



Advanced Network Management *a.k.a.*

AI Ops: **Autonomous IT Operations** or **AI for IT Operations**

course#: 80240663

Spring 2022



ANM2022



Instructor: Dan Pei

peidan@tsinghua.edu.cn

Teaching Assistant:

Qianyu Ouyang oyqy2019@mails.tsinghua.edu.cn



Please scan QR code to join class WeChat Group

Roadmap

- **Self-Introduction**
- Course Logistics
- What is AIOps?
- Course Coverage (Website)
- Assignments and Projects

How to pronounce “Pei”?

Just pronounce “Pei” as “Pay”

In fact, just call me “Dan”

Louvre Museum



I-M Pei: the famous architect who designed the glass pyramid




About the Instructor

- Tenured Associate Professor.
- UCLA Ph.D. Best Ph.D. Thesis Award in UCLA CS in 2005.
- Joined Tsinghua CS Department in December 2012
- Homepage: <http://netman.aiops.org/~peidan>
- Previously a Principal Researcher at AT&T Research, a co-founder and founding CEO of a mobile health company in Beijing, before joining Tsinghua.
- Supervised interns from CMU, Cornell, Princeton, UCLA, GaTech, Michigan, Northwestern etc. Now @ Google, MSR, IBM, Purdue, Northeastern, HKUST
- ACM/IEEE Senior Member

My Research Group @ Tsinghua: NetMan

<http://netman.aiops.org>

- Currently advising 11 Ph.D. and 6 M.S. students at Tsinghua.
- 2 affiliated assistant professors and 4 post-docs
- Graduated 15 PhDs & 5 Masters (4人拿到华为“天才少年”，三人拿到快手 Kuai Star, 4 went to MSRA, 5 became assistant/associate professors, 3 became CEO/CTO/director in a startup)



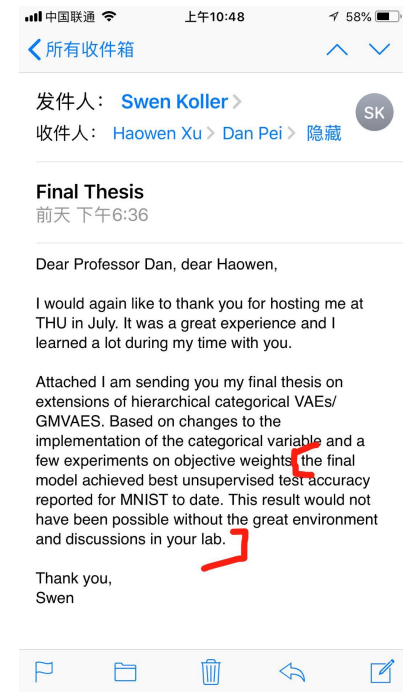
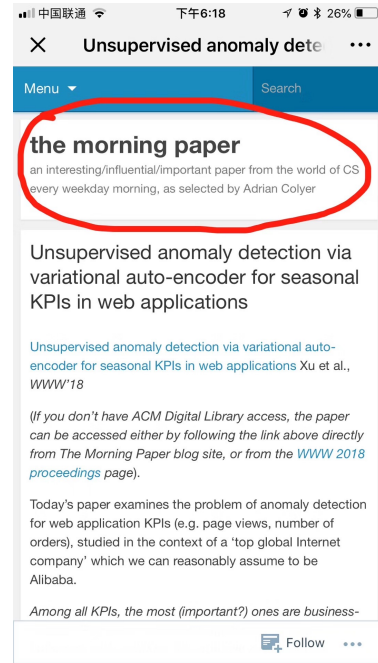
Industry Collaborators



Publications:

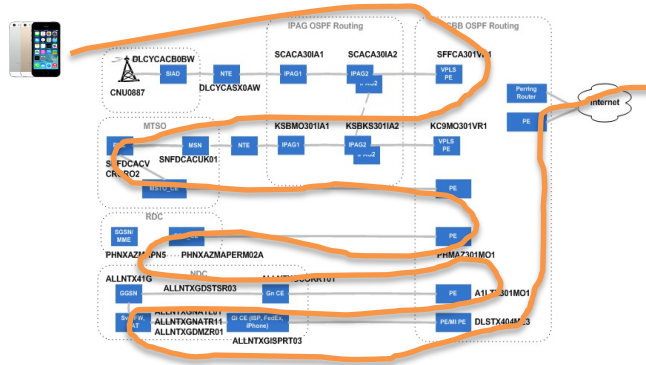
180+ AIOps papers and 20+ issued US Patents. Published in SIGCOMM、WWW、KDD、SIGMETRICS、TON、ATC、INFOCOM、FSE、ASE、IMC、CoNEXT、IJCAI etc.

Research results are covered by technology media such as MIT technology Review, Hacker News, Mother Board, Morning paper, and many Chinese media.




AIOps : Autonomous IT Operations through Machine Learning


Large & complex access network



Large & complex data center



Large & complex application software




2012 淘宝核心链路应用拓扑图

- Imagine that you are running an Internet-based service with hundreds of thousands of servers and many software modules, a large, complex, cross-layer, and rapidly evolving distributed system.
- You want to achieve 99.999% service reliability, but machine-generated monitoring data and hundreds of operators (IT operation engineers) alone won't get you there, because of the high complexity and sheer scale of the software/hardware system and the vast amount of machine-generated data.
- Machine learning is the direction to enable Autonomous IT Operations autonomous.



AIOps Architecture & Algorithms



- The major topics of AIOps often coincide with its more general counterparts in Machine Learning:
 1. Anomaly Detection in Time Series, Logs (semi-structured text), Traces (program execution trace), and Graphs
 2. Anomaly Localization
 3. Failure/Event Prediction
 4. Causal Inference and its application in Root Cause Analysis
- State-of-art Machine Learning Algorithms are applied to solve the unique challenges in AIOps:
 1. Deep Neural Networks for Time Series or Sequence
 2. Deep Generative Model (VAE, GAN)
 3. Deep Reinforcement Learning
 4. Natural Language Processing
 5. Causal Inference

PRINCIPLES RAY DALIO

“Ray Dalio has provided me with invaluable guidance and insights that are now available to you in *Principles*.”

—BILL GATES

“I found it to be truly extraordinary. Every page is full of so many principles of distinction and insights—and I love how Ray incorporates his history and his life in such an elegant way.”

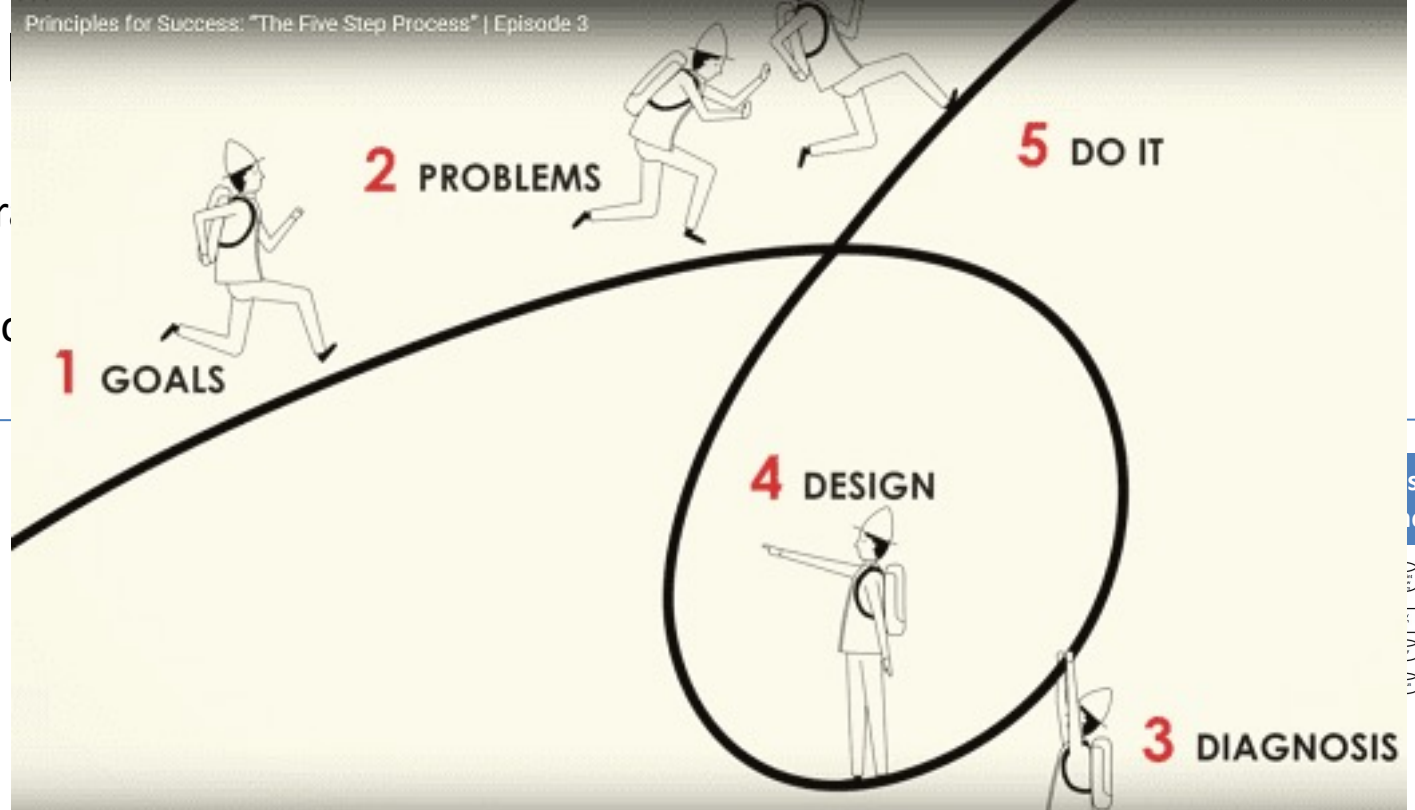
—TONY ROBBINS

#1 NEW YORK TIMES BESTSELLER

INFOCOM 2012

y:

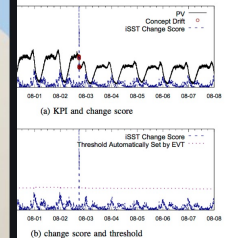

m Pr
lem
If no



EPISODE 03 PRINCIPLES FOR SUCCESS

Unsu
Learn

Transfer
Learning for
concept drift



SSRE 2018 Best
Paper

WW










Roadmap

- Self-Introduction
- **Course Logistics**
- What is AIOps?
- Course Coverage (Website)
- Assignments and Projects

Summary

- AI Ops is an **interdisciplinary** research field between **Machine Learning** and **Systems/Networking**, which is why this course had this historical title “Advanced Network Management”. If you are interested in learning how a large distributed system can be better run with the help of machine learning, this course is for you. If you want to learn **how machine learning can help solve challenging problems in a very complex system**, this course is for you.
- This course will cover the latest progress in major topics of AI Ops using **case studies** from recent research papers in top conferences in all major computer science fields, including Machine Learning, Data Mining, System/Networking, Software Engineering, Database, Multimedia, etc.
- Through these case studies, we will show how the latest Machine Learning Algorithms are applied to solve the unique challenges in AI Ops. The **basics of these Machine Learning** algorithms will be briefly **reviewed in an easy-to-understand way**, without going through the detailed theory behind them. Thus by the end of the course, you will be able to learn **roughly how these algorithms work, and how it can be applied to solve real-world problems**.





Course Requirements

- Course website: <http://course.aiops.org/>
- **Prerequisites:**
 - You are expected to be familiar with at least one programming language, preferably Python.
 - If not, please quickly learn one.
 - You don't have to have previous machine learning knowledge.
- Encourage interaction and discussion
 - stop me and ask questions at any time! (online students please “raise your hands” on Zhumu)
 - You get credits for interaction (Yes, our TA notices your interaction)


Course Info

- Time: Tuesday **7:20pm-9:45pm** (Weeks 1-6,7-16)
 - 15 classes, each has three 45-minute sessions.
 - Exception: Saturday May 7th **7:20pm-9:45pm**
- Session 1: Machine Learning Basics
- Sessions 2 & 3: AIOps case studies
- Office Hours:
 - After every class.


Grading:

- Attendance & Interaction: **10%**
 - Presence: 0.5% for each lecture in Weeks 3-14. 6% in total
 - 1% for each question asked or answered. 4% maximum.
- Assignments: **25%** (each student finishes each assignment alone)
 - Assignment 1: 10%
 - Assignment 2: 15%
- Project: **65%** (A project team of 2 or 3 students)
 - Ranking in algorithm competition & report: 55%
 - Presentation in week 16 (scored by all students, TA, and instructor): 10%
- The final grade will be in letter grading scale (e.g., A,B,C,D)

Project: Algorithm Competition



Invited Talks



Roadmap

- Self-Introduction
- Course Logistics
- **What is AIOps?**
- Course Coverage (Website)
- Assignments and Projects

Roadmap

- Self-Introduction
- Course Logistics
- What is AIOps?
- **Course Coverage (Website)**
 - <http://course.aiops.org>
- Assignments and Projects

Case Studies: Papers and Slides

- Already Posted on the course website
 - Maybe updated before each class, so please check the website weekly.
- Lecture Coverage
 - You were supposed to read these papers carefully
- Reading List
 - Relevant papers, read them if you have time

Why reading papers/case studies?

- Purpose: background knowledge, algorithm, methodology, writing, design principle, solution inspiration, evaluation methodology for the problem you are working on, and (sometimes) vision.
- My strongly biased personal opinion: Papers are not for finding topics for your next AI Ops paper.
 - Instead, if possible, “From Practice, into Practice”: E.g. try existing algorithms in papers in practice to discover where it does not work

What's a good paper

- A new and important problem, solid solution.
- Old and challenging problem, a new *simple but elegant* solution, with straightforward insight and intuition behind it.

Using Google Scholar to find relevant papers (demo)

- Browse latest proceedings of relevant conferences, find one relevant paper p . *Relevant paper set* $S=\{p\}$

for (each new p in S) {

 Browse p 's references, and put relevant ones into S ;

 Browse p 's citations in Google Scholar, and put relevant ones into S ;

}


<https://www.editage.com/insights/8-winning-hacks-to-use-google-scholar-for-your-research-paper>

Roadmap

- Self-Introduction
- Course Logistics
- What is AIOps?
- Course Coverage (Website)
- **Assignments and Projects**

Enjoy the course!

Spring 2017



Spring 2018



Fall 2018



Fall 2019



Fall 2020

