

ARTIFICIAL INTELLIGENCE 3E

FOUNDATIONS OF COMPUTATIONAL AGENTS


[Contents](#) [Index](#) [Home](#)
[12 Planning with Uncertainty](#)
[12.3 Sequential Decisions](#)
[12.5 Decision Processes](#)

12.4 The Value of Information and Control

Information can be valuable to agents if it helps them make better decisions.

Example 12.24. In [Example 12.22](#), the action *Check_smoke* provides information about fire. Checking for smoke costs 20 units and does not provide any direct reward; however, in an optimal policy, it is worthwhile to check for smoke when there is a report because the agent can condition its further actions on the information obtained. Thus, the information about smoke is valuable to the agent, even though smoke only provides imperfect information about whether there is fire.

One of the important lessons from this example is that an information-seeking action, such as *Check_smoke*, can be treated in the same way as any other action, such as *Call*. An optimal policy often includes actions whose only purpose is to find information, as long as subsequent actions can condition on some effect of the action. Most actions do not just provide information; they also have a more direct effect on the world.

If X is a random variable and D is a decision variable, the **value of information** about X for decision D is how much extra utility can be obtained by knowing the value for X when decision D is made. This depends on what is controlled and what else is observed for each decision, which is the information provided in a decision network.

The value of information about X for decision D in a no-forgetting decision network N is

- the value of decision network N with an arc added from X to D , and with arcs added from X to the decisions after D to ensure that the network remains a [no-forgetting decision network](#)
- minus the value of the decision network N where D does not have information about X , and the no-forgetting arcs are not added.

This is only defined when X is not a successor of D , because that would cause a cycle. (Something more sophisticated must be done when adding the arc from X to D causes a cycle.)

Example 12.25. In [Example 12.14](#), consider how much it could be worth to get a better forecast. The value of getting perfect information about the weather for the decision about whether to take an umbrella is the difference between the value of the network with an arc from *Weather* to *Umbrella* which, as calculated in [Example 12.23](#), is 91 and the original network, which, as computed in [Example 12.14](#), is 77. Thus, the value of information about *Weather* for the *Umbrella* decision is $91 - 77 = 14$.

The value of information has some interesting properties:

- The value of information is never negative. The worst that can happen is that the agent can ignore the information.
- If an optimal decision is to do the same thing no matter which value of X is observed, the value of information X is zero. If the value of information X is zero, there is an optimal policy that does not depend on the value of X (i.e., the same action can be chosen no matter which value of X is observed).

The value of information is a bound on the amount the agent should be willing to pay (in terms of loss of utility) for information X for decision D . It is an upper bound on the amount that **imperfect information** about the value of X at decision D would be worth. Imperfect information is the information available from a noisy sensor of X . It is not worth paying more for a sensor of X than the value of information about X for the earliest decision that could use the information of X .

Example 12.26. In the fire alarm problem of [Example 12.22](#), the agent may be interested in knowing whether it is worthwhile to try to detect tampering. To determine how much a tampering sensor could be worth, consider the value of information about tampering.

The following are the values (the expected utility of the optimal policy, to one decimal point) for some variants of the network. Let N_0 be the original network.

- The network N_0 has a value of -22.6 .
- Let N_1 be the same as N_0 but with an arc added from *Tampering* to *Call*. N_1 has a value of -21.3 .
- Let N_2 be the same as N_1 except that it also has an arc from *Tampering* to *Check_smoke*. N_2 has a value of -20.9 .
- Let N_3 be the same as N_2 but without the arc from *Report* to *Check_smoke*. N_3 has the same value as N_2 .

The difference in the values of the optimal policies for the first two decision networks, namely 1.3, is the value of information about *Tampering* for the decision *Call* in network N_0 . The value of information about *Tampering* for the decision *Check_smoke* in network N_0 is 1.7. Therefore, installing a tampering sensor could at most give an increase of 1.7 in expected utility.

In the context N_3 , the value of information about *Tampering* for *Check_smoke* is 0. In the optimal policy for the network with both arcs, the information about *Alarm* is ignored in the optimal decision function for *Check_smoke*; the agent never checks for smoke when deciding whether to call in the optimal policy when *Alarm* is a parent of *Call*.

The **value of control** specifies how much it is worth to control a variable. In its simplest form, it is the change in value of a decision network where a random variable is replaced by a decision variable, and arcs are added to make it a no-forgetting network. If this is done, the change in utility is non-negative; the resulting network always has an equal or higher expected utility than the original network.

Example 12.27. In the fire alarm decision network of [Figure 12.11](#), you may be interested in the value of controlling *Tampering*. This could, for example, be used to estimate how much it is worth to add security guards to prevent tampering. To compute this, compare the value of the decision network of [Figure 12.11](#) to the decision network where *Tampering* is a decision node and a parent of the other two decision nodes.

To determine the value of control, turn the *Tampering* node into a decision node and make it a parent of the other two decisions. The value of the resulting network is -20.7 . This can be compared to the value of N_3 in [Example 12.26](#) (which has the same arcs, and differs in whether *Tampering* is a decision or random node), which was -20.9 . Notice that control is more valuable than information.

The previous description assumed the parents of the random variable that is being controlled become parents of the decision variable. In this case, the value of control is never negative. However, if the parents of the decision node do not include all of the parents of the random variable, it is possible that control is less

valuable than information. In general, one must be explicit about what information will be available when controlling a variable.

Example 12.28. Consider controlling the variable *Smoke* in the network of [Figure 12.11](#). If *Fire* is a parent of the decision variable *Smoke*, it has to be a parent of *Call* to make it a no-forgetting network. The expected utility of the resulting network with *Smoke* coming before *Check_smoke* is -2.0 . The value of controlling *Smoke* in this situation is due to observing *Fire*. The resulting optimal decision is to call if there is a fire and not call otherwise.

Suppose the agent were to control *Smoke* without observing *Fire*. That is, the agent can decide to make smoke or prevent smoke, and *Fire* is not a parent of any decision. This situation can be modeled by making *Smoke* a decision variable with no parents. In this case, the expected utility is -23.20 , which is worse than the initial decision network, because blindly controlling *Smoke* loses its ability to act as a sensor for *Fire*.

[12.3 Sequential Decisions](#)

[ReferencesIndex](#)

[12.5 Decision Processes](#)

Generated on Sat Jun 28 18:40:00 2025 by [L^AT_EX](#)ML

[Artificial Intelligence: Foundations of Computational Agents, Poole & Mackworth](#)

Copyright © 2023, [David L. Poole](#) and [Alan K. Mackworth](#). 

This work is licensed under a [Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License](#).