Online Process Optimization
Grand Challenges and Opportunities

Manfred Morari

Workshop 08 Nov. 2019

NTNU Trondheim

A Practitioner's Perspective

- Chemical Process Control
  - Shell, BP, Exxon, DuPont, ICI PLC
- Building Climate/Energy Control (HVAC)
  - Siemens, Carrier
- Automotive Systems
  - Ford, Daimler-Chrysler
- Aircraft Systems
  - United Technologies
- Power Electronics, Electrical Power Systems
  - ABB

Chemical Engineering Science Vol. 36, pp. 653–660
Pergamon Press Ltd., 1981. Printed in Great Britain

0009–2509/81/0401–0653/$02.00/0

MODELLING OF FLUIDIZED BED REACTORS—VI(a)

AN ISOTHERMAL BED WITH STOCHASTIC BUBBLES

JOHN ROBERT LIGON
Department of Chemical Engineering and Materials Science, University of Minnesota, 151 Amundson Hall, 421 Washington Avenue S.E., Minneapolis, MN 55455, U.S.A.

and

NEAL R. AMUNDSON*
Department of Chemical Engineering, University of Houston, Houston, TX 77004, U.S.A.

(*Received* 30 *April* 1980; *accepted* 10 *July* 1980)

Abstract—A two-phase stochastic isothermal fluidized bed reactor model with first order reaction in the dense phase is developed to investigate the significance of the fluctuating nature of fluidized beds on reactor performance. Several stochastic processes are employed as the overall mass transfer coefficient between phases. Analytical moment solutions are obtained for white noise coefficients while hybrid computer simulation was used for correlated stochastic coefficients. Results indicate that a gamma distributed coefficient is preferred over white noise and Gaussian correlated coefficients. When compared with the deterministic model, randomness in the mass transfer coefficient is seen to lead to a decrease in reactor performance. Deviation from the deterministic model increases with increasing variance and decreasing fluctuation frequency of the correlated stochastic coefficients.

School of
Mathematics

(1929-2014)

Charlotte Striebel

Charlotte Striebel passed away on Wednesday March 12, 2014 at the age of 84. She was born on July 30, 1929 in Columbus, OH and attended Ohio State University as an undergraduate, where she obtained an M.A. degree in 1952. She went on to University of California Berkeley for graduate school and got her Ph.D. in 1960, working on stochastic processes. This remained the main area of her published research throughout her career. Between 1958 and 1964 she worked at Lockheed Missiles and Space Company in Sunnyvale, CA where she developed the initial workings of the GPS system in wide use today. She also analyzed satellite tracking data as well as statistics associated with Polaris missiles and the recovery of manned space capsules from the Gemini space program. After some

Condensed Directory (/directory/condensed-directory)

Administration (/about/directory?keys=&field_role_type_value=administration)

Faculty

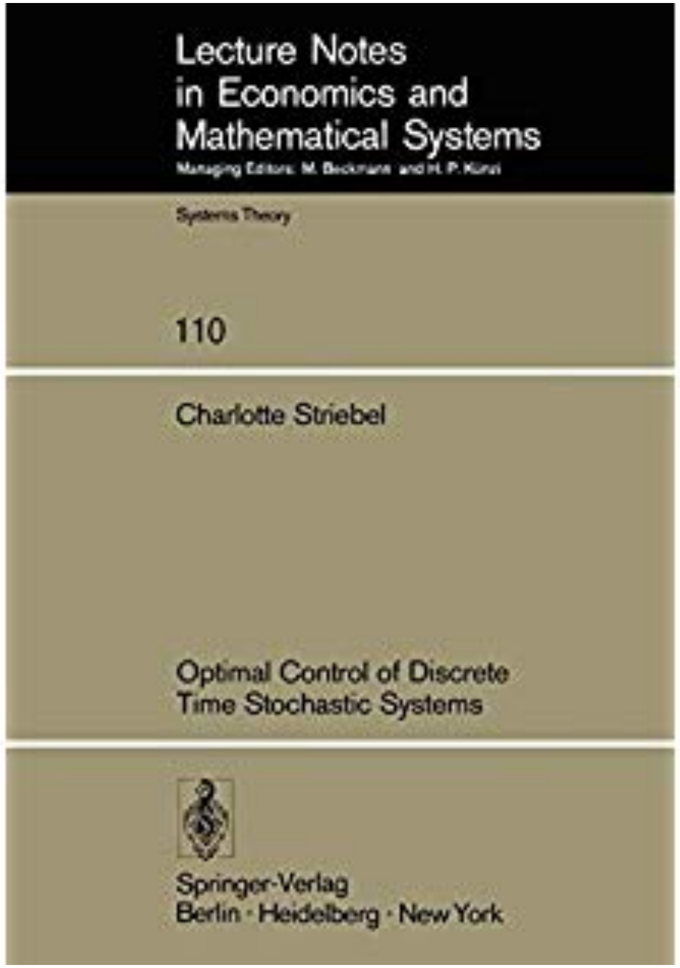

Proceedings of the
47th IEEE Conference on Decision and Control
Cancun, Mexico, Dec. 9-11, 2008

A Tractable Approximation of Chance Constrained Stochastic MPC based on Affine Disturbance Feedback

Frauke Oldewurtel, Colin N. Jones, Manfred Morari

Abstract— **This paper deals with model predictive control of uncertain linear discrete-time systems with polytopic constraints on the input and chance constraints on the states. When having polytopic constraints and bounded disturbances, the robust problem with an open-loop prediction formulation is known to be conservative. Recently, a tractable closed-loop prediction formulation was introduced, which can reduce the conservatism of the robust problem. We show that in the presence of chance constraints and stochastic disturbances, this closed-loop formulation can be used together with a tractable approximation of the chance constraints to further increase the performance while satisfying the chance constraints with the predefined probability.**

D. Bertsimas, M. Sim, "Tractable Approximations to Robust Conic Optimization Problems", *Math. Program.*, ser. B 107, 2006, pp. 5-36.

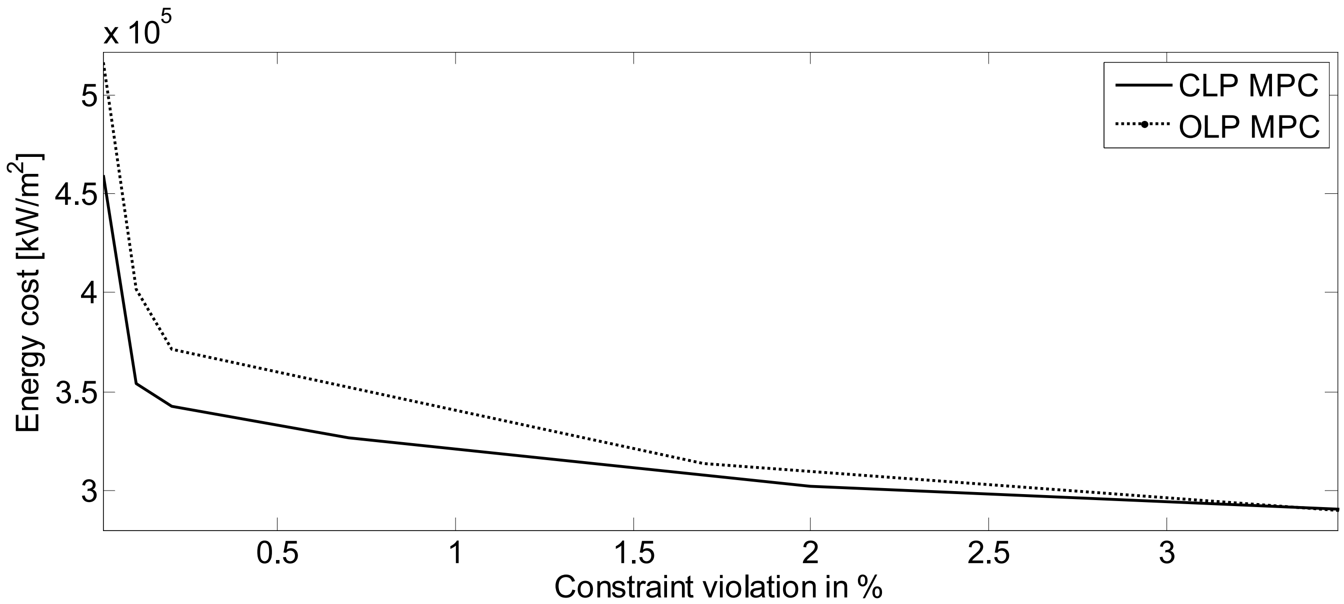

Fig. 3. **Tradeoff curve for energy consumption and constraint violation.** The curve depicts the tradeoff between a low energy consumption and a high degree of constraint satisfaction.

Idea: Get rid of “Modeling” …
…in Model-Based-Design

- Kalman (1958): Design of a self-optimizing control system. Trans. ASME
- Bellman (1961): Adaptive Control Processes
- Åström & Wittenmark (1973): On Self-Tuning Regulators. Automatica
- Landau (1974): A survey of model reference adaptive techniques, Automatica
- Narendra & Valavani (1976): Stable adaptive observers and controllers. Proc IEEE
- Åström, Borisson, Ljung, Wittenmark (1977): Theory and applications of self-tuning regulators. Automatica
- …..

ASEA Novatune introduced in 1982…

…and mostly abandoned by 1995

"Even if Novatune in many cases provides very good control, the experience is that the effort it takes, to make it work that well, is discouraging. ***It is worth the effort in some cases, but not as a general tool. What is needed is a tool that is much easier to use.*** You shouldn't be required to set any parameters, except to state what kind of result you desire."

Per Erik Maden (1995) Experiences with Adaptive Control since 1982. CDC Proc.

Embracing the Machine Learning and Artificial Intelligence Contributions

MIT IDSS LIDS SOCIOTECHNICAL SYSTEMS RESEARCH CENTER

Learning for Dynamics and Control (L4DC)

May 30 & 31, 2019 at the Ray and Maria Stata Center
Massachusetts Institute of Technology, Cambridge, MA

Over the next decade, the biggest generator of data is expected to be devices which sense and control the physical world.

This explosion of real-time data that is emerging from the physical world requires a rapprochement of areas such as machine learning, control theory, and optimization. While control theory has been firmly rooted in tradition of model-based design, the availability and scale of data (both temporal and spatial) will require rethinking of the foundations of our discipline. From a machine learning perspective, one of the main challenges going forward is to go beyond pattern recognition and address problems in data driven control and optimization of dynamical processes.

Why did Adaptive Control “fail”?
--- It was not appropriate

Anything works

- no specs
- model simple
- general solution

Learning Control

Nothing works

- tight specs
- model complex
- specific solution

Why did Adaptive Control "fail"?
--- tuning all the time not needed

- Åström & Hägglund (2000). Supervision of adaptive control algorithms
- PID Autotuner: Tune on demand only

Commercial Autotuners

- One-button autotuning
- Three settings: fast, slow, delay dominated
- Automatic generation of gain schedules
- Adaptation of feedback gains
- Adaptation of feedforward gain
- Many versions
  - Single loop controllers
  - DCS systems
- Robust
- Excellent industrial experience
- Large numbers

Thanks: Karl Aström

Learning Controllers

- Model-based vs. model-free
  - If you do not have a model, how can you verify the performance of the closed-loop control system?
  - If you do have a model, why would you use a model-free learning method?
- Policy learning based on reward function
  - Curse of dimensionality
  - Specification guarantees via definition of reward function

Learning Controllers

- Model-based vs. model-free
  - If you do not have a model, how can you verify the performance of the closed-loop control system?
  - If you do have a model, why would you use a model-free learning method?
- Policy learning based on reward function
  - Curse of dimensionality
  - Specification guarantees via definition of reward function

Design ≠ Optimization
Design ≈ Constraint Satisfaction

Propositional Logic Control Specifications for Refrigeration Cycle

Manipulated Inputs $u_{i,min} \leq u_i \leq u_{i,max}$ i=1,2,3

Controlled / Monitored Outputs $y_i = y_{i,set}$ i=1,2,3 / $z_i = z_{i,set}$ i=1,..4

Prioritized Objectives

$z_1 > z_{1,min}$	$z_3 < z_{3,max}$	$z_3 < z_{3,max}$	$y_3 = y_{3,set}$
$z_2 < z_{2,max}$	$z_4 < z_{4,max}$	$z_4 < z_{4,max}$	$y_2 = y_{2,set}$
$y_1 = y_{1,set}$	$z_4 > z_{4,min}$	$z_2 > z_{2,min}$	
	$y_2 = y_{2,set}$	$y_3 = y_{3,set}$	

Specification guarantees via definition of reward function?

GARY MARCUS BUSINESS 08.14.19 09:00 AM

WIRED

DEEPMIND'S LOSSES AND THE FUTURE OF ARTIFICIAL INTELLIGENCE

"…In some ways, deep reinforcement learning is a kind of turbocharged memorization; systems that use it are capable of awesome feats, but they have only a shallow understanding of what they are doing. As a consequence, current systems lack flexibility, and thus are unable to compensate if the world changes, sometimes even in tiny ways."

IAS
2019-10-16

Energy-Based Approaches To Representation Learning

Yann LeCun
New York University
Facebook AI Research
http://yann.lecun.com

facebook
Artificial Intelligence Research

Supervised Learning works but requires many labeled samples

- **Training a machine by showing examples instead of programming it**
- **When the output is wrong, tweak the parameters of the machine**
- **Works well for:**
  - Speech→words
  - Image→categories
  - Portrait→ name
  - Photo→caption
  - Text→topic
  - ….

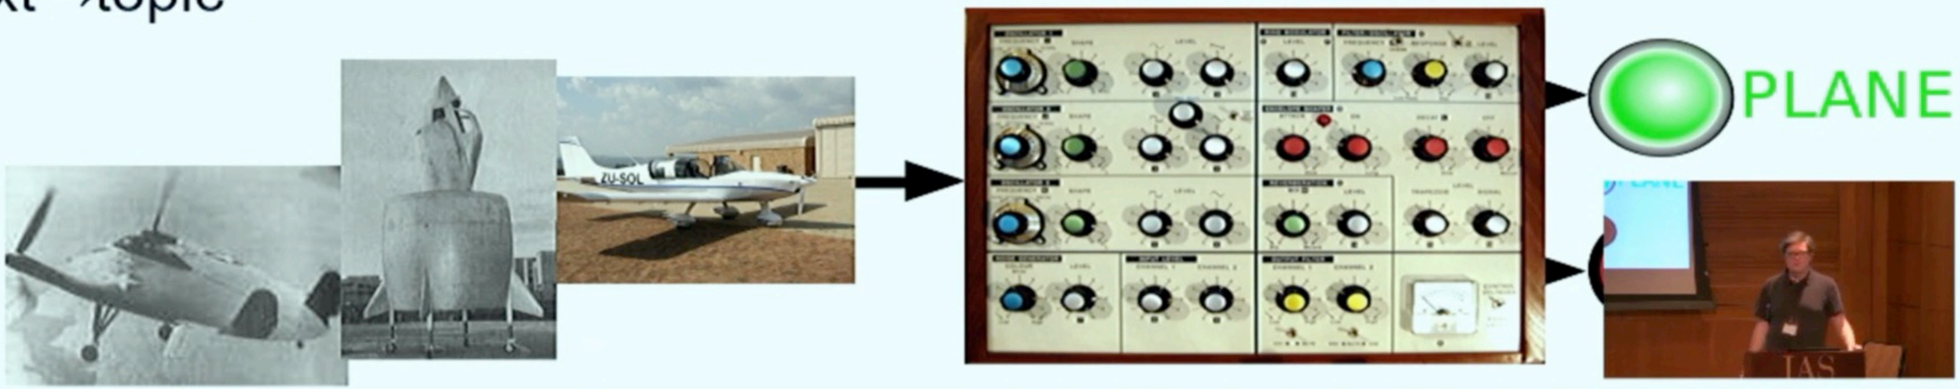

Y. LeCun

Reinforcement Learning: works great for games and simulations.

- **57 Atari games: takes 83 hours equivalent real-time (18 million frames) to reach a performance that humans reach in 15 minutes of play.**
  - [Hessel ArXiv:1710.02298]
- **Elf OpenGo v2: 20 million self-play games. (2000 GPU for 14 days)**
  - [Tian arXiv:1902.04522]
- **StarCraft: AlphaStar 200 years of equivalent real-time play**
  - [Vinyals blog post 2019]
- **OpenAI single-handed Rubik's cube**
  - 10,000 years of simulation

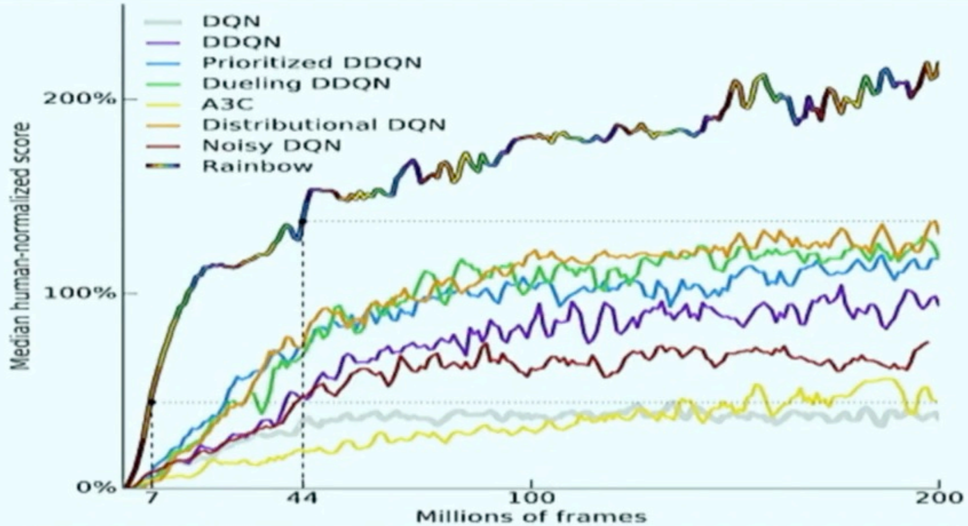

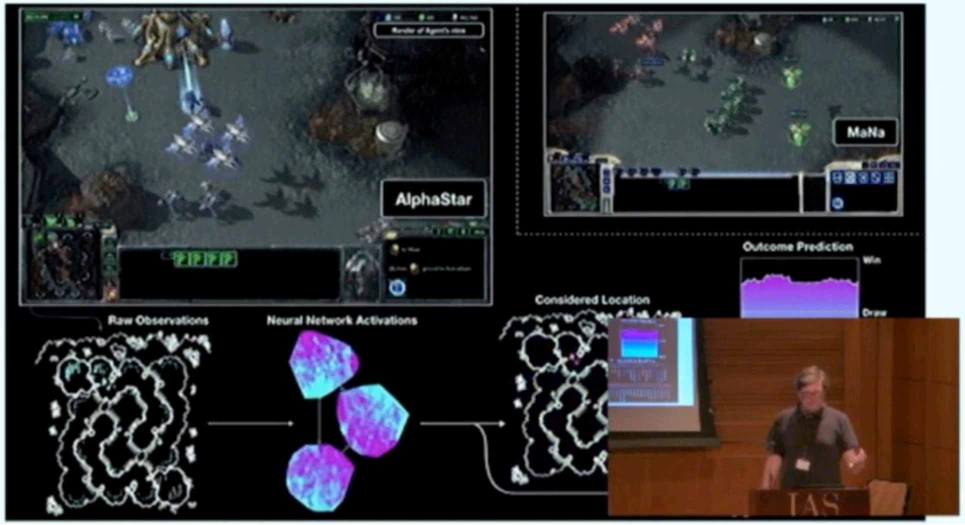

Y. LeCun

But RL Requires too many trials in the real world

- **Pure RL requires too many trials to learn anything**
  - it's OK in a game
  - it's not OK in the real world
- **RL works in simple virtual world that you can run faster than real-time on many machines in parallel.**

- **Anything you do in the real world can kill you**
- **You can't run the real world faster than real time**

Some Research Directions

- MPC Approximation via Neural Networks
- Robustness Analysis of Learning Enabled Components
- Gaussian-Process based Model Predictive Control

Example: Oscillating Masses

18 oscillating masses [1]
State dim: 36
Action dim: 9
Horizon: 50

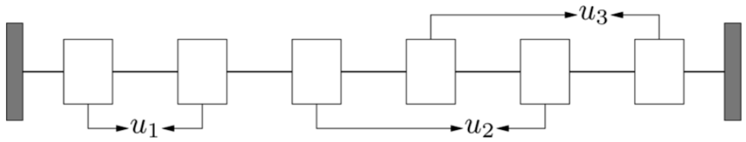

6 mass version

Training Parameter	Value
Training Set Size	2,500,000
Testing Set Size	250,000
# Training Epochs	200 Epochs (~40 hour)
Neural Network Depth	7 layers
Neural Network Hidden Width	128-512
# Neural Network Parameters	1,668,554

[1] Y. Wang and S. Boyd. Fast Model Predictive Control Using Online Optimization. IEEE Transactions on Control Systems Technology. Vol 18, No. 2, 2010.

DARPA Project: Assured Autonomy Unmanned Underwater Vehicle (UUV)

- Sonar data
- NN to locate pipeline on sea floor
- Steering control loop

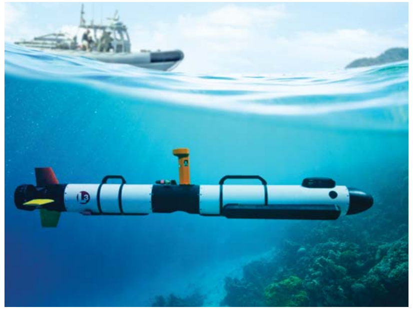

Kothare, Morari, Automatica (1999)

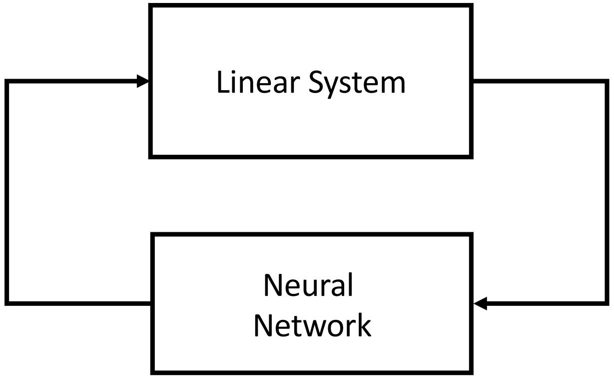

General Interconnection of Linear System and Quadratically-Constrained Nonlinearity