Sequential Experimentation by Evolutionary Algorithms

Ofer M. Shir
Tel-Hai College and Migal Research Institute, Upper Galilee, Israel

Thomas Bäck
Natural Computing Group, Leiden University, The Netherlands

Joshua Knowles
School of Computer Science, University of Birmingham, United Kingdom

Richard Allmendinger
Alliance Manchester Business School, University of Manchester, United Kingdom

GECCO '17 Companion, July 15-19, 2017, Berlin, Germany

ACM ISBN 978-1-4503-4939-0/17/07.
http://dx.doi.org/10.1145/3067695.3067701

Instructors

- **Ofer Shir** is a Senior Lecturer at the Computer Science Department in Tel-Hai College, and a Principal Investigator at Migal-Galilee Research Institute, where he heads the Scientific Informatics and Experimental Optimization group – both located in the Upper Galilee, Israel.
- **Thomas Bäck** is Professor of Computer Science at the Leiden Institute of Advanced Computer Science (LIACS), Leiden University, The Netherlands, where he is head of the Natural Computing group since 2002.
- **Joshua Knowles** is Professor of Natural Computation at the School of Computer Science, University of Birmingham, UK, and honorary professor at Alliance Manchester Business School.
- **Richard Allmendinger** is a Lecturer at the Alliance Manchester Business School, University of Manchester, UK. He was a Postdoctoral Research Associate and Honorary Lecturer at the Biochemical Engineering Department, University College London, UK, following a PhD degree in Computer Science from the University of Manchester.




Agenda

- What do we mean by “Sequential Experimentation”?
- Examples of what has been done
- Potential Application Areas
- Case-Study: Quantum Control Experiments
- Case-Study: Biological Experiments
- Case-Study: Instrument Setup Experiments
- Discussion: Conclusions and Open Questions

What do we mean by …

SEQUENTIAL EXPERIMENTATION

"Typical" Characteristics

- Experiments are time-consuming.
- Experiments are expensive.
- Only few experiments are possible.
- There are exceptions as well!

Quantum Control: Case-Study

- Evolution "in the loop"
- Thousands of experiments possible ("*kHz regime*")

Further Challenges

- Noise and uncertainty of measurements
- Multiple objectives
- Dynamically changing requirements of experimentalists/stakeholders!
- Dynamically changing (resource) constraints
- Cost choices during optimization
  ➔ Some experiments may cost more than others
- Unusual constraints on population sizes and other hyperparameters

➔ Biological Experiments: Case-Study

1960s 2000s

Examples:
- Flow Plate
- Bended Pipe
- Nozzle
- Nutrient Solutions
- Coffee Formulations
- Quantum Control
- Biological Experiments

EXAMPLE APPLICATIONS

Early Experiments I: Flow Plate

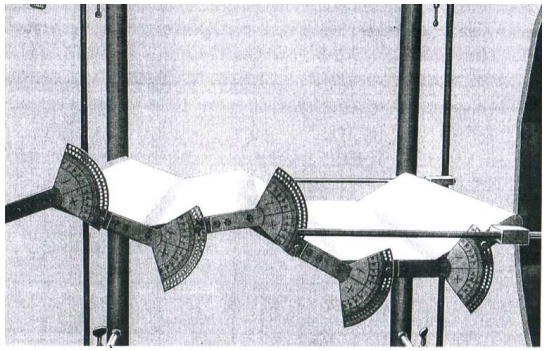

- A plate with 5 controllable angle brackets
- Measurable air flow drag (by a pitot tube)

Figure from: I. Rechenberg, Evolutionsstrategie '73, frommann-holzboog, Stuttgart 1973

Early Experiments I: Flow Plate

Experiment 1:
- Left / right supporting point at same y-coordinate.
- Horizontal flow.
- Minimize drag.

Experiment 2:
- Left supporting point 25% lower than right one.
- Horizontal flow.
- Minimize drag.

Plate drag

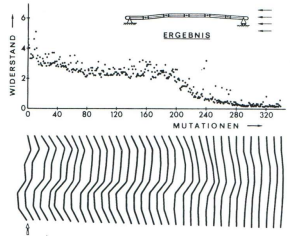

Number of mutations and selected plate shapes

Start	-30	-40	40	-30	40
End	0	4	0	6	-6

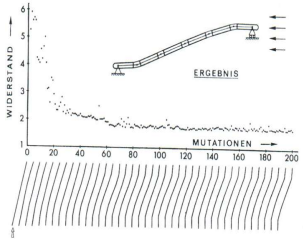

Number of mutations and selected plate shapes

Start	0	0	0	0	0
End	16	6	2	0	-18

Figures from: I. Rechenberg, Evolutionsstrategie ´73, frommann-holzboog, Stuttgart 1973

9

Early Experiments II: Bended Pipe

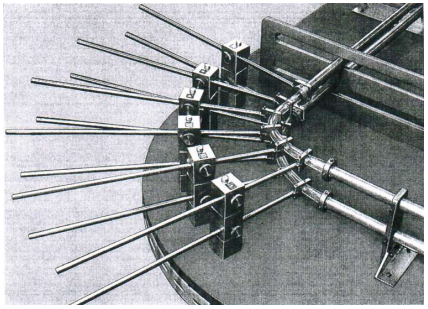

- A flexible pipe with 6 controllable bending devices
- Minimize bend losses of liquid flow
- Measure drag by pitot tube

Figure from: I. Rechenberg, Evolutionsstrategie ´73, frommann-holzboog, Stuttgart 1973

10

Early Experiments II: Bended Pipe

Pipe drag

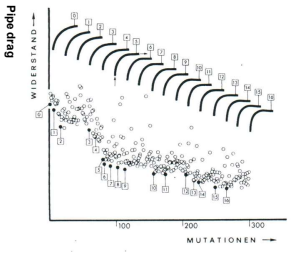

Number of mutations and selected pipe shapes

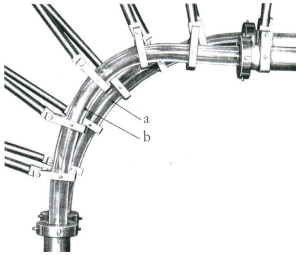

Initial (a) and optimized (b) pipe shape

- Bend loss of final form reduced by 10%
- Including drag a total reduction of 2%

Figure from: I. Rechenberg, Evolutionsstrategie ´73, frommann-holzboog, Stuttgart 1973

11

Early Experiment III: Nozzle

- What can be done if physics, (bio-) chemistry, … of process unkown?
- No model or simulation program available!
- Idea: Optimize with the real object
- “Hardware in the loop”
- Example: Supersonic nozzle, turbulent flow, physical model not available.

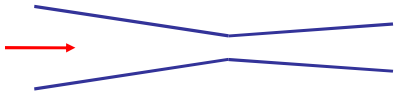

12

Experimental Setup: Nozzle

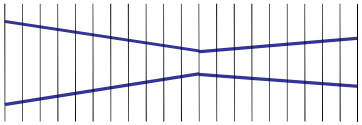

- Production of differently formed conic nozzle parts (pierced plates).
- Form of nozzle part is value of decision variable.

choosing conic nozzle parts (by EA)
clamping of conic nozzle parts (manually)
steam under high pressure passed into nozzle
degree of efficiency is **measured!**

„simulator replacement“

Nozzle Experiment (I)

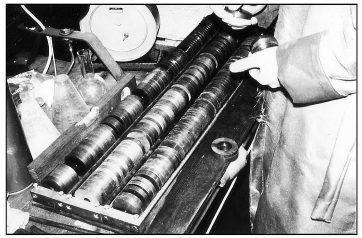

collection of conical nozzle parts

device for clamping nozzle parts

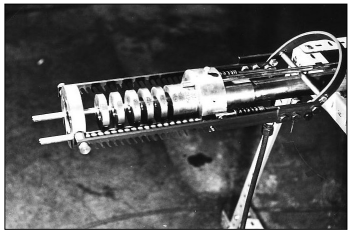

Figures courtesy of Hans-Paul Schwefel

Nozzle Experiment (II)

Hans-Paul Schwefel
while changing nozzle parts

Figures courtesy of Hans-Paul Schwefel

Nozzle Experiment (III)

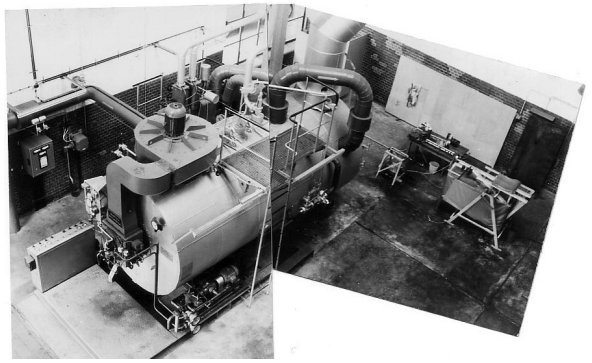

Figures courtesy of Hans-Paul Schwefel

steam plant / experimental setup

Nozzle Experiment (IV)

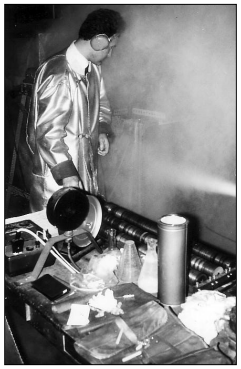

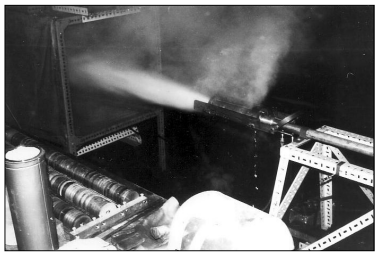

the nozzle in operation …

… while measuring degree of efficiency

Figures courtesy of Hans-Paul Schwefel

17

Nozzle Results (I)

- Illustrative Example: Optimize Efficiency
  - Initial:
  - Evolution:

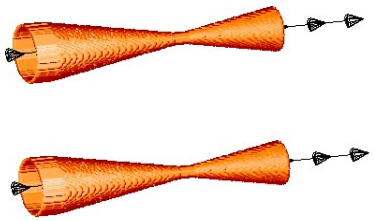

- 32% Improvement in Efficiency !

18

Nozzle Results (II)

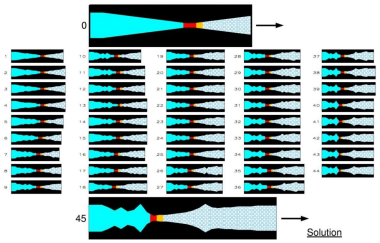

- 250 experiments were made.
- 45 improvements found.
- Discrete ring segments, variable-dimensional optimisation
- Gene duplication and deletion as additional operators.

J. Klockgether and H.-P. Schwefel, "Two-phase nozzle and hollow core jet experiments," in Proceedings of the 11th Symposium on Engineering Aspects of Magneto-Hydrodynamics, Caltech, Pasadena, California, USA, 1970.

19

Experiment: Coffee Formulations

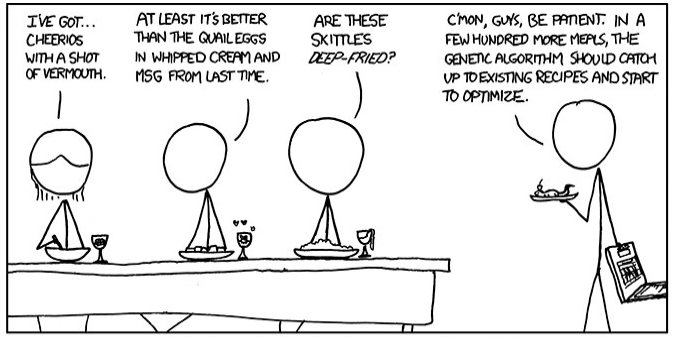

M. Herdy: Beiträge zur Theorie und Anwendung der Evolutionsstrategie, PhD Thesis, Technical University of Berlin, Germany, 2000.

20

Coffee Formulations: Results

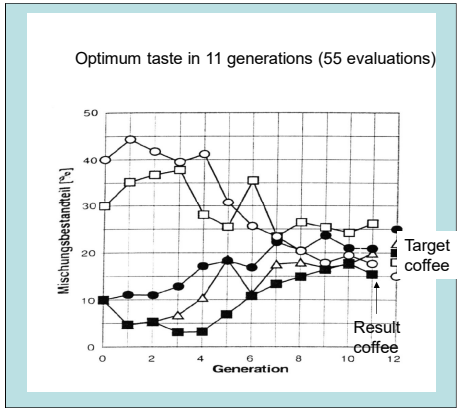

M. Herdy: Beiträge zur Theorie und Anwendung der Evolutionsstrategie, PhD Thesis, Technical University of Berlin, Germany, 2000.

Coffee Formulations: Results

- Coffee mixture differs a lot from target coffee !
- Taste is identical !
- Multiple realizations, but cost optimal !
- Approximation of cubic polynomial: 35 evals.

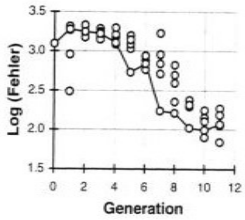

M. Herdy: Beiträge zur Theorie und Anwendung der Evolutionsstrategie, PhD Thesis, Technical University of Berlin, Germany, 2000.

EXPERIMENTAL OPTIMIZATION: FUNDAMENTALS

Experimental Requirements (for an Optimizer)

1. Speed: fast convergence is required
2. Reliability: reproducibility of results within a margin
   - Environmental parameters often hidden (temperature, pressure, …)
3. Robustness: manufacturing feasibility
4. Reference solution (recommended):

   pre-designed reference item, robust and stable, having a known objective function value

Convergence Speed

- Experiments are typically expensive:

Goal: Drive the system towards finding large improvements with as few experiments as possible.

- Practical solutions: “greedy” variants of evolutionary algorithms, e.g.,
  - Derandomized evolution strategies
  - ParEGO
  - Often “stochastic gradient search”
  - Need to support parallel execution!

See e.g. Bäck, Foussette, Krause: *Contemporary Evolution Strategies*, Springer 2013, for a comparison of evolution strategies when very few function evaluations are possible.

25

Reliability of Results

- Mostly *algorithm-dependent*
- Attained results must be reproducible
- Scenarios of recording *experimental outliers* must be avoided (elitism is tricky…)
- Perceived result versus *a posteriori* result
- Possible solutions:
  - Employing comma (non-elitist) strategies
  - In ES, the recombination operator assists in treating noise (The Genetic Repair (GR) Hypothesis, Beyer)
  - Increasing sampling rate of measurements (“signal averaging”)

26

Environmental Parameters

- As many as possible physical conditions should be recorded during the experiment
- Ideally, sensitivity of the system to the environment should be assessed
- Basic starting points: recording Signal/Noise, extracting power spectrum of the noise, etc.

27

Manufacturing Feasibility

- Mostly *system-dependent*
- Realization of the prescribed decision parameters of the experiment to equivalent systems, e.g., in a manufacturing stage
- To this end, sensitivity of the system must be assessed (electronics, for instance)
- Upon obtaining reproducible results, they should be verified on equivalent systems

28

Noise Colors

- White Noise ($1/f^0$ -noise)
- Pink Noise (flicker noise, or $1/f$ -noise)
- Red (Brownian) Noise ($1/f^2$ -noise)

Tip: Assess the stability of your system by extracting the Power Spectral Density of its signal-free state.

M. Roth, J. Roslund, and H. Rabitz, "Assessing and managing laser system stability for quantum control experiments", *Rev. Sci. Instrum.* **77**, *083107 (2006)*

29

POTENTIAL APPLICATION AREAS

30

A Classification

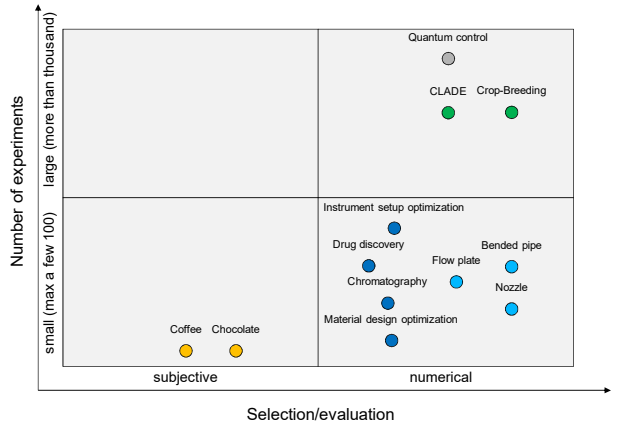

31

Potential Application Areas

- Cosmetics / Detergent Formulation Optimization
- Catalyst Formulation Optimization (Cost, Effectiveness, …)
- Subjective Evaluation Applications based on Human Taste or other Senses
- Engineering Applications Requiring Real-World Experiments for Measurement
- Concrete Formulation Optimization
- Glue Formulation Optimization
- Plant Startup Process
- Chemical Compound Synthesis Processes (e.g., Drugs)
- Instrument Setup Optimization

32

MANCHESTER 1824

The University of Manchester

Tel-Hai College

Case-Study:

QUANTUM CONTROL EXPERIMENTS

33

Altering the Course of Quantum Phenomena

Rabitz et al.
"Electric Field Design"
Quantum Control Theory

$$i\hbar\frac{\partial}{\partial t}|\psi(t)\rangle = \mathcal{H}(t)|\psi(t)\rangle$$
$$\mathcal{H}(t) = \mathcal{H}_0 - \vec{\mu}\cdot\vec{\epsilon}(t)$$

Find optimal $\vec{\epsilon}(t)$ s.t.
$|\langle\psi_{\text{target}}|\,\psi(T)\rangle|^2 \longrightarrow \max$

Hamiltonian required

PRA **37**, 4950 (1988)

Judson and Rabitz
"Teaching Lasers to Control Molecules"
Quantum Control Experiments

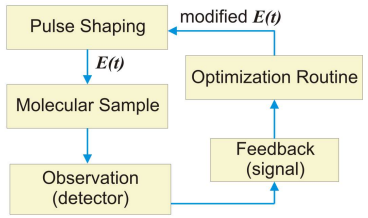

Hamiltonian **not** required

PRL **68**, 1500 (1992)

34

Quantum Control Experiments

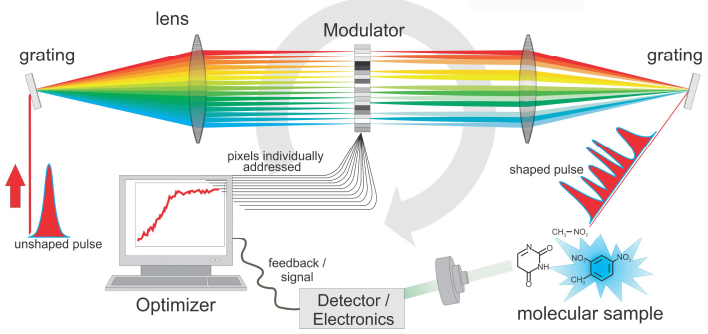

35

The QCE Arena: The Optical Table

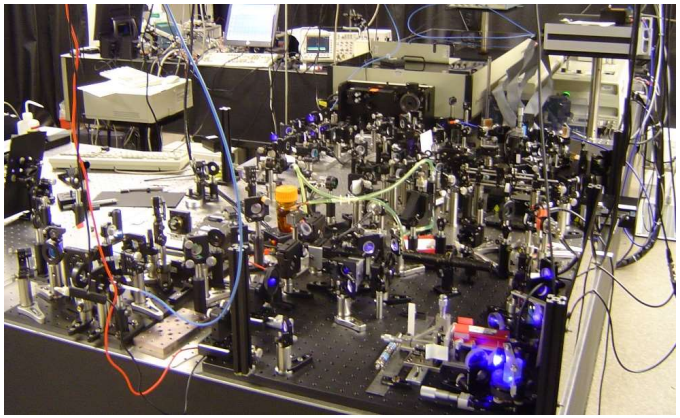

Figure courtesy of Jonathan Roslund

36

The Optical Table: Shaping the Pulse

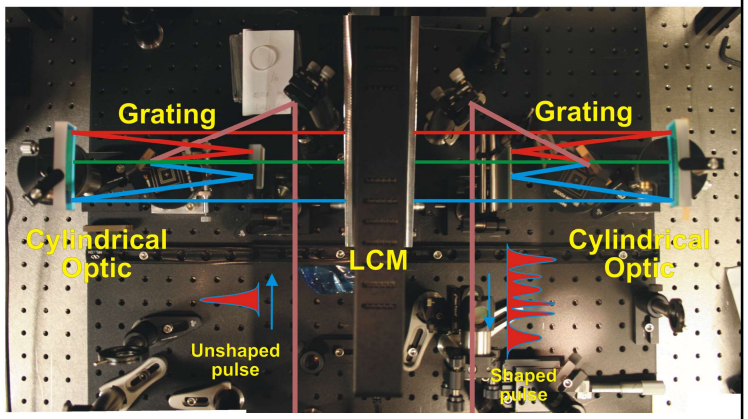

Figure courtesy of Jonathan Roslund

37

QCE: Sources of Noise/Uncertainty

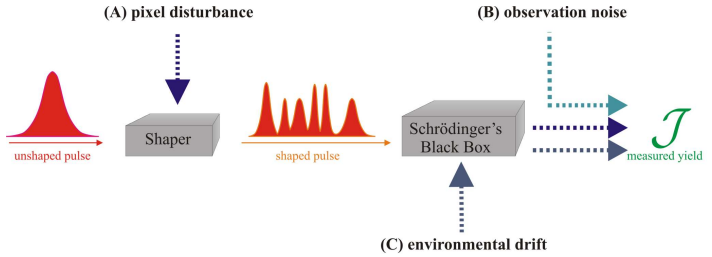

(A) $\tilde{\phi}(\omega) = \left(\phi(\omega_1) + \mathcal{N}_1\left(0, \epsilon_S^2\right), \ldots, \phi(\omega_n) + \mathcal{N}_n\left(0, \epsilon_S^2\right)\right)$

(B) $\tilde{\mathcal{J}} = \mathcal{J} + \mathcal{N}\left(0, \epsilon_{\mathcal{J}}^2\right)$ Signal Averaging: $\left\langle \tilde{\mathcal{J}} \right\rangle = \mathcal{J}$, $\mathrm{VAR}\left[\tilde{\mathcal{J}}\right] = \frac{\epsilon_{\mathcal{J}}^2}{k}$

(C) $\hat{\mathcal{J}}(t) = \tilde{\mathcal{J}} + \xi(t)$

38

Single-Objective QCE

- CMA-ES was observed to perform extremely well with small population sizes
- Recombination is indeed necessary (GR, Beyer)
- Robust, reproducible, reliable solutions

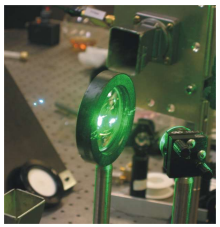

Figure courtesy of Jonathan Roslund

Roslund, J., Shir, O.M., Bäck, T., Rabitz, H.: Accelerated Optimization and Automated Discovery with Covariance Matrix Adaptation for Experimental Quantum Control. Physical Review A (Atomic, Molecular, and Optical Physics) **80**(4) (2009) 043415

39

Multi-Observable QCE

(a) Experimental Pareto frontier for the Total Ionization problem approximated by MO-CMA-ES, displaying the perceived frontier of a single experiment, the **reference frontier** of the intensity based non-shaped pulse, as well as a sampling of the Pareto optimal set.

(b) Experimental Pareto frontier for the Molecular Plasma Generation problem approximated by MO-CMA-ES remedied with occasional re-evaluation, displaying the perceived frontier, the reference frontier, and the reproduction of the Pareto optimal set.

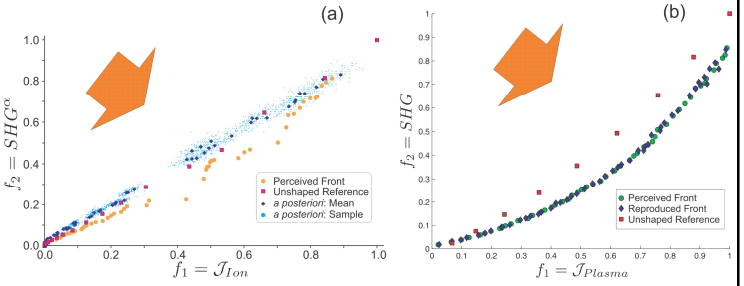

Shir, O.M., Roslund, J., Leghtas, Z., Rabitz, H.: Quantum Control Experiments as a Testbed for Evolutionary Multi-Objective Algorithms. Genetic Programming and Evolvable Machines **13**(4) (2012) 445—491

40

Radar Optimal Dynamic Discrimination

- **Competition: maximizing free electron number vs. minimizing SHG**
- **Pay-off over unshaped (TL) reference (HV ratio): 24.5%**

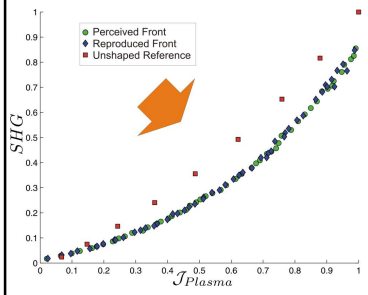

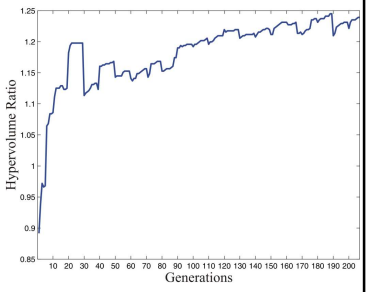

Shir, O.M., Roslund, J., Leghtas, Z., Rabitz, H.: Quantum Control Experiments as a Testbed for Evolutionary Multi-Objective Algorithms. Genetic Programming and Evolvable Machines **13**(4) (2012) 445—491

Extended Features: Statistical Learning (**FOCAL**)

- QCE and Derandomized ES enjoy a happy marriage
- However, the default CMA-ES does not learn a covariance matrix reflective of the inverse Hessian
- **FOCAL**, for experimental Hessian retrieval

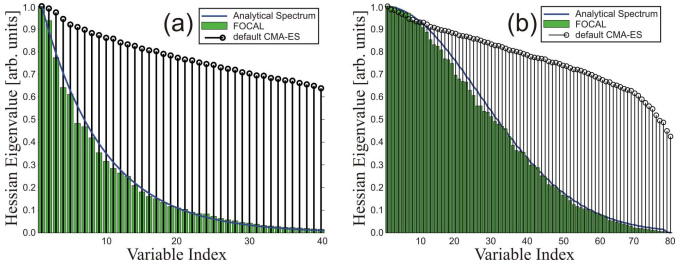

Shir, O.M., Roslund, J., Whitley, D., Rabitz, H.: Efficient retrieval of landscape Hessian: Forced optimal covariance adaptive learning. Physical Review E **89**(6) (2014) 063306

FOCAL: Experimental Results

(a) Retrieving the Hessian by FOCAL for rank-deficient atomic Rubidium
(b) 5 most important Hessian eigenvectors; Physical form is corroborated

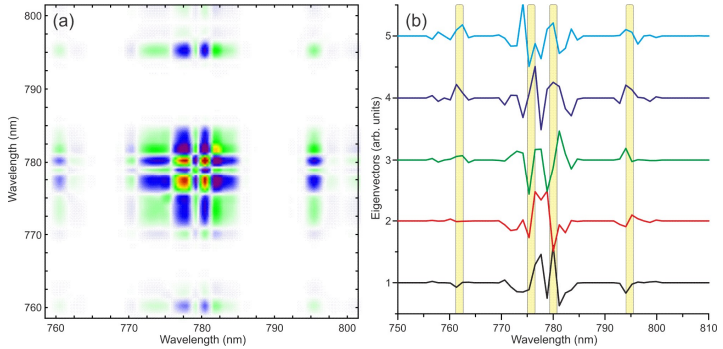

Shir, O.M., Roslund, J., Whitley, D., Rabitz, H.: Efficient retrieval of landscape Hessian: Forced optimal covariance adaptive learning. Physical Review E **89**(6) (2014) 063306

Case-Study:

BIOLOGICAL EXPERIMENTS

Experimental Optimization in Biology, Medicine and Food

- Evolution of real DNA on microarray chips
  - Dealing with a very large population size
  - Landscape analysis / optimizing evolution
- Optimizing the design of an analytical instrument
  - Handling *ephemeral resource constraints*
- Crop Breeding
  - How to use genotypic information effectively

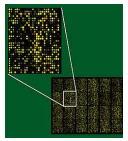

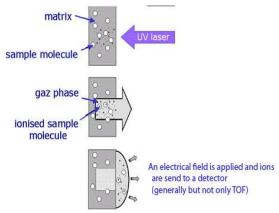

45

Challenges in Experimental Optimization

- Noise and uncertainty of measurements
- Multiple objectives
- Dynamically changing requirements of biologists !
- Dynamically changing (resource) constraints
- Cost choices during optimization
  - Some experiments may cost more than others
- Unusual constraints on population sizes and other hyperparameters

46

Evolution of drugs / enzymes / biomarkers

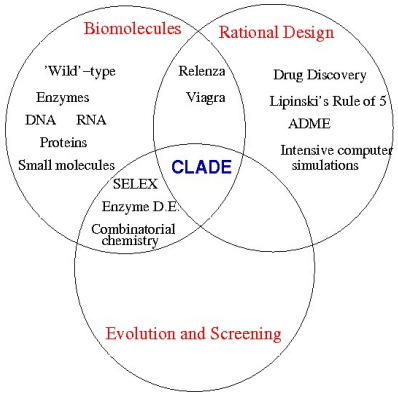

- Closed-loop aptameric directed evolution (CLADE)
- Molecules tested on a chip
- Evolution occurs on the computer
- Exquisite control of the evolution / easily incorporate rational design principles

47

CLADE Details

Microarray detail

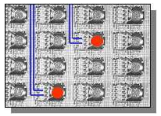

Binding of target proteins

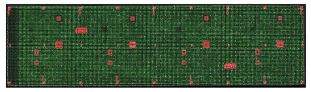

- Evolve short strands of DNA (up to 40 bases)
- Population size: ~90,000
- ~24 hr and £hundreds per cycle
- Evolve DNA to bind target proteins *specifically* and *tightly*
- 5-10 cycles used

48

CLADE – Looking at the aptamer fitness landscape

- *Measured fitness of 1m 10-base DNA strands in duplicate. Statistics below.*
- Number of local optima: 1809 (single-point mutations)
- Epistasis variance =0.532

$$\eta = \frac{\sum_{i \in I} (f_i - l_i)^2}{\sum_{i \in I} (f_i - \bar{f})^2},$$

- Autocorrelation correl'n length=6

$$r(s) = \frac{E[f(x^{t+s})f(x^t)] - E[f(x^t)]^2}{E[f(x^t)^2] - E[f(x^t)]^2}$$

- The information content sequence logo for the top 1000 aptamers

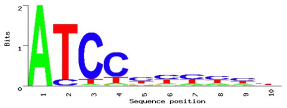

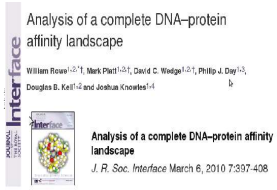

CLADE – How to exploit/handle large population sizes

Tested EAs on NK landscapes

- Population size=44000
- High selection pressure and high mutation rates generally favourable
- Resul
on rea

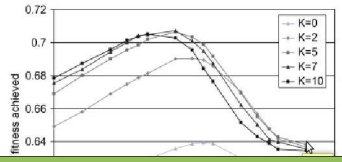

We would like to model and test more EAs on the aptamer landscape. We can do that with LSMs...

Fig. 4. Best fitnesses achieved by a (μ,λ) GA without crossover, with μ=4000 and λ=40,000.

Landscape State Machine Modelling Approach

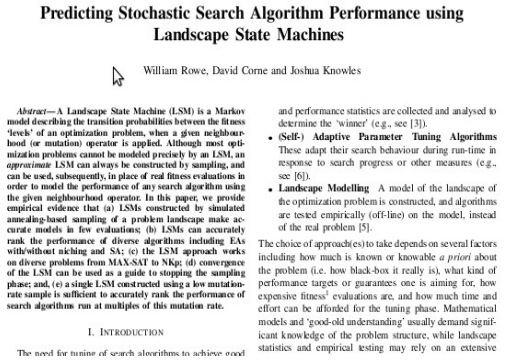

Predicting Stochastic Search Algorithm Performance using Landscape State Machines

William Rowe, David Corne and Joshua Knowles

Abstract— **A Landscape State Machine (LSM) is a Markov model describing the transition probabilities between the fitness 'levels' of an optimization problem, when a given neighbourhood (or mutation) operator is applied. Although most optimization problems cannot be modeled precisely by an LSM, an *approximate* LSM can always be constructed by sampling, and can be used, subsequently, in place of real fitness evaluations in order to model the performance of any search algorithm using the given neighbourhood operator. In this paper, we provide empirical evidence that (a) LSMs constructed by simulated annealing-based sampling of a problem landscape make accurate models in few evaluations; (b) LSMs can accurately rank the performance of diverse algorithms including EAs with/without niching and SA; (c) the LSM approach works on diverse problems from MAX-SAT to NKp; (d) convergence of the LSM can be used as a guide to stopping the sampling phase; and, (e) a single LSM constructed using a low mutation-rate sample is sufficient to accurately rank the performance of search algorithms run at multiples of this mutation rate.**

I. Introduction

The need for tuning of search algorithms to achieve good

and performance statistics are collected and analysed to determine the 'winner' (e.g., see [3]).

- **(Self-) Adaptive Parameter Tuning Algorithms** These adapt their search behaviour during run-time in response to search progress or other measures (e.g., see [6]).
- **Landscape Modelling** A model of the landscape of the optimization problem is constructed, and algorithms are tested empirically (off-line) on the model, instead of the real problem [5].

The choice of approach(es) to take depends on several factors including how much is known or knowable *a priori* about the problem (i.e. how black-box it really is), what kind of performance targets or guarantees one is aiming for, how expensive fitness[1] evaluations are, and how much time and effort can be afforded for the tuning phase. Mathematical models and 'good-old understanding' usually demand significant knowledge of the problem structure, while landscape statistics and empirical testing may rely on an extensive

LSM results on CLADE

- LSM predicts optimal mutation rate with high accuracy for all EA settings
- It is better than (or at least as good as) the 1/L heuristic on every condition tested

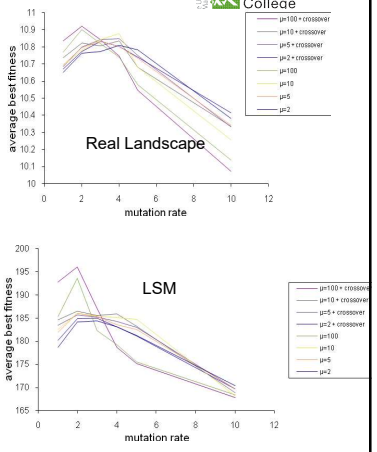

References on CLADE

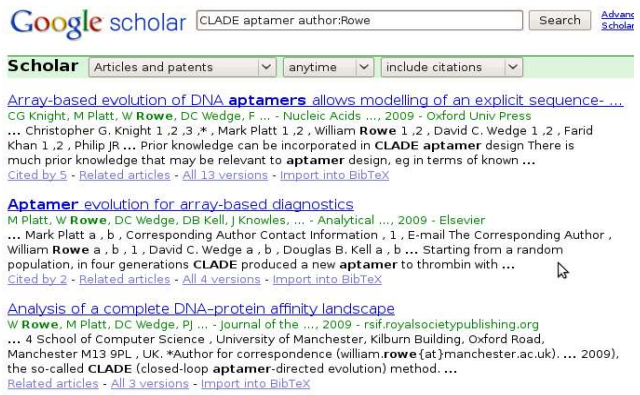

Motivation: automation of science experiments

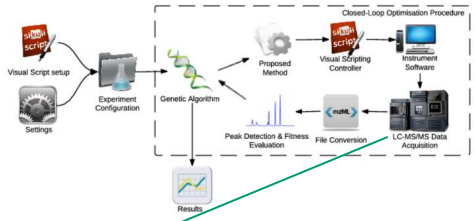

Figure S1: MUSCLE system diagram

Mass spectrometers optimized by ParEGO were used in the HUSERMET project, a large study of human blood serum in health and disease with over 800 patient subjects and performed in collaboration with GlaxoSmithKline, AstraZeneca, Stockport NHS Trust and others (Bradbury et al, 2014)

Bradbury, James, Grégory Genta-Jouve, J. William Allwood, Warwick B. Dunn, Royston Goodacre, Joshua D. Knowles, Shan He, and Mark R. Viant.. **MUSCLE: automated multi-objective evolutionary optimisation of targeted LC-MS/MS analysis,** Bioinformatics (2014): btu740.

LDI Spectroscopy: Optimizing "Flyability"

- In LDI, a laser is fired at a sample, e.g. blood serum, to analyse it
- The sample may be placed on a silicon matrix
- Properties of the matrix determine how much/which compounds in the sample will "fly" into the mass-spectrometer when the laser hits them

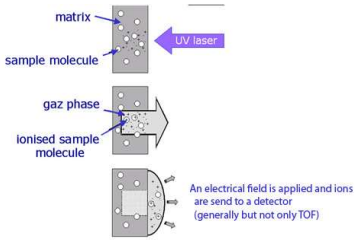

Optimize the Silicon Wafer Matrix

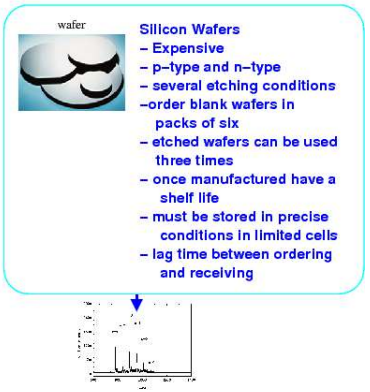

Optimize the Silicon Wafers

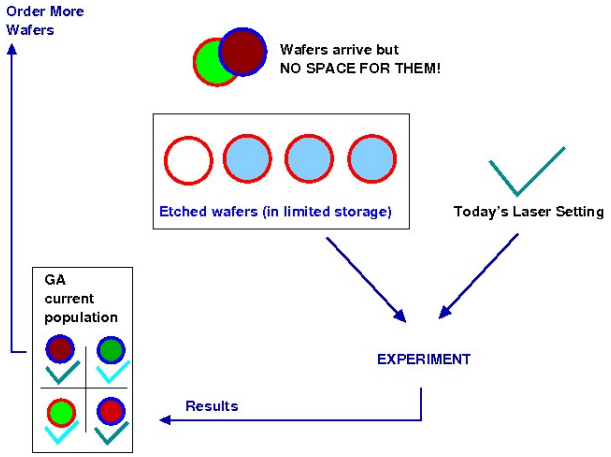

Ephemeral Resource Constraints (ERCs)

- Variables are not free during optimization – limited by current resources (E_t is defined by ERCs)

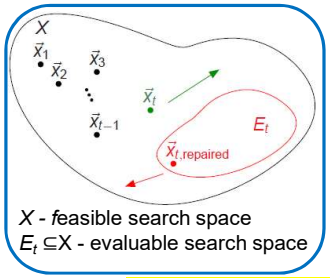

X - feasible search space
$E_t \subseteq X$ - evaluable search space

$$\text{maximize } y = f(\vec{x})$$
$$\text{subject to } \vec{x} \in X,$$
$$y_t = \begin{cases} f(\vec{x}_t) & \text{if } \vec{x}_t \in E_t(\sigma_t) \subseteq X \\ \text{null} & \text{otherwise,} \end{cases}$$

E_t can be defined in terms of (constraint) schemata, e.g. H = (*1**0)

R. Allmendinger and J. Knowles (2012): **On Handling Ephemeral Resource Constraints in Evolutionary Search**. *Evolutionary Computation*, 21(3): 497-531.
R. Allmendinger (2012): **Tuning Evolutionary Search for Closed-Loop Optimization**. PhD thesis. Department of Computer Science, University of Manchester, UK.

Ephemeral Resource Constraints (ERCs)

- ERCs raise many questions:
  - Who/what is in control ?
  - How to mesh scheduling with optimization
  - How to prevent diversity loss/ drift effects
  - How to reduce wastage of materials
- Many examples of ERCs available due to the complex nature of resources

R. Allmendinger and J. Knowles (2012): **On Handling Ephemeral Resource Constraints in Evolutionary Search**. *Evolutionary Computation*, 21(3): 497-531.
R. Allmendinger (2012): **Tuning Evolutionary Search for Closed-Loop Optimization**. PhD thesis. Department of Computer Science, University of Manchester, UK.

Ephemeral Resource Constraints (ERCs)

Periodic ERC: Models availability of a specific resource (defined by the constraint schema) at regular time intervals, e.g. *"In an optimization problem requiring skilled engineers to operate instruments, on Mondays, only engineer eng$_i$ is available."*

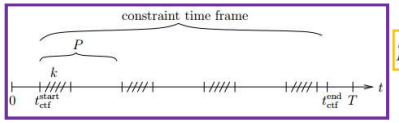

$perERC(t_{\text{ctf}}^{\text{start}}, t_{\text{ctf}}^{\text{end}}, k, P, H)$

Commitment relaxation ERC: Commits an optimizer to a specific variable value combination for some period of time whenever it uses this particular combination, e.g. *"In an optimization problem involving the selection of instrument settings, the configuration, c, once set, cannot be changed during the remainder of the working day."*

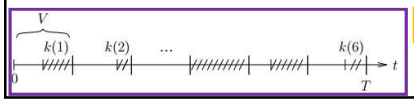

$commRelaxERC(t_{\text{ctf}}^{\text{start}}, t_{\text{ctf}}^{\text{end}}, V, H)$

Ephemeral Resource Constraints (ERCs)

Commitment composite ERC:

- Decision variables define a composite (a constraint schema) that requires resources to be available to evaluate solution.
- Limited storage cells
- Time-lag between ordering and receiving resources
- Resources have shelf life and limited reuses
- *Silicon wafer example*

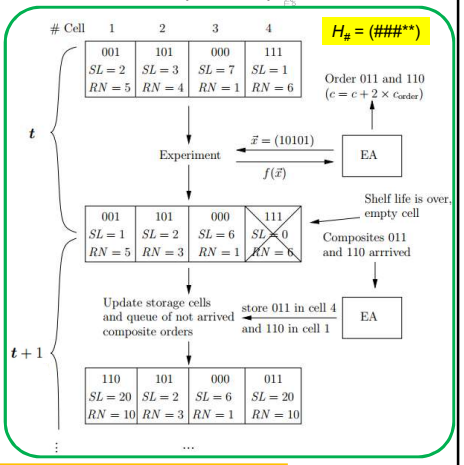

$commCompERC(H_{\#}, \#SC, TL, RN, SL)$

Ephemeral Resource Constraints (ERCs)

- Various approaches can be used for dealing with ERCs, e.g.:
  - Static constraint-handling strategies (Allmendinger and Knowles, 2012)
  - Learn policies (via reinforcement learning) to switch between constraint-handling strategies during optimization (Allmendinger and Knowles, 2011)
  - Optimization combined with online resource purchasing strategies (Allmendinger and Knowles, 2010)
- Theoretical studies (e.g. using Markov chains) on ERCs can guide choice of EA configurations (Allmendinger and Knowles, 2015)
- More work needed on combining ERCs with other experimental challenges e.g. noise, multiple objectives, non-homogeneous experimental costs.

R. Allmendinger and J. Knowles (2010): **On-Line Purchasing Strategies for an Evolutionary Algorithm Performing Resource-Constrained Optimization.** Parallel Problem Solving in Nature - PPSN XI , pp 161-170.
R. Allmendinger and J. Knowles (2011): **Policy Learning in Resource-Constrained Optimization**. Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation (GECCO '11), pp 1971-1978.
R. Allmendinger and J. Knowles (2012): On Handling **Ephemeral Resource Constraints in Evolutionary Search**. *Evolutionary Computation*, 21(3): 497-531.
R. Allmendinger and J. Knowles (2015): **Ephemeral Resource Constraints in Optimization**. In R. Datta and K. Deb (Eds.) Evolutionary Constrained Optimization, Springer, pp 95-134.

Further resourcing issues

Change of variables during the optimization

- E.g. Identify most effective drug cocktails from a non-stationary drug library

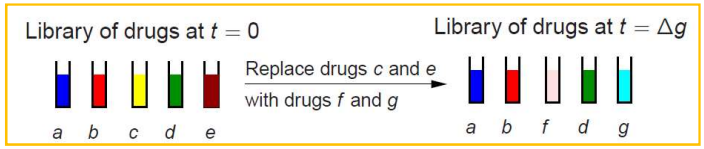

- *Fair mutation:* Perform usual optimization alongside rapid exploration of the space of solutions using any of the new variables

R. Allmendinger and J. Knowles (2010): **Evolutionary Optimization on Problems Subject to Changes of Variables**. *Parallel Problem Solving in Nature - PPSN XI* , pp 151-160.
B.G. Small, B.W. McColl, R. Allmendinger, J. Pahle, G. Lopez-Castejon, N.J. Rothwell, J. Knowles, P. Mendes, D. Brough, and D.B. Kell (2011): **Efficient discovery of anti-inflammatory small molecule combinations using evolutionary computing**. *Nature Chemical Biology*, 7: 902-908.

Further resourcing issues

Objectives exhibit non-uniform latencies

- Multiobjective optimization problems where at least one of the objective functions requires a relatively longer time to be evaluated than the cheapest/quickest of the objective functions → at any given time, fitness estimates of some solutions may only be partial
- E.g. identify most potent drug cocktails that is also most economical to manufacture
  - *Potency of drug cocktails:* laborious process involving creating the cocktails and then testing them
  - *Cost of manufacture:* relatively quickly to compute by looking up the drug amounts and costs and summing them up (or running a simulation tool).

R. Allmendinger and J. Knowles (2013): **'Hang On a Minute': Investigations on the Effects of Delayed Objective Functions in Multiobjective Optimization**. Evolutionary Multi-Criterion Optimization (EMO 2013), LNCS 7811, pp 6-20.
R. Allmendinger, J. Handl, and J. Knowles (2015): **Multiobjective Optimization: When Objectives Exhibit Unequal Latencies.** European Journal of Operations Research, 243(2): 497-513.

2 objectives, one fast and one slow one

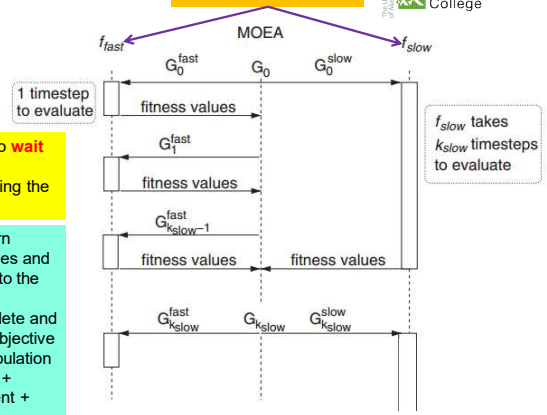

Standard approach is to **wait** until all objectives are evaluated before updating the population

Don't wait but e.g. return available objective values and assign **pseudovalues** to the other objectives
→ Solutions with complete and partial information on objective values co-exist in a population
Challenges: Selection + Pseudovalue assignment + Ranking

R. Allmendinger and J. Knowles (2013): **'Hang On a Minute': Investigations on the Effects of Delayed Objective Functions in Multiobjective Optimization**. Evolutionary Multi-Criterion Optimization (EMO 2013), LNCS 7811, pp 6-20.
R. Allmendinger, J. Handl, and J. Knowles (2015): **Multiobjective Optimization: When Objectives Exhibit Unequal Latencies.** European Journal of Operations Research, 243(2): 497-513.

Optimization performance affected by
- Correlation between objectives
- Length of delay period
- Search space characteristics

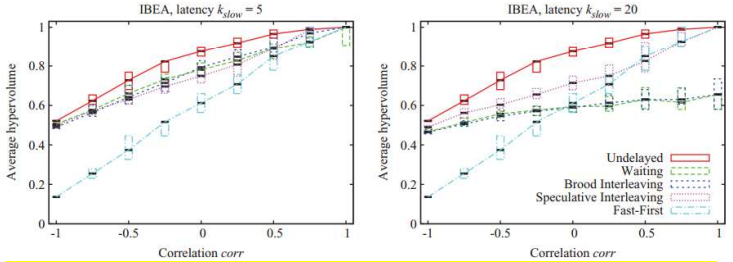

Average hypervolume as a function of the correlation level between two objectives (fixed delay period). Waiting performs well for short delay periods, and non-waiting strategies for positively correlated objectives and longer waiting periods.

R. Allmendinger, J. Handl, and J. Knowles (2015): **Multiobjective Optimization: When Objectives Exhibit Unequal Latencies.** European Journal of Operations Research, 243(2): 497-513.

Optimization performance affected by
- Correlation between objectives
- Length of delay period
- Search space characteristics

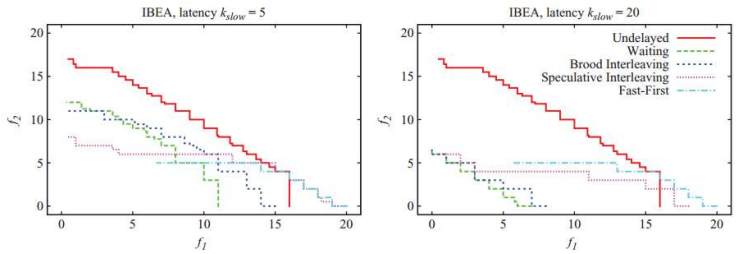

Median attainment surface for a fixed delay period. Objective f_1 is the quicker to evaluate objective. Some strategies suffer from bias towards f_1.

R. Allmendinger, J. Handl, and J. Knowles (2015): **Multiobjective Optimization: When Objectives Exhibit Unequal Latencies.** European Journal of Operations Research, 243(2): 497-513.

Average hypervolume as a function of the correlation level between two objectives and the delay period. Darker shading indicates better performance.

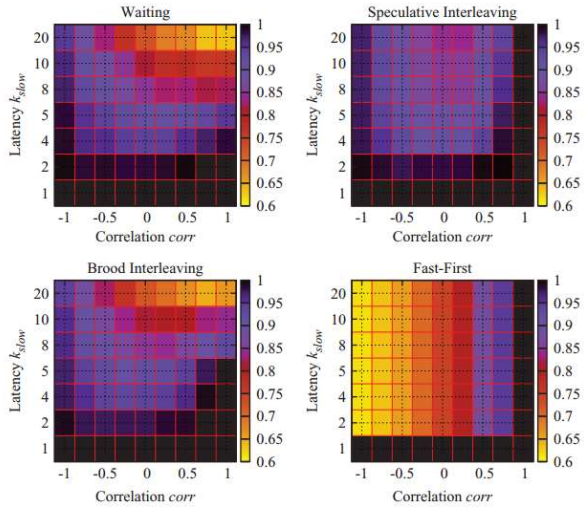

R. Allmendinger, J. Handl, and J. Knowles (2015): **Multiobjective Optimization: When Objectives Exhibit Unequal Latencies.** European Journal of Operations Research, 243(2): 497-513.

Crop-Breeding Strategies

- Food and energy crops in short supply (simplifying very complex global socio-political situation)
- Some traits in crops are *quantitative,* e.g. energy yield, low-temperature resistance
- Given new sequencing technologies, we can see the quantitative trait loci (QTLs) – the genotype. **Will this really help us breed faster for the traits we want?**

69

EAs that know where they are?

Fitness-only "F" Algorithms
- Standard EA
- FUSS
- Crowding

Genotypic "G" Algorithms
- Genotypic Niching
- Surrogate modeling / LEM
- EDAs

Which of these are really better, F or G?

Can we tell crop breeders how to use sequence information to do efficient search? Can we sell them the algorithms?

70

Simulation Details

Real Crop Breeding
- Crops take months to ~ 1 year to grow
- But population size can be 1,000 - 20,000
- Genome size ~10,000
- Non-viables do not grow – provide no feedback
- Starting population is evolved but diverse

Simulation
- *NKt* landscape. The *t* param controls number of traits to enhance
- Fitness below 0.65 → 0
- Start from diverse evolved population on *NK*
- Allow only 10 generations!
- Population size as for real world

71

Crop-Breeding: Some Results

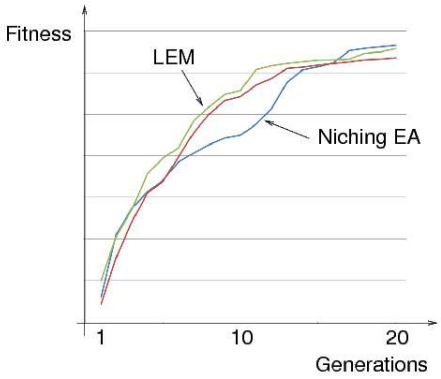

LEM = Michalski R.S., "LEARNABLE EVOLUTION MODEL: Evolutionary Processes Guided by Machine Learning," *Machine Learning* , 38, pp 9-40, 2000.

72

Other applications in biology / food / medicine

- Chocolate flavour/aroma
  - Subjective measurement and time lags
- Configuration of GC-MS instrument for looking at Human serum
  - Multiobjective
  - Used ParEGO - surrogate modeling approach
- Evolving nano-technologies or autonomous robots
  - Limited resources

73

The University of Manchester

Tel-Hai College

Case-Study:

INSTRUMENT SETUP OPTIMIZATION

74

Instruments

Precision tools used to create value from raw material or automate a certain operation

75

Instrument setup optimization

General goal: Configure operating conditions of instrument such that the instruments performs as "efficiently" as possible.

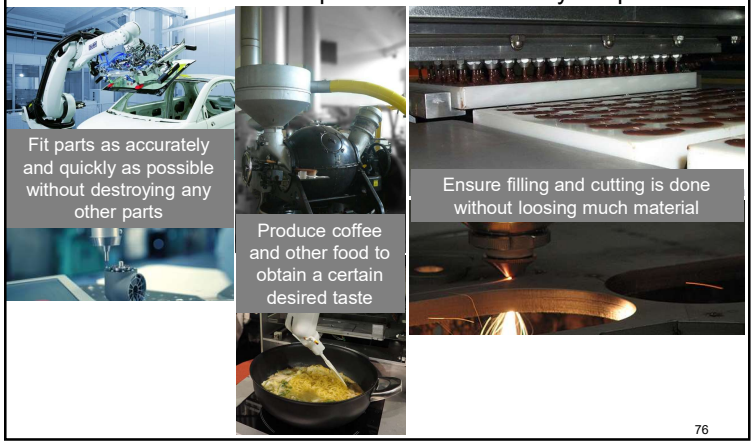

76

Chromatography

- Well-established approach for purifying proteins
- Relies on using expensive raw materials (resin)
- Time-consuming
- Requires operating conditions to be optimized
- Various chromatography techniques are available
- Combination of these need to be used in purification

77

Overview of a typical biologics manufacturing process

Upstream
Growing of cells

Downstream
Processing of cell mass from the upstream to meet purity and quality requirements

WPI @BETC_WPI [BETC Webinar Series]

78

Overview of a typical biologics manufacturing process

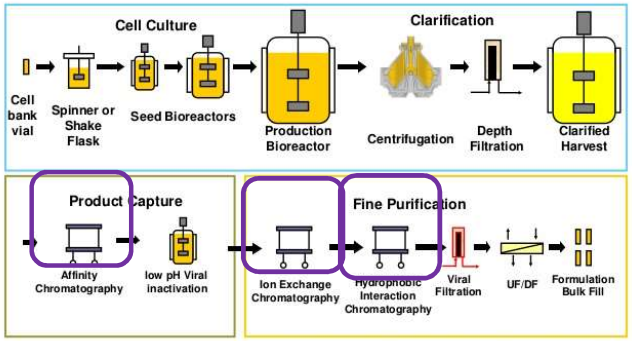

WPI @BETC_WPI [BETC Webinar Series]

79

Optimization of Chromatography Operating Conditions (Typically)

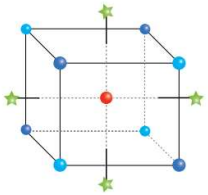

Central Composite Circumscribed (CCC) design

Operating conditions *x*

Liquid handling robot

Measure protein purity and yield via HPLC

Typically:
- DoE-driven
- Chromatography steps are considered independently (one step at the time)

80

Optimization of Chromatography Operating Conditions (Proposed)

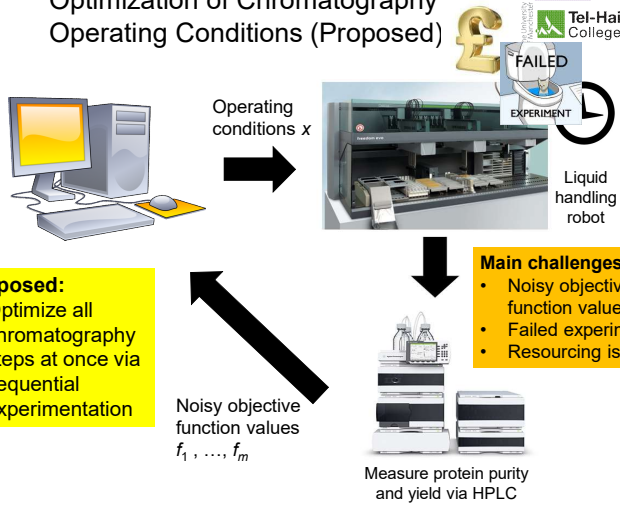

Proposed:
- Optimize all chromatography steps at once via sequential experimentation

Main challenges:
- Noisy objective function values
- Failed experiments
- Resourcing issues

R. Allmendinger, S. Gerontas, N.J. Titchener-Hooker, and S.S. Farid (2014): **Tuning Evolutionary Multiobjective Optimization for Closed-Loop Estimation of Chromatographic Operating Conditions**. *Parallel Problem Solving in Nature - PPSN XIII*, pp 741-750.

Challenges

- **Linkage constraints:** Operating condition setting at step i can dictate lowest possible setting of another operating condition at step $i+1$

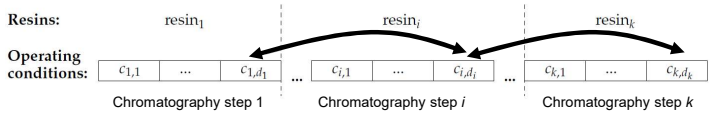

- **Multiple objectives:**
  - Yield: max $f_{yield} = Y_1 \text{x} \ldots \text{x}\ Y_k$, measured at each step $i = 1,\ldots,k$
  - Impurities: min $f_{impurity}$ = Total amount of impurities remaining

- **Resourcing issue:**
  Stop experiment if total yield $f_{yield} = Y_1 \text{x} \ldots \text{x}\ Y_p < 90\%$

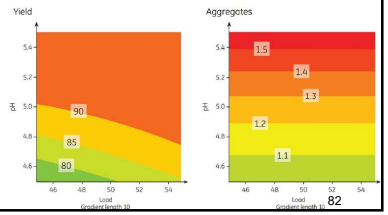

Challenges

- **Resourcing issue:**
  Stop experiment if total yield $f_{yield} = Y_1 \text{x} \ldots \text{x}\ Y_p < 90\%$

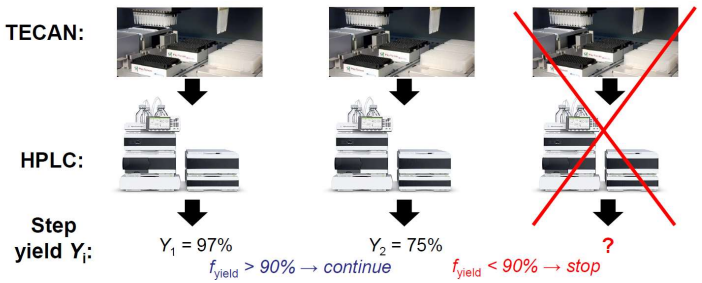

R. Allmendinger, S. Gerontas, N.J. Titchener-Hooker, and S.S. Farid (2014): **Tuning Evolutionary Multiobjective Optimization for Closed-Loop Estimation of Chromatographic Operating Conditions**. *Parallel Problem Solving in Nature - PPSN XIII*, pp 741-750.

Some results for dealing with the resourcing issue

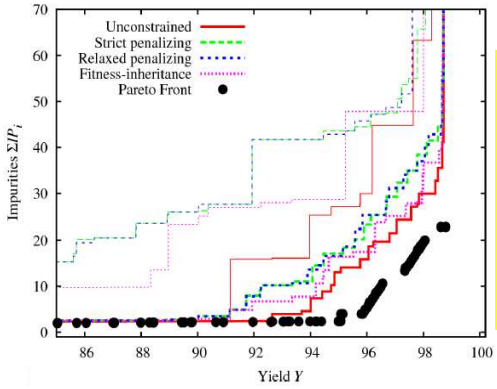

Fitness-inheritance strategy:
Fill missing yield measurements Y_p+1 ,…, Y_k with measurements obtained for the most similar (in terms of operating conditions) and fully evaluated multi-step chromatography process, and also copy $f_{impurity}$ of that process.

R. Allmendinger, S. Gerontas, N.J. Titchener-Hooker, and S.S. Farid (2014): **Tuning Evolutionary Multiobjective Optimization for Closed-Loop Estimation of Chromatographic Operating Conditions**. *Parallel Problem Solving in Nature - PPSN XIII*, pp 741-750.

DISCUSSION

Evolutionary Algorithms Used

•Nozzle Experiments: Two-Membered Evolution Strategy [Rechenberg; 1973]

•Quantum Control Experiments: Derandomized Evolution Strategies [Hansen et al.; 1994-2008]

•Biological Experiments: PESA-II [Corne et al, 2002], ParEGO [Knowles et al. 2006], (μ+λ)-ES, and Learnable evolution model [Michalski,2000]

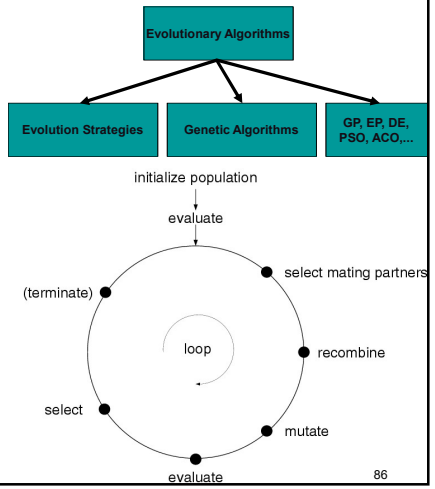

Some Practical Principles for Closed-Loop Optimization

- Keep experimentalists in the loop
- Understand the experimental platform
- Simulate the platform, and compare algorithms
- Do it for real – and get feedback

Keep experimentalists in the loop

- Explain EAs, manage expectations of outcomes.
- Understand the variables and objectives. Confirm 3 times at least.
- Still be prepared to change objectives half-way through!
- Enable them to use familiar software for viewing results.

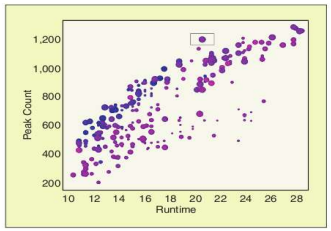

Objectives shown above were changed during optimization

Understand the experimental platform

- Variables, constraints, measurements, noise
- Financial costs, time lags
- Resource constraints
- Batch size of platform dictates/constrains population size of EA

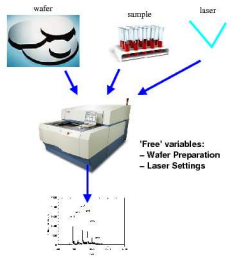

Simulations prior to the *real thing*

- Really helpful to manage expectations of stakeholders
- Tune your algorithms for weird and wonderful population sizes, constraints, budget limitations of real experimental platform
- If possible, use domain experts to design test problems that are similar to the real problem

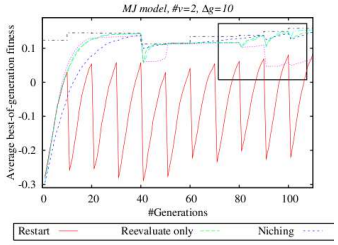

Conclusions

- Experimental Optimization is hard – but an Evolutionary approach is feasible!
- EAs should be given a chance in new application areas
- Fundamental research in EAs is much needed

Goals and Open Questions

- Given a budget of ***k*** experiments – what strategy should be taken?
- NFL holds more than ever – there will be no winner algorithm handling all experimental scenarios!
- How do statistical approaches perform in comparison?
  - *Design-of-Experiments*
- Holy Grail: A package of strategies to drive an experimental system to a reliable maximum with minimum experiments

Box, G. E. P.; Hunter, J. S.; Hunter, W. G.: Statistics for experimenters: design, innovation, and discovery; Wiley-Interscience, 2005.

Acknowledgments

- Hans-Paul Schwefel
- Herschel Rabitz
- Jonathan Roslund
- Douglas Kell
- Roy Goodacre and Seetharaman Vaidyanathan
- David Corne
- Will Rowe
- Chris Knight and the rest of the CLADE team

93