The Promise of Artificial Intelligence in Process Systems Engineering:
Is it here, finally?

Venkat Venkatasubramanian

Samuel Ruben-Peter G. Viele Professor of Engineering

Center for the Management of Systemic Risk

Columbia University

New York, NY 10027

2040 Visions of Process Systems Engineering
Symposium on the Occasion of George Stephanopoulos's
70th Birthday and Retirement from MIT, June 1-2, 2017

150 COLUMBIA | ENGINEERING
The Fu Foundation School of Engineering and Applied Science
1864–2014

CRIS.lab

Talk Philosophy

- **Objectives**
  - Review AI in PSE: 1980s to Present
  - Potential of AI in PSE: Present – 2040?
  - Identify the challenges: Intellectual, Implementational, Organizational
- **Broad overview**
  - Not a detailed, in-depth technical presentation
  - More details in these papers

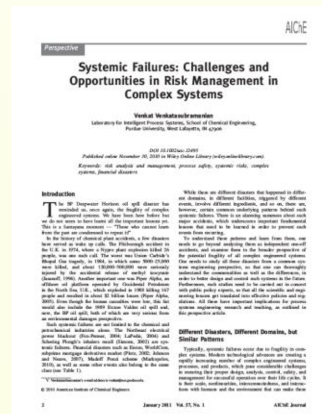

V. Venkatasubramanian, "Systemic Failures: Challenges and Opportunities for Risk Management in Complex Systems", Perspective Article, AIChE Journal, Jan 2011.

V. Venkatasubramanian, "Drowning in Data: Informatics and Modelling Challenges in a Data-Rich Networked World", Perspective Article, AIChE Journal, Jan 2009.

150 COLUMBIA | ENGINEERING
The Fu Foundation School of Engineering and Applied Science
1864–2014

CRIS.lab

Branches of AI

- Games - study of state space search, e.g., Chess, GO
- Automated reasoning and theorem proving, e.g., Logic Theorist
- Robotics and planning – e.g., driverless cars
- Vision – e.g., facial recognition
- Natural language understanding and semantic modeling, e.g. Siri
- **Expert Systems or Knowledge-based systems**
- **Machine Learning** – e.g., Bayesian classifiers, Deep neural nets
- Automatic programming
- Hardware for AI
- Distributed & Self-organizing AI – e.g., Drone swarms
- Artificial Life – e.g., cellular automata, agent-based modeling

Promise of AI in PSE

- In essence, AI is about problem-solving and decision-making under complex conditions
  - Ill-posed problems
  - Model and data uncertainties
  - Combinatorial search spaces
  - Nonlinearity and multiple local optima
  - Noisy data
  - Fast decisions are required – e.g., fight or flight responses
- But these are applicable to many PSE problems: Design, Control, Optimization
- So some of us went about developing AI approaches in the mid-80s
  - Davis, Kramer, Stephanopoulos, Ungar, Venkatasubramanian and Westerberg
- We expected significant impact from AI, much like Optimization and MPC
- But it did not happen – Why not?

AI in PSE:
Why very little impact?

Before I answer this question, let me first review the different phases of AI in PSE

AI in PSE: Different Phases

- **Phase I: Expert Systems in PSE (1983 – 1995)**
  - Davis, Kramer, Stephanopoulos, Ungar, Venkatasubramanian and Westerberg
  - CONPHYDE (1983), DECADE (1985), MODEX (1986), DESIGN-KIT (1987), MODEL.LA (1990), …
  - LISPE Consortium founded at MIT (1985)
  - First course on AI in PSE developed at Columbia (1986)

Fall 1986

CEE chemical engineering education

VOLUME XX • NUMBER 4 • FALL 1986

CHEMICAL ENGINEERING DIVISION OF AMERICAN SOCIETY FOR ENGINEERING EDUCATION

GRADUATE EDUCATION ISSUE

A Research Program On
ARTIFICIAL INTELLIGENCE IN PROCESS ENGINEERING

GEORGE STEPHANOPOULOS
Massachusetts Institute of Technology
Cambridge, MA 02139

THE REEMERGENCE of artificial intelligence as a viable and utilitarian discipline offers the potential of harvesting early promises on intelligent man-machine interaction. For process engineering, these promises have nurtured and disillusioned a generation of engineers. Presently, the mood is cautiously optimistic. The "novelty" of the technology has taken most by surprise and has found the large majority, even among the early devotees in artificial intelligence, unprepared for meaningful engineering applications. Nevertheless, idling skepticism has been replaced by a wide-spread activism, leading to a multitude of exploratory prototypes. But, what we observe as a feverish research and development activity in knowledge-based expert system is nothing else but a very serious effort in self-education. We will have to wait for the next phase of developments to see useful and practical products for process systems engineering purposes.

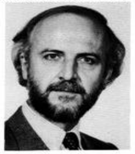

George Stephanopoulos was an undergraduate at the National Technical University of Athens, Greece, received his ME at McMaster University, Canada, and did his doctoral studies in chemical engineering at the University of Florida. In 1974 he joined the faculty at the University of Minnesota, and fr[...]
National Technical University of A[...]
he is presently the J.R. Mares prof[...]
the author of two books: *Chemica[...]*
Theory and Practice, and *Synthes[...]*
He has been a Dreyfus Scholar [...]
Award (1982) of AIChE and the C[...]
ASEE. His research interests are in [...]
ing, which he and his students [...]
methodologies from artificial int[...]
computers.

Existing prototypes of expert systems are interesting examples, and some of them have had significant economic impact in areas other than those related to chemical and biochemical engineering. They have provided certain paradigms which later efforts have tried to imitate. But, are these prototypes appropriate for process engineering?

- **Can they "model" the human activity during the conception of a chemical process, the design of a product, the development of a process flowsheet, or the synthesis of control configurations and operating procedures for complete processing plants?**
- **Can they support engineering activities, capitalizing on the innate "intelligence" of expert technologists and designers, as this intelligence is articulated within the context of the problem being solved?**
- **Do they provide high level, transparent communication between man and machine during the graphic generation of process flowsheets, or control configurations, or the analytic development of process models, the introduction of qualitative reasoning, or the formulation of design problems (assumptions, assertions, hypothesis testing, etc.)?**

It is our view that the existing paradigms cannot satisfy the above needs; after all, they were conceived to solve different problems. New prototypes are needed which should reflect the particularities of the process systems engineering problems.

THE M.I.T.-LISPE

The Laboratory for Intelligent Systems for Pro-

A course in
ARTIFICIAL INTELLIGENCE IN PROCESS ENGINEERING
Experiences From a Graduate Course

V. VENKATASUBRAMANIAN
Columbia University
New York, NY 10027

OVER THE RECENT past, notable advances have been made in the field of artificial intelligence (AI) that are poised to make important contributions to various engineering disciplines [2]. Chemical engineering, process engineering in particular, stands to make significant gains by the application of the AI methodology called "Knowledge-Based Expert Systems" (KBES). Briefly, AI is the study of understanding human information processing with the aid of computers and computational models. KBES is the first attempt towards this goal by concentrating on narrow, restricted domains of knowledge (such as those of experts), rather than tackling the entire spectrum of human intelligence. Such an attempt has resulted in some progress towards the understanding of the different facets of human cognition [13]. In this paper we discuss the organization and content of a new course that has been specifically designed for chemical engineers on the application of KBES methodology in process engineering.

MOTIVATION

It is becoming increasingly clear that areas such as process synthesis and design, process diagnosis and safety, intelligent computer-aided instruction and training, etc., will derive substantial benefits by integrating the KBES methodology into the existing predominantly algorithmic approaches. We are then faced with the question of how to go about doing this.

The current approach used in the application of the KBES methodology is the so-called *dialogue* approach, where one or more computer scientists trained in AI (called the "knowledge engineers") interact with one or more chemical engineers (called the "domain experts"), and together they develop the knowledge-based system for the given problem. This approach has the drawback that the knowledge engineer spends a considerable amount of time and effort in learning the problem domain (say, a given problem in process synthesis or diagnosis) in order to be able to design an appropriate system. Similarly, the domain engineer spends considerable time and effort in conveying the domain knowledge to the knowledge engineer as well as learning something about AI and KBES. It seems that a better approach would be to train chemical engineers in AI, let them develop the appropriate knowledge-based systems for their problems, and let the computer science expert (knowledge engineer) be involved only as an occasional consultant for some difficult AI related problems which are beyond the scope of our artificially intelligent chemical engineer. Such an approach is, in fact, similar in spirit to what chemical engineers have been doing for a long

In this paper we discuss the organization and content of a new course that has been specifically designed for chemical engineers on the application of KBES methodology in process engineering.

Venkat Venkatasubramanian is an assistant professor in chemical engineering at Columbia University. After receiving his doctoral degree from Cornell University in 1983, he worked as a research associate in the Department of Computer Science at Carnegie-Mellon University. At Columbia he is directing the research efforts in the "Intelligent Process Engineering Laboratory," and is currently working on developing knowledge-based expert systems for process diagnosis, design, and training.

First AI in PSE Meeting
Columbia University, March 1987

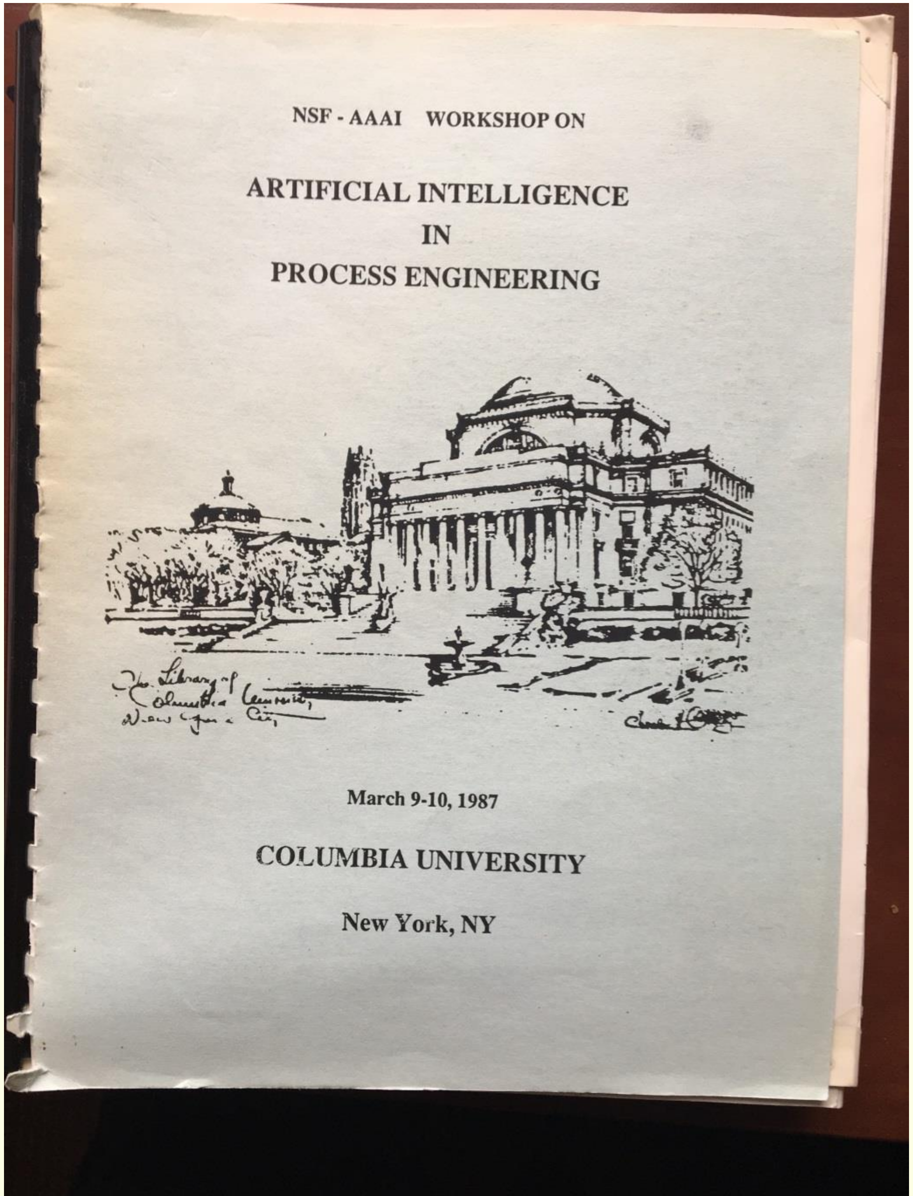
NSF - AAAI WORKSHOP ON

ARTIFICIAL INTELLIGENCE
IN
PROCESS ENGINEERING

March 9-10, 1987

COLUMBIA UNIVERSITY

New York, NY

Porto Carras, Greece, June 20-24, 1988

Stephanopoulos, Ungar, and Venkatasubramanian

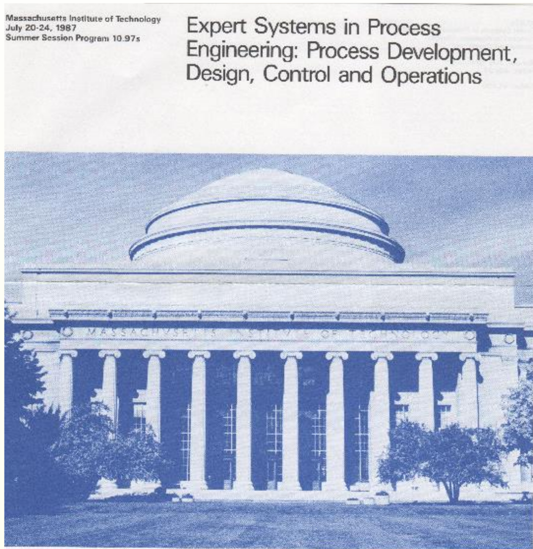

Same location for ESCAPE 21, June 2011

AI in PSE: Phase II

- **Phase II: Machine Learning I - Neural Networks (1990 – 2005)**
- **Backpropagation algorithm:** Rumelhart, Hinton and Williams (1986)
- Whitley and Davis (1993, 1994)
- Hoskins and Himmelblau (1988); Matsuura, Abe, Kubota, Himmelblau (1989)
- Kramer (1991); Leonard and Kramer (1991, 1992, 1993)
- Bhat and McAvoy (1990); Qin and McAvoy (1992)
- Bakshi and Stephanopoulos (1992, 1993)
- Ungar, Powell, and Kamens (1990); Psichogios and Ungar (1991, 1992)
- Venkatasubramanian (1985); Venkatasubramanian and Chan (1989); Kavuri and Venkatasubramanian (1993, 1994)
- Also progress in **Expert Systems** and **Genetic Algorithmic** methods
- **Most work was on process control and fault diagnosis**

ASM (1995 – 2000)

Ohio State
(Davis)

Purdue
(Venkatasubramanian)

University of Toronto
(Kim Vicente)

NIST AEGIS Program

USD $17,000,000 (49% matched back)

Fore-runner to the **Smart Manufacturing Initiative** (2016)

Diagnostic ToolKit (Dkit): 1995-2000

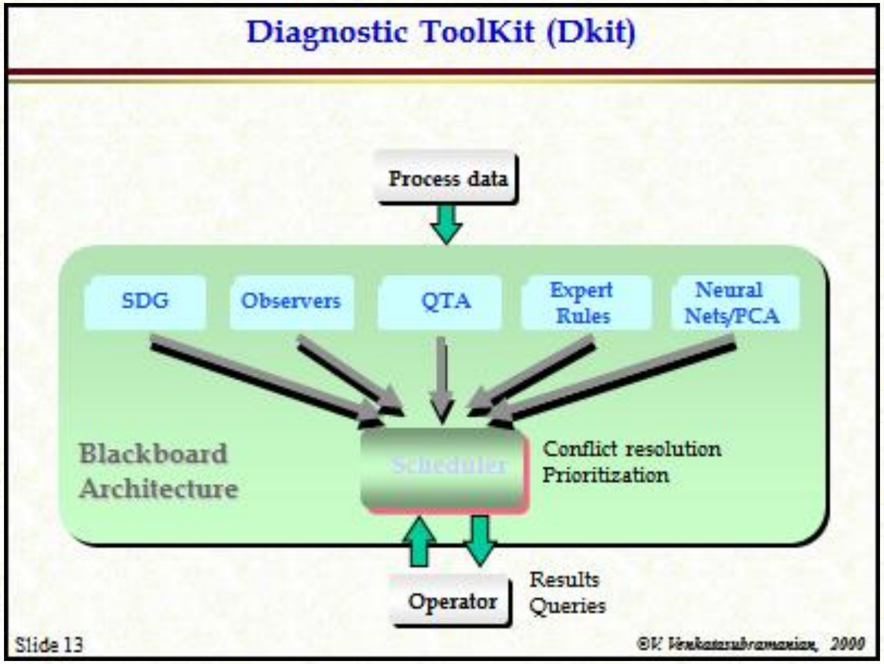

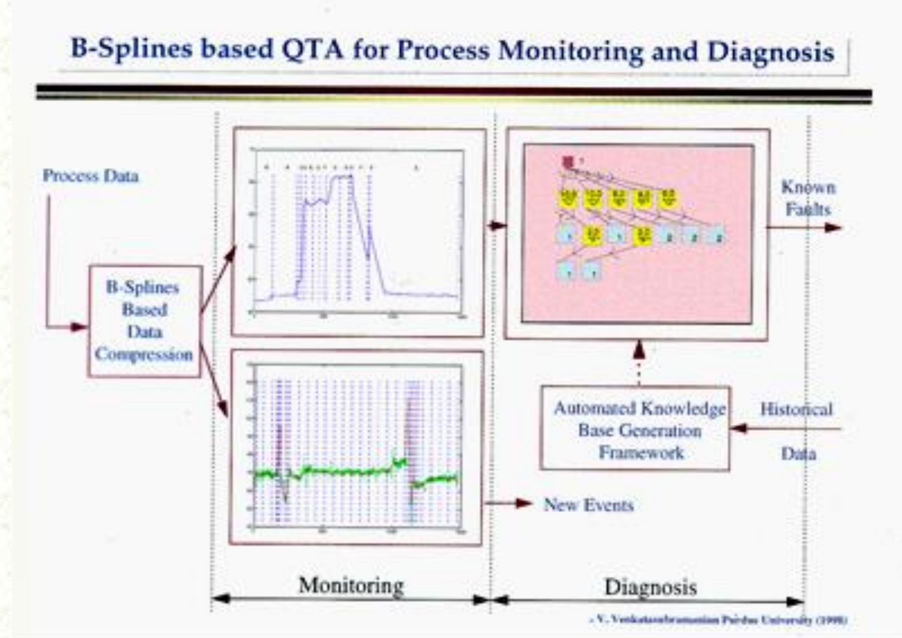

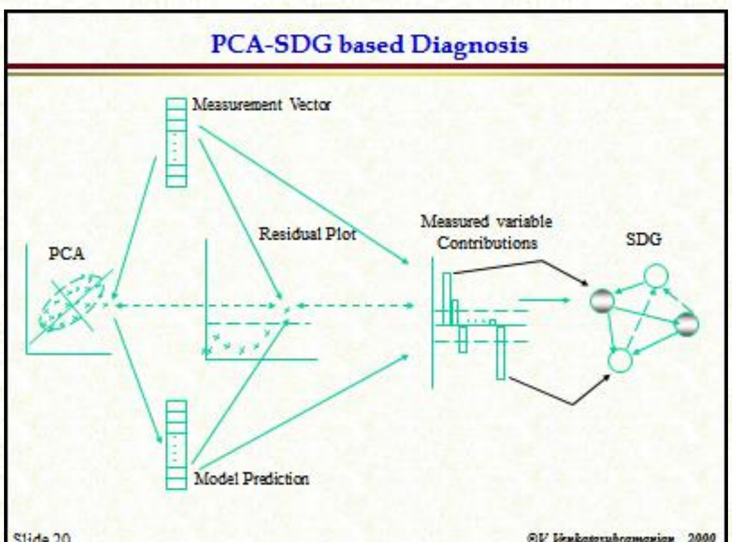

- Dkit successfully anticipated and diagnosed several failures even before the alarms went off (~1/2 – 2 hours ahead)
- Implemented in G2, tested at Exxon's BRCP
- Dkit was licensed to Honeywell in 1998
- **We were about 20-30 years too early to tackle this problem!**

So, why wasn't AI in PSE NOT impactful in Industry during (1985- 2015)?

- For the same reasons it was not impactful in other domains
  - Lack of computational power and computational storage
  - Lack of communication infrastructure – NO Internet, Wireless
  - Lack of convenient software environment
  - Lack of specialized hardware – e.g., NVIDIA GPU for simulations
  - Lack of data
  - Lack of acceptance of computer generated advice
  - Costs were prohibitive
- NO technology PUSH
- NO market PULL
  - Low-hanging fruits in optimization and control applications
  - No need to go after the more challenging AI applications
- Technology usually takes ~40-50 years to reach wide adoption – e.g., Aspen+, LP, MINLP, MPC, etc.

What is Different Now?

- **Cray-2 Supercomputer (1985)**
  - 1.9 GFLOPS
  - 244 MHz
  - 150 kW!
  - $32 Million! (2010 dollars)

Your text here

- **Apple Watch (2015)**
  - 3 GFLOPS
  - 1 GHz
  - 1 W!
  - $300!

- Performance/unit cost Gain ~150,000

So, what happened?

- **Basically Moore's Law happened over the last 30 years!**
- **All these metrics improved by orders of magnitude!**
  - Computational power
  - Computational storage
  - Communication infrastructure: Internet, Wireless
  - Convenient software infrastructure – Python, Java, OWL, …
  - Specialized hardware – graphics processors
  - Big Data
  - Trust & Acceptance – Google, Yelp, Trip Advisor, Tinder, …
- Technology **PUSH** is there now
- Market **PULL** is there now
  - Many low-hanging fruits in optimization and control applications have been picked in the last 30 years
  - Need to go after the more challenging tasks for further improvements
- There is **Great Convergence** now!

So, what happened?
Watson and AlphaGO

- **Deep Blue (IBM) vs Gary Kasparov**
  - May 11, 1997 – New York City
  - Score: 3.5 – 2.5
  - First computer program to defeat a world champion in a match under tournament regulations
- **Watson (IBM) wins *Jeopardy***
  - Feb 2011
  - Human Champs: Jennings (2nd) and Rutter (3rd)
- **AlphaGO (DeepMind) vs Lee Seedol**
  - Mar 2016
  - Score: 4-1
  - Deep Learning Neural Networks

Source: Wiki

AI in PSE:
Entered Phase III

- **Phase III: Machine Learning II - Data Science (2005 – Present)**
  - Deep Learning Neural Nets
  - Statistical Machine Learning
  - Reinforcement Learning
- Big impact on NLP, Robotics, Vision
  - Watson, AlphaGO, Self-driving cars

How about "Watson" for PSE?

- **What will it take to develop "Watson" for PSE?**
  - Not just qualitative facts
  - Quantitative
  - Math Models
  - Charts, Tables, Spectra
  - Heuristic Knowledge

"Watson" for Pharmaceutical Engineering (2005-2011)

User
Dissolution Problems

Modeling

- Multiscale models for synthesis mehtods
- Modeling and Analysis Crystalline Structure
- Unit Operation models
  DEM/CFD/FEM
- Process modeling
  First Principles Process Model

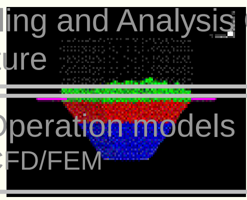

Ontological Informatics Infrastructure

- Knowledge relationships
- Process
  Crystallization, Granulation, Drying, Size Reduction etc.
- Equipment:
  Blender, tabletting press, roller compactor, fluid bed
- Organic composites:
  Physical, chemical and powder properties

Physical Properties

Powder Properties

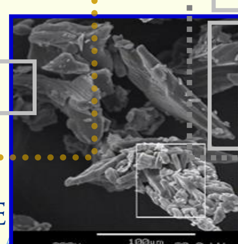

Material Science

- Material synthesis method development
- Systematic functionalization
- Control methodology for spatial structure of organic composites
- Determine structure-function-performance relationships

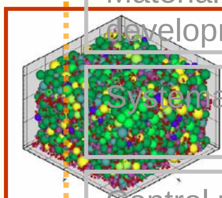

Spatial Structure

Tools for Design and Manufacturing

- Intelligent Systems for Preformulation
  Formulation
  Process Monitoring and Control
- Tools for Organic Solids Mfg.
  Optimization, Simulation, Safety Analysis, Scheduling

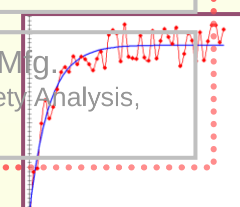

Experiments

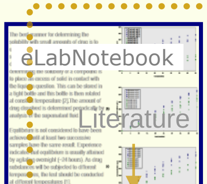

eLabNotebook

Literature

Experiments

Instruments

Experiment Report

Data

Venkatasubramanian, V., Zhao, C., Joglekar, G., Jain, A., Hailemariam, L., Sureshbabu, P., Akkisetti, P., Morris, K. and Reklaitis, G.V., "Ontological Informatics Infrastructure for Chemical Product Design and Process Development", *Comp. & Chem. Engg.*, 2006.

"Watson" for Pharmaceutical Engineering (2005-2011)

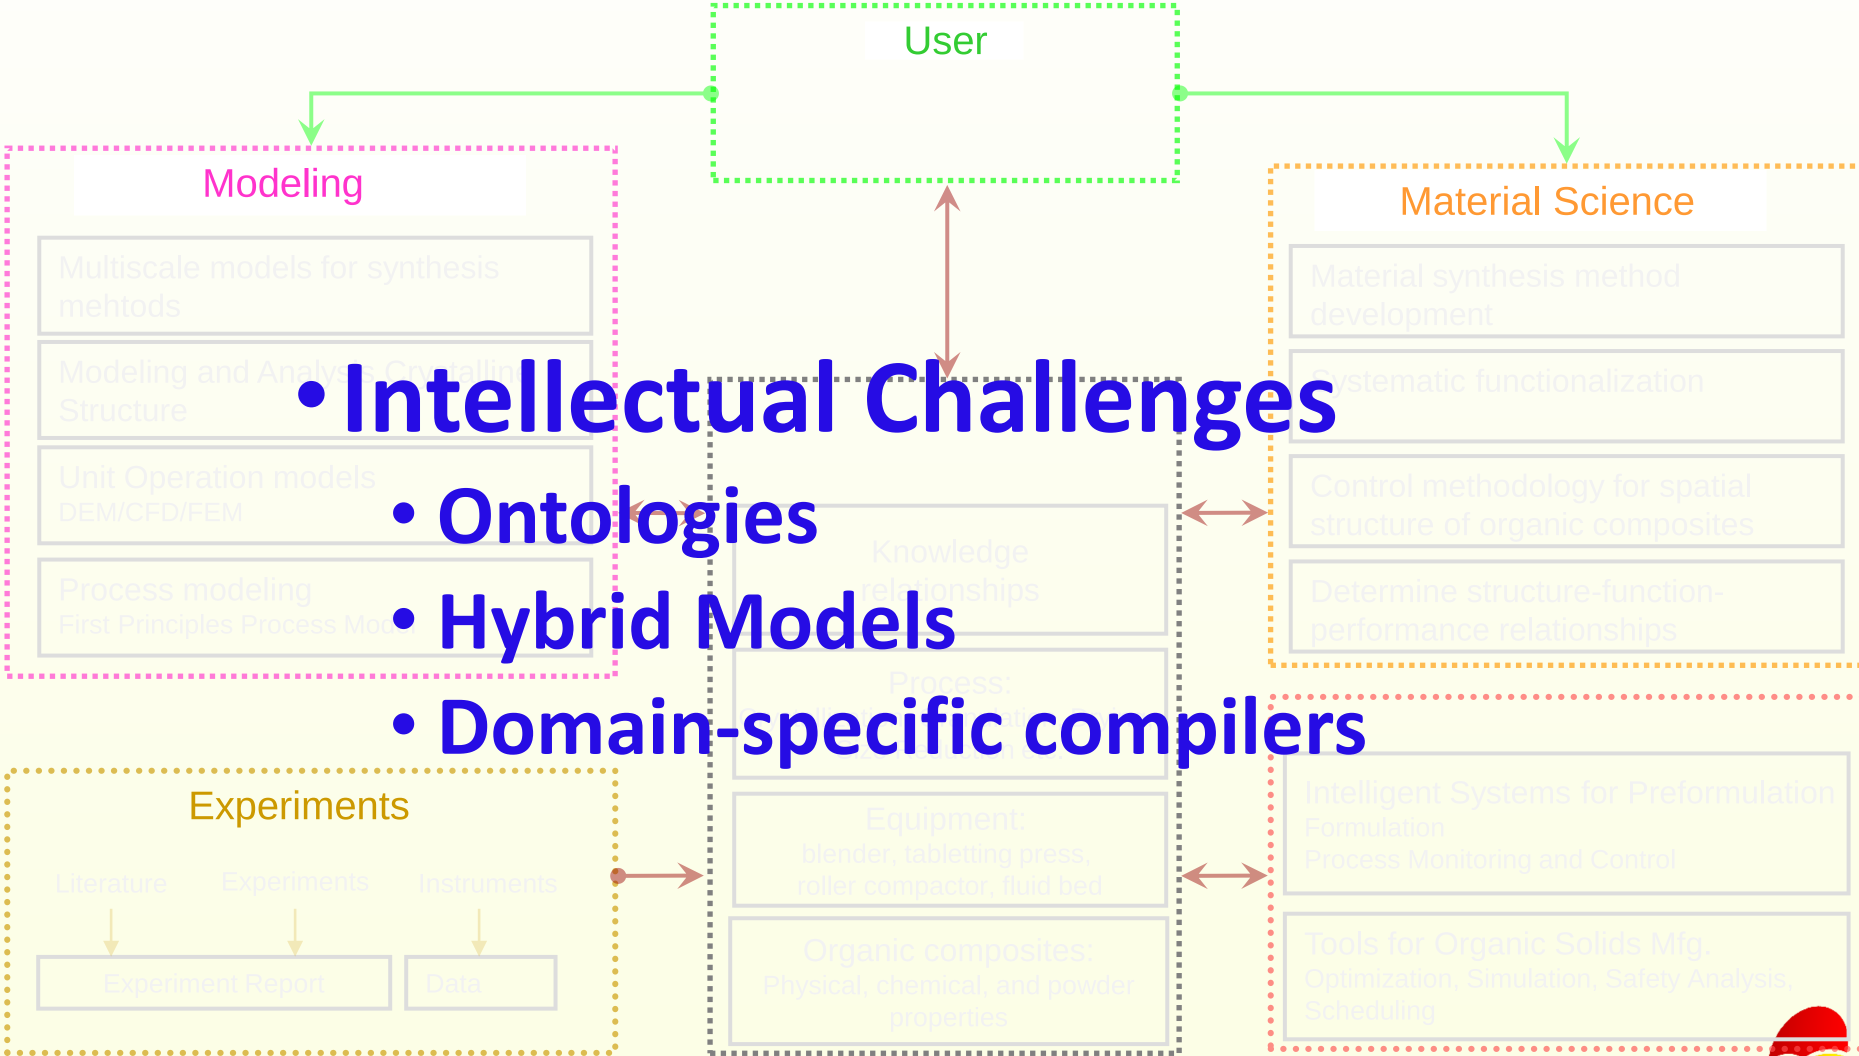

- **Intellectual Challenges**
  - **Ontologies**
  - **Hybrid Models**
  - **Domain-specific compilers**

Venkatasubramanian, V., Zhao, C., Joglekar, G., Jain, A., Hailemariam, L., Sureshbabu, P., Akkisetti, P., Morris, K. and Reklaitis, G.V., "Ontological Informatics Infrastructure for Chemical Product Design and Process Development", *Comp. & Chem. Engg.*, 2006.

HOLMES: Semantic Search Engine (2011-2017)

- HOLMES: Ontology-Learning Materials Engineering System
- Machine Learning
- anguage Processing (NLP)

Academic Journal Articles

Semantic Storage

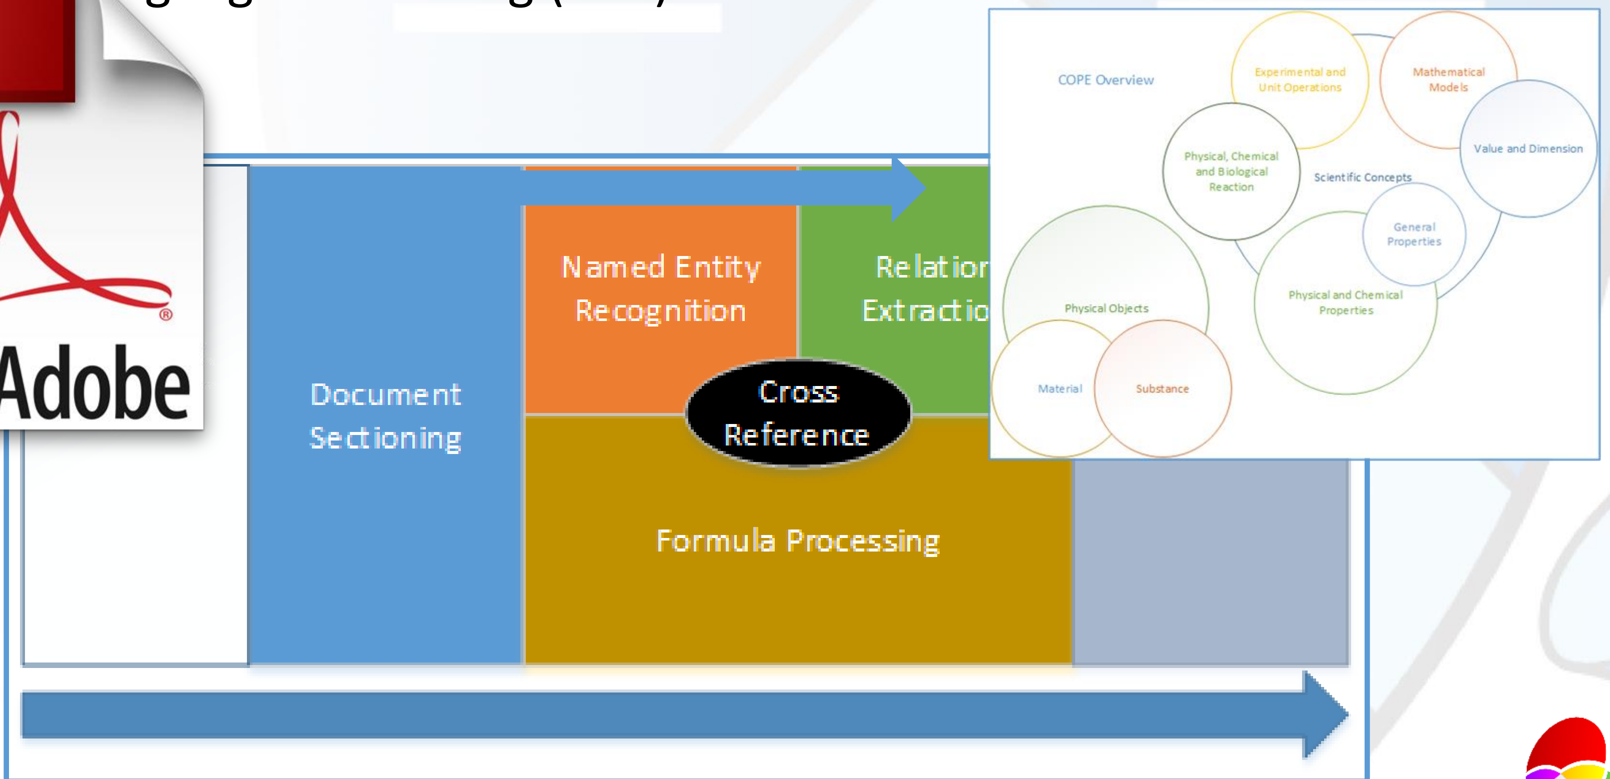

150 Columbia | Engineering
The Fu Foundation School of Engineering and Applied Science
1864–2014

HOLMES: Hybrid ontology-learning materials engineering system for pharmaceutical products: Multi-label entity recognition and concept detection, M. M. Remolona et. al., *Comp. & Chem. Engg*, in press, 2017.

CRIS.lab

AI in PSE: Phase III – Data Science

- **Challenges: Intellectual, Implementational and Organizational**
- **Smart Manufacturing Initiative**
  - Many relevant algorithms and knowledge modeling frameworks are **already known**
- **Implementational**
  - Computational power, storage, communication are here now!
  - **Integrating** Hardware, Software, Communication, and Models
  - **Managing** and updating data, knowledge and models
- **Organizational**
  - Personnel training
  - User acceptance and trust
  - System maintenance
- **These were the main limitations of the Honeywell ASM Program in 1995-2000**
- **Intellectual challenges**
  - Hybrid models
  - Domain-specific compilers
  - Ontologies
  - Custom languages and representations – e.g., Chemistry
  - Semantic search engines
  - Visualization

AI in PSE: Phase IV (2010 - ?)

- **Self-organizing Intelligent Systems**
- Modeling, predicting, and controlling the behavior a large population of self-organizing intelligent agents
  - Drone swarms, Driverless car fleets
  - Self-assembling nanostructures
- **Science of Emergence**
- **Grand conceptual challenges here**

Science of Self-organizing Systems

- **20th Century Science was largely Reductionist**
  - Quantum Mechanics and Elementary Particle Physics
  - Molecular Biology, Double Helix, Sequencing Human Genome

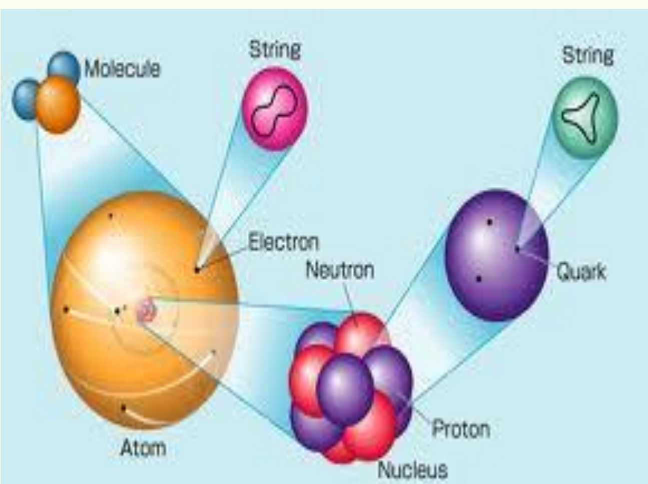

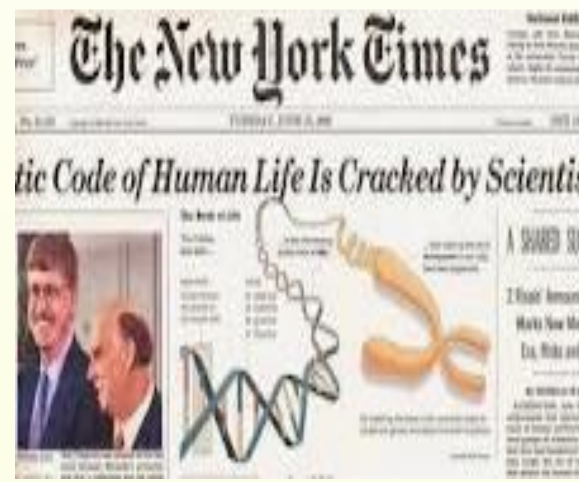

Complex Self-organizing Systems

- But can **reductionism** answer this question?
- Given the properties of a neuron, can we predict the behavior of a system of 100 billion neurons?

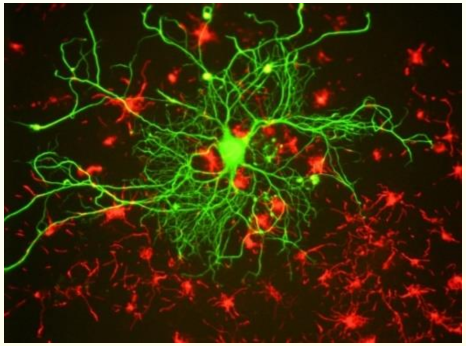

- From **Neuron** → **Brain** → **Mind**
- How do you go from Parts to System?
- **Reductionism cannot answer this!**

Two Small Clouds at the Dawn of 20^{th} Century

- Lord Kelvin's lecture, Royal Society, London, in April 1900
- "Nineteenth Century Clouds Over the Dynamic Theory of Heat and Light"
- "Physics knowledge is almost complete, except for two small "clouds" that remain over the horizon"
- These small "clouds" Revolutionized 20^{th} Century Physics
  - Blackbody Radiation: Quantum Mechanics
  - Michelson-Morley Null Experiment: Relativity

Max Planck

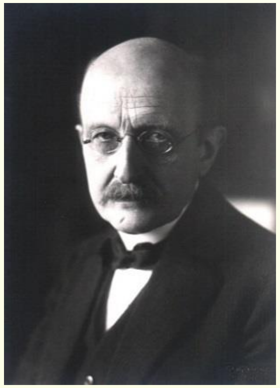

Albert Einstein

Lord Kelvin

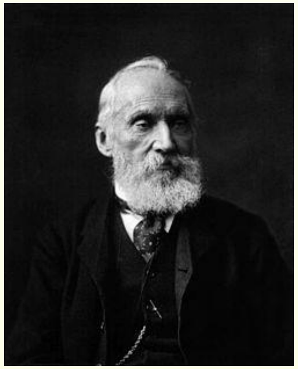

"Large Cloud" at the Dawn of 21st Century

- How do you go from **Parts to Whole**?

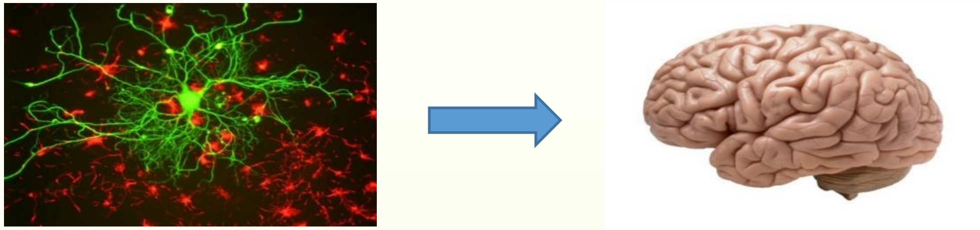

- Reductionism can't help here!

- Need an **Constructionist** Theory of **Emergent Behavior**

- Requires a NEW conceptual synthesis across AI, Systems Engineering, Statistical Mechanics, Game Theory, and Biology

- What might such a theory look like?

"Large Cloud" at the Dawn of 21st Century

- Individual agent properties ⟶ Emergent properties of millions of agents

- "Dumb" agents – e.g., Molecules
  - Classical Mechanics (Small) - e.g., Planetary motion
  - Statistical Mechanics (Large) - e.g., Gas

- "Intelligent" agents – e.g., People
  - Classical Mechanics ⟶ **Neoclassical Economics**
  - Statistical Mechanics ⟶ **???**

- Conceptual problem with **Entropy as Disorder**

"Large Cloud" at the Dawn of 21st Century

- True meaning of Entropy: Measure of Fairness in a Distribution
- Statistical Mechanics → Statistical Teleodynamics (4 Laws)
- Dynamics of Ideal Free Market
  - Proves equilibrium is reached by Maximizing Fairness
  - Proves equilibrium is both Statistical and Nash
  - Deep connection between Statistical Mechanics and Game Theory
  - Proves Existence, Uniqueness, Optimality, and Asymptotic Stability
  - Proves the Emergence of Income Distribution: Lognormal
  - Fairest Inequality
  - Guidelines for Tax Policy and Executive Compensation

Predictions for Different Countries

- Theory estimates lognormal-based income shares for Top 1%, Top 10-1%, and Bottom 90% for ideally fair societies
- Piketty's World Top Incomes Database (WTI)

Non-ideal Inequality Coefficient $\psi = \frac{Actual\ share}{Ideal\ share} - 100\%$

ψ = 0 Fairest Inequality; ψ ≠ 0 Unfair Inequality

Norway: Non-ideal Inequality ψ

Fairest Inequality Line at 0%

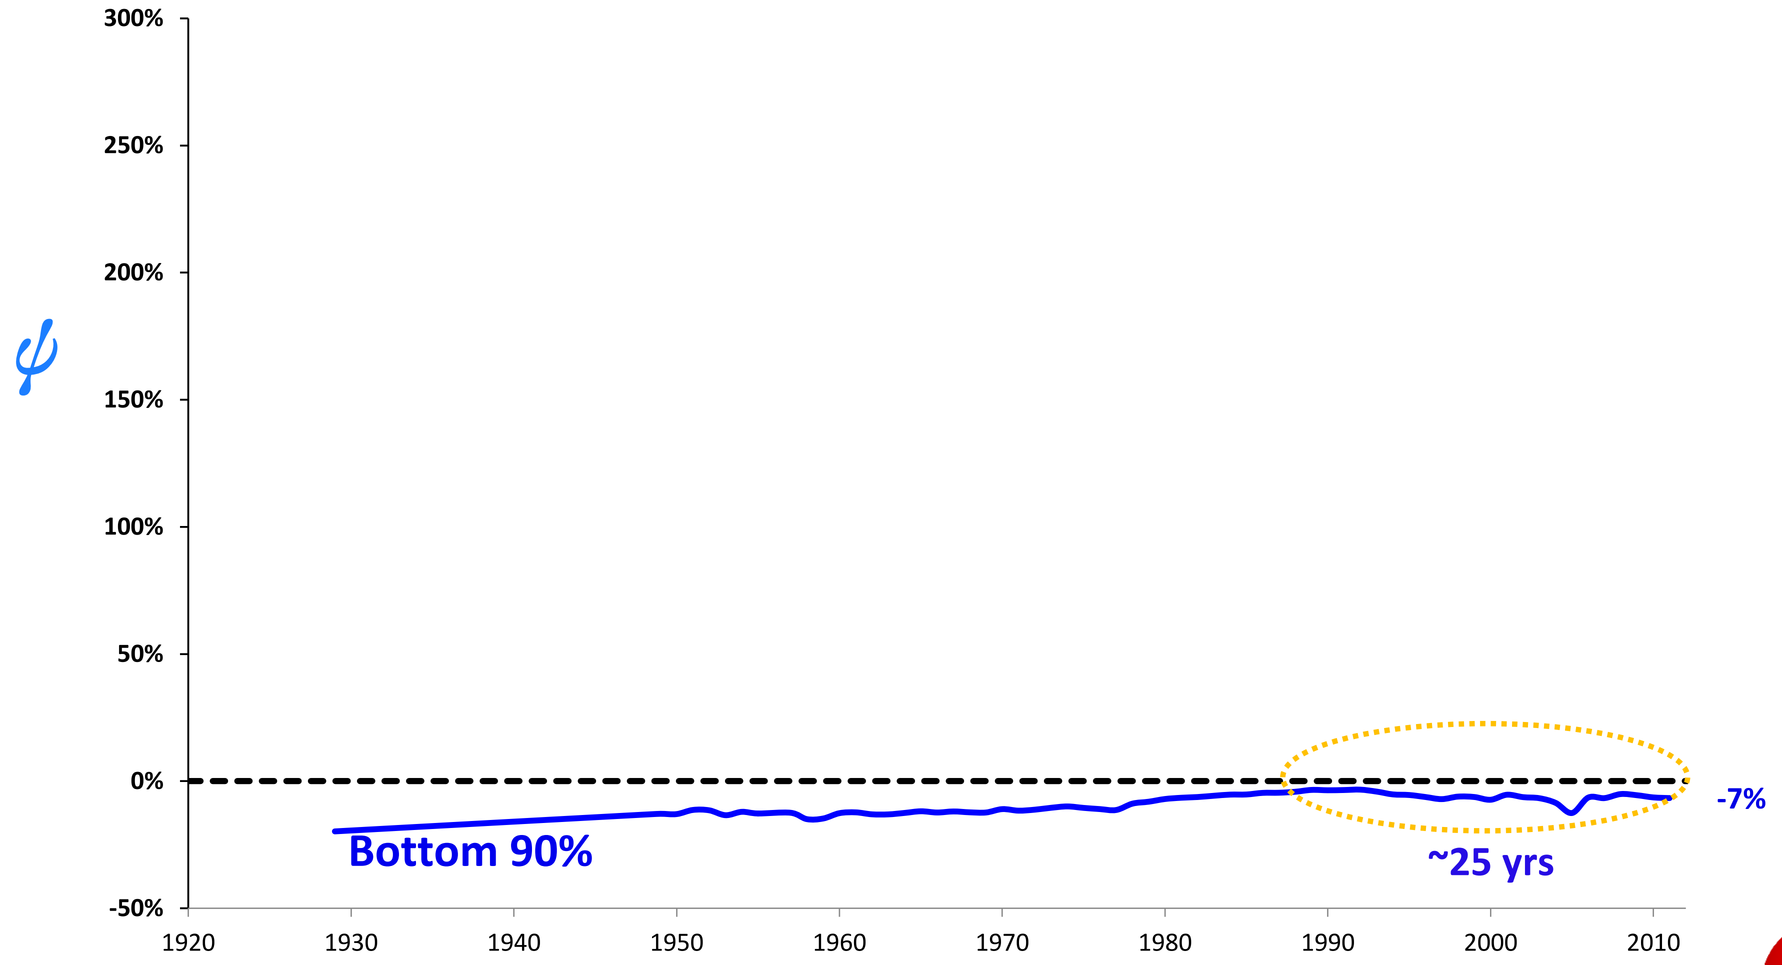

COLUMBIA | ENGINEERING
The Fu Foundation School of Engineering and Applied Science

CRIS.lab

Norway: Non-ideal Inequality ψ

Fairest Inequality Line at 0%

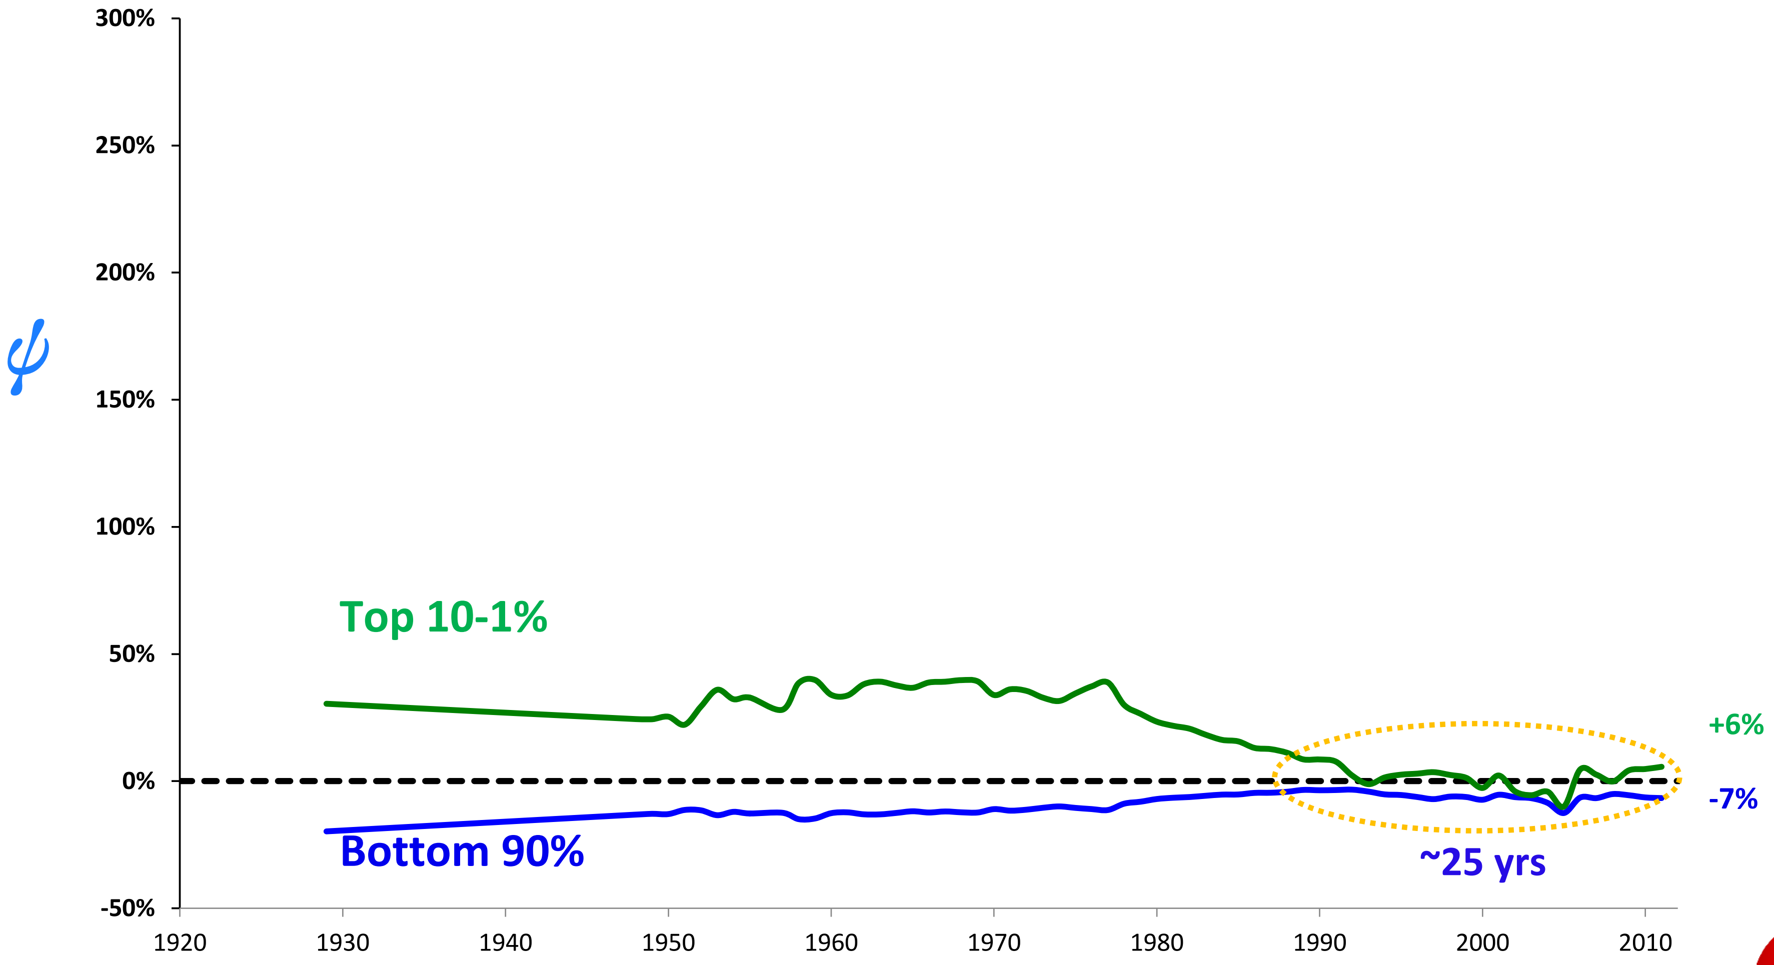

CRIS.lab

Norway: Non-ideal Inequality ψ

Fairest Inequality Line at 0%

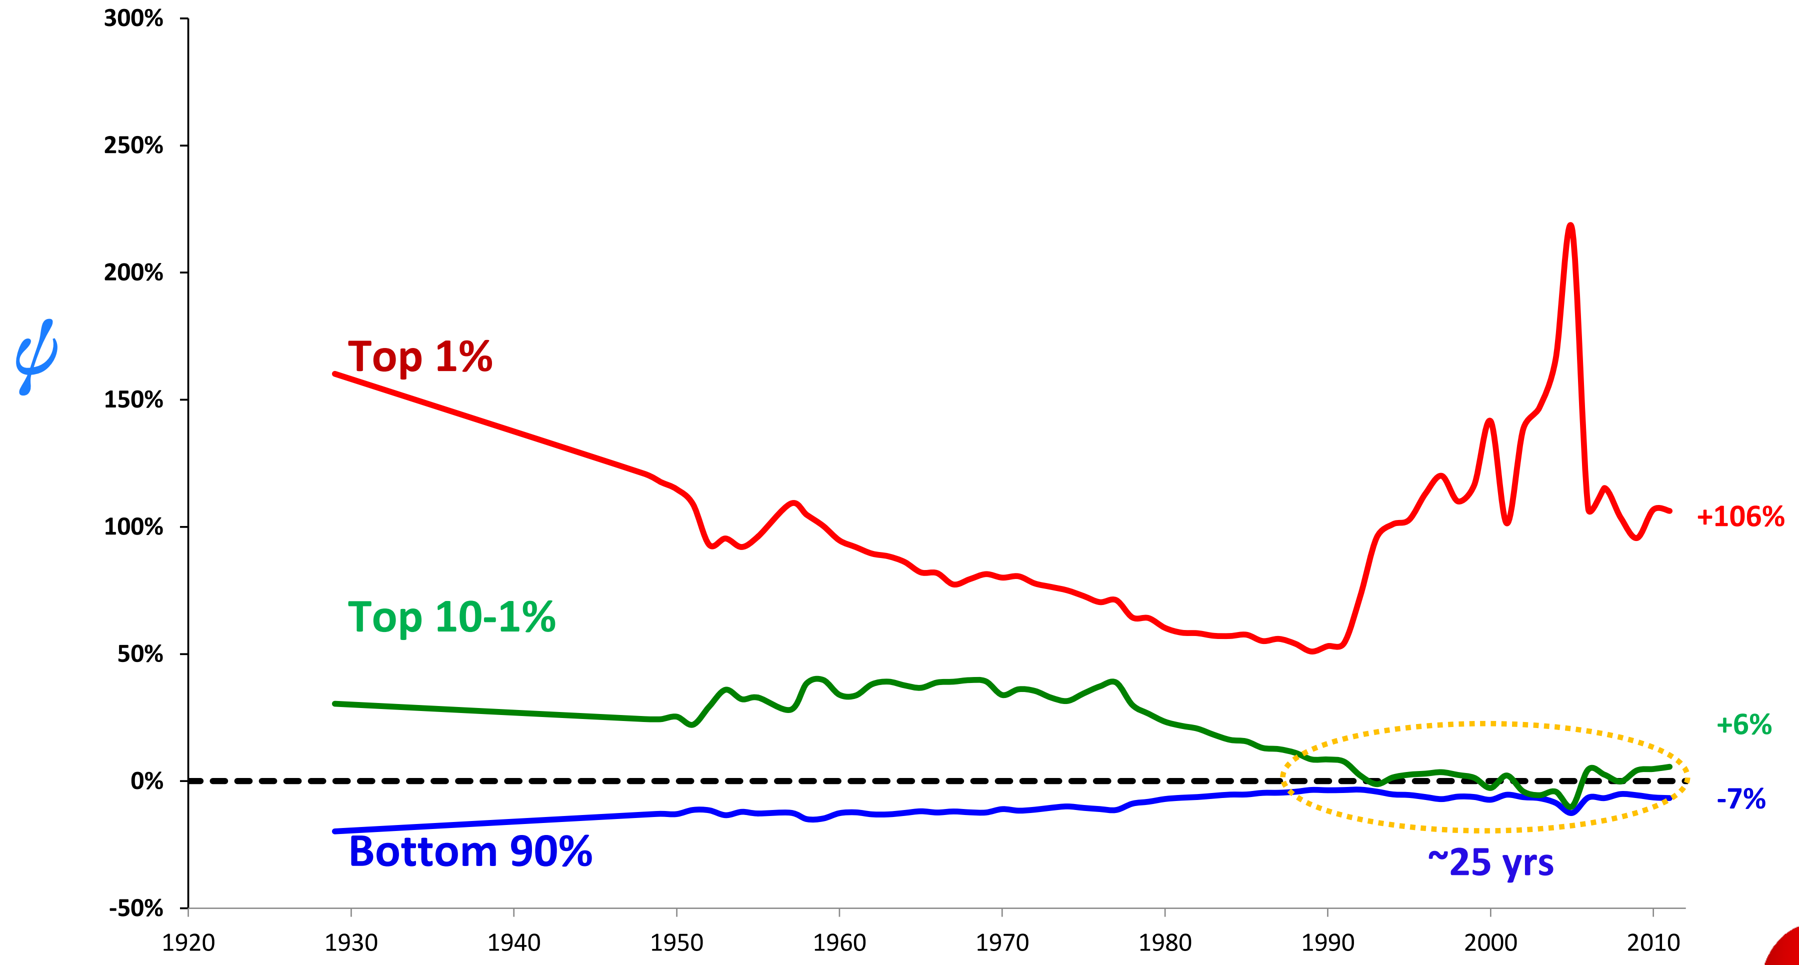

USA: Non-ideal Inequality ψ

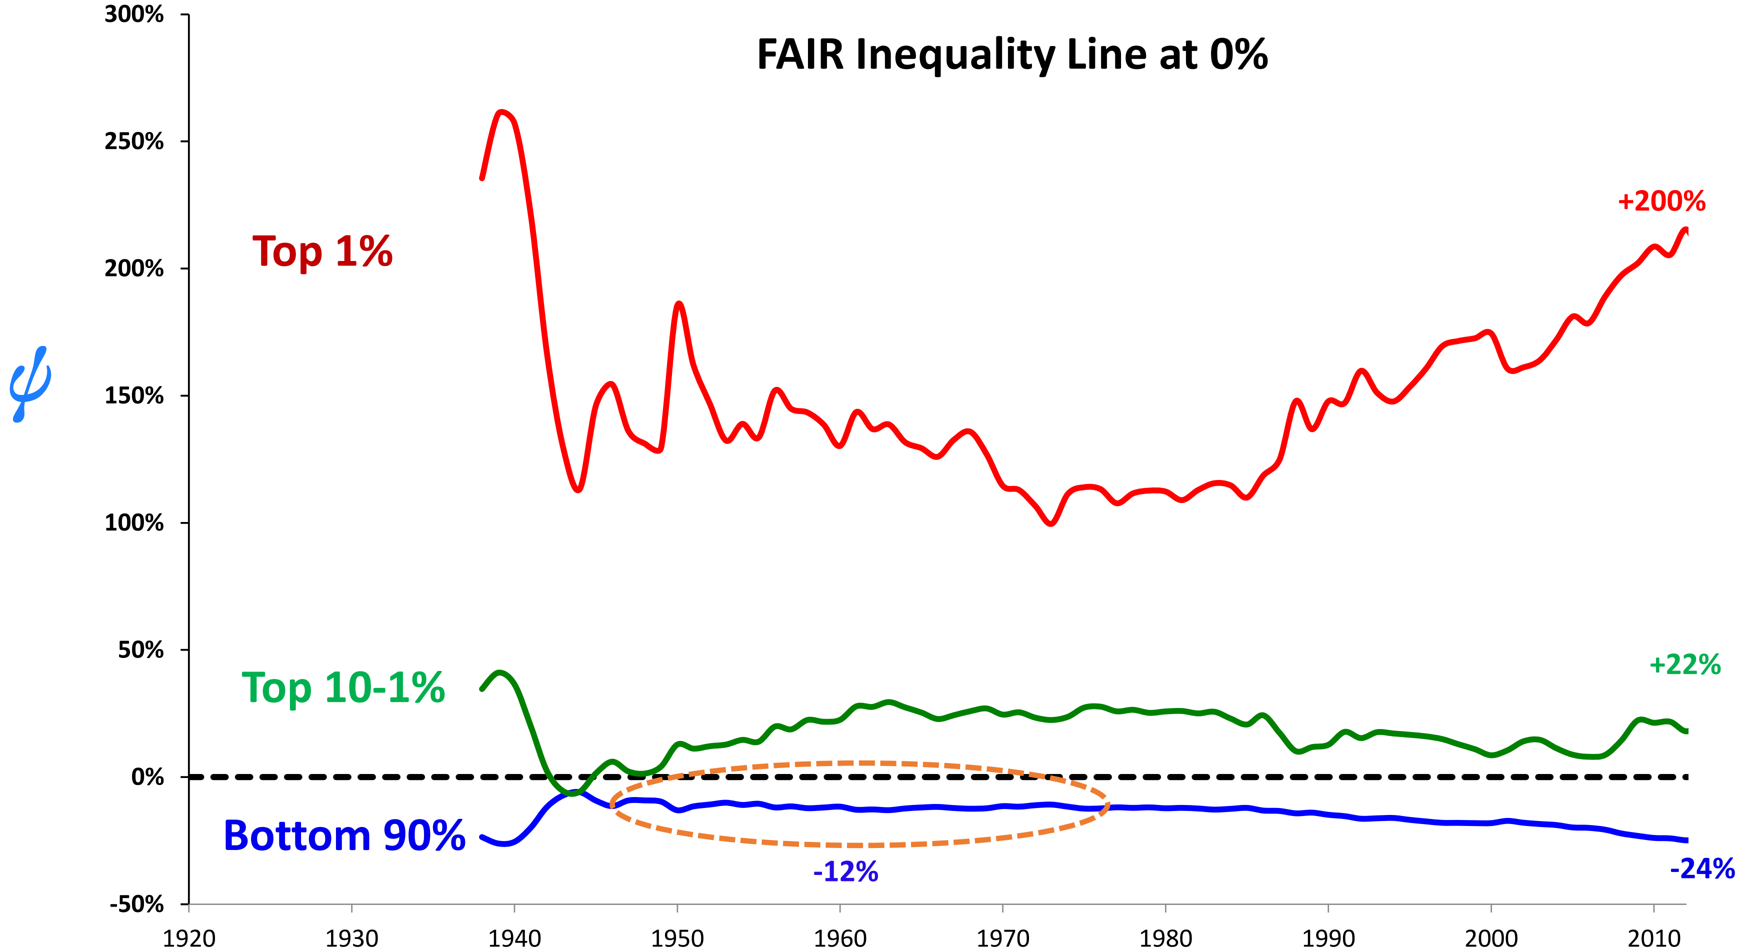

Data: Piketty's World Top Income Database

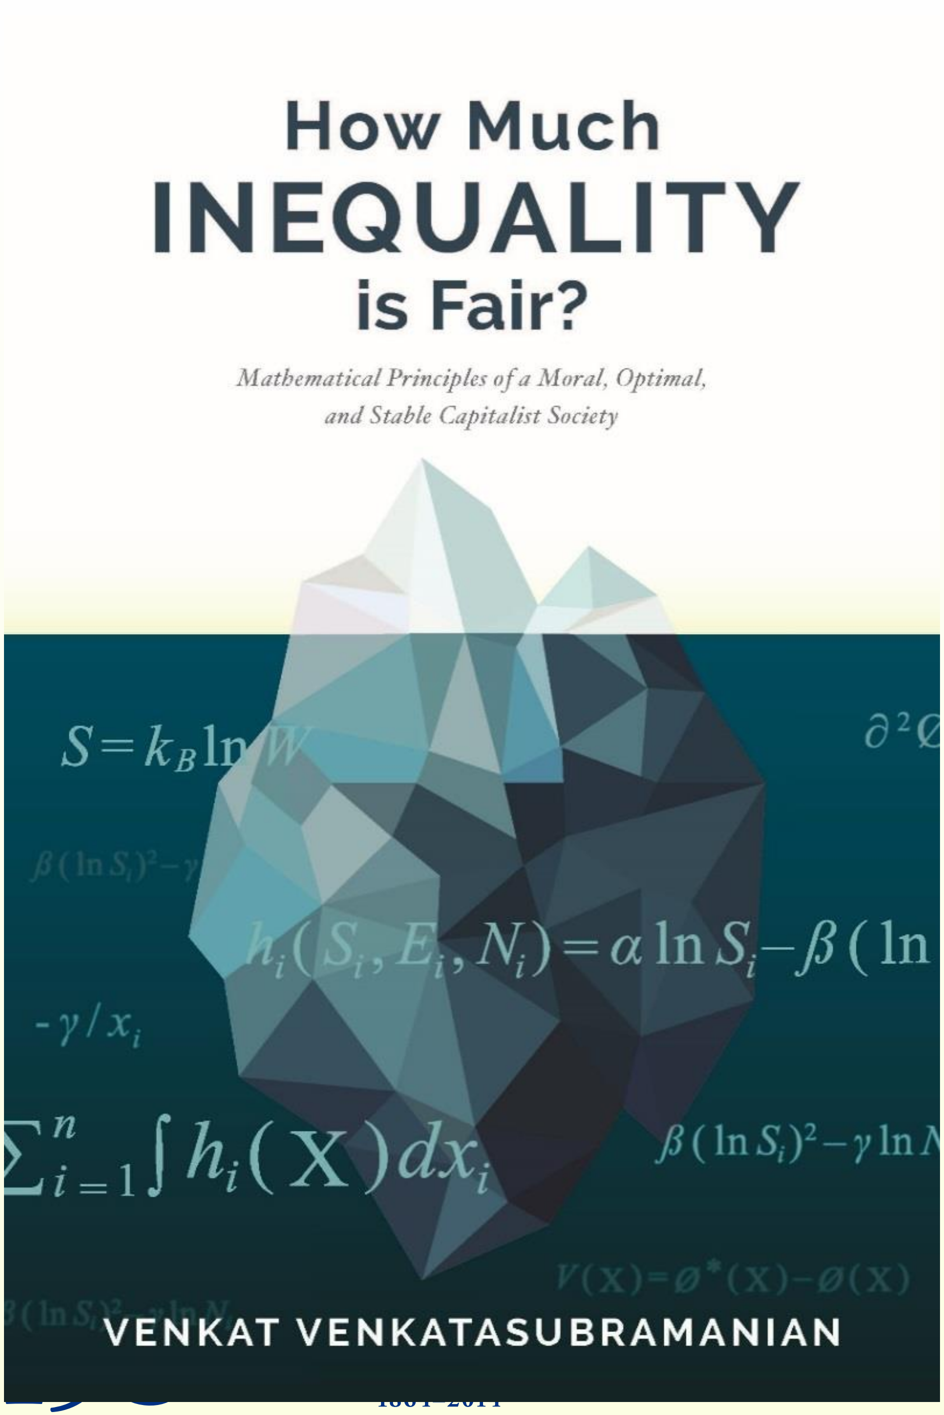

- Mathematical and Conceptual Foundations of Statistical Teleodynamics

- Synthesis of Concepts from Political Philosophy, Economics, Game Theory, Statistical Mechanics, Information Theory, and Systems Engineering

- Theory of Emergence of Income Distribution

Columbia University Press
Economics Series
July 2017

CRIS.lab

"Large Cloud" at the Dawn of 21st Century

- How do you go from **Parts to Whole**?

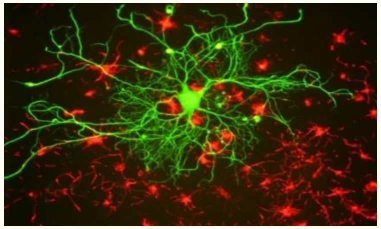

- Need an **Constructionist** Theory of **Emergent Behavior**

- Requires a NEW conceptual synthesis across AI, Systems Engineering, Statistical Mechanics, Game Theory, and Biology

AI in PSE: Dawn of a New Era

- **Grand Intellectual Challenges at the intersection of Complexity Science, AI and Systems Engineering**
  - Theory of Emergence
  - Design, Control, Optimization and Risk Management by Self-Organization
- **Impact of AI in PSE**
  - Hardware, software, communication, cost, acceptance are here
  - But will still take **20 – 30** years to reach significant impact
    - Hybrid models
    - Domain-specific Compilers
    - Ontologies
    - Custom languages and representations
    - Semantic search engines
    - Visualization
- **Revolutionize all aspects of PSE**
  - Energy, Sustainability, Materials, Pharmaceuticals, Healthcare, Systems Biology

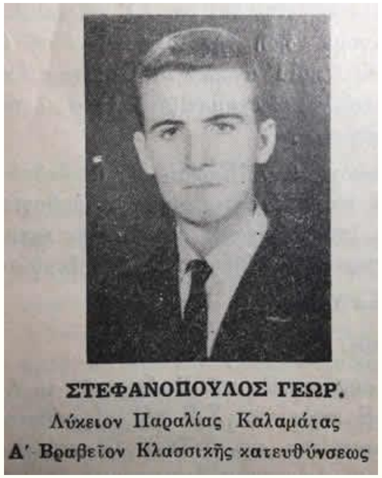

Thank You, George!

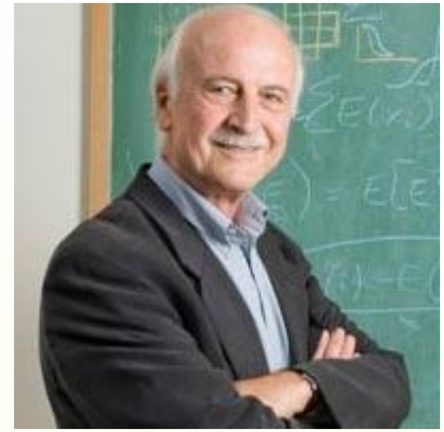

For your great contributions to PSE!

For your support in my career!

- **Happy 70th Birthday!**
- **Best wishes for a happy retired life!**

Image Credits

creative commons




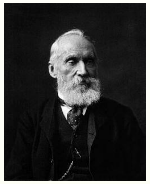


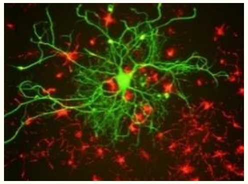



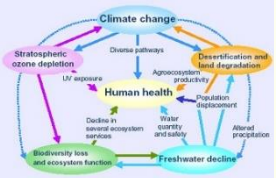

gettyimages

The New York Times


Other

http://link.springer.com/chapter/10.1007%2F978-3-540-72608-1_3

Valdis Krebs, http://orgnet.com/cdo.html

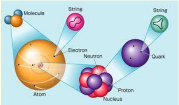

http://www.rikenresearch.riken.jp/eng/frontline/4584

http://www.sciencephotogallery.co.uk/

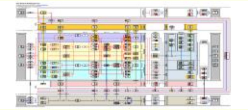

Shadow Banking
Zoltan Pozsar, Tobias Adrian, Adam Ashcraft, and Hayley Boesky
Federal Reserve Bank of New York Staff Reports, no. 458

Thank You for Your Attention!

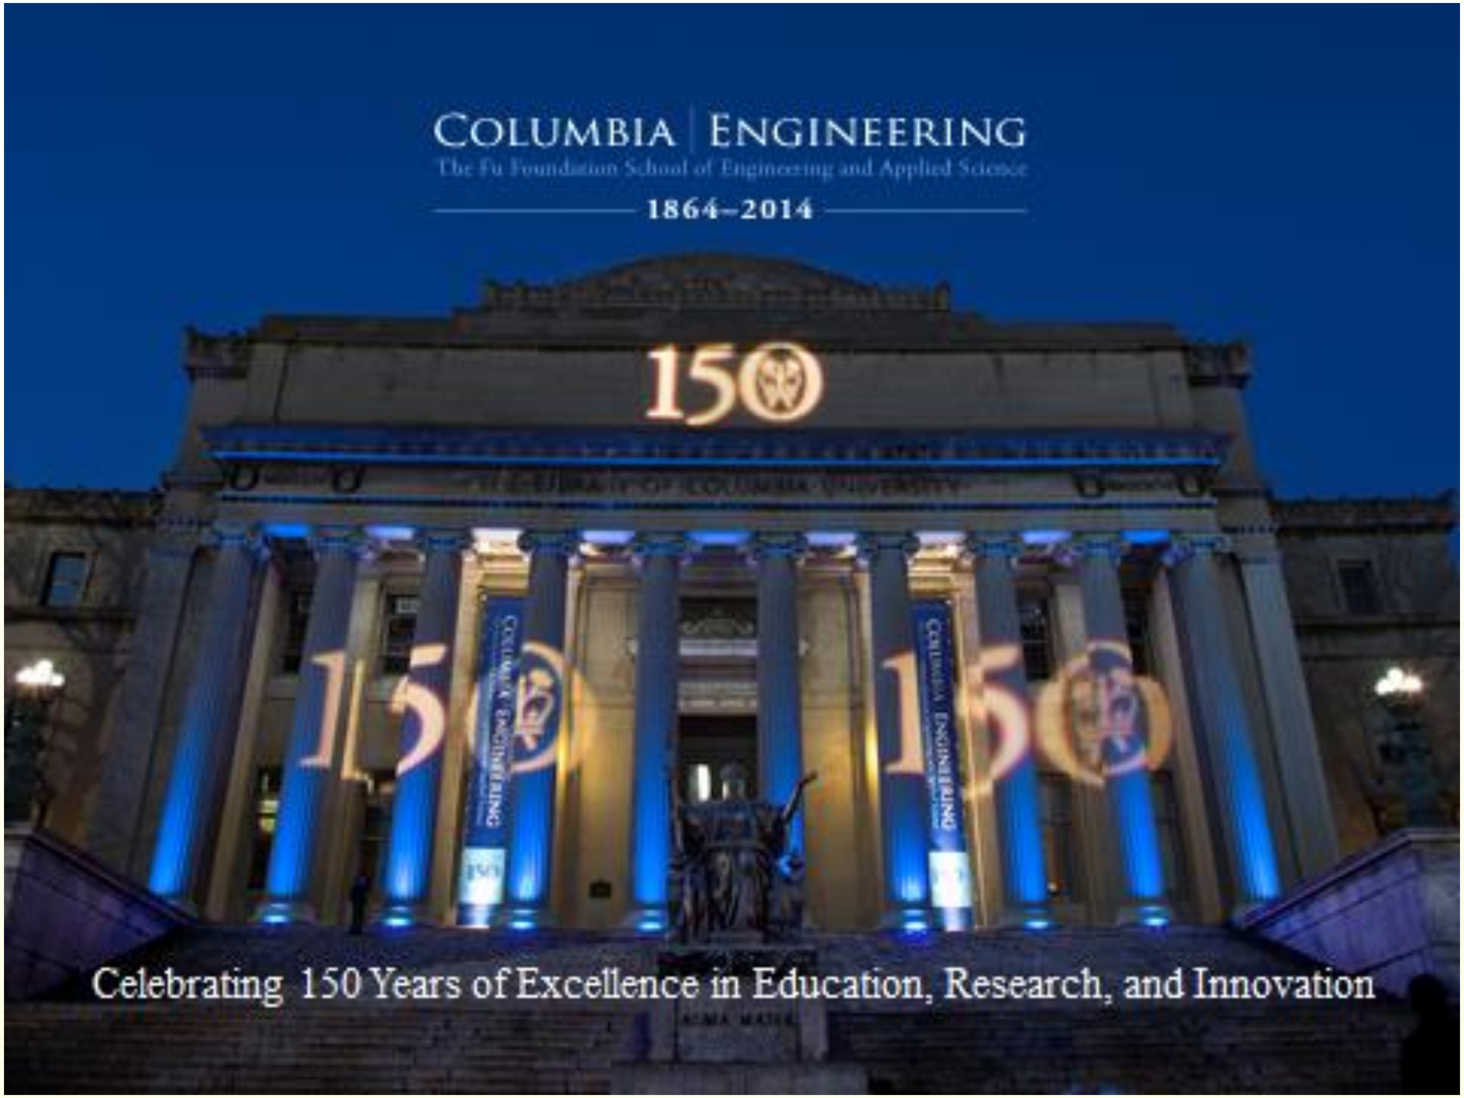

Questions?