Advanced Optimization
Lecture 1: Discrete Optimization

November 20, 2018

Master AIC

Université Paris-Saclay, Orsay, France

Anne Auger

INRIA Saclay – Ile-de-France

Dimo Brockhoff

INRIA Saclay – Ile-de-France

Typical scenario in the continuous case:

Optimize $f : \Omega \subset \mathbb{R}^n \mapsto \mathbb{R}^k$

$x \in \mathbb{R}^n$ → [black box] → $f(x) \in \mathbb{R}^k$

derivatives not available or not useful

Lecture Goals

As in introductory lecture: always examples and small exercises to learn "on-the-fly" the concepts and fundamentals

Overall goals:

❶ give more details on a few important aspects of blackbox optimization

❷ prepare you better for a potential Master's thesis (PhD thesis) on the topic

Hence, I will give later on some details on our available projects

Course Overview

	Date		Topic
1	Tue, 20.11.2018	Dimo	Randomized Algorithms for Discrete Problems
2	Tue, 27.11.2018	Dimo	Exercise: The Travelling Salesperson Problem
3	Tue, 4.12.2018	Dimo	Evolutionary Multiobjective Optimization I
4	Tue, 11.12.2018	Dimo	Evolutionary Multiobjective Optimization II
	vacation		
5	Tue, 8.1.2019	Dimo	Looking at Data
6	Tue, 15.1.2019	Anne	Continuous Optimization I
7	Tue, 29.1.2019	Anne	Continuous Optimization II
	???		oral presentations (individual time slots)

all lectures from 14h00 till 17h15

here in E107 in Nov/Dec and in E105 in January

No Exam...

Since the idea is to prepare you for your Master's thesis:

- we don't have a written exam ☺

- but instead work towards research:

    - each student is assigned a scientific paper

    - which is to be read, understood, critically questioned, and finally presented

    - summarize the paper in a short abstract in your own words

    - oral presentations in the end of the course (15min presentation + 15min oral "exam")

    - also the exercises are (closer to) research questions than before

Additional Offer: Solving COCO Issues

In addition, we plan to offer an upgrade of your grade (by 1 point max.) if you happen to solve an issue from the COCO issue tracker!

`https://github.com/numbbo/coco/issues/`

deadline:
day of last lecture: January 29, 2019

Details on the Paper Project/Oral Presentation

- no written exam but instead each student is assigned a scientific paper (list online and on next slide)

    - to be read, understood, critically questioned, and presented

    - maximal 3 students per paper

    - decision made until *November 27, 2018 (next lecture)*

- summary of the paper in a short abstract in your own words

    - handed in via a web form until *January 15, 2019, 23h59* (https://docs.google.com/forms/d/1E8t5beYzkFkxWU2gLAVbouTRrPUWtWfBcRrS8OYA6N8/edit)

    - 4000 characters max.

- individual oral presentations at the end of the course

    - 15min presentation + 15min oral "exam"

    - dates to be decided with you towards the end of the lecture

    - slides to be sent by email to us until last lecture (*Jan. 29, 2019*)

The List of Papers

All papers are relevant to current research in Randopt but the starred ones indicate possible *concrete research projects as follow-ups.

1*) Two-dimensional subset selection for hypervolume and epsilon-indicator

2*) RM-MEDA: A regularity model-based multiobjective estimation of distribution algorithm.

3*) A universal catalyst for first-order optimization.

4*) Optimized Approximation Sets for Low-Dimensional Benchmark Pareto Fronts.

5*) Covariance matrix adaptation for multi-objective optimization.

101) Efficient optimization of many objectives by approximation-guided evolution.

102) PISA - A Platform and Programming Language Independent Interface for Search Algorithms.

103) A Mean-Variance Optimization Algorithm.

104) Theoretical foundation for CMA-ES from information geometry perspective.

105) Population Size Adaptation for the CMA-ES Based on the Estimation Accuracy of the Natural Gradient.

106) CMA-ES with Optimal Covariance Update and Storage Complexity.

107) Exponential natural evolution strategies.

Today's Lecture

❶ present open projects

❷ randomized search heuristics in the discrete domain

❸ exercise: Pure Random Search (PRS) and the (1+1)EA

❶ present open projects

Potential Research Topics for Master's/PhD Theses

`http://randopt.gforge.inria.fr/thesisprojects/`

THESIS PROJECTS

[[start]]

Home

Welcome!

On this page, you will find various current technical and scientific projects in the field of stochastic blackbox optimization proposed by Anne Auger, Dimo Brockhoff, and Nikolaus Hansen at Inria. Depending on the subject, the projects can be Bachelor, Master's, or PhD theses, or related to internships and might be carried out in close relationship with external collaborators, including companies.

If you are interested in (stochastic) blackbox optimization but your favorite topic is not mentioned here, feel free to contact us personally. We might always have other topics in mind, which range from theoretical studies to algorithm design but which have not yet been formalized here.

Current Openings

- Stopping Criteria for Multiobjective Optimizers (Master's project)
- Various technical projects around the COCO platform (Internships/Bachelor)
- Large-scale Stochastic Black-box Optimization (Master's project)
- The Orbit Algorithm for Expensive Numerical Blackbox Problems (Bachelor/Master's project)

Previous Announcements

- Adaptive Stochastic Search Algorithms for Constrained Optimization (Master's thesis project)
- Data Mining Performance Results of Numerical Optimizers (Master's thesis project)
- General Constraint Handling in the Stochastic Numerical Optimization Algorithm CMA-ES (CIFRE PhD)
- Designing Variants of the Covariance Matrix Adaptation Evolution Strategy to Handle Multiobjective Blackbox

Projects without the involvement of companies:

- stopping criteria in multiobjective optimization
- large-scale variants of CMA-ES
- algorithms for expensive optimization based on CMA-ES

all above: relatively flexible between theoretical (e.g. proofs of convergence) and practical projects

Coco-related:

- implementing and benchmarking existing algorithms
  [new test suites for constraints & mixed-integer available]
- recommendations for noisy optimization

not all subject ideas online:
please contact us if you are interested!

❷ randomized search heuristics in the discrete domain

[mainly a recap of the introductory lecture]

Reminder: Discrete Optimization

Context discrete optimization:

- discrete variables
- or optimization over discrete structures (e.g. graphs)
- search space often finite, but typically too large for enumeration
- → need for smart algorithms

Algorithms for discrete problems:

- typically problem-specific
- but some general concepts are repeatedly used:
  - greedy algorithms
  - dynamic programming
  - randomized search heuristics
  - branch and bound

Remark: Coping with Difficult Problems

Exact

- brute-force often too slow
- better strategies such as dynamic programming & branch and bound
- still: often exponential runtime

Approximation Algorithms

- guarantee of low run time
- guarantee of high quality solution
- obstacle: difficult to prove these guarantees

Heuristics

- intuitive algorithms
- guarantee to run in short time
- often no guarantees on solution quality

Motivation General Search Heuristics

- often, problem complicated and not much time available to develop a problem-specific algorithm
- search heuristics are a good choice:
  - relatively easy to implement
  - easy to adapt/change/improve
    - e.g. when the problem formulation changes in an early product design phase
    - or when slightly different problems need to be solved over time
    - remember blackbox scenario
- search heuristics are also often "any-time", i.e. give a feasible solution early on which is then improved throughout the algorithm run → might be important in practice

Idea behind (Randomized) Local Search:

- explore the local neighborhood of the current solution (randomly)

Pure Random Search:

- go to randomly chosen neighbor

First Improvement Local Search:

- go to first (randomly) chosen neighbor which is better

Best Improvement strategy:

- always go to the best neighbor
- not random anymore
- computationally expensive if neighborhood large

One class of (bio-inspired) stochastic optimization algorithms:
Evolutionary Algorithms (EAs)

1859

- Class of optimization algorithms originally inspired by the idea of biological evolution
- selection, mutation, recombination

Metaphors

Classical Optimization	Evolutionary Computation
variables or parameters	variables or chromosomes
candidate solution vector of decision variables / design variables / object variables	individual, offspring, parent
set of candidate solutions	population
objective function loss function cost function error function	fitness function
iteration	generation

initialization

evaluation

potential parents

best individual

stop?

parents

evaluation

offspring
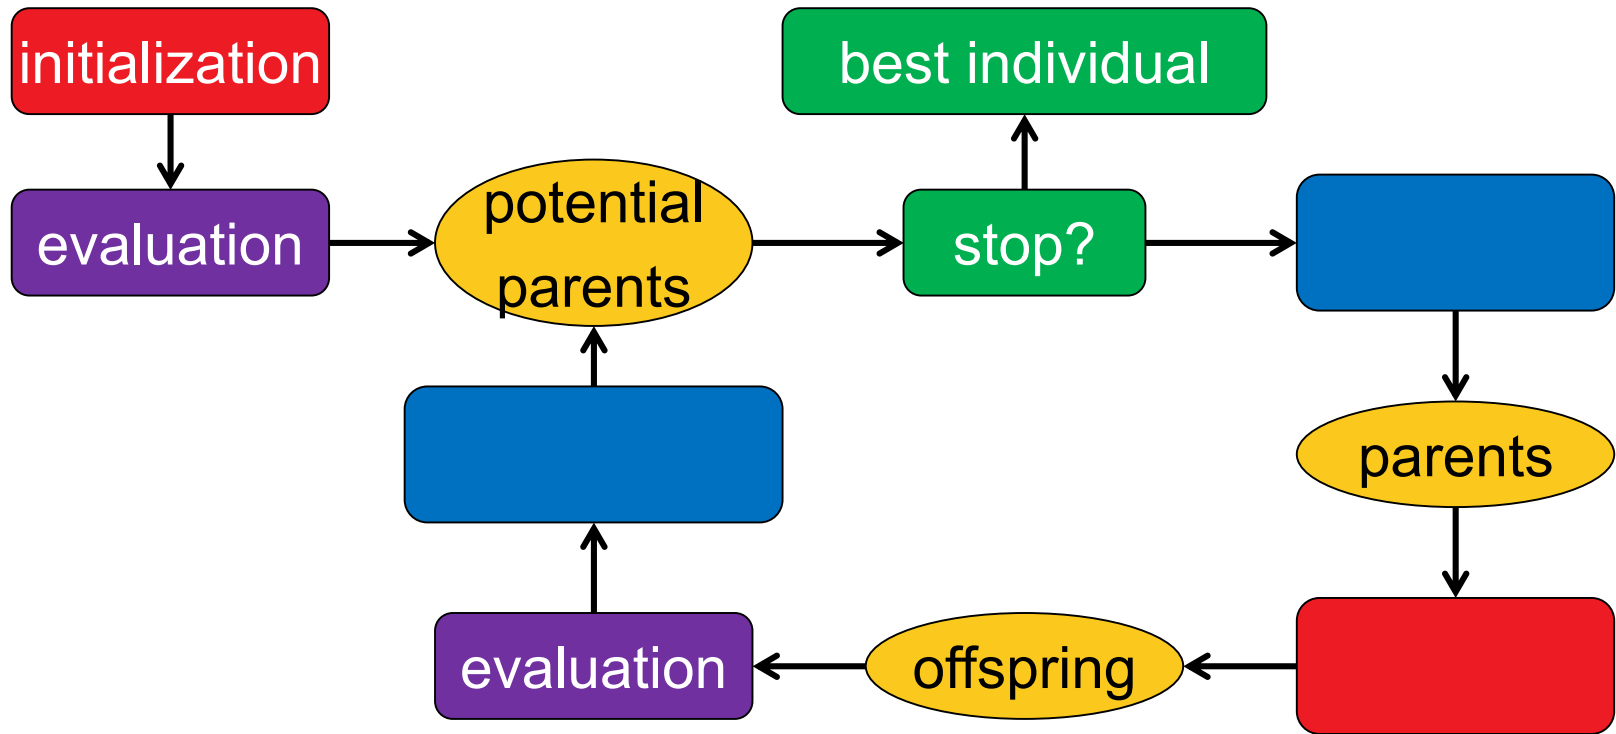
Question:
which pieces are missing?
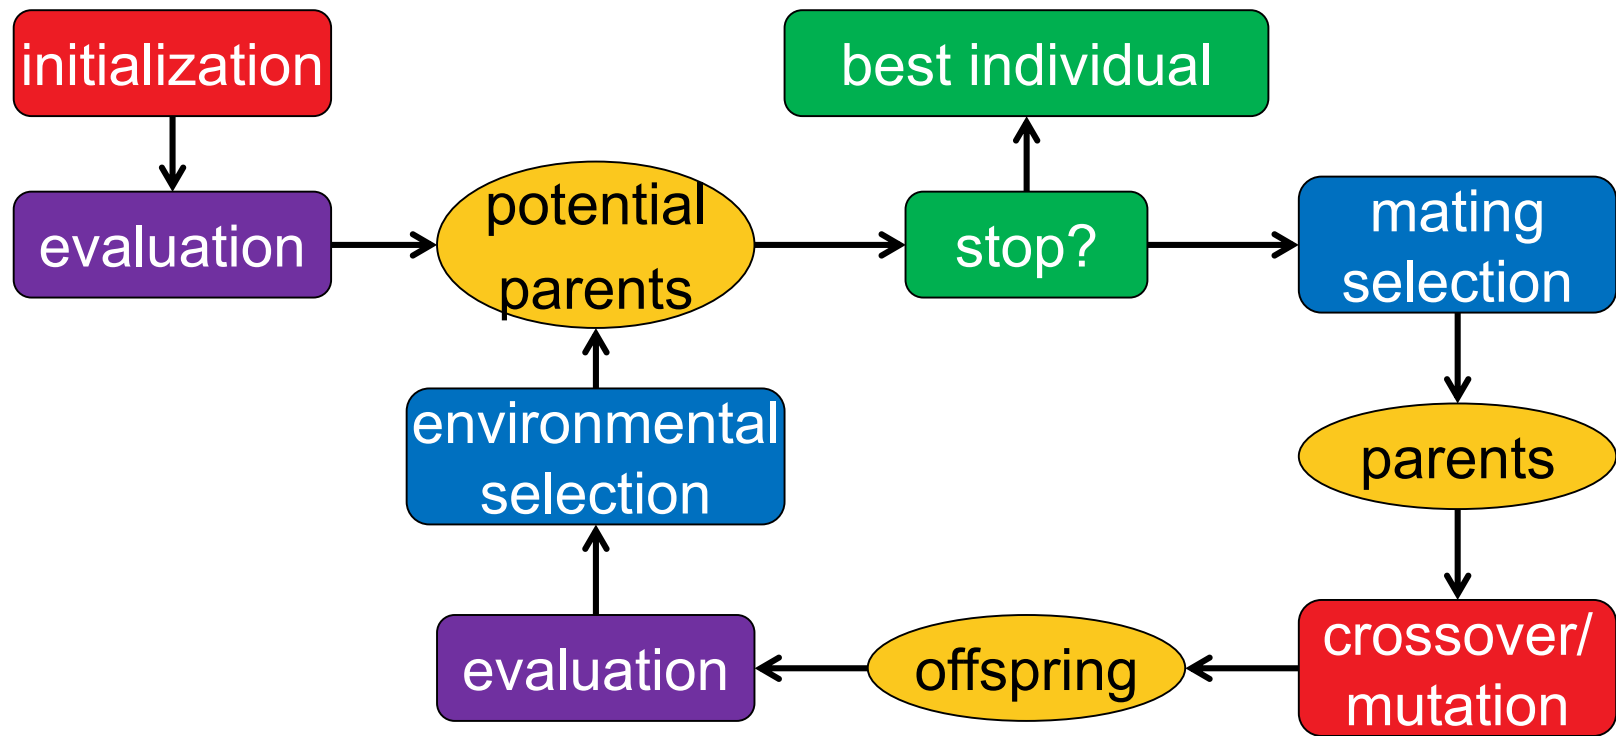
Genetic Algorithms (GA)

J. Holland 1975 and D. Goldberg (USA)

$$\Omega = \{0, 1\}^n$$

Evolution Strategies (ES)

I. Rechenberg and H.P. Schwefel, 1965 (Berlin)

$$\Omega = \mathbb{R}^n$$

Evolutionary Programming (EP)

L.J. Fogel 1966 (USA)

Genetic Programming (GP)

J. Koza 1990 (USA)

$$\Omega = \text{space of all programs}$$

nowadays one umbrella term: evolutionary algorithms

The genotype – phenotype mapping

- related to the question: how to come up with a fitness of each individual from the representation?

- related to DNA vs. actual animal (which then has a fitness)

fitness of an individual not always = f(x)

- include constraints

- include diversity

- others

- but needed: always a total order on the solutions

  - even in the case of multiobjective optimization!

Reminder of some EA parts

Selection is the major determinant for specifying the trade-off between exploitation and exploration

Selection is either

stochastic or deterministic

e.g. fitness proportional

$$Q_i = \frac{f(x_i)}{\sum_{j=1}^{\mu} f(x_j)}$$

Disadvantage: depends on scaling of f

e.g. via a tournament

Selection is the major determinant for specifying the trade-off between exploitation and exploration

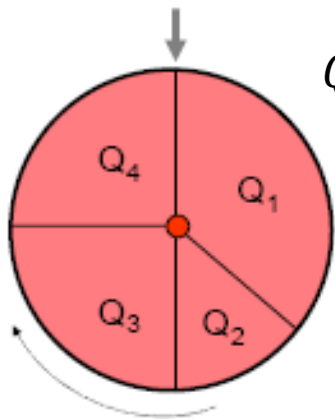Selection is either

stochastic   or   deterministic

e.g. fitness proportional   e.g. (μ+λ), (μ,λ)

$$Q_i = \frac{f(x_i)}{\sum_{j=1}^{\mu} f(x_j)}$$

e.g. via a tournament

Question: what is plus and what is comma selection?

Selection is the major determinant for specifying the trade-off between exploitation and exploration

Selection is either

stochastic   or   deterministic

e.g. fitness proportional

$$Q_i = \frac{f(x_i)}{\sum_{j=1}^{\mu} f(x_j)}$$

e.g. (μ+λ), (μ,λ)

best μ from offspring *and* parents

best μ from offspring only

e.g. via a tournament

Mating selection (selection for variation): usually stochastic
Environmental selection (selection for survival): often deterministic

Variation aims at generating new individuals on the basis of those individuals selected for mating

Variation = Mutation and Recombination/Crossover

mutation: *mut:* $\Omega \rightarrow \Omega$

recombination: *recomb:* $\Omega^r \rightarrow \Omega^s$ where $r \geq 2$ and $s \geq 1$

- choice always depends on the problem and the chosen representation
- however, there are some operators that are applicable to a wide range of problems and tailored to standard representations such as vectors, permutations, trees, etc.

Variation Operators: Guidelines

Two desirable properties for mutation operators:

- every solution can be generation from every other with a probability greater than 0 ("exhaustiveness")

- $d(x, x') < d(x, x'') => Prob(\text{mut}(x) = x') > Prob(\text{mut}(x) = x'')$ ("locality")

Desirable property of recombination operators ("in-between-ness"):

$$x'' = \text{recomb}(x, x') \Rightarrow d(x'', x) \leq d(x, x') \wedge d(x'', x') \leq d(x, x')$$

Swap:

1 2 3 4 5 6 → 1 4 3 2 5 6

swap

Scramble:

1 2 3 4 5 6 → 1 3 4 2 5 6

rearrange

Invert:

1 2 3 4 5 6 → 1 4 3 2 5 6

reverse

Insert:

1 2 3 4 5 6 → 1 4 2 3 5 6

move

1-point crossover

n-point crossover

uniform crossover

choose each bit independently from one parent or another

Assumptions:

- search space Ω = set of all bitstrings of length n ($\Omega = \{0,1\}^n$)
- minimization of objective function $f: \{0,1\}^n \to \mathbb{R}$

Algorithm:

- init: sample a point x uniformly at random in Ω
- while not happy:
  - $y \leftarrow \text{mutate}(x)$
  - if $f(y) \leq f(x)$:
    - $x \leftarrow y$

Variants:

- Randomized Local Search (RLS): mutate = sample a Hamming neighbor uniformly at random (= flip exactly one bit)
- (1+1)-ES: mutate = flip each bit with probability $1/n$

Conclusions

- EAs are generic algorithms (randomized search heuristics, meta-heuristics, ...) for black box optimization

  no or almost no assumptions on the objective function

- They allow for an easy and rapid implementation and therefore to find good solutions fast

  we will see this next week

  also: easy to incorporate (recommended!)
  problem-specific knowledge to improve the algorithm

❸ exercise: Pure Random Search (PRS) and the (1+1)EA

Exercise:
Pure Random Search and the (1+1)EA

```
http://www.cmap.polytechnique.fr/~dimo.brockhoff/
advancedOptSaclay/2018/exercises.php
```

$$\Omega = \{0, 1\}^n$$

$$f(x) = \sum_{i=1}^{n} x_i$$

$$P_i = \{x | f(x) = i\}$$

$$\Omega = \bigcup_{i=0}^{n} P_i$$

1^n

$1^i 0^{n-i}$

i ones

$0^{n-i} 1^i$

fitness level "i"

0^n

$$T = \inf\{t \in \mathbb{N}, X_t = (1, \ldots, 1)\}$$

X_t estimate of solution at iteration t

T_i time to leave fitness level "i"

$$E(T) \leq \sum_{i=0}^{n-1} E(T_i)$$

$\text{Prob}(\text{leave}P_i) \geq \frac{1}{n}(1 - \frac{1}{n})^{n-1} \times (n - i)$ (proba to flip one and only one of the $(n - i)$ remaining 0)

$$(1 - \frac{1}{n})^{n-1} \geq \frac{1}{e}$$

$\text{Prob}(\text{leave}P_i) \geq \frac{n-i}{en}$

$$E(T_i) \leq \frac{en}{n-i}$$

$$E(T) \leq \sum_{i=0}^{n-1} \frac{en}{n-i} \leq e.n(\log n + 1)$$