LOST BETWEEN TWO SHORES

JASPER A. VRUGT

Department of Civil and Environmental Engineering
University of California, Irvine, USA

Institute for Biodiversity and Ecosystem Dynamics (IBED)
University of Amsterdam, The Netherlands

THE NETHERLANDS (HOLLAND)

Known for:

1. Beer (Heineken, Amstel)
2. Soccer (Cruijff, Sneijder)
3. Cheese (Gouda, Edam)
4. Beauty of its capital (??)
5. Flowers
6. Non-interesting geology

ALMA MATER: UNIVERSITY OF AMSTERDAM

CAMPUS AREA

LOTS OF BIKES

TRANSITION TO THE WORLD

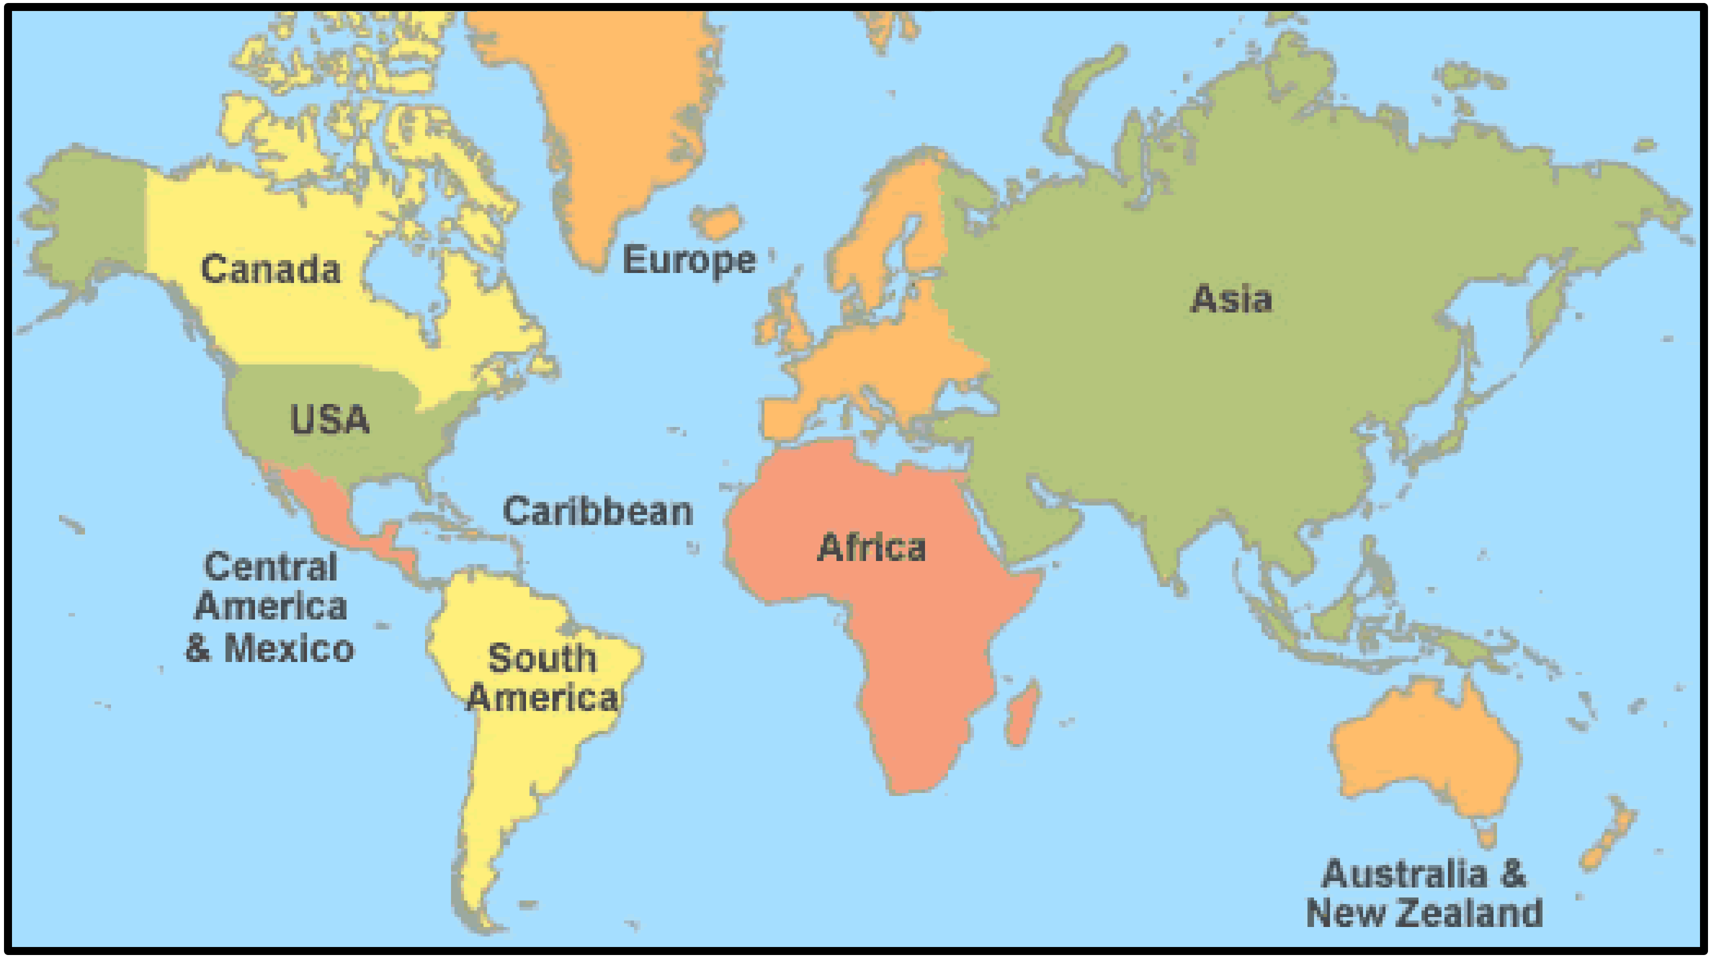

NOW UC - IRVINE

CAMPUS AREA

NEWPORT BEACH

WHAT DO I ACTUALLY DO?

NORMAN BEAN

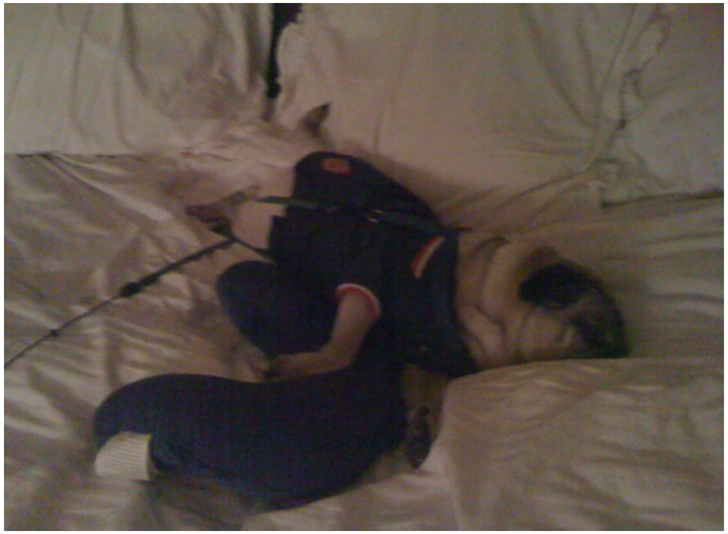

ENVIRONMENTAL MODELING

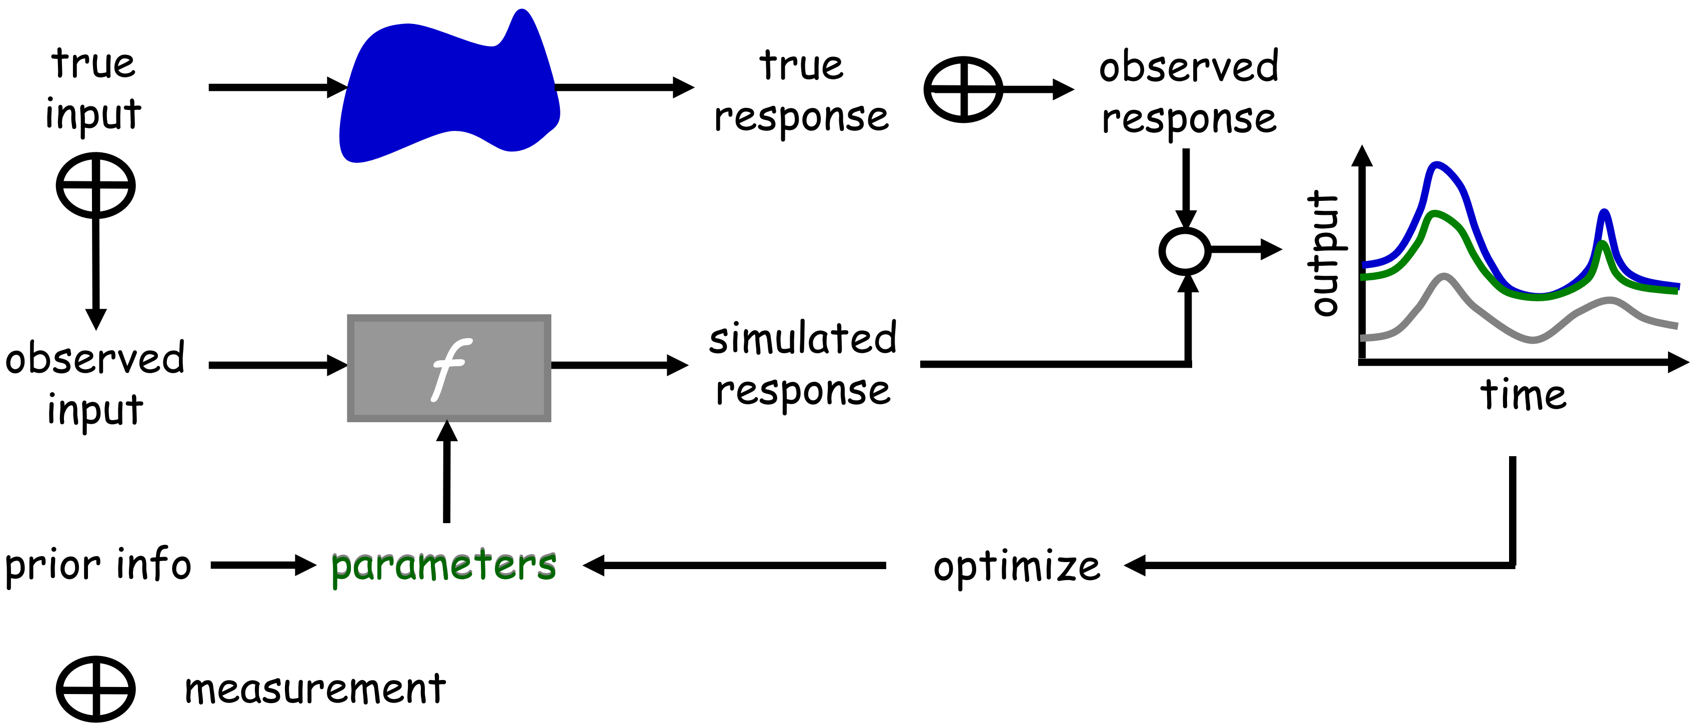

TUNING THE PARAMETERS SO THAT CLOSEST FIT TO THE OBSERVED SYSTEM RESPONSE IS OBTAINED

CLASSICAL ESTIMATION: A SINGE SOLUTION

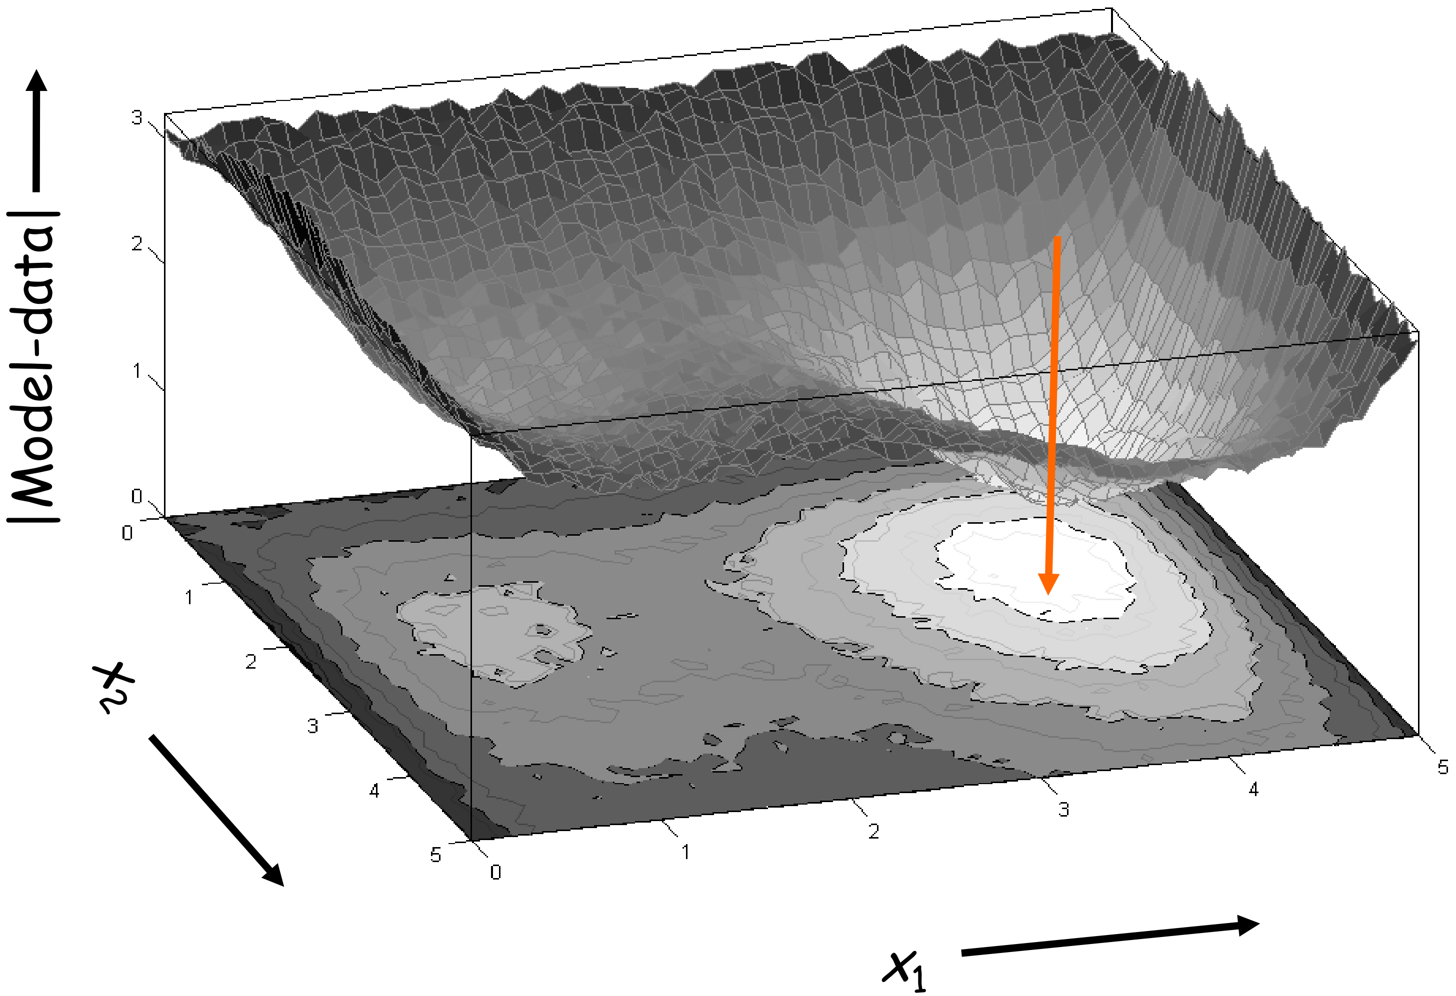

MODELING FRAMEWORK WITH UNCERTAINTY

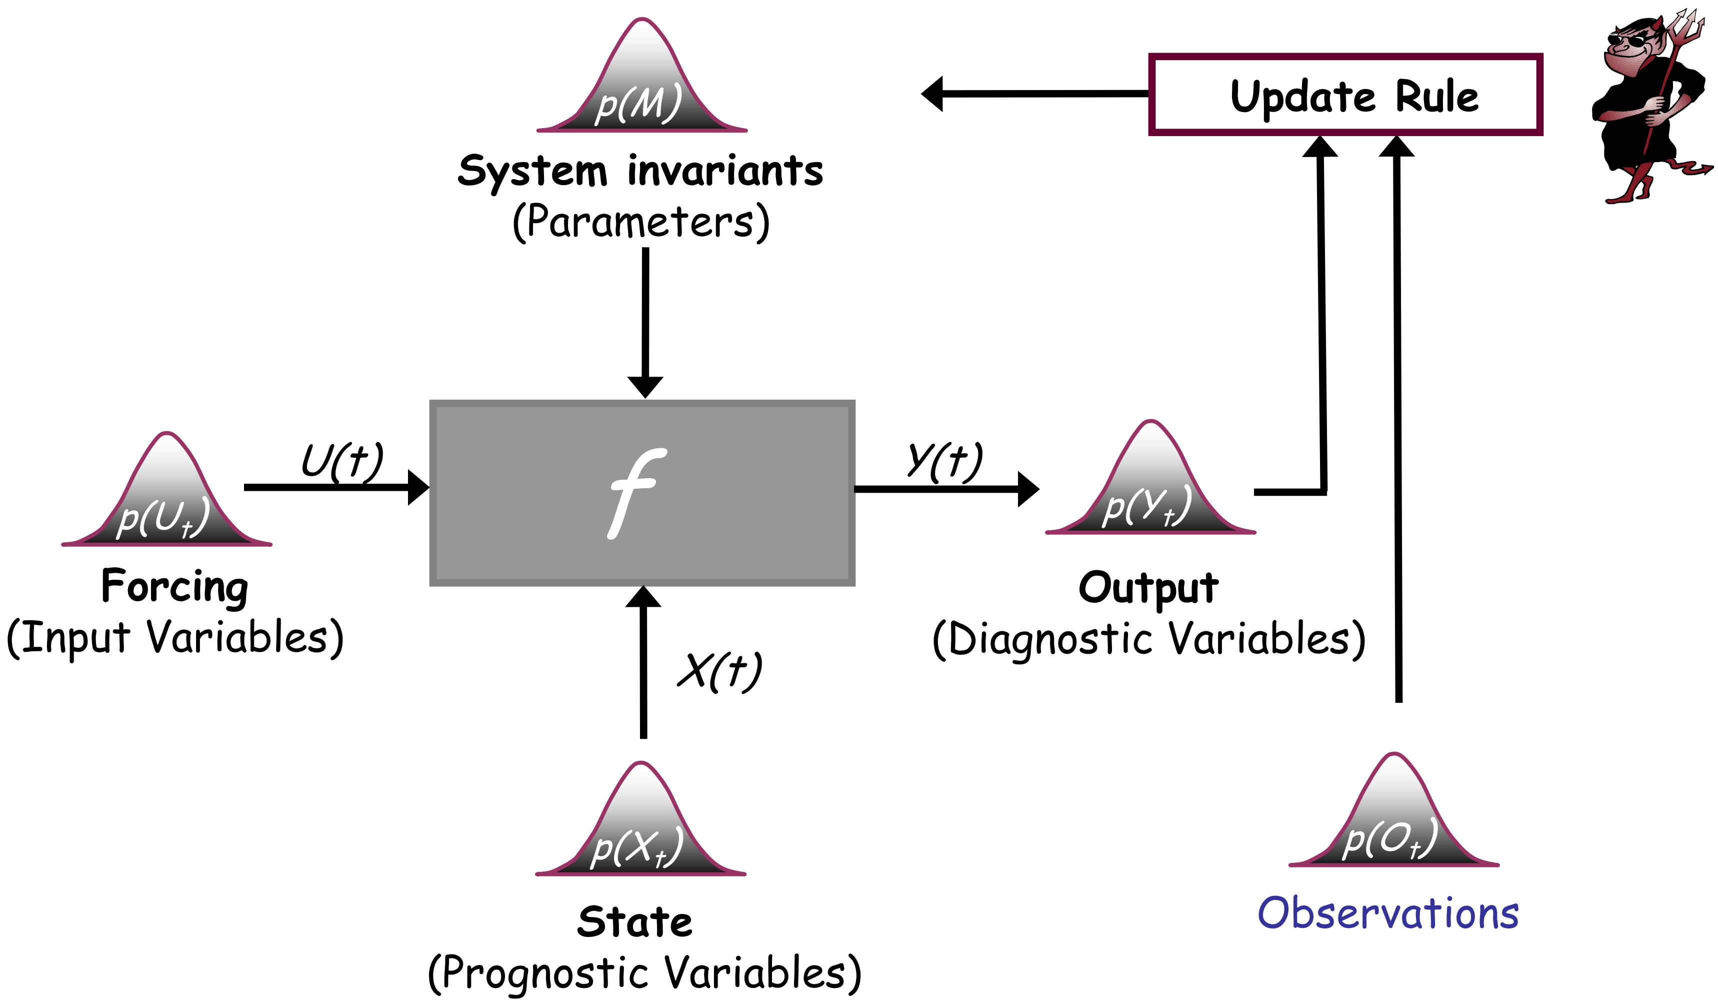

I. Pareto optimization

Multi-objective optimization

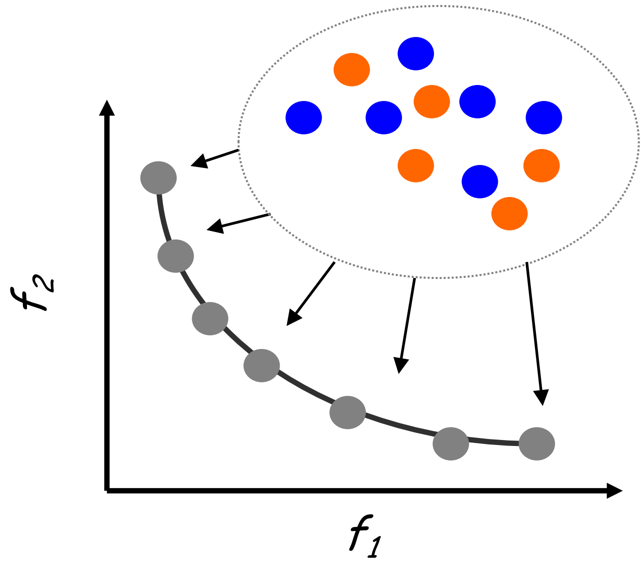

AMALGAM optimization

II. Stochastic Optimization

Markov Chain Monte Carlo sampling

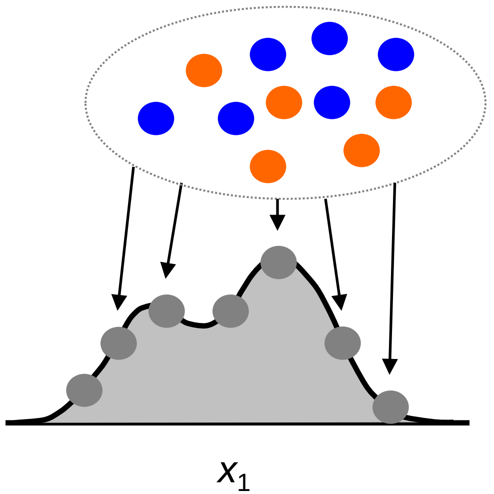

DREAM - MCMC

OTHER MODEL - DATA SYNTHESIS METHODS

III. Sequential Data Assimilation

Combined Parameter & State Estimation

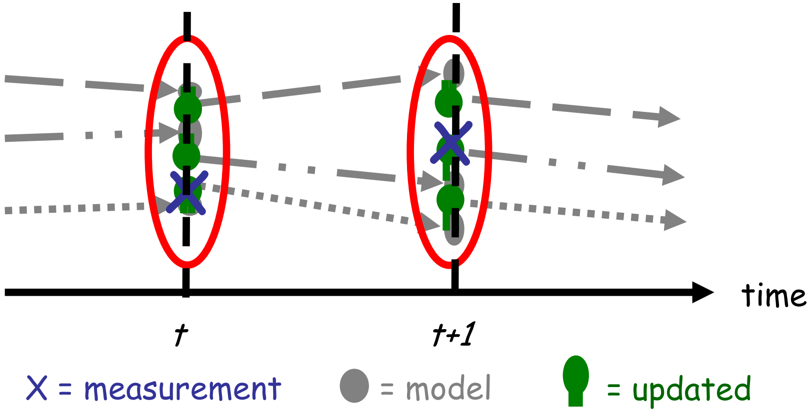

IV. Model Averaging

Bayesian / Mallows Model Averaging

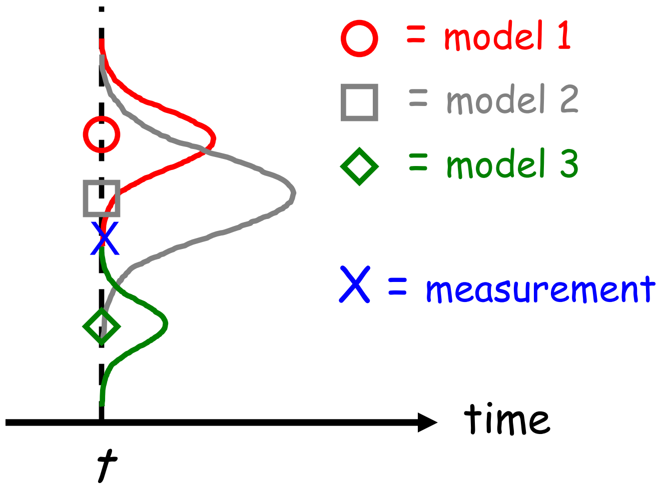

P-PREAM: Differential Evolution Particle Filter
Sequential Monte Carlo + MCMC with DREAM

Quantile Regression
PDF sharper than BMA

DIFFERENTIAL EVOLUTION ADAPTIVE METROPOLIS

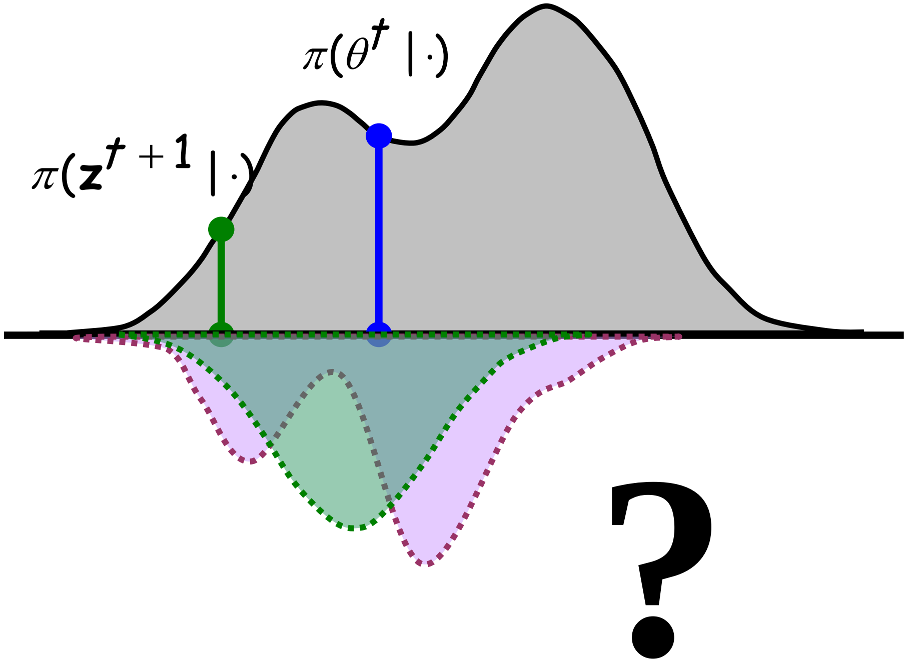

DREAM: Continuously Updates the Scale and Orientation of the Proposal Distribution

Maintains Detailed Balance and is Ergodic
Handles Multimodality Efficiently
High-dimensionality

ESPECIALLY DESIGNED FOR PARALLEL COMPUTING

Accelerating Markov Chain Monte Carlo Simulation By Self-Adaptive Differential Evolution with Randomized Subspace Sampling

Jasper A. Vrugt*, Cajo J.F. ter Braak ‡, Cees G.H. Diks §, Bruce A. Robinson ¶, James M. Hyman ‖, and Dave Higdon **

† Center for NonLinear Studies (CNLS), Los Alamos National Laboratory, Los Alamos, NM 87545, USA, ‡ Biometris, Wageningen University and Research Centre, 6700 AC, Wageningen, The Netherlands, § Center for Nonlinear Dynamics in Economics and Finance, University of Amsterdam, Amsterdam, The Netherlands, ¶ Civilian Nuclear Program Office (SPO-CNP), Los Alamos National Laboratory, Los Alamos, NM 87545, USA, ‖ Mathematical Modeling and Analysis group (T-7), Los Alamos National Laboratory, Los Alamos, NM 87545, USA, and ** Statistical Sciences (CCS-6), Los Alamos National Laboratory, Los Alamos, NM 87545, USA

Submitted to Proceedings of the National Academy of Sciences of the United States of America

Markov chain Monte Carlo (MCMC) methods have found widespread use in many fields of study to estimate the average properties of complex systems, and for posterior inference in a Bayesian framework. Existing theory and experiments prove convergence of well constructed MCMC schemes to the appropriate limiting distribution under a variety of different conditions. In practice, however this convergence is often observed to be disturbingly slow. This is frequently caused by an inappropriate selection of the proposal distribution used to generate trial moves in the Markov Chain. Here we show that significant improvements to the efficiency of MCMC simulation can be made by using a self-adaptive Differential Evolution learning strategy within a population-based evolutionary framework. This scheme, entitled DiffeRential Evolution Adaptive Metropolis or DREAM, runs multiple different chains simultaneously for global exploration, and automatically tunes the scale and orientation of the proposal distribution during the search. Ergodicity of the algorithm is proved, and various examples involving nonlinearity, high-dimensionality, and multimodality show that DREAM is generally superior to other adaptive MCMC sampling approaches. The DREAM scheme significantly enhances the applicability of MCMC simulation to complex, multi-modal search problems.

Markov chain Monte Carlo | adaptive proposal | randomized subspace sampling | differential evolution | delayed rejection |

In 1953, Metropolis et al. [28] introduced the Markov chain Monte Carlo (MCMC) scheme to estimate $E_\pi f(x)$, the expectation of a function f with respect to a distribution π. The basis of this method is a Markov chain that generates a random walk through the search space and successively visits solutions with stable frequencies stemming from a fixed probability distribution. The MCMC estimator is approximated as the unweighted mean of f along the last M elements of the realized path of the chain, $\frac{1}{M}\sum_{t=1}^{M} f(x_t)$, that is, after a burn-in period to allow the chain to explore the search space and reach its stationary regime. This algorithm has been used extensively in statistical physics, and appeared also in spatial statistics and statistical image analysis. In [11] the MCMC method was extended for posterior inference in a Bayesian framework. Ever since, the method has found widespread use in many different fields ranging from physics and chemistry, to finance, economics, genetics, statistical inference, biology and bioinformatics [9, 10, 39, 6, 30, 38, 1, 23, 21].

To visit configurations with a stable frequency, an MCMC algorithm generates trial moves from the current ("*old*") position of the Markov chain x_{t-1} to a new state z. The earliest and most general MCMC approach is the random walk Metropolis (RWM) algorithm. Assume that we have already sampled points $\{x_0, \ldots, x_{t-1}\}$ this algorithms proceeds in the following three steps. First, a candidate point z is sampled from a proposal distribution q that depends on the present location, x_{t-1} and is symmetric, $q(x_{t-1}, z) = q(z, x_{t-1})$. Next, the candidate point is either accepted or rejected using the Metropolis acceptance probability:

$$\alpha(x_{t-1}, z) = \begin{cases} \min\left(\frac{\pi(z)}{\pi(x_{t-1})}, 1\right) & \text{if } \pi(x_{t-1}) > 0 \\ 1 & \text{if } \pi(x_{t-1}) = 0 \end{cases} \qquad [1]$$

where $\pi(\cdot)$ denotes the probability density function (pdf) of the target distribution. Finally, if the proposal is accepted the chain moves to z, otherwise the chain remains at its current location x_{t-1}.

The original RWM scheme is constructed to maintain detailed balance with respect to $\pi(\cdot)$ at each step in the chain:

$$p(x_{t-1})p(x_{t-1} \to z) = p(z)p(z \to x_{t-1}) \qquad [2]$$

where $p(x_{t-1})$ ($p(z)$) denotes the probability of finding the system in state $x_{t-1}(z)$, and $p(x_{t-1} \to z)$($p(z \to x_{t-1})$) denotes the conditional probability to perform a trial move from x_{t-1} to z (z to x_{t-1}). The result is a Markov chain which under some regularity conditions has a unique stationary distribution with pdf $\pi(\cdot)$. Hastings [20] extended Eq. 1 to include non-symmetrical proposal distributions, i.e. $q(x_{t-1}, z) \neq q(z, x_{t-1})$ in which a proposal jump to z and the reverse jump do not have equal probability. This extension is called the Metropolis Hastings algorithm (MH), and has become the basic building block of many existing MCMC sampling schemes.

The simplicity of the original MH algorithm and theoretically sound statistical basis of the method has lead to widespread implementation and use. However, in practice the MH algorithm requires tuning of some internal variables before the MCMC simulator works properly. The efficiency of the method is essentially determined by the scale and orientation of the proposal distribution, $q(x_{t-1}, \cdot)$ used to generate trial moves (transitions) in the Markov Chain. When the

The authors declare no conflict of interest.

Abbreviations: MCMC, markov chain monte carlo; DREAM, diffeRential evolution adaptive metropolis; DRAM, delayed rejection adaptive metropolis; DE-MC, differential evolution markov chain

* To whom correspondence should be addressed. E-mail: vrugt@lanl.gov

Vrugt et al., *WRR*, (2008); Vrugt et al., *IJNSNS*, (2009)

HYDROGEOLOGIC MODEL CALIBRATION

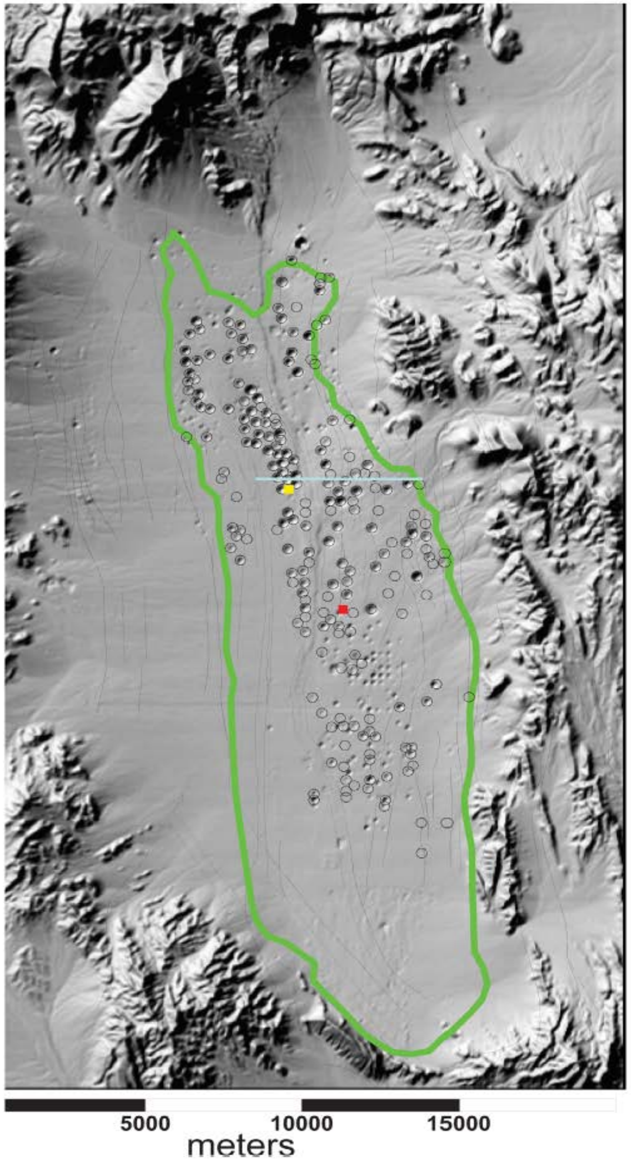

NEVADA TEST SITE

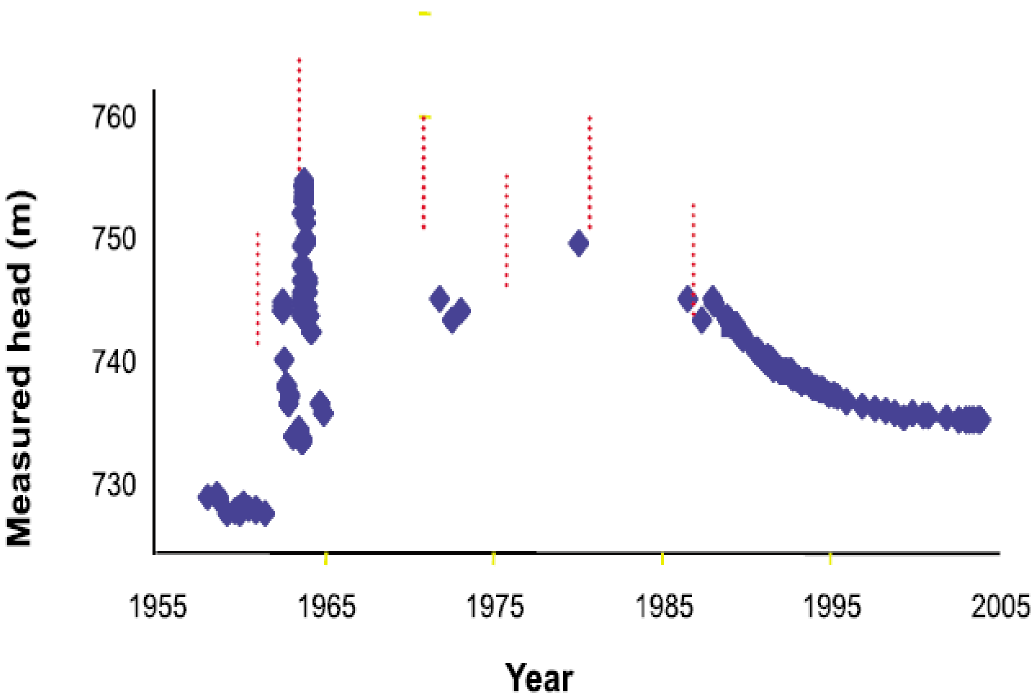

241 parameters

Keating et al., *WRR* (2010)

MT-DREAM$_{ZS}$: CALIBRATION RESULTS

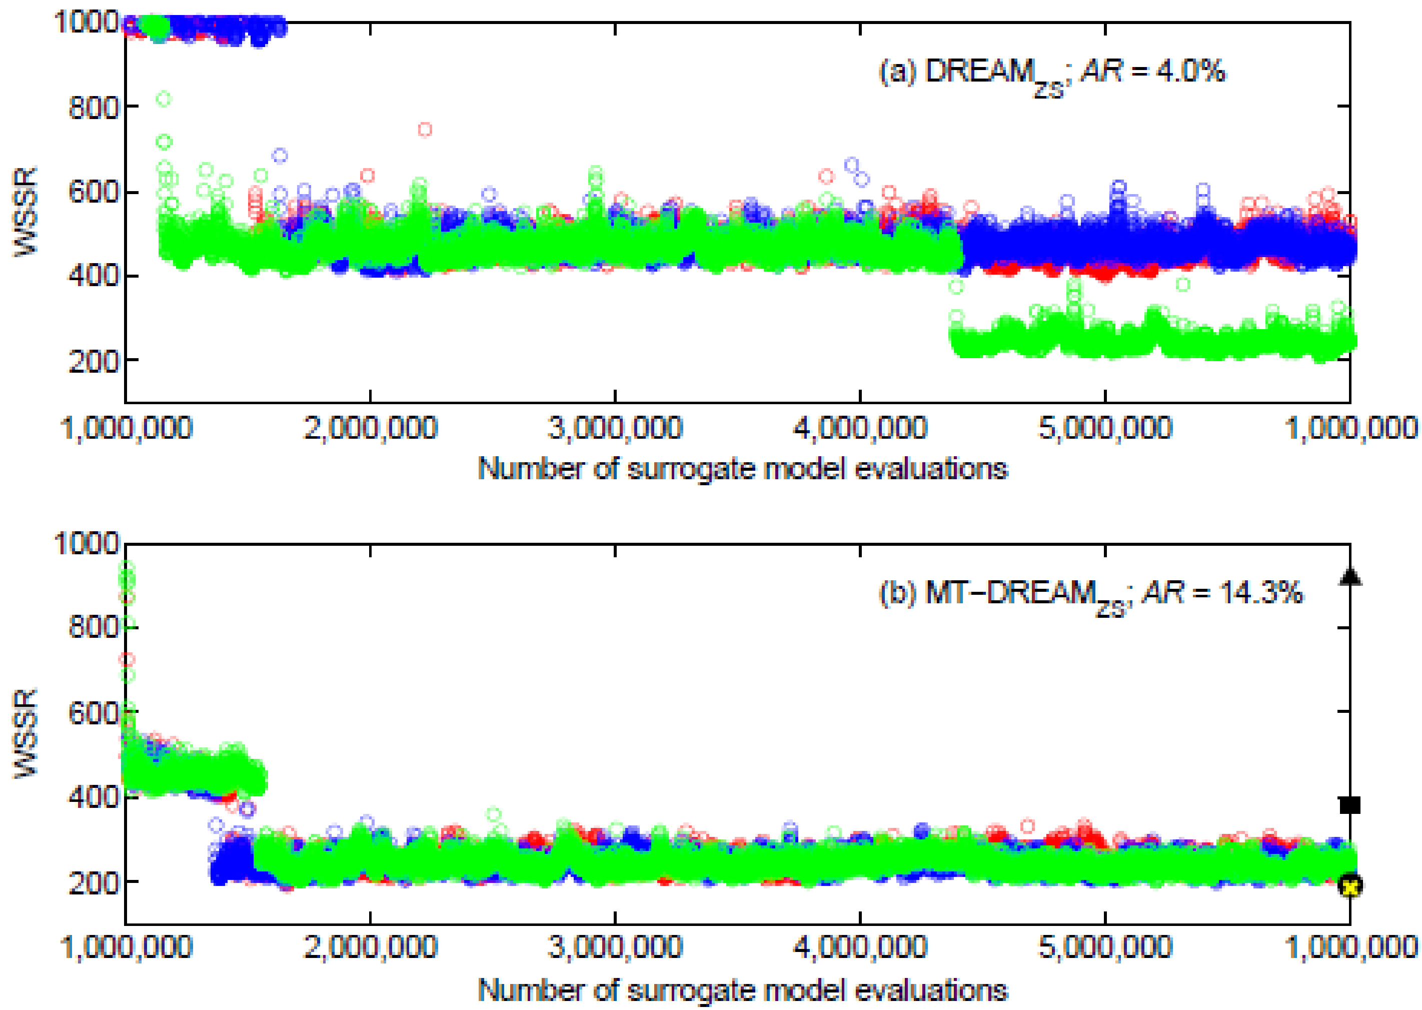

Laloy and Vrugt, *WRR* (2011)

SOFTWARE: JASPER.ENG.UCI.EDU

Software

<< Home | MATLAB Packages | R Packages | Python Packages | C Packages | Other Languages | Parallel Computation | Forum »

MATLAB Packages

DREAM: Differential Evolution Adaptive Metropolis (DREAM) Markov Chain Monte Carlo (MCMC) sampling of the posterior probability density function. This code runs multiple different chains simultaneously for global exploration, and automatically tunes the scale and orientation of the proposal distribution using differential evolution. The algorithm maintains detailed balance and ergodicity and is generally superior to other adaptive MCMC sampling approaches, especially in the presence of high-dimensionality and multimodality. This algorithm is a follow up on the SCEM-UA global optimization algorithm (which can be obtained upon request) and is especially designed to take full advantage of the power of distributed computer networks.

DREAM$_{(ZS)}$: Differential Evolution Adaptive Metropolis (DREAM) Markov Chain Monte Carlo (MCMC) of the posterior probability density function. DREAM$_{(ZS)}$ is based on the original DREAM algorithm, but uses sampling from an archive of past states to generate candidate points in each individual chain. Sampling from the past circumvents the need for a large number of parallel chains, designed to accelerate convergence for high-

Request Software Packages

It's easy – just fill out this form :)

What software packages are you interested in?
Please choose your package(s) below.

MATLAB Packages *

- DREAM
- DREAM_(ZS)
- DREAM_(D)
- AMALGAM
- AMALGAM-SO
- SODA
- BMA
- None – I want something else

Other Packages

- R Packages
- Python Packages
- C Packages
- Parallel Computation
- Other Languages

About
How are you planning on using this software?

How did you hear about this software? *

WHAT AM I PLANNING TO DO?

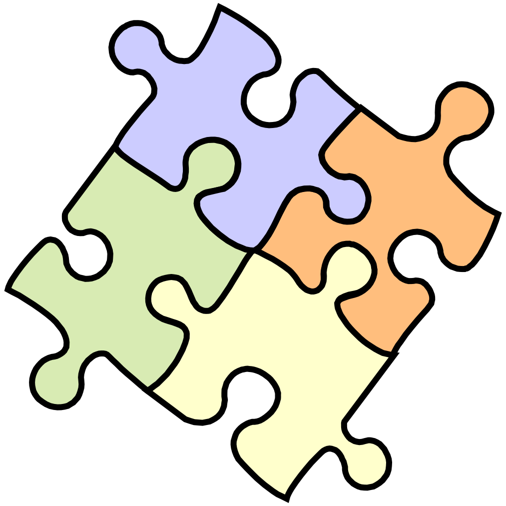

LEAST SQUARES REVISITED

$$\min_{\theta \epsilon \Theta} F(\theta) = \sum_{t=1}^{n} e_t(\theta)^2$$

IS THIS REALLY THE WAY TO GO?

ARTIFICIALLY INTRODUCES EQUIFINALITY

TREMENDOUS INFORMATION LOSS OF DATA

COMPLICATES FINDING BEST SOLUTIONS

NEW RULES – AFTER BILL MAHER

- TRY TO BE THE FIRST TO PUBLISH ON A NEW SUBJECT – LESS IMPORTANT TO BE RIGHT
- DON'T READ TOO MUCH LITERATURE; IT WILL KILL YOUR CREATIVITY AND ABILITY TO THINK OUTSIDE THE BOX
- A GOOD PROFESSOR CAN EXPLAIN A CONCEPT TO A 15 YEAR OLD KID. I CAN'T.
- ALWAYS QUESTION CONVENTIONAL WISDOM. DO NOT TAKE PROFESSORS TOO SERIOUSLY
- HIRING IS EASY, BUT HOME GROWN EXCELLENCE IS THE ULTIMATE STAMP OF A GOOD PROGRAM
- OPEN SOURCE DATA, MODELS, CODES. MANY PEOPLE BENEFIT FROM THIS