# Introduction to Optimization <br> Lecture 4: Gradient-based Optimization 

October 5, 2018<br>TC2 - Optimisation<br>Université Paris-Saclay

Dimo Brockhoff
Inria Saclay - Ile-de-France

## Course Overview



## Course Overview

| 1 | Mon, 17.9.2018 | Monday's lecture: introduction, example problems, problem types <br> Thu, 20.9.2018 |
| :--- | :--- | :--- |
| groups defined via wiki <br> everybody went (actively!) through the Getting Started part of <br> github.com/numbbo/coco |  |  |
| 2 | Fri, 21.9.2018 | lecture "Benchmarking", final adjustments of groups everybody can run <br> and postprocess the example experiment (~1h for final questions/help <br> during the lectur) |
| 3 | Fri, 28.9.2018 | lecture "Introduction to Continuous Optimization" |
| 4 | Fri, 5.10.2018 | lecture "Gradient-Based Algorithms" + DFO |

## Details on Continuous Optimization Lectures

Introduction to Continuous Optimization

- examples (from ML / black-box problems)
- typical difficulties in optimization

Mathematical Tools to Characterize Optima

- reminders about differentiability, gradient, Hessian matrix
- unconstraint optimization
- first and second order conditions
- convexity
- constraint optimization

Gradient-based Algorithms

- quasi-Newton method (BFGS)
- [DFO trust-region method]

Learning in Optimization / Stochastic Optimization

- CMA-ES (adaptive algorithms / Information Geometry)
- PhD thesis possible on this topic
method strongly related to ML / new promising research area interesting open questions


## Constrained Optimization

## Equality Constraint

## Objective:

Generalize the necessary condition of $\nabla f(x)=0$ at the optima of f when $f$ is in $\mathcal{C}^{1}$, i.e. is differentiable and its differential is continuous

Theorem:
Be $U$ an open set of $(E,\| \|)$, and $f: U \rightarrow \mathbb{R}, g: U \rightarrow \mathbb{R}$ in $\mathcal{C}^{1}$. Let $a \in E$ satisfy

$$
\left\{\begin{array}{c}
f(a)=\inf \left\{f(x) \mid x \in \mathbb{R}^{n}, g(x)=0\right\} \\
g(a)=0
\end{array}\right.
$$

i.e. $a$ is optimum of the problem

If $\nabla g(a) \neq 0$, then there exists a constant $\lambda \in \mathbb{R}$ called Lagrange multiplier, such that
$\nabla f(a)+\lambda \nabla g(a)=0 \quad$ Euler - Lagrange equation
i.e. gradients of $f$ and $g$ in $a$ are colinear

## Geometrical Interpretation Using an Example

## Exercise:

Consider the problem

$$
\inf \left\{f(x, y) \mid(x, y) \in \mathbb{R}^{2}, g(x, y)=0\right\}
$$

$$
f(x, y)=y-x^{2} \quad g(x, y)=x^{2}+y^{2}-1=0
$$

1) Plot the level sets of $f$, plot $g=0$
2) Compute $\nabla f$ and $\nabla g$
3) Find the solutions with $\nabla f+\lambda \nabla g=0$
equation solving with 3 unknowns ( $x, y, \lambda$ )
4) Plot the solutions of 3 ) on top of the level set graph of 1 )

## Interpretation of Euler-Lagrange Equation

Intuitive way to retrieve the Euler-Lagrange equation:

- In a local minimum $a$ of a constrained problem, the hypersurfaces (or level sets) $f=f(a)$ and $g=0$ are necessarily tangent (otherwise we could decrease $f$ by moving along $g=0$ ).
- Since the gradients $\nabla f(a)$ and $\nabla g(a)$ are orthogonal to the level sets $f=f(a)$ and $g=0$, it follows that $\nabla f(a)$ and $\nabla g(a)$ are colinear.


## Generalization to More than One Constraint

## Theorem

- Assume $f: U \rightarrow \mathbb{R}$ and $g_{k}: U \rightarrow \mathbb{R}(1 \leq k \leq p)$ are $\mathcal{C}^{1}$.
- Let $a$ be such that

$$
\left\{\begin{array}{r}
f(a)=\inf \left\{f(x) \mid x \in \mathbb{R}^{n}, \quad g_{k}(x)=0, \quad 1 \leq k \leq p\right\} \\
g_{k}(a)=0 \text { for all } 1 \leq k \leq p
\end{array}\right.
$$

- If $\left(\nabla g_{k}(a)\right)_{1 \leq k \leq p}$ are linearly independent, then there exist $p$ real constants $\left(\lambda_{k}\right)_{1 \leq k \leq p}$ such that

$$
\nabla f(a)+\sum_{k=1 \uparrow}^{p} \lambda_{k} \nabla g_{k}(a)=0
$$

again: $a$ does not need to be global but local minimum

## The Lagrangian

- Define the Lagrangian on $\mathbb{R}^{n} \times \mathbb{R}^{p}$ as

$$
\mathcal{L}\left(x,\left\{\lambda_{k}\right\}\right)=f(x)+\sum_{k=1}^{p} \lambda_{k} g_{k}(x)
$$

- To find optimal solutions, we can solve the optimality system
$\left\{\right.$ Find $\left(x,\left\{\lambda_{k}\right\}\right) \in \mathbb{R}^{n} \times \mathbb{R}^{p}$ such that $\nabla f(x)+\sum_{k=1}^{p} \lambda_{k} \nabla g_{k}(x)=0$

$$
g_{k}(x)=0 \text { for all } 1 \leq k \leq p
$$

$$
\Leftrightarrow\left\{\begin{array}{c}
\text { Find }\left(x,\left\{\lambda_{k}\right\}\right) \in \mathbb{R}^{n} \times \mathbb{R}^{p} \text { such that } \nabla_{x} \mathcal{L}\left(x,\left\{\lambda_{k}\right\}\right)=0 \\
\nabla_{\lambda_{k}} \mathcal{L}\left(x,\left\{\lambda_{k}\right\}\right)(x)=0 \text { for all } 1 \leq k \leq p
\end{array}\right.
$$

## Inequality Constraint: Definitions

Let $\mathcal{U}=\left\{x \in \mathbb{R}^{n} \mid g_{k}(x)=0\right.$ (for $k \in E$ ), $g_{k}(x) \leq 0$ (for $k \in I$ ) $\}$.

Definition:
The points in $\mathbb{R}^{n}$ that satisfy the constraints are also called feasible points.

## Definition:

Let $a \in \mathcal{U}$, we say that the constraint $g_{k}(x) \leq 0$ (for $k \in I$ ) is active in $a$ if $g_{k}(a)=0$.

## Inequality Constraint: Karush-Kuhn-Tucker Theorem

Theorem (Karush-Kuhn-Tucker, KKT):
Let $U$ be an open set of $(E,\| \|)$ and $f: U \rightarrow \mathbb{R}, g_{k}: U \rightarrow \mathbb{R}$, all $\mathcal{C}^{1}$
Furthermore, let $a \in U$ satisfy
$\int f(a)=\inf \left(f(x) \mid x \in \mathbb{R}^{n}, g_{k}(x)=0\right.$ (for $\left.k \in E\right), g_{k}(x) \leq 0$ (for $k \in \mathrm{I}$ )

$$
\begin{array}{cc}
g_{k}(a)=0(\text { for } k \in E) & \text { also works again for } a \\
g_{k}(a) \leq 0(\text { for } k \in I) & \text { being a local minimum }
\end{array}
$$

Let $I_{a}^{0}$ be the set of constraints that are active in $a$. Assume that $\left(\nabla g_{k}(a)\right)_{k \in E \cup I_{a}^{0}}$ are linearly independent.
Then there exist $\left(\lambda_{k}\right)_{1 \leq k \leq p}$ that satisfy

$$
\left\{\begin{array}{c}
\nabla f(a)+\sum_{k=1}^{p} \lambda_{k} \nabla g_{k}(a)=0 \\
g_{k}(a)=0(\text { for } k \in E) \\
g_{k}(a) \leq 0(\text { for } k \in I) \\
\left.\lambda_{k} \geq 0 \text { (for } k \in I_{a}^{0}\right) \\
\lambda_{k} g_{k}(a)=0(\text { for } k \in E \cup I)
\end{array}\right.
$$

## Inequality Constraint: Karush-Kuhn-Tucker Theorem

Theorem (Karush-Kuhn-Tucker, KKT):
Let $U$ be an open set of $(E,\| \|)$ and $f: U \rightarrow \mathbb{R}, g_{k}: U \rightarrow \mathbb{R}$, all $\mathcal{C}^{1}$
Furthermore, let $a \in U$ satisfy
$\left\{\begin{array}{c}f(a)=\inf \left(f(x) \mid x \in \mathbb{R}^{n}, g_{k}(x)=0(\text { for } k \in E), g_{k}(x) \leq 0(\text { for } k \in \mathrm{I})\right. \\ g_{k}(a)=0(\text { for } k \in E) \\ g_{k}(a) \leq 0(\text { for } k \in I)\end{array}\right.$
Let $I_{a}^{0}$ be the set of constraints that are active in $a$. Assume that $\left(\nabla g_{k}(a)\right)_{k \in E \cup I_{a}^{0}}$ are linearly independent.
Then there exist $\left(\lambda_{k}\right)_{1 \leq k \leq p}$ that satisfy

$$
\left\{\begin{array}{c}
\nabla f(a)+\sum_{k=1}^{p} \lambda_{k} \nabla g_{k}(a)=0 \\
g_{k}(a)=0(\text { for } k \in E) \\
g_{k}(a) \leq 0(\text { for } k \in I) \\
\lambda_{k} \geq 0\left(\text { for } k \in I_{a}^{0}\right) \\
\lambda_{k} g_{k}(a)=0(\text { for } k \in E \cup I)
\end{array}\right.
$$

either active constraint or $\lambda_{k}=0$

## Descent Methods

## Descent Methods

## General principle

(1) choose an initial point $x_{0}$, set $t=0$
(2) while not happy

- choose a descent direction $\boldsymbol{d}_{t} \neq 0$
- line search:
- choose a step size $\sigma_{t}>0$
- set $\boldsymbol{x}_{t+1}=\boldsymbol{x}_{t}+\sigma_{t} \boldsymbol{d}_{t}$
- set $t=t+1$


## Remaining questions

- how to choose $\boldsymbol{d}_{t}$ ?
- how to choose $\sigma_{t}$ ?


## Gradient Descent

Rationale: $\boldsymbol{d}_{t}=-\nabla f\left(\boldsymbol{x}_{t}\right)$ is a descent direction indeed for $f$ differentiable

$$
\begin{aligned}
f(x-\sigma \nabla f(x)) & =f(x)-\sigma\|\nabla f(x)\|^{2}+o(\sigma\|\nabla f(x)\|) \\
< & f(x) \text { for } \sigma \text { small enough }
\end{aligned}
$$

## Step-size

- optimal step-size: $\sigma_{t}=\operatorname{argmin} f\left(\boldsymbol{x}_{t}-\sigma \nabla f\left(\boldsymbol{x}_{t}\right)\right)$
- Line Search: total or partial optimization w.r.t. $\sigma$ Total is however often too "expensive" (needs to be performed at each iteration step)
Partial optimization: execute a limited number of trial steps until a loose approximation of the optimum is found. Typical rule for partial optimization: Armijo rule (see next slides)


## Typical stopping criterium:

norm of gradient smaller than $\epsilon$

## The Armijo-Goldstein Rule

Choosing the step size:

- Only to decrease $f$-value not enough to converge (quickly)
- Want to have a reasonably large decrease in $f$


## Armijo-Goldstein rule:

- also known as backtracking line search
- starts with a (too) large estimate of $\sigma$ and reduces it until $f$ is reduced enough
- what is enough?
- assuming a linear $f$ e.g. $m_{k}(x)=f\left(x_{k}\right)+\nabla f\left(x_{k}\right)^{T}\left(x-x_{k}\right)$
- expected decrease if step of $\sigma_{k}$ is done in direction $\boldsymbol{d}$ : $\sigma_{k} \nabla f\left(x_{k}\right)^{T} \boldsymbol{d}$
- actual decrease: $f\left(x_{k}\right)-f\left(x_{k}+\sigma_{k} \boldsymbol{d}\right)$
- stop if actual decrease is at least constant times expected decrease (constant typically chosen in [0, 1])


## The Armijo-Goldstein Rule

## The Actual Algorithm:

Input: descent direction d, point $\mathbf{x}$, objective function $f(\mathbf{x})$ and its gradient $\nabla f(\mathbf{x})$, parameters $\sigma_{0}=10, \theta \in[0,1]$ and $\beta \in(0,1)$
Output: step-size $\sigma$
Initialize $\sigma: \sigma \leftarrow \sigma_{0}$
while $f(\mathbf{x}+\sigma \mathbf{d})>f(\mathbf{x})+\theta \sigma \nabla f(\mathbf{x})^{T} \mathbf{d}$ do
$\sigma \leftarrow \beta \sigma$
end while

Armijo, in his original publication chose $\beta=\theta=0.5$.
Choosing $\theta=0$ means the algorithm accepts any decrease.

## The Armijo-Goldstein Rule

## Graphical Interpretation


linear approximation
(expected decrease)

## The Armijo-Goldstein Rule

## Graphical Interpretation


linear approximation
(expected decrease)

## The Armijo-Goldstein Rule

## Graphical Interpretation


linear approximation
(expected decrease)

## Newton Algorithm

## Newton Method

- descent direction: $-\left[\nabla^{2} f\left(x_{k}\right)\right]^{-1} \nabla f\left(x_{k}\right)$ [so-called Newton direction]
- The Newton direction:
- minimizes the best (locally) quadratic approximation of $f$ :

$$
\tilde{f}(x+\Delta x)=f(x)+\nabla f(x)^{T} \Delta x+\frac{1}{2}(\Delta x)^{T} \nabla^{2} f(x) \Delta \mathrm{x}
$$

- points towards the optimum on $f(x)=\left(x-x^{*}\right)^{T} A\left(x-x^{*}\right)$
- however, Hessian matrix is expensive to compute in general and its inversion is also not easy
quadratic convergence

$$
\text { (i.e. } \left.\lim _{k \rightarrow \infty} \frac{\left|x_{k+1}-x^{*}\right|}{\left|x_{k}-x^{*}\right|^{2}}=\mu>0\right)
$$

## Remark: Affine Invariance

Affine Invariance: same behavior on $f(x)$ and $f(A x+b)$ for $A \in$
$\operatorname{GLn}(\mathbb{R})=$ set of all invertible $n \times n$ matrices over $\mathbb{R}$

- Newton method is affine invariant
see http: //users.ece.utexas.edu/~cmcaram/EE381V_2012F/ Lecture_6_Scribe_Notes.final.pdf
- same convergence rate on all convex-quadratic functions
- Gradient method not affine invariant


## Quasi-Newton Method: BFGS

$x_{t+1}=x_{t}-\sigma_{t} H_{t} \nabla f\left(x_{t}\right)$ where $H_{t}$ is an approximation of the inverse Hessian

## Key idea of Quasi Newton:

successive iterates $x_{t}, x_{t+1}$ and gradients $\nabla f\left(x_{t}\right), \nabla f\left(x_{t+1}\right)$ yield second order information

$$
\begin{gathered}
q_{t} \approx \nabla^{2} f\left(x_{t+1}\right) p_{t} \\
\text { where } p_{t}=x_{t+1}-x_{t} \text { and } q_{t}=\nabla f\left(x_{t+1}\right)-\nabla f\left(x_{t}\right)
\end{gathered}
$$

Most popular implementation of this idea: Broyden-Fletcher-Goldfarb-Shanno (BFGS)

- default in MATLAB's fminunc and python's scipy.optimize.minimize


## Conclusions

I hope it became clear...
...what are the difficulties to cope with when solving numerical optimization problems
in particular dimensionality, non-separability and ill-conditioning
...what are gradient and Hessian
...what is the difference between gradient and Newton direction
...and that adapting the step size in descent algorithms is crucial.

## Derivative-Free Optimization

## Derivative-Free Optimization (DFO)

DFO = blackbox optimization


## Why blackbox scenario?

- gradients are not always available (binary code, no analytical model, ...)
- or not useful (noise, non-smooth, ...)
- problem domain specific knowledge is used only within the black box, e.g. within an appropriate encoding
- some algorithms are furthermore function-value-free, i.e. invariant wrt. monotonous transformations of $f$.


## Derivative-Free Optimization Algorithms

- (gradient-based algorithms which approximate the gradient by finite differences)
- coordinate descent
- pattern search methods, e.g. Nelder-Mead
- surrogate-assisted algorithms, e.g. NEWUOA or other trustregion methods
- other function-value-free algorithms
- typically stochastic
- evolution strategies (ESs) and Covariance Matrix Adaptation Evolution Strategy (CMA-ES)
- differential evolution
- particle swarm optimization
- simulated annealing


## Downhill Simplex Method by Nelder and Mead

While not happy do:
[assuming minimization of $f$ and that $x_{1}, \ldots, x_{n+1} \in \mathbb{R}^{n}$ form a simplex]

1) Order according to the values at the vertices: $f\left(x_{1}\right) \leq f\left(x_{2}\right) \leq \cdots \leq f\left(x_{n+1}\right)$
2) Calculate $x_{o}$, the centroid of all points except $x_{n+1}$.
3) Reflection

Compute reflected point $x_{r}=x_{o}+\alpha\left(x_{o}-x_{n+1}\right)(\alpha>0)$
If $x_{r}$ better than second worst, but not better than best: $x_{n+1}:=x_{r}$, and go to 1 )
4) Expansion

If $x_{r}$ is the best point so far: compute the expanded point

$$
x_{e}=x_{o}+\gamma\left(x_{r}-x_{o}\right)(\gamma>0)
$$

If $x_{e}$ better than $x_{r}$ then $x_{n+1}:=x_{e}$ and go to 1)
Else $x_{n+1}:=x_{r}$ and go to 1)
Else (i.e. reflected point is not better than second worst) continue with 5)
5) Contraction (here: $\left.f\left(x_{r}\right) \geq f\left(x_{n}\right)\right)$

Compute contracted point $x_{c}=x_{o}+\rho\left(x_{n+1}-x_{o}\right)(0<\rho \leq 0.5)$
If $f\left(x_{c}\right)<f\left(x_{n+1}\right): x_{n+1}:=x_{c}$ and go to 1)
Else go to 6)
6) Shrink
$x_{i}=x_{1}+\sigma\left(x_{i}-x_{1}\right)$ for all $i \in\{2, \ldots, n+1\}(\sigma<1)$ and go to 1$)$
J. A Nelder and R. Mead (1965). "A simplex method for function minimization".

Computer Journal. 7: 308-313. doi:10.1093/comjn//7.4.308

## Nelder-Mead: Reflection

2) Calculate $x_{o}$, the centroid of all points except $x_{n+1}$.
3) Reflection

Compute reflected point $x_{r}=x_{o}+\alpha\left(x_{o}-x_{n+1}\right)(\alpha>0)$
If $x_{r}$ better than second worst, but not better than best: $x_{n+1}:=x_{r}$, and go to 1)


## Nelder-Mead: Reflection

2) Calculate $x_{o}$, the centroid of all points except $x_{n+1}$.
3) Reflection

Compute reflected point $x_{r}=x_{o}+\alpha\left(x_{o}-x_{n+1}\right)(\alpha>0)$
If $x_{r}$ better than second worst, but not better than best: $x_{n+1}:=x_{r}$, and go to 1)


## Nelder-Mead: Reflection

2) Calculate $x_{o}$, the centroid of all points except $x_{n+1}$.
3) Reflection

Compute reflected point $x_{r}=x_{o}+\alpha\left(x_{o}-x_{n+1}\right)(\alpha>0)$
If $x_{r}$ better than second worst, but not better than best: $x_{n+1}:=x_{r}$, and go to 1)


## Nelder-Mead: Expansion

2) Calculate $x_{o}$, the centroid of all points except $x_{n+1}$.
3) Expansion

If $x_{r}$ is the best point so far: compute the expanded point

$$
x_{e}=x_{o}+\gamma\left(x_{r}-x_{o}\right)(\gamma>0)
$$

If $x_{e}$ better than $x_{r}$ then $x_{n+1}:=x_{e}$ and go to 1)
Else $x_{n+1}:=x_{r}$ and go to 1)
Else (i.e. reflected point is not better than second worst) continue with 5)


## Nelder-Mead: Expansion

2) Calculate $x_{o}$, the centroid of all points except $x_{n+1}$.
3) Expansion

If $x_{r}$ is the best point so far: compute the expanded point

$$
x_{e}=x_{o}+\gamma\left(x_{r}-x_{o}\right)(\gamma>0)
$$

If $x_{e}$ better than $x_{r}$ then $x_{n+1}:=x_{e}$ and go to 1)
Else $x_{n+1}:=x_{r}$ and go to 1)
Else (i.e. reflected point is not better than second worst) continue with 5)


## Nelder-Mead: Expansion

2) Calculate $x_{o}$, the centroid of all points except $x_{n+1}$.
3) Expansion

If $x_{r}$ is the best point so far: compute the expanded point

$$
x_{e}=x_{o}+\gamma\left(x_{r}-x_{o}\right)(\gamma>0)
$$

If $x_{e}$ better than $x_{r}$ then $x_{n+1}:=x_{e}$ and go to 1)
Else $x_{n+1}:=x_{r}$ and go to 1)
Else (i.e. reflected point is not better than second worst) continue with 5)


## Nelder-Mead: Expansion

2) Calculate $x_{o}$, the centroid of all points except $x_{n+1}$.
3) Expansion

If $x_{r}$ is the best point so far: compute the expanded point

$$
x_{e}=x_{o}+\gamma\left(x_{r}-x_{o}\right)(\gamma>0)
$$

If $x_{e}$ better than $x_{r}$ then $x_{n+1}:=x_{e}$ and go to 1)
Else $x_{n+1}:=x_{r}$ and go to 1)
Else (i.e. reflected point is not better than second worst) continue with 5)


## Nelder-Mead: Expansion

2) Calculate $x_{o}$, the centroid of all points except $x_{n+1}$.
3) Contraction (here: $\left.f\left(x_{r}\right) \geq f\left(x_{n}\right)\right)$

Compute contracted point $x_{c}=x_{o}+\rho\left(x_{n+1}-x_{o}\right)(0<\rho \leq 0.5)$ If $f\left(x_{c}\right)<f\left(x_{n+1}\right): x_{n+1}:=x_{c}$ and go to 1) Else go to 6)


## Nelder-Mead: Expansion

2) Calculate $x_{o}$, the centroid of all points except $x_{n+1}$.
3) Contraction (here: $\left.f\left(x_{r}\right) \geq f\left(x_{n}\right)\right)$

Compute contracted point $x_{c}=x_{o}+\rho\left(x_{n+1}-x_{o}\right)(0<\rho \leq 0.5)$ If $f\left(x_{c}\right)<f\left(x_{n+1}\right): x_{n+1}:=x_{c}$ and go to 1) Else go to 6)


## Nelder-Mead: Expansion

2) Calculate $x_{o}$, the centroid of all points except $x_{n+1}$.
3) Shrink
$x_{i}=x_{1}+\sigma\left(x_{i}-x_{1}\right)$ for all $i \in\{2, \ldots, n+1\}$ and go to 1)


## Nelder-Mead: Expansion

2) Calculate $x_{o}$, the centroid of all points except $x_{n+1}$.
3) Shrink
$x_{i}=x_{1}+\sigma\left(x_{i}-x_{1}\right)$ for all $i \in\{2, \ldots, n+1\}$ and go to 1)

## Nelder-Mead: Standard Parameters

- reflection parameter : $\alpha=1$
- expansion parameter: $\gamma=2$
- contraction parameter: $\rho=\frac{1}{2}$
- shrink paremeter: $\sigma=\frac{1}{2}$
some visualizations of example runs can be found here: https://en.wikipedia.org/wiki/Nelder\�\�\�Mead_method


## stochastic algorithms

## Stochastic Search Template

A stochastic blackbox search template to minimize $f: \mathbb{R}^{\boldsymbol{n}} \rightarrow \mathbb{R}$ Initialize distribution parameters $\theta$, set population size $\lambda \in \mathbb{N}$
While happy do:

- Sample distribution $P(\boldsymbol{x} \mid \theta) \rightarrow \boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{\lambda} \in \mathbb{R}^{n}$
- Evaluate $\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{\lambda}$ on $f$
- Update parameters $\theta \leftarrow F_{\theta}\left(\theta, \boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{\lambda}, f\left(\boldsymbol{x}_{1}\right), \ldots, f\left(\boldsymbol{x}_{\lambda}\right)\right)$
- All depends on the choice of $P$ and $F_{\theta}$
deterministic algorithms are covered as well
- In Evolutionary Algorithms, $P$ and $F_{\theta}$ are often defined implicitly via their operators.


## Generic Framework of an Evolutionary Algorithm


stochastic operators
"Darwinism"
stopping criteria

Nothing else: just interpretation change

## CMA-ES in a Nutshell

## The CMA-ES

Input: $m \in \mathbb{R}^{n}, \sigma \in \mathbb{R}_{+}, \lambda$
Initialize: $\mathbf{C}=\mathbf{I}$, and $p_{\mathrm{c}}=\mathbf{0}, p_{\sigma}=\mathbf{0}$,
Set: $c_{\mathrm{c}} \approx 4 / n, c_{\sigma} \approx 4 / n, c_{1} \approx 2 / n^{2}, c_{\mu} \approx \mu_{w} / n^{2}, c_{1}+c_{\mu} \leq 1, d_{\sigma} \approx 1+\sqrt{\frac{\mu_{w}}{n}}$, and $w_{i=1 \ldots \lambda}$ such that $\mu_{w}=\frac{1}{\sum_{i=1}^{\mu} w_{i}^{2}} \approx 0.3 \lambda$
While not terminate

$$
\begin{aligned}
& \boldsymbol{x}_{i}=m+\sigma \boldsymbol{y}_{i}, \quad \boldsymbol{y}_{i} \sim \mathcal{N}_{i}(\mathbf{0}, \mathrm{C}), \quad \text { for } i=1, \ldots, \lambda \quad \text { sampling } \\
& m \leftarrow \sum_{i=1}^{\mu} w_{i} \boldsymbol{x}_{i: \lambda}=m+\sigma \boldsymbol{y}_{w} \quad \text { where } \boldsymbol{y}_{w}=\sum_{i=1}^{\mu} w_{i} \boldsymbol{y}_{i: \lambda} \quad \text { update mean } \\
& p_{\mathrm{c}} \leftarrow\left(1-c_{\mathrm{c}}\right) p_{\mathrm{c}}+1_{\left\{\left\|p_{c}\right\|<1.5 \sqrt{n}\right\}} \sqrt{1-\left(1-c_{\mathrm{c}}\right)^{2}} \sqrt{\mu_{w}} \boldsymbol{y}_{w} \quad \text { cumulation for } \mathrm{C} \\
& p_{\sigma} \leftarrow\left(1-c_{\sigma}\right) p_{\sigma}+\sqrt{1-\left(1-c_{\sigma}\right)^{2}} \sqrt{\mu_{w}} \mathbf{C}^{-\frac{1}{2}} \boldsymbol{y}_{w} \quad \text { cumulation for } \sigma \\
& \mathrm{C} \leftarrow\left(1-c_{1}-c_{\mu}\right) \mathrm{C}+c_{1} p_{\mathrm{c}} p_{\mathrm{c}}{ }^{\mathrm{T}}+c_{\mu} \sum_{i=1}^{\mu} w_{i} \boldsymbol{y}_{i: \lambda} \boldsymbol{y}_{i: \lambda}^{\mathrm{T}} \\
& \sigma \leftarrow \sigma \times \exp \left(\frac{c_{\sigma}}{d_{\sigma}}\left(\frac{\left\|p_{\sigma}\right\|}{\mathrm{E}\|\mathbb{N}(0,1)\|}-1\right)\right) \quad \text { update of } \sigma
\end{aligned}
$$

Not covered on this slide: termination, restarts, useful output, boundaries and encoding

## CMA-ES in a Nutshell

## The CMA-ES

Input: $m \in \mathbb{R}^{n}, \sigma \in \mathbb{R}_{+}, \lambda$
Initialize: $\mathbf{C}=\mathbf{I}$, and $p_{\mathrm{c}}=\mathbf{0}, p_{\sigma}=\mathbf{0}$,
Set: $c_{\mathrm{c}} \approx 4 / n, c_{\sigma} \approx 4 / n, c_{1} \approx 2 / n^{2}, c_{\mu} \approx \mu_{w} / n^{2}, c_{1}+c_{\mu} \leq 1, d_{\sigma} \approx 1+\sqrt{\frac{\mu_{w}}{n}}$, and $w_{i=1 \ldots \lambda}$ such that $\mu_{w}=\frac{1}{\sum_{i=1}^{\mu} w_{i}} \approx 0.3 \lambda$
While not terminate

$$
\begin{array}{rlr}
\boldsymbol{x}_{i} & =m+\sigma \boldsymbol{y}_{i}, \quad \boldsymbol{y}_{i} \sim \mathcal{N}_{i}(\mathbf{0}, \mathrm{C}), \quad \text { for } i=1, \ldots, \lambda & \text { sampling } \\
m & \leftarrow \sum_{i=1}^{\mu} w_{i} \boldsymbol{x}_{i: \lambda}=m+\sigma \boldsymbol{y}_{w} \quad \text { where } \boldsymbol{y}_{w}=\sum_{i=1}^{\mu} w_{i} \boldsymbol{y}_{i: \lambda} & \text { update mean } \\
p_{\mathrm{c}} & \leftarrow\left(1-c_{\mathrm{c}}\right) p_{\mathrm{c}}+\mathbb{1}_{\left\{\left\|p_{\sigma}\right\|<1.5 \sqrt{n}\right\}} \sqrt{1-\left(1-c_{\mathrm{c}}\right)^{2}} \sqrt{\mu_{w}} \boldsymbol{y}_{w} & \text { cumulation for } \mathrm{C} \\
p_{\sigma} & \leftarrow\left(1-c_{\sigma}\right) p_{\sigma}+\sqrt{1-\left(1-c_{\sigma}\right)^{2}} \sqrt{\mu_{w}} \mathrm{C}^{-\frac{1}{2}} \boldsymbol{y}_{w} & \text { cumulation for } \sigma \\
\mathrm{C} & \leftarrow\left(1-c_{1}-c_{\mu}\right) \mathrm{C}+c_{1} p_{c} p_{\mathrm{c}}^{\mathrm{T}}+ & \sim^{\mu} \\
\sigma & \leftarrow \sigma \times \exp \left(\frac{c_{\sigma}}{d_{\sigma}}\left(\frac{\left\|\mid p_{\sigma}\right\|}{\mathrm{E} \| \mathcal{N}, \mathbf{1}) \|}-1\right)\right. & \text { Goal of next lecture: } \\
\text { Got covered on this slide: terminatic } & \text { Understand the main principles } \\
\text { of this state-of-the-art algorithm. }
\end{array}
$$

