

Heuristic Algorithms

Master's Degree in Computer Science/Mathematics

Roberto Cordone

DI - Università degli Studi di Milano



- Schedule: **Wednesday 13.30 - 16.30 in classroom Alfa**
Thursday 09.30 - 12.30 in classroom Alfa
- Office hours: **on appointment**
- E-mail: **roberto.cordone@unimi.it**
- Web page: **<https://homes.di.unimi.it/cordone/courses/2026-ae/2026-ae.html>**
- Ariel site: **<https://myariel.unimi.it/course/view.php?id=7439>**

Constructive metaheuristics

The constructive algorithms have strong limitations on many problems
What can be done, without abandoning the general scheme?

Iterate the scheme to generate many (potentially) different solutions

- the efficiency decreases: the computational times are summed
- the effectiveness increases: the best solution is returned

The trade-off must be carefully tuned

The iterated scheme can apply

- multi-start, that is different algorithm at each iteration $l = 1, \dots, \ell$
(this requires to define multiple \mathcal{F}_{A_l} and φ_{A_l})

but it is more flexible to apply metaheuristics, that exploit

- randomisation (operations based on a random seed), as in the case of semigreedy algorithms, *GRASP* and *Ant System* (partly, *ART*)
- memory (operations based on the solutions of previous iterations), as in the case of *ART*, cost perturbation and *Ant System*

Termination condition

The iterated scheme can ideally proceed for an infinite time

In practice, one uses termination conditions that can be “absolute”

- 1 a given **total number of iterations of the basic scheme**
- 2 a given **total execution time**
- 3 a given **target value of the objective**

or “relative” to the profile of f^*

- 1 a given **number of iterations of the basic scheme without improving f^***
- 2 a given **execution time without improving f^***
- 3 a given **minimum ratio between the improvement of f^* and the number of iterations of the basic scheme or the execution time**
(e.g.: f^* improves less than 1% in the last 1 000 iterations)

Fair comparisons require absolute conditions

Multi-start (or **restart**) is a classical, very simple and natural approach:

- define **different search spaces** $\mathcal{F}_{A^{[l]}}$ and **selection criteria** $\varphi_{A^{[l]}}(i, x)$
- apply each resulting algorithm $A^{[l]}$ to obtain $x^{[l]}$
- return the **best solution** $x = \arg \min_{l=1, \dots, \ell} f(x^{[l]})$

A typical case is to tune $\varphi_A(i, x)$ with numerical parameters μ

The construction graph can model this situation

- including all nodes and arcs admitted by at least one algorithm $A^{[l]}$:

$$\mathcal{F}_A = \bigcup_{l=1}^{\ell} \mathcal{F}_{A^{[l]}}$$

- setting arc weights depending on l : $\varphi_A(i, x, l) = \varphi_{A^{[l]}}(i, x)$
- setting an infinite arc weight for the arcs that are forbidden in a specific algorithm $A^{[l]}$: $\varphi_A(i, x, l) = +\infty$

Example

A whole family of heuristics for the *TSP* can be obtained setting:

- insertion criterium:

$$i_k^* = \arg \min_{i \in \{1, \dots, |x|\}} \gamma_{i,k} = \mu_1 (c_{s_i,k} + c_{k,s_{i+1}}) + (1 - \mu_1) (-c_{s_i,s_{i+1}})$$

where $\mu_1 \in [0; 1]$ tunes the relative strength of the

- increase in cost due to the added node k
 - decrease in cost due to the removed edge (s_i, s_{i+1})
- selection criterium:

$$k^* = \arg \min_{k \in N \setminus N_x} \varphi_A(k, x) = \mu_2 \gamma_{i_k^*, k} + \mu_3 d(x, k) + (1 - \mu_2 - \mu_3) (-d(x, k))$$

where $\mu_2, \mu_3 \in [0; 1]$ tune the relative strength (and sign) of the

- increase in cost due to the added node k
- distance of the added node k from the current circuit x

to get *CI* for $\mu = (1/2, 1, 0)$, *NI* for $\mu = (1/2, 0, 1)$, *FI* for $\mu = (1/2, 0, 0)$

Constructive metaheuristics

The main constructive metaheuristics are

- 1 **Adaptive Research Technique (ART)** or Tabu Greedy:
forbid some moves based on the solutions of the previous iterations

$$\min_{i \in \Delta_{A^{[l]}}^+(x)} \varphi_A(i, x)$$

with $\Delta_{A^{[l]}}^+(x) = \left\{ i \in B \setminus x : x \cup \{i\} \in \mathcal{F}^{[l]}(x_A^{[1]}, \dots, x_A^{[l-1]}) \subseteq \mathcal{F} \right\}$
This is much less popular than the other two

- 2 **semigreedy** and **GRASP**: use a randomised selection criterium

$$\min_{i: x \cup \{i\} \in \mathcal{F}} \varphi_A^{[l]}(i, x, \omega^{[l]})$$

- 3 **Ant System (AS)**: use a randomised selection criterium
depending on the solutions of the previous iterations

$$\min_{i: x \cup \{i\} \in \mathcal{F}} \varphi_A^{[l]}(i, x, \omega^{[l]}, x_A^{[1]}, \dots, x_A^{[l-1]})$$

New information on the arcs of the construction graph guides the search

The **ART** uses memory, the **GRASP** randomisation, the **AS** both

Adaptive Research Technique

It was proposed by Patterson et al. (1998) for the *CMSTP*

When deceptively good elements are included in the first steps
the final solution can be quite bad

Aiming to avoid that

- the roll-out approach makes a look-ahead on each possible element
(*but a single step can be insufficient to identify the misleading ones*)
- the *ART* forbids some elements to drive subset x on the right path
in the search space

(*how to identify the misleading elements?*)

The prohibitions are temporary, with an expiration time of L iterations;
otherwise, building feasible solutions would become impossible

Adaptive Research Technique

Define a basic constructive heuristic A

Let T_i be the **starting iteration of the prohibition** for each element $i \in B$ and x^* be the best solution found

Set $T_i = -\infty$ for all $i \in B$ to indicate that no element is forbidden

At each iteration $l \in \{1, \dots, \ell\}$

- 1 **apply heuristic A forbidding all elements i such that $l \leq T_i + L$**
(all prohibitions older than L iterations automatically expire);
let $x^{[l]}$ be the resulting solution
- 2 if $x^{[l]}$ is better than x^* , set $x^* := x^{[l]}$ and save $T_i - l$ for all $i \in B$
- 3 **decide which elements to forbid and set $T_i = l$ for them:**
each element is forbidden with probability π (any better ideas?)
- 4 make **minor tweaks to L , π or T_i**

At the end, return x^*

Example: ART for the SCP

$$c \quad \begin{array}{|c|c|c|c|c|} \hline 25 & 6 & 8 & 24 & 12 \\ \hline \end{array}$$

$$A \quad \begin{array}{|c|c|c|c|c|} \hline 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 \\ \hline \end{array}$$

Let $L = 2$, $\pi = 0.15$, pseudorandom numbers $0.1, 0.9, 0.4, 0.5, 0.1, 0.2, \dots$

- 1 the basic heuristic finds solution $x^{[1]} = \{2, 3, 5, 4\}$ of cost $f(x^{[1]}) = 50$;
forbid column 2 (because $0.1 \leq \pi < 0.9$, 0.4 and 0.5)
- 2 the basic heuristic finds solution $x^{[2]} = \{3, 1\}$ (2 is forbidden)
of cost $f(x^{[2]}) = 33$; forbid column 3 (because $0.1 \leq \pi < 0.2$)
- 3 the basic heuristic finds solution $x^{[3]} = \{1\}$ (3 and 2 are forbidden) of cost
 $f(x^{[3]}) = 25$, that is optimal
- 4 ...

An unlucky sequence could forbid column 1 at step 2

Parameter tuning

The ART has three basic parameters

- the **total number of iterations** ℓ (*tuned mainly by the available time*)
- the **length** L of the prohibition
- the **probability** π of the prohibition

How to assign effective values to the parameters?

The experimental comparison of different values is necessary but complex

- ① it requires **long experimental campaigns**, because **the number of configurations grows combinatorially with**
 - the number of parameters
 - the number of tested values for each parameter
(*the more sensitive the result, the more values must be tested*)
- ② it risks **overfitting**, that is **labelling as absolutely good values which are good only on the benchmark instances considered**

The excess of parameters is an undesirable aspect, and often reveals an **insufficient study of the problem and of the algorithm**

More on this point later

Diversification and intensification

Diversification aims to obtain different solutions at every iteration

The *ART* achieves it forbidding elements of the previous solutions

An excessive diversification can hinder the discovery of the optimum

Intensification aims to focus the search on the more promising subsets

Diversification and intensification play complementary roles

Their relative strength can be tuned through the parameters based on

- **problem data:** assign
 - a smaller probability π_i to be forbidden
 - a shorter expiration time L_i of the prohibitionto promising elements (e. g., cheaper ones)
- **memory:**
 - assign a smaller π_i or L_i (i is never forbidden when $\pi_i = 0$ or $L_i = 0$) to promising elements (e. g., appearing in the best known solutions)
 - periodically restart the algorithm with the $T_i - l$ values associated with the best known solution, instead of $T_i = -\infty$

Semi-greedy heuristics

A nonexact constructive algorithm has at least one step t which builds a subset $x^{(t)}$ not included in any optimal solution

Since the element selected is the best according to the selection criterium

$$i^* = \arg \min_{i \in \Delta_A^+(x)} \varphi_A(i, x)$$

necessarily $\varphi_A(i, x)$ is incorrect, but probably not completely wrong

The **semi-greedy algorithm** (Hart and Shogan, 1987) assumes that elements that lead to the optimum are very good for $\varphi_A(i, x)$, even if not strictly the best

How to know which one?

If it is not possible to refine $\varphi_A(i, x)$

- define a suitable probability distribution on $\Delta_A^+(x)$ favouring the elements with the best values of $\varphi_A(i, x)$
- select $i^*(\omega)$ according to the distribution function

Semi-greedy heuristics

Since the set of alternative choices is finite, this means to assign

- probability $\pi_A(i, x)$ to arc $(x, x \cup \{i\})$ of the construction graph (with a sum equal to 1 for the outgoing arcs of each node)

$$\sum_{i \in \Delta_A^+(x)} \pi_A(i, x) = 1 \quad \text{for all } x \in \mathcal{F}_A : \Delta_A^+(x) \neq \emptyset$$

- higher probabilities to the better elements for the selection criterium

$$\varphi_A(i, x) \leq \varphi_A(j, x) \Leftrightarrow \pi_A(i, x) \geq \pi_A(j, x)$$

for each $i, j \in \Delta_A^+(x), x \in \mathcal{F}_A$

This heuristic approach has important properties

- it can reach an optimal solution if there is a path from \emptyset to X^*
(this is a basic condition)
- it can be reapplied several times obtaining different solutions and the probability to reach the optimum grows gradually
(each iteration decreases the probability of missing an optimal path)

Convergence to the optimum

The probability of

- following a path γ is the product of the probabilities on the arcs

$$\prod_{(y, y \cup \{i\} \in \gamma)} \pi_A(i, y)$$

- obtaining a solution x is the sum of those of the paths Γ_x reaching x

$$\sum_{\gamma \in \Gamma_x} \prod_{(y, y \cup \{i\} \in \gamma)} \pi_A(i, y)$$

This implies that the probability to reach the optimum:

- 1 is nonzero if and only if there exists a path of nonzero probability from \emptyset to X^*
- 2 increases as $\ell \rightarrow +\infty$
(the probability of not reaching it decreases gradually)

It tends to 1 for probabilistically approximatively complete algorithms

Convergence to the optimum

In this context, a *random walk* is a constructive metaheuristic in which all the arcs going out of the same node have equal probability

- it finds a path to the optimum with probability 1 (if one exists)
- the time required can be extremely long

The exhaustive algorithm is exact and requires finite time

A deterministic constructive heuristic sets all probabilities to zero except for those on the arcs of a single path

- it finds the optimum only if it enjoys specific properties
- it finds the optimum in a single run

Randomised heuristics that favour promising arcs and penalise the others

- accelerate the average convergence time
- decrease the guarantee of convergence in the worst case

There is a trade-off between expected and worst result

Arcs with zero probability can block the path to the optimum

Arcs with probability converging to zero reduce the probability to find it

GRASP, that is **Greedy Randomised Adaptive Search Procedure** (Feo and Resende, 1989) is a sophisticated variant of the semi-greedy heuristic

- *Greedy* indicates that it uses a constructive basic heuristic
- *Randomised* indicates that the basic heuristic makes random steps
- *Adaptive* indicates that the heuristic uses an adaptive selection criterium $\varphi_A(i, x)$, depending also on x (not strictly necessary)
- *Search* indicates that it alternates the constructive heuristic and an exchange heuristic (differently from the semi-greedy approach)

The use of auxiliary exchange heuristics allows strongly better results

This aspect will be investigated in the following lessons

What probability function?

Several functions $\pi_A(i, x)$ are monotonous with respect to $\varphi_A(i, x)$


$$\varphi_A(i, x) \leq \varphi_A(j, x) \Leftrightarrow \pi_A(i, x) \geq \pi_A(j, x)$$

- **uniform probability**: each arc going out of x has the same $\pi_A(i, x)$; the algorithm performs a **random path** in \mathcal{F}_A (**random walk**)
- **Heuristic-Biased Stochastic Sampling (HBSS)**:
 - sort the arcs going out of x by nonincreasing values of $\varphi_A(i, x)$
 - assign a decreasing probability according to the position in the order based on a simple scheme (linear, exponential, ecc...)
- **Restricted Candidate List (RCL)**:
 - sort the arcs going out of x by nonincreasing values of $\varphi_A(i, x)$
 - insert the best arcs in a list (*How many?*)
 - assign uniform probability to the arcs of the list, zero to the others

The most common strategy is the *RCL*, even if the zero probability arcs potentially cancel the global convergence to the optimum

Common probability functions

Suppose that at the current step $|\Delta_A^+(x)| = 10$ elements can be added



Definition of the RCL

Two main strategies are used to define the *RCL*

- **cardinality**: the RCL includes the best μ elements of $\Delta_A^+(x)$, where $\mu \in \{1, \dots, |\Delta_A^+(x)|\}$ is a parameter fixed by the user
 - $\mu = 1$ yields the constructive basic heuristic
 - $\mu = |B|$ (i. e., $|\Delta_A^+(x)|$ for each x) yields the *random walk*
- **value**: the RCL includes all the elements of $\Delta_A^+(x)$ whose value is between φ_{\min} and $(1 - \mu)\varphi_{\min} + \mu\varphi_{\max}$ where

$$\varphi_{\min}(x) = \min_{i \in \Delta_A^+(x)} \varphi_A(i, x) \quad \varphi_{\max}(x) = \max_{i \in \Delta_A^+(x)} \varphi_A(i, x)$$

and $\mu \in [0; 1]$ is a parameter fixed by the user

- $\mu = 0$ yields the constructive basic heuristic
- $\mu = 1$ yields the *random walk*

General scheme of GRASP

Algorithm GRASP(l)

$x^* := \emptyset$; $f^* := +\infty$; { Best solution found so far }

For $l = 1$ to ℓ *do*

{ Constructive heuristic with random steps }

$x := \emptyset$;

While $\Delta_A^+(x) \neq \emptyset$ *do*

$\varphi_i := \varphi_A(i, x)$ for each $i \in \Delta_A^+(x)$

$\pi := \text{AssignProbabilities}(\Delta_A^+(x), \varphi, \mu)$;

$i := \text{RandomExtract}(\Delta_A^+(x), \pi)$;

$x := x \cup \{i\}$;

EndWhile;

$x := \text{Search}(x)$; { Exchange heuristic }

If $x \in X$ and $f(x) < f^*$ *then* $x^* := x$; $f^* := f(x)$;

EndFor;

Return (x^*, f^*) ;

Example: GRASP for the SCP

$$c \quad \begin{array}{|c|c|c|c|c|} \hline 25 & 6 & 8 & 24 & 12 \\ \hline \end{array}$$

$$A \quad \begin{array}{|c|c|c|c|c|} \hline 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 1 \\ \hline \end{array}$$

Let $\mu = 2$ and the pseudorandom sequence be: 0.6, 0.8, ...

- 1 start with the empty subset: $x^{(0)} = \emptyset$
- 2 build the *RCL* with columns 2 ($\varphi_2 = 2$) and 3 ($\varphi_3 = 4$);
select column 3 (because $0.6 > 1/2$);
- 3 build the *RCL* with columns 2 ($\varphi_2 = 3$) and 1 ($\varphi_3 = 6.25$);
select column 1 (because $0.8 > 1/2$);
- 4 the solution obtained is $x = \{3, 1\}$ of cost $f(x) = 33$

With $\mu = 2$, the optimal solution cannot be obtained; with $\mu = 3$ it can

The optimum is found with $\mu = 2$ if a destructive phase is applied

Reactive parameter tuning

Once again there are parameters to tune:

- the number of iterations ℓ
- the value μ determining the size of the *RCL*

An idea to exploit memory is to **learn from the previous results**

- 1 select m configurations of parameters μ_1, \dots, μ_m and set $\ell_r = \ell/m$
- 2 run each configuration μ_r for ℓ_r iterations
- 3 evaluate the mean $\bar{f}(\mu_r)$ of the results obtained with μ_r
- 4 **update the number of iterations** ℓ_r for each μ_r based on $\bar{f}(\mu_r)$

$$\ell_r = \frac{\frac{1}{\bar{f}(\mu_r)}}{\sum_{s=1}^m \frac{1}{\bar{f}(\mu_s)}} \ell \quad \text{for } r = 1, \dots, m$$

increasing it for the more effective configurations

- 5 repeat the whole process, going back to point 2, for R times

Other schemes use scores based on the number of best known results

Cost perturbation methods

Instead of forbidding/forcing some choices, or modifying their probability, it is possible to **modify the appeal of the available choices**

Given a basic constructive heuristic A , at each step of iteration l

- **tune the selection criterium $\varphi_A(i, x)$ with a factor $\tau_A^{[l]}(i, x)$**

$$\psi_A^{[l]}(i, x) = \frac{\varphi_A(i, x)}{\tau_A^{[l]}(i, x)}$$

- **update $\tau_A^{[l]}(i, x)$ based on the previous solutions $x^{[1]}, \dots, x^{[l-1]}$**


The elements with a better $\varphi_A(i, x)$ tend to be favoured, but $\tau_A^{[l]}(i, x)$ tunes this effect, promoting

- **intensification** if $\tau_A^{[l]}(i, x)$ increases for the most frequent elements; this favours **solutions similar to the previous ones**
- **diversification** if $\tau_A^{[l]}(i, x)$ decreases for the most frequent elements; this favours **solutions different from the previous ones**

Ant Colony Optimization

It was devised by Dorigo, Maniezzo and Colorni in 1991 drawing inspiration from the social behaviour of ants

Stigmergy = indirect communication among different agents who are influenced by the results of the actions of all agents



Each agent is an application of the basic constructive heuristic

- **it leaves a trail on the data** depending on the solution generated
- **it performs choices influenced by the trails** left by the other agents

The choices of the agent have also a **random component**

As in the semi-greedy heuristic

- a basic constructive heuristic A is given
- each step performs a partially random choice

Differently from the semi-greedy heuristic

- each iteration l runs h times heuristic A (population)
- all the choices of $\Delta_A^+(x)$ are feasible (there is no RCL)
- the probability $\pi_A(i, x)$ depends on
 - ① the selection criterium $\varphi_A(i, x)$
 - ② auxiliary information $\tau_A(i, x)$ denoted as **trail** produced in previous iterations (sometimes by other agents in the same iteration)

The trail is **uniform at first** ($\tau_A(i, x) = \tau_0$), and later tuned

- increasing it to favour promising choices
- decreasing it to avoid repetitive choices

For the sake of simplicity, the trail $\tau_A(i, x)$ is not associated to each arc $(x, x \cup \{i\})$, but is the same for blocks of arcs (e.g., depending only on i)

Random choice

Instead of selecting the best element according to criterium $\varphi_A(i, x)$, i is extracted from $\Delta_A^+(x)$ with probability

$$\pi_A(i, x) = \frac{\tau_A(i, x)^{\mu_\tau} \eta_A(i, x)^{\mu_\eta}}{\sum_{j \in \Delta_A^+(x)} \tau_A(j, x)^{\mu_\tau} \eta_A(j, x)^{\mu_\eta}}$$

where

- the denominator normalizes the probability
- the **visibility** is the auxiliary function

$$\eta_A(i, x) = \begin{cases} \varphi_A(i, x) & \text{for maximisation problems} \\ \frac{1}{\varphi_A(i, x)} & \text{for minimisation problems} \end{cases}$$

The promising choices have larger visibility

- the parameters μ_τ and μ_η tune the weights of the two terms

Balancing given and learned information

The original Ant System tunes the probabilities

$$\pi_A(i, x) = \frac{\tau_A(i, x)^{\mu_\tau} \eta_A(i, x)^{\mu_\eta}}{\sum_{j \in \Delta_A^+(x)} \tau_A(j, x)^{\mu_\tau} \eta_A(j, x)^{\mu_\eta}}$$

with parameters μ_η and μ_τ that control the amount of randomness

- $\mu_\eta \approx 0$ and $\mu_\tau \approx 0$ push towards **randomness**
- **large values of μ_η and μ_τ** push towards **determinism**
(favour $\arg \max_{i \in \Delta^+(x)} \tau_A(i, x)^{\mu_\tau} \eta_A(i, x)^{\mu_\eta}$)

and the relative weight of the data and of memory

- $\mu_\eta \gg \mu_\tau$ **favours the data**, simulating the basic constructive heuristic which makes sense when the known solutions are not very significant
- $\mu_\eta \ll \mu_\tau$ **favours memory**, keeping close to the previous solutions which makes sense when the known solutions are very significant

Balancing given and learned information

The **Ant Colony System** variant splits the selection into two phases

- 1 decide the selection procedure to use
 - with probability q , choose i **deterministically**
 - with probability $(1 - q)$, choose i **stochastically**

where parameter q tunes the randomness

- $q \approx 0$ favours **random choices**
 - $q \approx 1$ favours **deterministic choices**
- 2 apply the selection procedure
 - the deterministic one selects the best element

$$i^* = \arg \max_{i \in \Delta^+(x)} \tau_A(i, x) \eta_A(i, x)^{\mu_\eta}$$

- the stochastic one select a random element with probabilities

$$\pi_A(i, x) = \frac{\tau_A(i, x) \eta_A(i, x)^{\mu_\eta}}{\sum_{j \in \Delta_A^+(x)} \tau_A(j, x) \eta_A(j, x)^{\mu_\eta}}$$

where parameter μ_η tunes the relative weight of data and memory

- $\mu_\eta \gg 1$ favours the **data**
- $\mu_\eta \ll 1$ favours **memory**

(setting $\mu_\tau = 1$ as a form of normalisation)

Trail update

At each iteration ℓ

- 1 run h instances of the basic heuristic A
- 2 select a subset $\tilde{X}^{[\ell]}$ of the solutions obtained, in order to favour their elements in the following iterations
- 3 update the trail according to the formula

$$\tau_A(i, x) := (1 - \rho) \tau_A(i, x) + \rho \sum_{y \in \tilde{X}^{[\ell]}: i \in y} F_A(y)$$

where

- $\rho \in [0; 1]$ is an oblivion parameter
- $F_A(y)$ is a fitness function expressing the quality of solution y (such that $F > \tau$: e.g., $F(y) = Q/f(y)$ for a suitable constant Q)

The purpose of the update is to

- 1 increase the trail on the elements of specific solutions ($y \in \tilde{X}^{[\ell]}$)
- 2 decrease the trail on the other elements

The oblivion parameter

$$\tau_A(i, x) := (1 - \rho) \tau_A(i, x) + \rho \sum_{y \in \tilde{X}^{[i]: i \in y}} F_A(y)$$

The oblivion parameter $\rho \in [0; 1]$ tunes the behaviour of the algorithm:

- **diversification**: a **high oblivion** ($\rho \approx 1$) **cancels the current trail** based on the intuition that
 - the solutions obtained are not trustworthy
 - different solutions should be explored
- **intensification**: a **low oblivion** ($\rho \approx 0$) **preserves the current trail** based on the intuition that
 - the solutions obtained are trustworthy
 - similar solutions should be explored

Selection of the influential solutions

$\tilde{X}^{[l]}$ collects the solutions around which the search will be intensified

- the classical Ant System considers all the solutions of iteration $l - 1$
- the elitist methods consider the best known solutions
 - the best solution of iteration $l - 1$
 - the best solution of all iterations $< l$

The elitist methods

- find better results in shorter time
- require additional mechanisms to avoid premature convergence

Some variants of the Ant System

- **MAX – MIN Ant System**: imposes on the trail a limited range of values $[\tau_{\min}; \tau_{\max}]$, experimentally tuned
- **HyperCube Ant Colony Optimization (HC-ACO)**: normalizes the trail between 0 and 1
- **Ant Colony System**: updates the trail on two levels
 - the **global update** (already seen) modifies it at each iteration ℓ
The purpose is to intensify the search
 - the **local update** updates the trail at each application g of the basic heuristic in order to discourage identical choices in the following

$$\tau_A(i, x) := (1 - \rho) \tau_A(i, x) \quad \text{for each } i \in x_A^{[l, g]}$$

The purpose is to diversify the search

General scheme of the *Ant System*

Algorithm AntSystem(l)

$x^* := \emptyset$; $f^* := +\infty$; { Best solution found so far }

$\tau_A := \tau_0$;

For $l = 1$ to ℓ do

$\tilde{X}^{[l]} := \emptyset$;

 For $g = 1$ to h do

$x := A(l, \tau_A)$; { Basic heuristic with random steps and memory }

$x := \text{Search}(x)$; { Exchange heuristic }

$\tau_A := \text{LocalUpdate}(\tau_A, x)$; { Local update of the trail }

 If $f(x) < f^*$ then $x^* := x$; $f^* := f(x)$;

$\tilde{X}^{[l]} := \text{Update}(\tilde{X}^{[l]}, x)$;

 EndFor;

$\tau_A := \text{GlobalUpdate}(\tau_A, \tilde{X}^{[l]})$; { Global update of the trail }

EndFor;

Return (x^*, f^*) ;

Convergence to the optimum

Some variants of the Ant System converge to the optimum with probability 1 (Gutjahr, 2002)

The analysis is based on the construction graph

- the trail $\tau_A(i, x)$ is laid down on the arcs $(x, x \cup \{i\})$
- no information from the data is used, that is $\eta_A(i, x) \equiv 1$ (*this strange assumption simplifies the proof, but is not necessary*)
- $\tau^{[l]}$ is the trail function at the beginning of iteration l
- $\gamma^{[l]}$ is the best path on the graph at the end of iteration l ,
- $(\tau^{[l]}, \gamma^{[l-1]})$ is the state of a nonhomogeneous Markov process:
 - the probability of each state depends only on the previous iteration
 - the process is nonhomogeneous because the dependency varies with l

The proof concludes that for $l \rightarrow +\infty$, with probability 1

- 1 the best path found γ is one of the optimal paths in \mathcal{F}
- 2 the trail τ tends to a maximum along γ , to zero on the other arcs

provided that a suitable parameter tuning is adopted

First variant with global convergence

The trail is updated with a variable coefficient of oblivion

$$\tau^{[l]}(i, x) := \begin{cases} (1 - \rho^{[l-1]}) \tau^{[l-1]}(i, x) + \rho^{[l-1]} \frac{1}{|\gamma^{[l-1]}|} & \text{if } (x, x \cup \{i\}) \in \gamma^{[l-1]} \\ (1 - \rho^{[l-1]}) \tau^{[l-1]}(i, x) & \text{otherwise} \end{cases}$$

where $\gamma^{[l-1]}$ is the best path found in the graph up to iteration $l - 1$ and $|\gamma^{[l-1]}|$ is the number of its arcs
(to normalise the trail)

If the oblivion decreases slowly enough

$$\rho^{[l]} \leq 1 - \frac{\log l}{\log(l+1)} \quad \text{and} \quad \sum_{l=0}^{+\infty} \rho^{[l]} = +\infty$$

then with probability 1 the state converges to (τ^*, γ^*) , where

- γ^* is an optimal path in the construction graph
- $\tau^*(i, x) = \frac{1}{|\gamma^*|}$ for $(x, x \cup \{i\}) \in \gamma^*$, 0 otherwise

Second variant with global convergence

Alternatively, if the oblivion ρ remains constant,
but the trail is forced a slowly decreasing minimum threshold

$$\tau(i, x) \geq \frac{c_l}{\log(l+1)} \quad \text{and} \quad \lim_{l \rightarrow +\infty} c_l \in (0; 1)$$

then with probability 1 the state converges to (τ^*, γ^*)

Here the oblivion is restricted by the minimum threshold

In practice, all algorithms proposed so far in the literature

- associate the trail to groups of arcs $(x, x \cup \{i\})$
(e.g., to element i)
- use constant values for parameters ρ and τ_{\min}

therefore do not guarantee convergence

The trail τ , and therefore π , can tend to zero on every optimal path