Multi-Agent Systems

B. Nebel, F. Lindner, T. Engesser
Winter Semester 2018/19

University of Freiburg
Department of Computer Science

Exercise Sheet 13

Due: February 6, 2019, 16:00

Exercise 13.1 (But-For Causality, 2+2+2)

Consider the situation in which an autonomous, self-aware war robot has to decide to either kill a person or to let her live. However, unknown to the robot there exists a safeguard in its programming: If the robot's decision is to kill that particular person, then the safeguard will activate causing the robot to be unable to carry out its decision.

(a) Represent this situation in a causal model.

(b) Check if the decision to let the person live is a but-for cause of the person's survival.

(c) Check if the decision to let the person live is a cause of the person's survival according to the modified HP definition.

Exercise 13.2 (Normality, 2+2+2)

Consider the situation that an Assassin puts poison in a Victim's coffee and the Victim's Bodyguard puts an antidote in the coffee such that the Victim survives. However, the Assassin would not have put the poison in the coffee if the Bodyguard hadn't put antidote in it in the first place.

(a) Represent the situation as a causal model.

(b) First check that the Bodyguard putting the antidote in the coffee is a cause of the Victim's survival according to the modified HP definition.

(c) Extend your causal model with a normality relation, and show that now the unintuitive result from (b) does not hold anymore.