



AI:Constraint Satisfaction

ABOUT THE COURSE:

Human beings solve problems in many different ways. Problem solving in artificial intelligence (AI) is inspired from these diverse approaches. AI problem solvers may be based on search, on memory, or on knowledge representation and reasoning. An approach to problem solving is to pose problems as constraint satisfaction problems (CSP), and employ general methods to solve them. The task of a user then is only to pose a problem as a CSP, and then call an off-the-shelf solver. CSPs are amenable to combining search based methods with reasoning. In this 2 credit course we will look at general approaches to solving finite domain CSPs, and explore how search can be combined with constraint propagation to find solutions.

This course is a companion to the course “Artificial Intelligence: Search Methods for Problem Solving” that was offered recently and “Artificial Intelligence: Knowledge Representation & Reasoning” that is being offered concurrently. The lectures for both courses are available online.

COURSE LAYOUT:

Module 1: Constraint satisfaction problems (CSP), examples.

Module 2: Constraint networks, equivalent and projection networks.

Module 3: Constraint propagation, arc consistency, path consistency, i-consistency.

Module 4: Directional consistency and graph ordering, backtrack free search, adaptive consistency.

Module 5: Search methods for solving CSPs, lookahead methods, dynamic variable and value ordering.

Module 6: Lookback methods, Gaschnig's backjumping, graph based backjumping, conflict directed backjumping. Combining lookahead with lookback, learning.

Module 7: Model based systems, model based diagnosis, truth maintenance systems, planning as CSP.

Wrapping up.