

Constraint Satisfaction Methods

Daniel Tauritz, PhD

Associate Professor
Department of Computer Science and Software Engineering
Auburn University

September 23, 2024

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints

Ignore Constraints

Ignore the constraints under the motto: all is well that ends well.

Ignore Constraints

Ignore the constraints under the motto: all is well that ends well.

Pros Simplest approach by far (requires no additional programming) and least computationally expensive (incurs no overhead).

Ignore Constraints

Ignore the constraints under the motto: all is well that ends well.

Pros Simplest approach by far (requires no additional programming) and least computationally expensive (incurs no overhead).

Cons Allows invalid solutions to flourish.

Ignore Constraints

Ignore the constraints under the motto: all is well that ends well.

Pros Simplest approach by far (requires no additional programming) and least computationally expensive (incurs no overhead).

Cons Allows invalid solutions to flourish.

When to use Only if invalid solutions have an inherently lower fitness than valid solutions which means that the fitness function implicitly accounts for the constraints.

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions

Kill Infeasible Solutions

Upon generating an infeasible solution, immediately kill it and generate a new solution; repeat this step until a feasible solution is generated.

Kill Infeasible Solutions

Upon generating an infeasible solution, immediately kill it and generate a new solution; repeat this step until a feasible solution is generated.

Pros Simple to implement (requires only a validity check) and guarantees only valid solutions propagate.

Kill Infeasible Solutions

Upon generating an infeasible solution, immediately kill it and generate a new solution; repeat this step until a feasible solution is generated.

Pros Simple to implement (requires only a validity check) and guarantees only valid solutions propagate.

Cons Computational overhead is proportional to the ratio of invalid to total solutions; if that ratio gets too high, then the overhead makes this approach infeasible. Also, the largest diameter of invalid space imposes a typically unknown lower bound on mutation rate to guarantee reachability of the global optimum.

Kill Infeasible Solutions

Upon generating an infeasible solution, immediately kill it and generate a new solution; repeat this step until a feasible solution is generated.

Pros Simple to implement (requires only a validity check) and guarantees only valid solutions propagate.

Cons Computational overhead is proportional to the ratio of invalid to total solutions; if that ratio gets too high, then the overhead makes this approach infeasible. Also, the largest diameter of invalid space imposes a typically unknown lower bound on mutation rate to guarantee reachability of the global optimum.

When to use Only if (a) the ratio of invalid to total solutions is sufficiently low to make the overhead of generating solutions and checking their validity lower than the overhead of generating guaranteed valid solutions or imposing a penalty function, and (b) global maximum reachability is guaranteed.

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness

Assign Arbitrarily Low Fitness

Infeasible solutions are assigned arbitrarily low fitness.

Assign Arbitrarily Low Fitness

Infeasible solutions are assigned arbitrarily low fitness.

Assign Arbitrarily Low Fitness

Infeasible solutions are assigned arbitrarily low fitness.

Pros Simple to implement (requires only a validity check) and may slightly increase the gene pool in case of stochastic selection.

Assign Arbitrarily Low Fitness

Infeasible solutions are assigned arbitrarily low fitness.

Pros Simple to implement (requires only a validity check) and may slightly increase the gene pool in case of stochastic selection.

Cons Allows invalid offspring.

Assign Arbitrarily Low Fitness

Infeasible solutions are assigned arbitrarily low fitness.

Pros Simple to implement (requires only a validity check) and may slightly increase the gene pool in case of stochastic selection.

Cons Allows invalid offspring.

When to use Only if the ratio of invalid to total solutions is sufficiently low to make the overhead of generating solutions and checking their validity lower than the overhead of generating guaranteed valid solutions or imposing a penalty function.

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness
- Penalty Function

Penalty Function

Employ a penalty function that reduces the fitness of infeasible solutions, preferably so that the fitness is reduced in proportion to the number of constraints violated, or to the distance from the feasible region.

Penalty Function

Employ a penalty function that reduces the fitness of infeasible solutions, preferably so that the fitness is reduced in proportion to the number of constraints violated, or to the distance from the feasible region.

Pros Supports a user controllable balance between ridding the population of invalid solutions and allowing exploration of invalid regions of the search space in order to make the global maximum (easier) reachable.

Penalty Function

Employ a penalty function that reduces the fitness of infeasible solutions, preferably so that the fitness is reduced in proportion to the number of constraints violated, or to the distance from the feasible region.

- Pros** Supports a user controllable balance between ridding the population of invalid solutions and allowing exploration of invalid regions of the search space in order to make the global maximum (easier) reachable.
- Cons** Requires problem dependent design of a penalty function, adds the penalty coefficient to the strategy parameters that need to be tuned for good performance, and uses potentially expensive evaluations to determine the fitness of invalid solutions.

Penalty Function

Employ a penalty function that reduces the fitness of infeasible solutions, preferably so that the fitness is reduced in proportion to the number of constraints violated, or to the distance from the feasible region.

Pros Supports a user controllable balance between ridding the population of invalid solutions and allowing exploration of invalid regions of the search space in order to make the global maximum (easier) reachable.

Cons Requires problem dependent design of a penalty function, adds the penalty coefficient to the strategy parameters that need to be tuned for good performance, and uses potentially expensive evaluations to determine the fitness of invalid solutions.

When to use If you cannot guarantee global maximum reachability from the starting population and a high quality decoder function cannot be designed.

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness
- Penalty Function

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness
- Penalty Function
- Repair Function

Repair Function

Employ a repair function that takes infeasible solutions and “repairs” them by transforming them into a related feasible solution, typically as close as possible to the infeasible one.

Repair Function

Employ a repair function that takes infeasible solutions and “repairs” them by transforming them into a related feasible solution, typically as close as possible to the infeasible one.

Pros Guarantees valid solutions and reduces the effective search space.

Repair Function

Employ a repair function that takes infeasible solutions and “repairs” them by transforming them into a related feasible solution, typically as close as possible to the infeasible one.

Pros Guarantees valid solutions and reduces the effective search space.

Cons Requires problem dependent design of a repair function and the largest diameter of invalid space imposes a typically unknown lower bound on mutation rate to guarantee reachability of the global optimum.

Repair Function

Employ a repair function that takes infeasible solutions and “repairs” them by transforming them into a related feasible solution, typically as close as possible to the infeasible one.

Pros Guarantees valid solutions and reduces the effective search space.

Cons Requires problem dependent design of a repair function and the largest diameter of invalid space imposes a typically unknown lower bound on mutation rate to guarantee reachability of the global optimum.

When to use If you can guarantee global maximum reachability from the starting population and a computationally efficient repair algorithm is available.

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness
- Penalty Function

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness
- Penalty Function
- Repair Function

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness
- Penalty Function
- Repair Function
- Closed Feasible Solution Space

Closed Feasible Solution Space

Employ a closed feasible solution space which guarantees that the initial population consists of feasible solutions only and all evolutionary operations on feasible solutions are guaranteed to result in feasible solutions. Typically a combination of custom representation, initialization, recombination, and mutation is employed to achieve this.

Closed Feasible Solution Space

Employ a closed feasible solution space which guarantees that the initial population consists of feasible solutions only and all evolutionary operations on feasible solutions are guaranteed to result in feasible solutions. Typically a combination of custom representation, initialization, recombination, and mutation is employed to achieve this.

Pros Guarantees valid solutions and reduces the effective search space.

Closed Feasible Solution Space

Employ a closed feasible solution space which guarantees that the initial population consists of feasible solutions only and all evolutionary operations on feasible solutions are guaranteed to result in feasible solutions. Typically a combination of custom representation, initialization, recombination, and mutation is employed to achieve this.

- Pros** Guarantees valid solutions and reduces the effective search space.
- Cons** Requires problem dependent design of closed evolutionary operators and the largest diameter of invalid space imposes a typically unknown lower bound on mutation rate to guarantee reachability of the global optimum.

Closed Feasible Solution Space

Employ a closed feasible solution space which guarantees that the initial population consists of feasible solutions only and all evolutionary operations on feasible solutions are guaranteed to result in feasible solutions. Typically a combination of custom representation, initialization, recombination, and mutation is employed to achieve this.

Pros Guarantees valid solutions and reduces the effective search space.

Cons Requires problem dependent design of closed evolutionary operators and the largest diameter of invalid space imposes a typically unknown lower bound on mutation rate to guarantee reachability of the global optimum.

When to use If you can guarantee global maximum reachability from the starting population and a computationally efficient repair algorithm is not available.

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness
- Penalty Function

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness
- Penalty Function
- Repair Function

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness
- Penalty Function
- Repair Function
- Closed Feasible Solution Space

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness
- Penalty Function
- Repair Function
- Closed Feasible Solution Space
- Feasible Decoder

Feasible Decoder

Employ a decoder function that maps genotype space to phenotype space such that the phenotypes are guaranteed to be feasible even when the genotypes are infeasible. Typically this involves mapping multiple different genotypes to the same phenotype.

Feasible Decoder

Employ a decoder function that maps genotype space to phenotype space such that the phenotypes are guaranteed to be feasible even when the genotypes are infeasible. Typically this involves mapping multiple different genotypes to the same phenotype.

Pros Guarantees valid solutions while imposing no limitations on the search of genotype space.

Feasible Decoder

Employ a decoder function that maps genotype space to phenotype space such that the phenotypes are guaranteed to be feasible even when the genotypes are infeasible. Typically this involves mapping multiple different genotypes to the same phenotype.

- Pros** Guarantees valid solutions while imposing no limitations on the search of genotype space.
- Cons** Requires problem dependent design of a high quality decoder function and for the typical case of a non-injective decoder function the effective search space (genotype space) can be far larger than what is implied by phenotype space.

Feasible Decoder

Employ a decoder function that maps genotype space to phenotype space such that the phenotypes are guaranteed to be feasible even when the genotypes are infeasible. Typically this involves mapping multiple different genotypes to the same phenotype.

Pros Guarantees valid solutions while imposing no limitations on the search of genotype space.

Cons Requires problem dependent design of a high quality decoder function and for the typical case of a non-injective decoder function the effective search space (genotype space) can be far larger than what is implied by phenotype space.

When to use If you cannot guarantee global maximum reachability from the starting population and a high quality decoder function can be designed.

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness

Constraint Satisfaction Techniques

There are a variety of techniques to solve Constraint Satisfaction Problems with EAs. Seven such techniques are:

- Ignore Constraints
- Kill Infeasible Solutions
- Assign arbitrarily low fitness
- Penalty Function
- Repair Function
- Closed Feasible Solution Space
- Feasible Decoder