300 \rightarrow 311 \rightarrow 110 \rightarrow 221 \rightarrow 020 \rightarrow 031 \rightarrow 010 \rightarrow 021 \rightarrow 000. \tag{1}$$

If this representation is used the state space has 32 elements some of which are forbidden and two others are unreachable. It is an elementary student exercise to write a program to search the space and get the above sequence of states, and people are always solving it without a computer or without even a pencil. Amarel [Ama71] pointed out that this representation has fewer states than a representation with named missionaries and cannibals. The Amarel representation is preferred if one has this one problem MCP, is free to choose the representation for MCP only and plans an *ad hoc* computer program for solving it. This is just the situation of an exercise in programming tree search.

What more does this problem offer AI?

If one indeed begins with the Amarel representation, the problem is indeed trivial. However, suppose we want a program that begins, as people do, with a natural language presentation of the problem. It is still trivial if the program need only solve the missionaries and cannibals problem. The programmer can then cheat as much as he likes by making his program exactly suited to MCP. The extreme of cheating is to make a program that merely prints

$$331 \rightarrow 310 \rightarrow 321 \rightarrow 300 \rightarrow 311 \rightarrow 110 \rightarrow 221 \rightarrow 020 \rightarrow 031 \rightarrow 010 \rightarrow 021 \rightarrow 000.$$

Of course, the readers will complain, but it won't be clear what does and doesn't count as cheating.

The way to disallow cheating is to demand a program that can solve any problem in a suitable set of problems. To illustrate this we consider a large set of elaborations of MCP0. It won't be trivial to make a program that can solve all of them unless the human sets up each of them as a state space search analogous to the original MCP0. We demand that the program use background common sense knowledge like that about rivers and boats that is used by a human solver.

We will skip the part about going from an English statement of the problem to a logical statement. The problem is then to make a program that will solve any of the problems using logically expressed background knowledge. The background knowledge should be described in a general way, not specifically oriented to MCP and related problems.

We have two reasons for skipping the translation from English in this paper. First, we don't have anything new to say about parsing English. Second, we don't yet have the logical target language that the parsing program should aim at. Progress toward establishing this language is the goal of the paper.

This much was already proposed in **Programs with Common Sense**[1] [McC58]. What is new in the present paper is spelling out the idea of *elaboration tolerance* that was perhaps implicit in the 1959 paper. We require a formulation of MCP that readily tolerates elaborations of the problem and preferably allows them to be described by sentences added to the statement of the problem rather than by surgery on the problem. English language formulations allow this, but the Amarel-type formulations do not. AI requires a logical language that allows elaboration tolerant formulations.

We begin with examples of English language formulations of elaborations of MCP0. Each of them involves adding sentences (sometimes in a metalanguage) to the basic description of MCP. For each of them, we discuss the requirements on the logical formalization of MCP0 that will permit the elaboration in question. The goal (not achieved in this paper) is to make a logical formalization that will admit all of them.

2 Examples of Elaboration Tolerance in English

Consider the following elaborations of MCP.

MCP1: an oar on each bank Add to the statement of MCPS0

> There is an oar on each bank of the river. The boat can carry one person with only one oar but requires two oars if it is to carry two people.

[1] http://www-formal.stanford.edu/jmc/mcc59.html

It is a *Gricean implicature* [2] that the boat is a rowboat, because the proposer would be misleading us if it were not.

The same nonmonotonic reasoning that reduces MCP0 to a state space search problem should do the same for MCP2.

[We send a cannibal to get the oar on the far side, and we are reduced to the previous problem. However, an unsophisticated program might have to solve the basic missionaries and cannibals problem with an enlarged search space, because there will still be the irrelevant actions of picking up and depositing oars.]

MCP2: a bridge on which two can walk This should make the problem trivial. MCP2 is superficially like the original MCP0. However, common sense knowledge about bridges leads to a simpler solution. Namely, they cross in missionary-cannibal pairs. This works, because there is no requirement analogous to the requirement, part of common sense knowledge, that the boat crosses the river along with its users. With a bridge, an arbitrary number can cross, but proving this requires some kind of mathematical induction. It would be interesting to know whether the solution is *obvious* in the sense of [McA91].

At one level of detail, a boat and a bridge are similar. Each is a tool that can be used to cross a river. Part of the planning can treat them the same and decide $use(bridge)$ or $use(boat)$. As described above, the detailed plans are different. Human planners use successive levels of detail. From an *ad hoc* AI point of view, it is simpler to put the full level of detail into the original description of the facts. We may do that in this paper for simplicity, but then our system will not be a proper prototype of a system that has to use many levels of detail. For example, a planner may decide to use a boat before knowing the specific requirements for that boat.

MCP3: bad boat The boat is defective and must be repaired. This elaboration to allow inferring that repairing the boat is (a) required, and (b) will lead to a solution. This should be inferrable without a method for repairing the boat being provided.

[2] The philosopher Paul Grice [Gri89] studied what may be inferred from an utterance beyond what is stated under the assumption that the speaker is trying to inform the hearer and to avoid misleading him.

MCP4: four missionaries and four cannibals The problem is now unsolvable, but this also requires an induction. However, this isn't obvious in the McAllister sense.

MCP5: The missionaries can't row. It must then follow that the problem is unsolvable. Nonmonotonic reason will be needed to exclude unmentioned ways of crossing the river and induction may be needed to show that no sequence of actions improves the situation.

MCP6: Only one missionary and one cannibal can row. The problem is still solvable. The Amarel formulation is readily modified to include the numbers of rowing missionaries and cannibals on the initial bank as part of the state, which then has 5 components instead of 3. The key question is how much the system needs to know to perform the modification.

MCP7: big cannibal eats small missionary if alone together A solution to the basic problem can be specialized to avoid this contingency. In English, the existence of unique biggest missionary and smallest cannibal is a Gricean implicature.

MCP8: biggest cannibal fills boat Therefore, he can cross only by himself. This additional precondition on the use of the boat still allows a solution with a few more steps.

MCP9: biggest missionary fills boat Therefore, he can cross only by himself. This is not ok. *Lemma: It is never ok for a missionary to be in the boat by himself.* This is obvious to a human, because if so, the cannibals will outnumber the missionaries elsewhere. I doubt it is obvious to many present theorem provers.

MCP10: There is an island. Now the problem can be solved with any number of missionaries and cannibals (assuming the same number of each. Reaching this conclusion should not require a complete reformulation of the facts about crossing rivers.

The obvious traditional mathematical approach suggests a mathematical induction on the number missionaries. However, reaching this conclusion requires only a special common sense case of mathematical induction. We should first make the nonmonotonic inference that if each cannibal can be ferried to the island, then they all can. We certainly shouldn't have to enumerate the missionaries and cannibals in order to do the inference. (This tractable mathematical induction resembles concepts treated by David McAllester [McA91].

MCP11: Jesus Christ One of the missionaries is Jesus Christ who can walk on water. It is interesting to ask what cultural literacy a program would need to understand this elaboration and how it ought to be expressed.

Now four missionaries and four cannibals can cross.

MCP12: a leak The boat has a leak, so it is necessary to bail. We can treat bailing as a parallel operation even if one person is rowing and also bailing. He stops rowing to bail, but we don't need to take into account the sequence. We can take into account the need for a bucket or can or we can assume that the requirements are met. We can worry about the leak being too big. The elaboration requires a formalism that admits concurrent events.

MCP13: irrelevant event in boat Some event can occur in the boat, so we have to introduce states in which some missionaries and cannibals are in the boat. The simplest case is an event irrelevant to crossing the river, e.g. buttoning and unbuttoning a jacket. Common sense knowledge of boats can make further elaborations, e.g. untying the boat, relevant.

MCP14: continuous action Consider an elaboration that requires considering continuous actions and other events. There are four cannibals and four missionaries, but if the strongest of the missionaries rows fast enough, the cannibals won't have gotten so hungry that they will eat the missionaries.

A missionary can deposit a cannibal at an arbitrary distance up the river except that fuel for the boat is limited. The missionaries then outnumber the cannibals until the marooned cannibal can walk back. Depending on the fuel available, this can make the problem solvable with more missionaries and cannibals.

MCP15: simple probabilities The probability is 1/2 that the largest missionary cannot fit into the boat with another person. Then the probability is 1/2 that the problem is solvable. Clearly we don't want to go to a logic based on probabilities in order to handle this small amount of probabilistic reasoning.

MCP16: details, details Another direction of elaboration is to give more detail. Crossing the river involves some missionaries and cannibals getting into the boat, propelling the boat across the river, and some missionaries and cannibals getting out. Crossing the river can be further elaborated to take into account the order of entering and leaving the boat and auxiliary actions like tying and untying the boat to the bank or dock. When a formalism that elaborates unnecessarily is used, the formalization should admit collapsing sequences of actions to single actions so as to restore simplicity.

MCP17: time The obedience of a cannibal will last only for the time of two back-and-forth trips in the absence of a missionary.

MCP18: Answering questions Here are some. Which missionaries actually rowed? When a missionary and a cannibal row together, which gets into the boat first and which gets out first? Perhaps it is safer if the cannibal gets into the boat first and the missionary gets out When two row, does one do all the rowing or do they take turns?

MCP19: increasing hunger Cannibals are initially not hungry, but rowing makes them hungry.

MCP20: There is more than one boat.

MCP21: elaboration requiring a strategy In all the previous elaborations, the solution is a determinate sequence of actions. Suppose we have two boats, only one of which works and an action that allows observing which one it is. The solution is then to observe and then use the good boat.

This is the first step in the direction of problems whose solutions are strategies, i.e. programs, rather than sequences of actions. MCP has many elaborations of this kind.

MCP22: conversion Three missionaries alone with a cannibal can convert him into a missionary. Whether a person is a missionary or a cannibal is then situation dependent. However, if the possibility of conversion is not mentioned, then nonmonotonic reasoning from the initial statement of the problem should produce a compact formalization in which being a missionary is not situation dependent.[3]

As exemplified above, the English language version of missionaries and cannibals can be elaborated by various single sentences. Each of the above elaborations except a few, which aren't well defined as stated, gives rise to a problem definite enough so that people will agree on whether it is solvable.

To get any definite problem out of the English or its sentence-by-sentence translation into mathematical logical sentences, one has to do nonmonotonic reasoning.

References

[Ama71] Saul Amarel. On representation of problems of reasoning about action. In Donald Michie, editor, *Machine Intelligence 3*, pages 131–171. Edinburgh University Press, 1971.

[Gri89] Paul Grice. *Studies in the Way of Words*. Harvard University Press, 1989.

[McA91] David McAllester. Some observations on cognitive judgements. In *AAAI-91*, pages 910–915. Morgan Kaufmann Publishers, July 1991. internet file ftp.ai.mit.edu:/pub/dam/aaai91a.ps.

[McC58] John McCarthy. Programs with Common Sense[4]. In *Mechanisation of Thought Processes, Proceedings of the Symposium of the National Physics Laboratory*, pages 77–84, London, U.K., 1958. Her Majesty's Stationery Office. Reprinted in McC90.

/@steam.stanford.edu:/u/ftp/jmc/missionaries1.tex: begun Thu Jul 31 13:49:44 1997, latexed August 11, 1997 at 4:39 p.m.

[3]A missionary who goes without food for a very long time becomes a cannibal.
[4]http://www-formal.stanford.edu/jmc/mcchay69.html