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

GECCO '18 Companion, July 15–19, 2018, Kyoto, Japan

ACM ISBN 978-1-4503-5764-7/18/07.
https://doi.org/10.1145/3205651.3207885

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.

Contributors and former-instructors:

- **Joshua Knowles**, University of Birmingham, UK.
- **Richard Allmendinger**, University of Manchester, UK.

Agenda

- What do we mean by "Sequential Experimentation"?
- Examples of what has been done
- Reference: Statistical Design of Experiments
- Potential Application Areas
- Case-Study: Quantum Control Experiments
- Hot off the lab-bench: Protein Expression
- 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

1960s … … … 2010s

Examples:
- Flow Plate
- Bended Pipe
- Nozzle
- Nutrient Solutions
- Coffee Formulations
- Quantum Control
- Protein Expression

EXAMPLE APPLICATIONS

Early Experiments I: Flow Plate

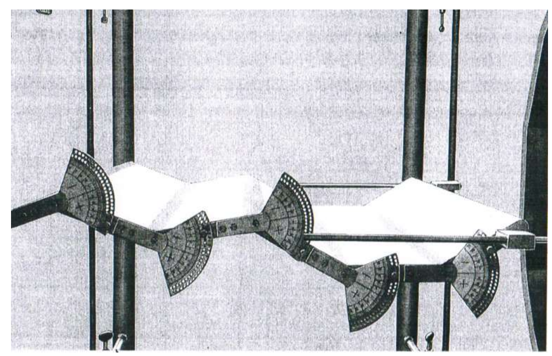

- A plate with 5 controllable angle brackets
- Measurable air flow drag (by a pitot tube)

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

Early Experiments I: Flow Plate

Experiment 1:	Experiment 2:
• Left / right supporting point at same y-coordinate. • Horizontal flow. • Minimize drag.	• Left supporting point 25% lower than right one. • Horizontal flow. • Minimize drag.

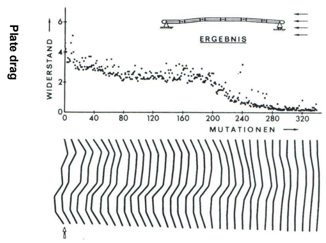

Number of mutations and selected plate shapes

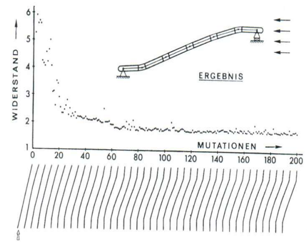

Number of mutations and selected plate shapes

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

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

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

Early Experiments II: Bended Pipe

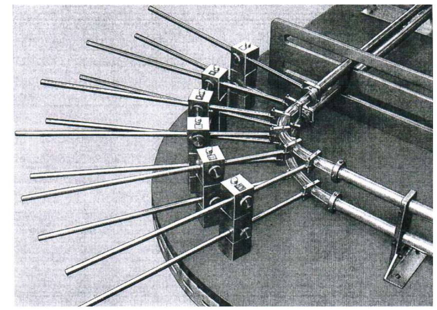

- A flexible pipe with 6 controllable bending devices
- Minimize bend losses of liquid flow
- Measure drag by pitot tube

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

Early Experiments II: Bended Pipe

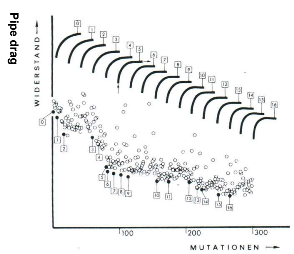

Number of mutations and selected pipe shapes

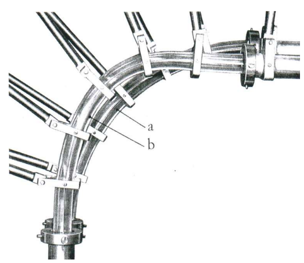

Initial (a) and optimized (b) pipe shape

- 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

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.

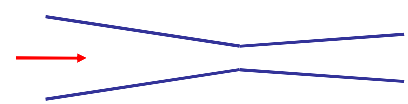

Experimental Setup: Nozzle

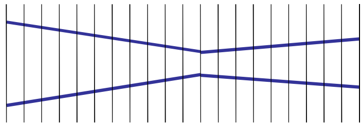

- Production of differently formed conic nozzle parts (pierced plates).
- Form of nozzle part is value of decision variable.

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)

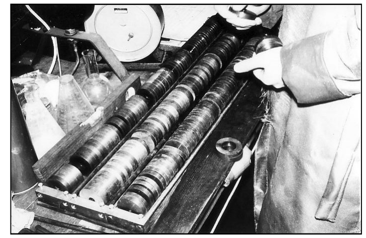

collection of conical nozzle parts

device for clamping nozzle parts

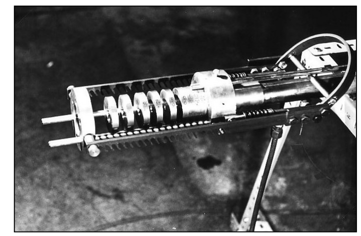

Figures courtesy of Hans-Paul Schwefel

Nozzle Experiment (II)

Hans-Paul Schwefel
while changing nozzle parts

Figures courtesy of Hans-Paul Schwefel

Nozzle Experiment (III)

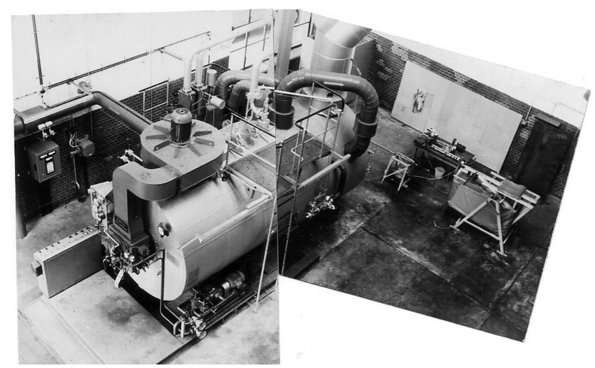

Figures courtesy of Hans-Paul Schwefel

steam plant / experimental setup

Nozzle Experiment (IV)

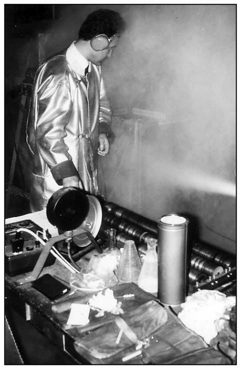

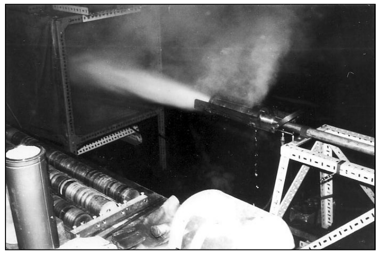

the nozzle in operation …

… while measuring degree of efficiency

Figures courtesy of Hans-Paul Schwefel

Nozzle Results (I)

- Illustrative Example: Optimize Efficiency
  - Initial:

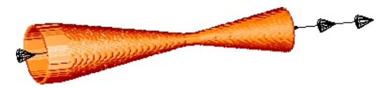

  - Evolution:

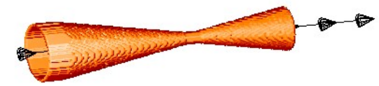

- 32% Improvement in Efficiency !

Nozzle Results (II)

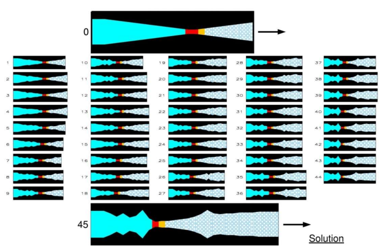

- 250 experiments were made.
- 45 improvements found.
- Discrete ring segments, variable-dimensional optimisation
- Gene duplication and deletion as additional operators.

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.

Experiment: Coffee Formulations

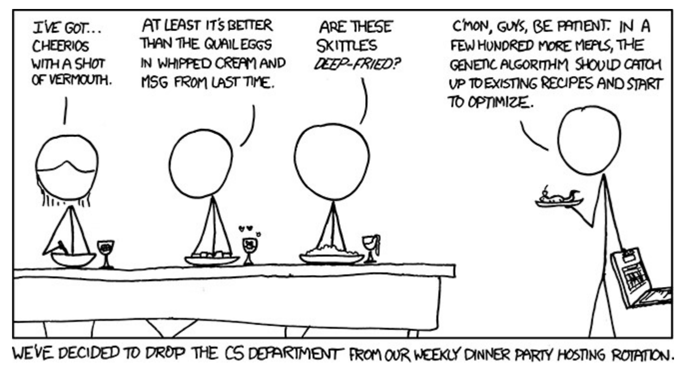

M. Herdy: Beiträge zur Theorie und Anwendung der Evolutionsstrategie, PhD Thesis, Technical University of Berlin, Germany, 2000.

Coffee Formulations: Results

Optimum taste in 11 generations (55 evaluations)

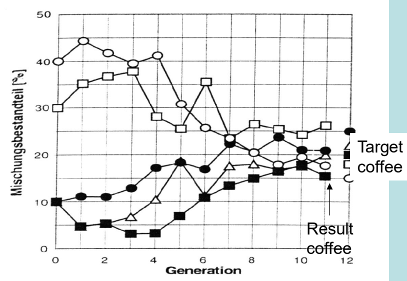

M. Herdy: Beiträge zur Theorie und Anwendung der Evolutionsstrategie, PhD Thesis, Technical University of Berlin, Germany, 2000.

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.

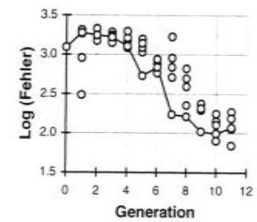

M. Herdy: Beiträge zur Theorie und Anwendung der Evolutionsstrategie, PhD Thesis, Technical University of Berlin, Germany, 2000.

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!

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”)

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.

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

Noise “Colors”

Autocorrelation of the noise spectrum indicates the “memory property” of the disturbance –

- White Noise: $^1/_{f^0} \rightarrow \delta(t)$ (no correlation)
- Pink (Flicker) Noise: $^1/_{f^1} \rightarrow \text{unknown}$
- Red (Brownian) Noise: $^1/_{f^2} \rightarrow e^{-\lambda t}$ (exp. distribution)

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)*

APPLICATION AREAS

Basic Science: Discoveries as Combinatorial Optimization Problems

- A problem shared by scientists is to achieve optimal behavior of their systems and arrive at new discoveries while searching over an array of parameters
- It is commonly visualized in terms of a ‘*landscape*’, in which a candidate solution is mapped onto a ‘*position*’, its quality onto an ‘*altitude*’
- The task is translated into *efficiently navigating within this search-space,* which **scales exponentially with the number of variables**

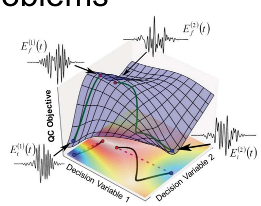

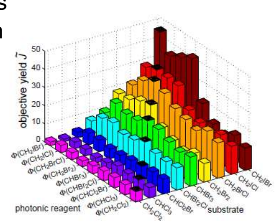

Kell, D.B., *Scientific discovery as a combinatorial optimisation problem: How best to navigate the landscape of possible experiments?* BioEssays, 2012. 34(3): p. 236-244.

A Classification

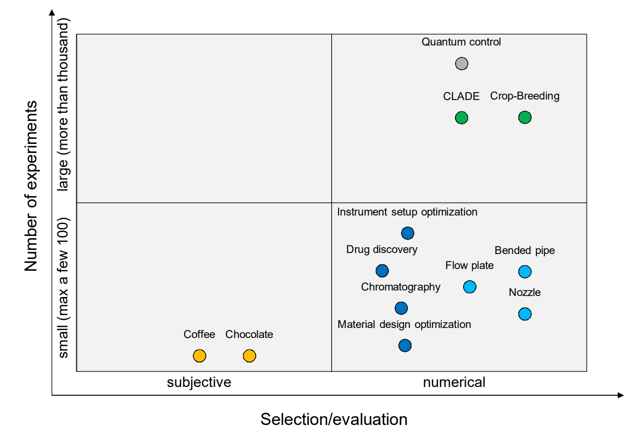

Potential Application Areas

- 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

Reference/State-of-the-Art:

STATISTICAL DESIGN OF EXPERIMENTS

Charts are based on Joshua Knowles' Jyvaskyla Summer School Course (2016)..

DoE: The Industry's Golden Standard

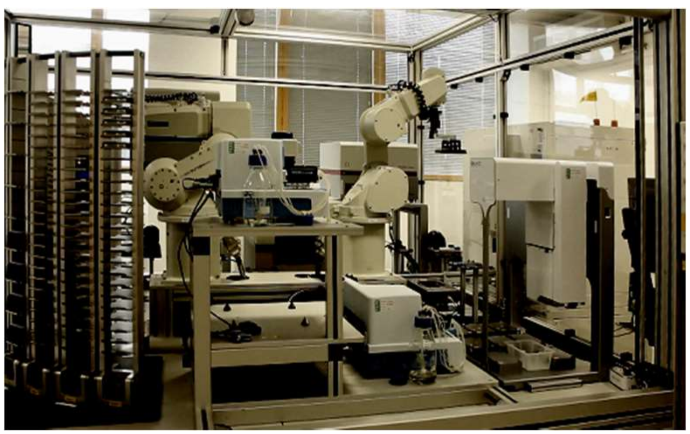

Experimentation terminology

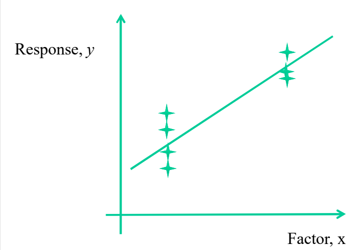

DoE
Response: also known as *effect*
Factor = independent variable
Factors have *levels*
A Factor at a particular level is a *treament*
The regression line is a *model, fit* or *response surface*

Machine Learning
Factors are *features*
The response is the *class* or *output*

Optimization
Factors are *decision variables*
Response is *objective value, cost, benefit, utility* or *fitness*

From classical to modern

- Classical

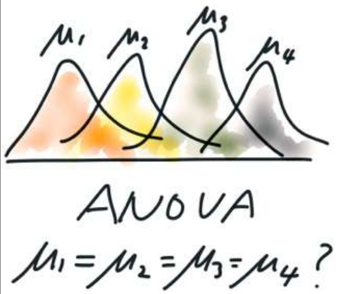

- Modern

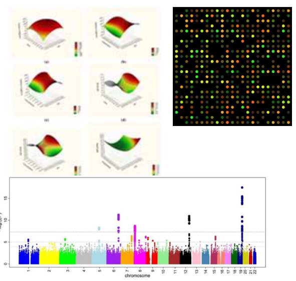

➢ Are modern experiments asking too much?

Classical statistical question

- **One factor**, also known as a treatment, at two different levels, giving *two groups*
- **One effect**, also known as dependent variable, outcome, objective value
- Question: does the factor influence the effect?

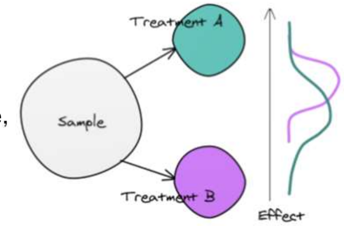

The NIST Engineering Handbook calls this a *comparative design.*

Classical method: ANOVA (Fisher)

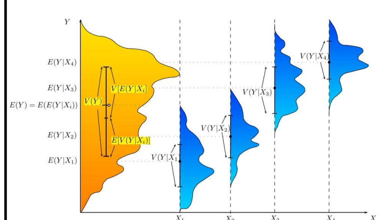

Figure 1: ANOVA : Fair fit

- ANOVA compares the variance between groups and within groups, with total variance
- If variance between groups is dominant then the statistic is significant

Modern experiment 1

- *N* factors, *N*>>2, e.g. genes
- *M* effects, *M*>1, e.g. disease, + other effects
- *P*>1 nuisance factors, ages, gender, etc

- Question: which genes are *most responsible* for the disease, which groups of genes *work together*, and *are other effects involved* in explaining the disease?

Modern experiment 2

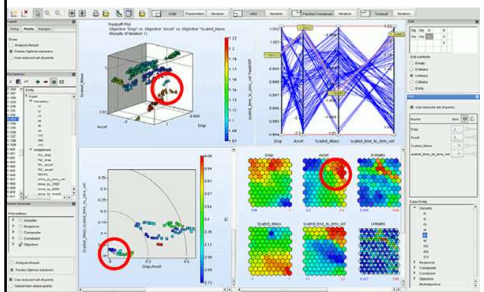

- Many factors
- Several effects
- Several nuisance variables
- Limited number of samples
- Noise (variance)
- Purpose: **Optimize** the effect

Classical DoE: topics

- Statistical models and fitting
- Control experiments and Blocking
- Randomization and Balance
- Bias and Blinding
- OFAT
- Full factorial
- Fractional factorial
- Optimal design, e.g. D-optimal
- Other designs

Variation

- Experiments involve measurements of quantities that are assumed to be *random variables*
- To determine if any observed effects are due to a factor (independent variable), we must account for the natural variation

Wikimedia Commons: Ron, S.; Venegas, P. J.; Toral, E.; Read, V. M.; Ortiz, D.; Manzano, A. (2012)

Statistical models

- To estimate model parameters from a set of experimental results is often a regression problem:

$y = \alpha + \beta.x1 + \gamma.x2 + \varepsilon$ (Solve system of equations)

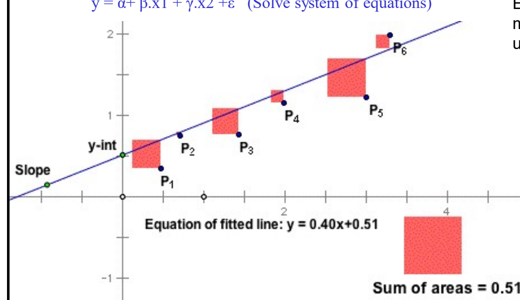

Even assuming a linear model as here, there is no unique solution.

The method of Least Squares has statistical support. It is the best linear unbiased estimator (blue)

(Laplace, Legendre, Gauss)

Statistical models

- To estimate model parameters from a set of experimental results - a regression problem

$y = \alpha + \beta.f(x1) + \gamma.g(x2) + \varepsilon$ (Solve system of equations)

NOTICE: We can fit nonlinear surfaces using linear regression. We just need to assume the functions f g, etc

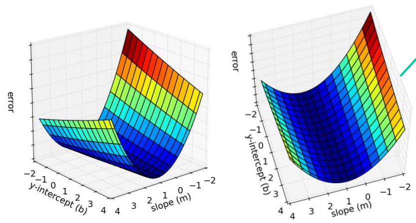

The method of Least Squares has statistical support. It is the best linear unbiased estimator (blue)

Replication

- The bedrock of experimentation is repetition or replication
  - In the simplest case this entails repeating a measurement, keeping all factors the same and taking the *mean* of the response
  - **More generally**, we wish to replicate some factors and let others vary. One way to control this is called *blocking*...[see later]

Randomization

- Several unwanted effects can be reduced by **shuffling** *before* assigning to treatment groups
- Let's say we have three treatments *A,B,C* and *N*=21 trials to do
  - **Nonrandom:** *AAAAAAABBBBBBBCCCCCCC*
  - **Complete Randomization**: *ABAABCCCBABBCCABBCCCB*
  - **Permuted Block Randomization:** *ABC|CAB|BCA|ACB|ACB|BAC|CBA*

Blocking

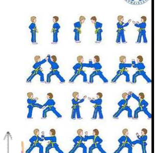

- Randomization is ok, but **blocking** can help even more
- We mitigate the disruptive effects of each potentially confounding factor by splitting the data into groups (or blocks) by that factor

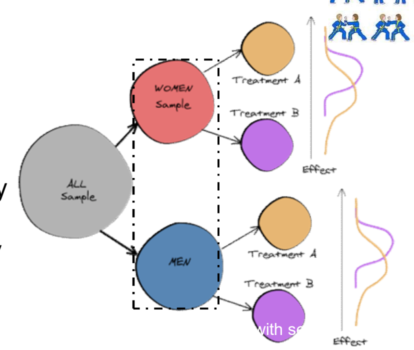

Handling multiple factors

- It is typical that we have N>1 factors to control
- The high-school solution to this is called

OFAT

(or one-factor-at-a-time)

- You hold all but one factor constant and vary that. Then you go onto the second factor ... and so on

OFAT

An OFAT design in two variables

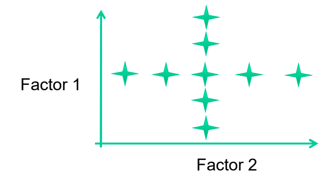

Weaknesses of OFAT

1. OFAT requires more* runs for the same precision in effect estimation
2. OFAT cannot estimate interactions between factors
3. OFAT can miss optimal settings of factors

*compared with experimental designs like Plackett-Burman

Why OFAT is poor: correlations

With several factors, the factors can **confound** each other, which means that one factor causes an increase while others cause a decrease.

The way to overcome this efficiently (in number of trials) is to *change several variables at once!*

Consider the 12 coins problem:

You have 12 coins, one of which is heavier or lighter than the others.
Determine which it is, and whether it is heavy or light.
You have only 3 weighings on a balance!!!

SOLUTION

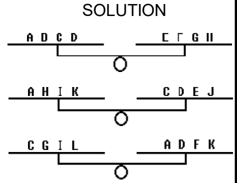

*compared with experimental designs like Plackett-Burman

Full Factorial design

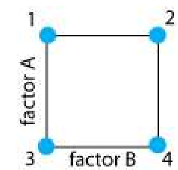

factor levels

trial	A	B
1	+	–
2	+	+
3	–	–
4	–	+

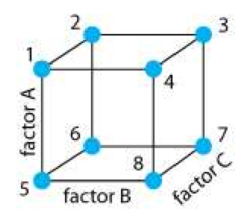

factor levels

trial	A	B	C
1	+	–	–
2	+	–	+
3	+	+	+
4	+	+	–
5	–	–	–
6	–	–	+
7	–	+	+
8	–	+	–

Fractional Factorial Design

A fractional factorial design is not unique. 2^p different designs

Another 2^{4-1} Experimental Design

Expt No.	A	B	C	D	AC	BC	ABC
1	-1	-1	-1	1	1	1	-1
2	1	-1	-1	-1	-1	1	1
3	-1	1	-1	-1	1	-1	1
4	1	1	-1	1	-1	-1	-1
5	-1	-1	1	1	-1	-1	1
6	1	-1	1	-1	1	-1	-1
7	-1	1	1	-1	-1	1	-1
8	1	1	1	1	1	1	1

Confoundings:

I=ABD, A=BD, B=AD, C=ABCD, D=AB, AC=BCD, BC=ACD, ABC=CD

Other fractional designs

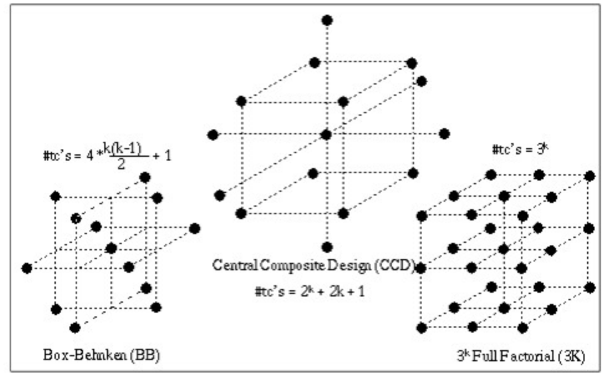

Latin hypercube screening design

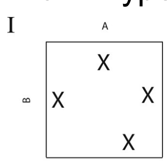

Random design

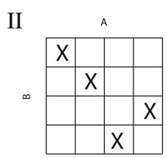

Latin Hypercube

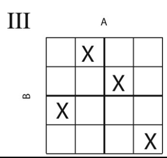

Orthogonal

EAs *versus* DoE comparison

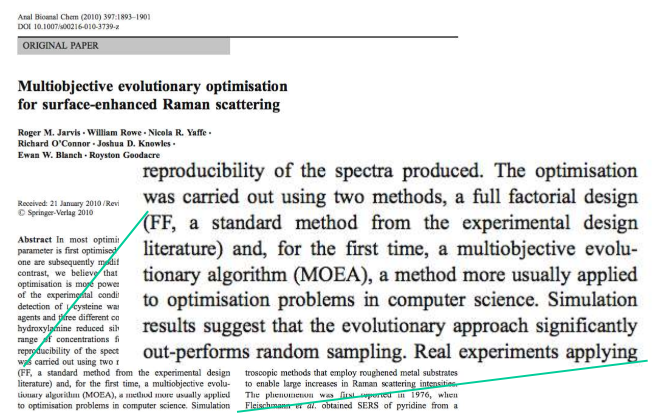
Anal Bioanal Chem (2010) 397:1893–1901
DOI 10.1007/s00216-010-3739-z

ORIGINAL PAPER

Multiobjective evolutionary optimisation for surface-enhanced Raman scattering

Roger M. Jarvis · William Rowe · Nicola R. Yaffe · Richard O'Connor · Joshua D. Knowles · Ewan W. Blanch · Royston Goodacre

Received: 21 January 2010 /Revi
© Springer-Verlag 2010

Abstract In most optimi… parameter is first optimised… one are subsequently modif… contrast, we believe that… optimisation is more power… of the experimental condi… detection of L-cysteine was… agents and three different co… hydroxylamine reduced silv… range of concentrations f… reproducibility of the spect… was carried out using two r… (FF, a standard method from the experimental design literature) and, for the first time, a multiobjective evolutionary algorithm (MOEA), a method more usually applied to optimisation problems in computer science. Simulation

reproducibility of the spectra produced. The optimisation was carried out using two methods, a full factorial design (FF, a standard method from the experimental design literature) and, for the first time, a multiobjective evolutionary algorithm (MOEA), a method more usually applied to optimisation problems in computer science. Simulation results suggest that the evolutionary approach significantly out-performs random sampling. Real experiments applying

troscopic methods that employ roughened metal substrates to enable large increases in Raman scattering intensities. The phenomenon was first reported in 1976, when Fleischmann et al. obtained SERS of pyridine from a

Comparing LHS and (1+1)-Cholesky-CMA-ES

- Comparison on BBOB, 5-d and 20-d, 80 repetitions, ECDF plots
- Above 10d FEs, (1+1)-Cholesky-CMA-ES always outperforms LHS

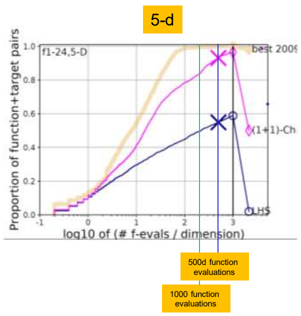

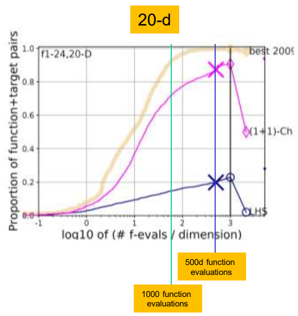

Case-Study:

QUANTUM CONTROL EXPERIMENTS

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. $|\langle\psi_{\mathrm{target}}|\,\psi(T)\rangle|^2 \longrightarrow \max$

Hamiltonian required

PRA **37**, 4950 (1988)

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

Pulse Shaping → $E(t)$ → Molecular Sample → Observation (detector) → Feedback (signal) → Optimization Routine → modified $E(t)$ → Pulse Shaping

Hamiltonian **not** required

PRL **68**, 1500 (1992)

Quantum Control Experiments

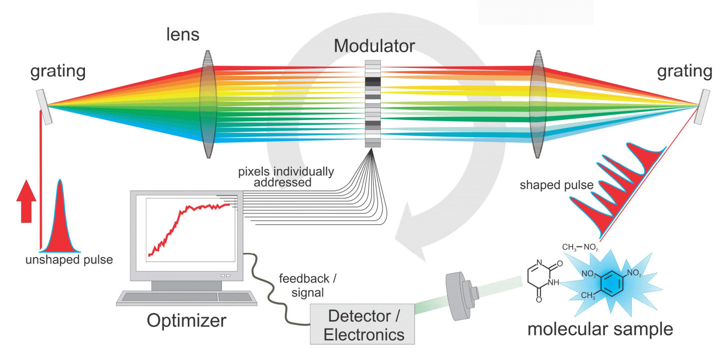

Tel-Hai College

The QCE Arena: The Optical Table

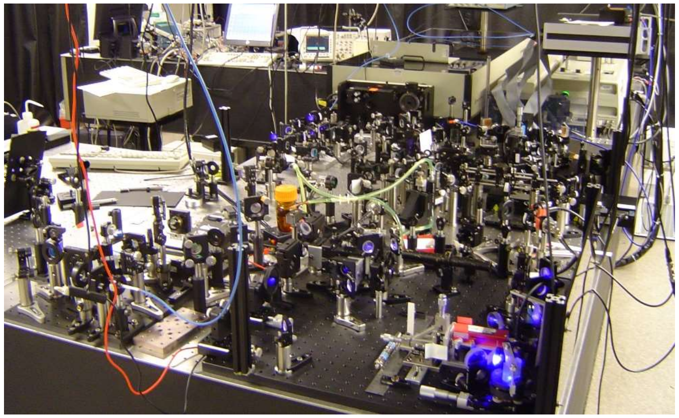

Figure courtesy of Jonathan Roslund

Tel-Hai College

The Optical Table: Shaping the Pulse

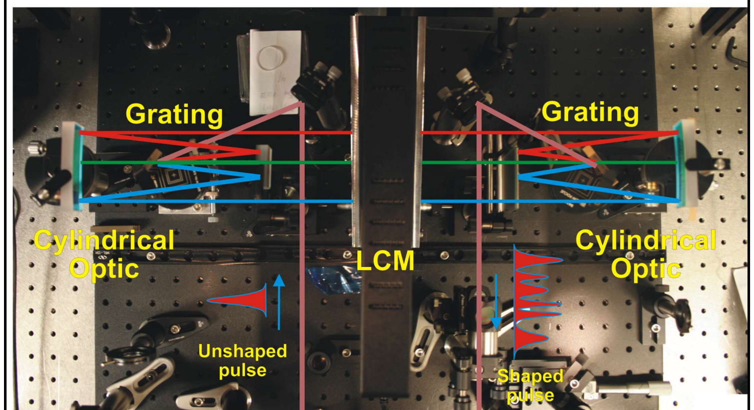

Figure courtesy of Jonathan Roslund

Tel-Hai College

QCE: Sources of Noise/Uncertainty

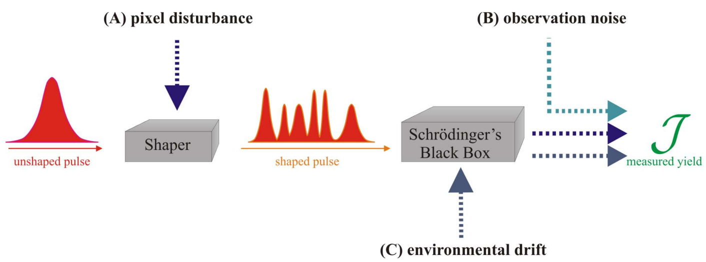

(A) $\tilde{\phi}(\omega) = \left(\phi(\omega_1) + \mathcal{N}_1\left(0, \epsilon_S^2\right), \ldots, \phi(\omega_n) + \mathcal{N}_n\left(0, \epsilon_S^2\right)\right)$

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

Tel-Hai College

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

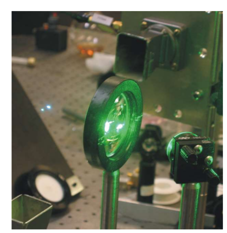

Figure courtesy of Jonathan Roslund

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

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.

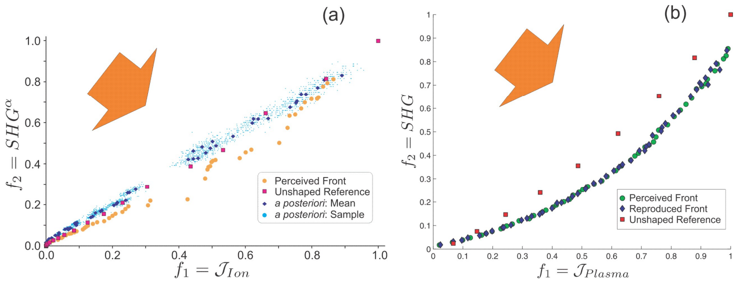

Shir, O.M., Roslund, J., Leghtas, Z., Rabitz, H.: Quantum Control Experiments as a Testbed for Evolutionary Multi-Objective Algorithms. Genetic Programming and Evolvable Machines **13**(4) (2012) 445—491

Radar Optimal Dynamic Discrimination

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

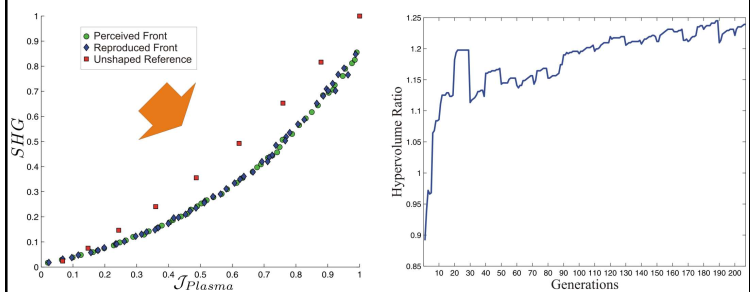

Shir, O.M., Roslund, J., Leghtas, Z., Rabitz, H.: Quantum Control Experiments as a Testbed for Evolutionary Multi-Objective Algorithms. Genetic Programming and Evolvable Machines **13**(4) (2012) 445—491

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

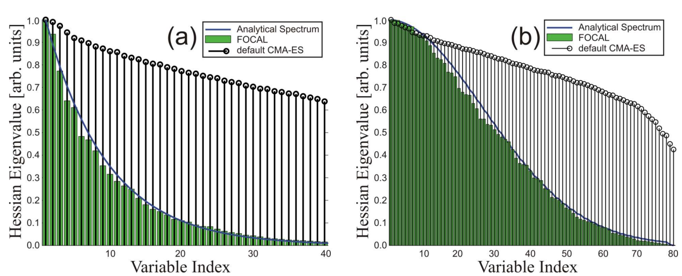

Shir, O.M., Roslund, J., Whitley, D., Rabitz, H.: Efficient retrieval of landscape Hessian: Forced optimal covariance adaptive learning. Physical Review E **89**(6) (2014) 063306

FOCAL: Experimental Results

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

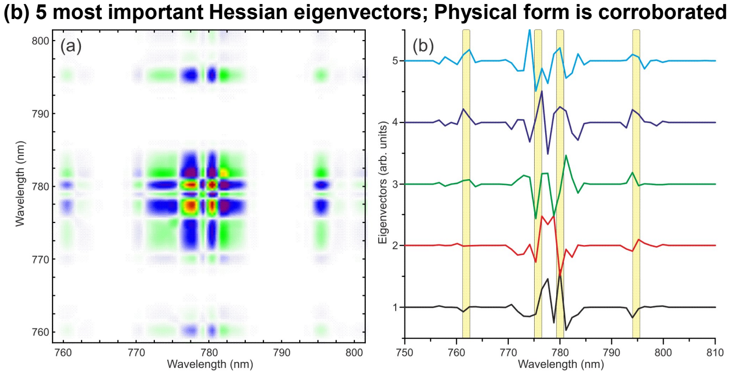

Shir, O.M., Roslund, J., Whitley, D., Rabitz, H.: Efficient retrieval of landscape Hessian: Forced optimal covariance adaptive learning. Physical Review E **89**(6) (2014) 063306

Hot-off-the-lab-bench

PROTEIN EXPRESSION

Heterologous Protein Expression

Four genes are required for ApcA heterologous expression in E. coli.

Goal: maximize the heterologous expression level

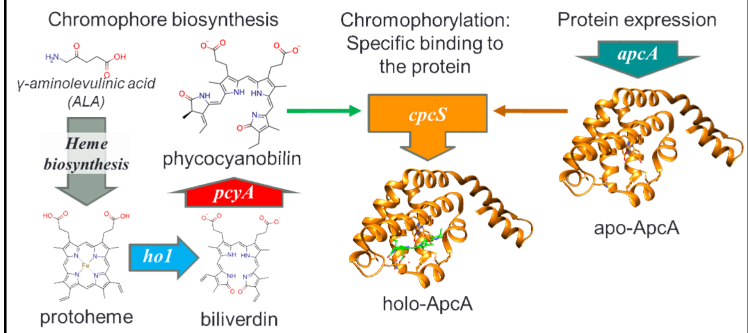

ApcA Expression System: Four Genes in Two Plasmids

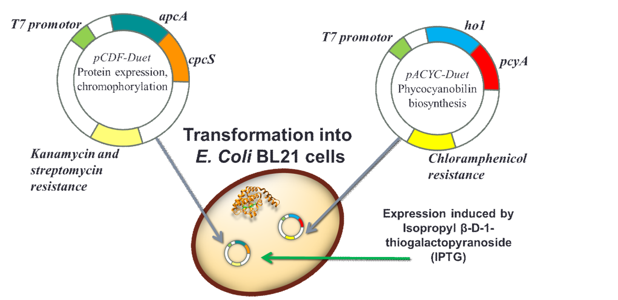

- Controls (10 categorical decision variables):
  4 growth temp., 5 expression temp., 3 growth volumes,
  6 IPTG concent., 5 O.D. values, 4 induction durations,
  7 gamma-ALA concent., 2 gamma-ALA timings,
  11 FeCl_3 concent., 3 Medium types
- Search-space cardinality: ~3×10^6 possible combinations

Closed Feedback Loop

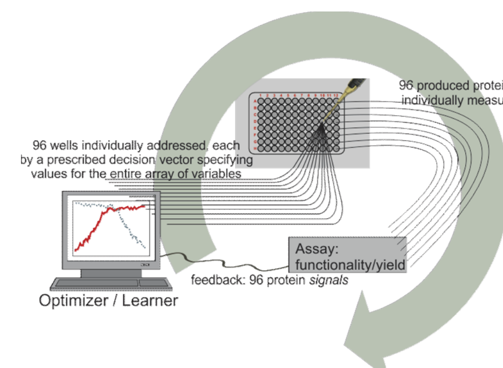

A scheme of a typical single iteration (step) in a proposed optimization run of a production system. The input variables for each well are prescribed by the algorithm, whereas the output (feedback) of each produced molecule is provided by the assay - altogether closing a feedback loop.

ApcA Expression in E. Coli: Assay

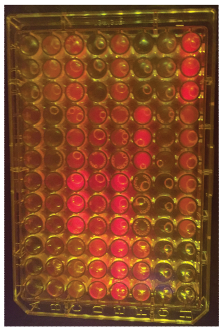

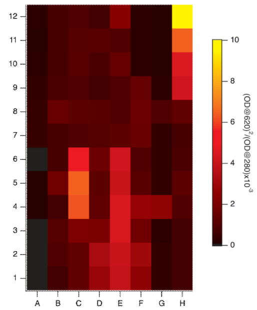

ApcA Expression in E. Coli: 4 generations

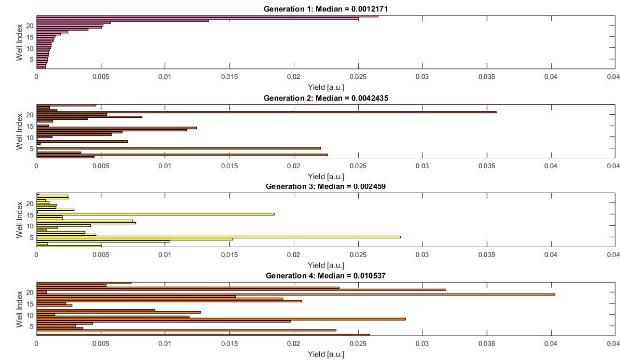

DISCUSSION

Evolutionary Algorithms Used

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

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

•Protein Expression Experiments: Categorical ES [unpublished]

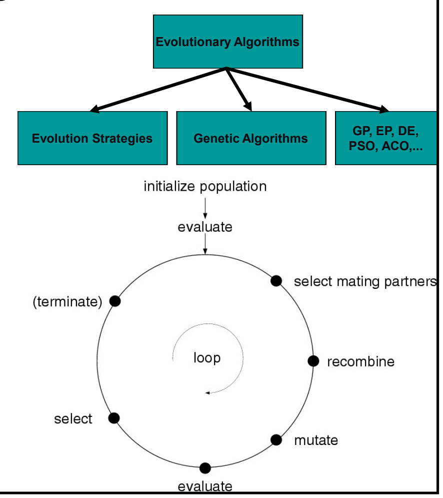

Some Practical Principles for Closed-Loop Optimization

- 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.

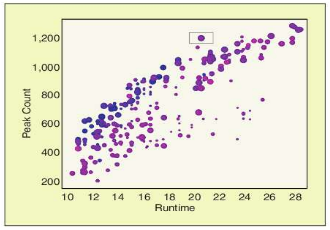

Objectives shown above were changed during optimization

Understand the experimental platform

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

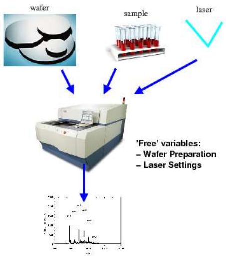

Simulations prior to the *real thing*

- 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

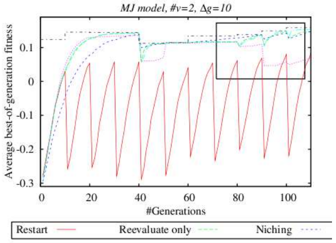

Conclusions

- Experimental Optimization is hard – but an Evolutionary approach is feasible!
- EAs should be given a chance in new application areas
- The human/psychological factor among the experimentalists plays a dominant role in making a decision on starting a campaign
- 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?
  - Especially *DoE*
- **The preliminary comparison presented earlier is a fine starting point** (slide #57)
- Holy Grail: A package of strategies to drive an experimental system to a reliable maximum with minimum experiments

Acknowledgments

- Joshua Knowles
- Richard Allmendinger
- Hans-Paul Schwefel
- Herschel Rabitz
- Jonathan Roslund
- Dror Noy
- Chen Erlich
- Hao Wang

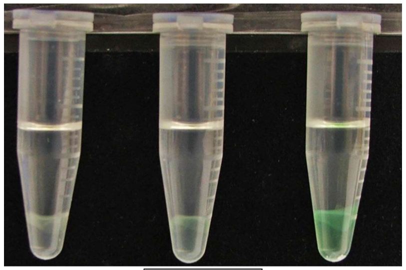

Figure courtesy of Dror Noy

Bibliography

- Allmendinger, R. and Knowles, J. (2013). On handling ephemeral resource constraints in evolutionary search. *Evolutionary Computation*, **21**(3):497–531.
- Bäck, T., Foussette, C., and Krause, P. (2013). *Contemporary Evolution Strategies.* Natural Computing Series. Springer-Verlag Berlin Heidelberg.
- Box, G. E., Hunter, J. S., and Hunter, W. G. (2005). *Statistics for Experimenters: Design, Innovation and Discovery*. John Wiley and Sons, Hoboken, NJ, USA, second edition.
- Eiben, A. E. and Smith, J. (2015). From evolutionary computation to the evolution of things. *Nature*, 521:476–482.
- Kell, D. B. (2012). Scientific discovery as a combinatorial optimisation problem: How best to navigate the landscape of possible experiments? *BioEssays*, **34**(3):236–244.
- Knowles, J. (2006). ParEGO: A Hybrid Algorithm with On-Line Landscape Approximation for Expensive Multiobjective Optimization Problems. *IEEE Transactions on Evolutionary Computation*, **10**(1):50–66.
- Knowles, J. (2009). Closed-loop evolutionary multiobjective optimization. *IEEE Computational Intelligence Magazine*, **4**(3):77–91.
- Rechenberg, I. (2000). Case studies in evolutionary experimentation and computation. *Computer Methods in Applied Mechanics and Engineering*, **186**(24):125–140.
- Roslund, J., Shir, O. M., Bäck, T., and Rabitz, H. (2009). Accelerated Optimization and Automated Discovery with Covariance Matrix Adaptation for Experimental Quantum Control. *Physical Review A*, **80**(4):043415.
- Shir, O. M., Roslund, J., Leghtas, Z., and Rabitz, H. (2012). Quantum control experiments as a testbed for evolutionary multi-objective algorithms. *Genetic Programming and Evolvable Machines*, **13**(4):445–491.
- Shir, O. M., Roslund, J., Whitley, D., and Rabitz, H. (2014). Efficient retrieval of landscape hessian: Forced optimal covariance adaptive learning. *Physical Review E*, **89**:063306.