Computational Sustainability: Sustainable Data Centers

CS 325: Topics in Computational Sustainability, Spring 2016

Manish Marwah
Senior Research Scientist
Hewlett Packard Labs
manish.marwah@hpe.com

Sustainability

"sustainable development is development that meets the needs of the present without compromising the ability of future generations to meet their own needs"

the Brundtland Commission of the United Nations, 1987

What is “sustainability”?

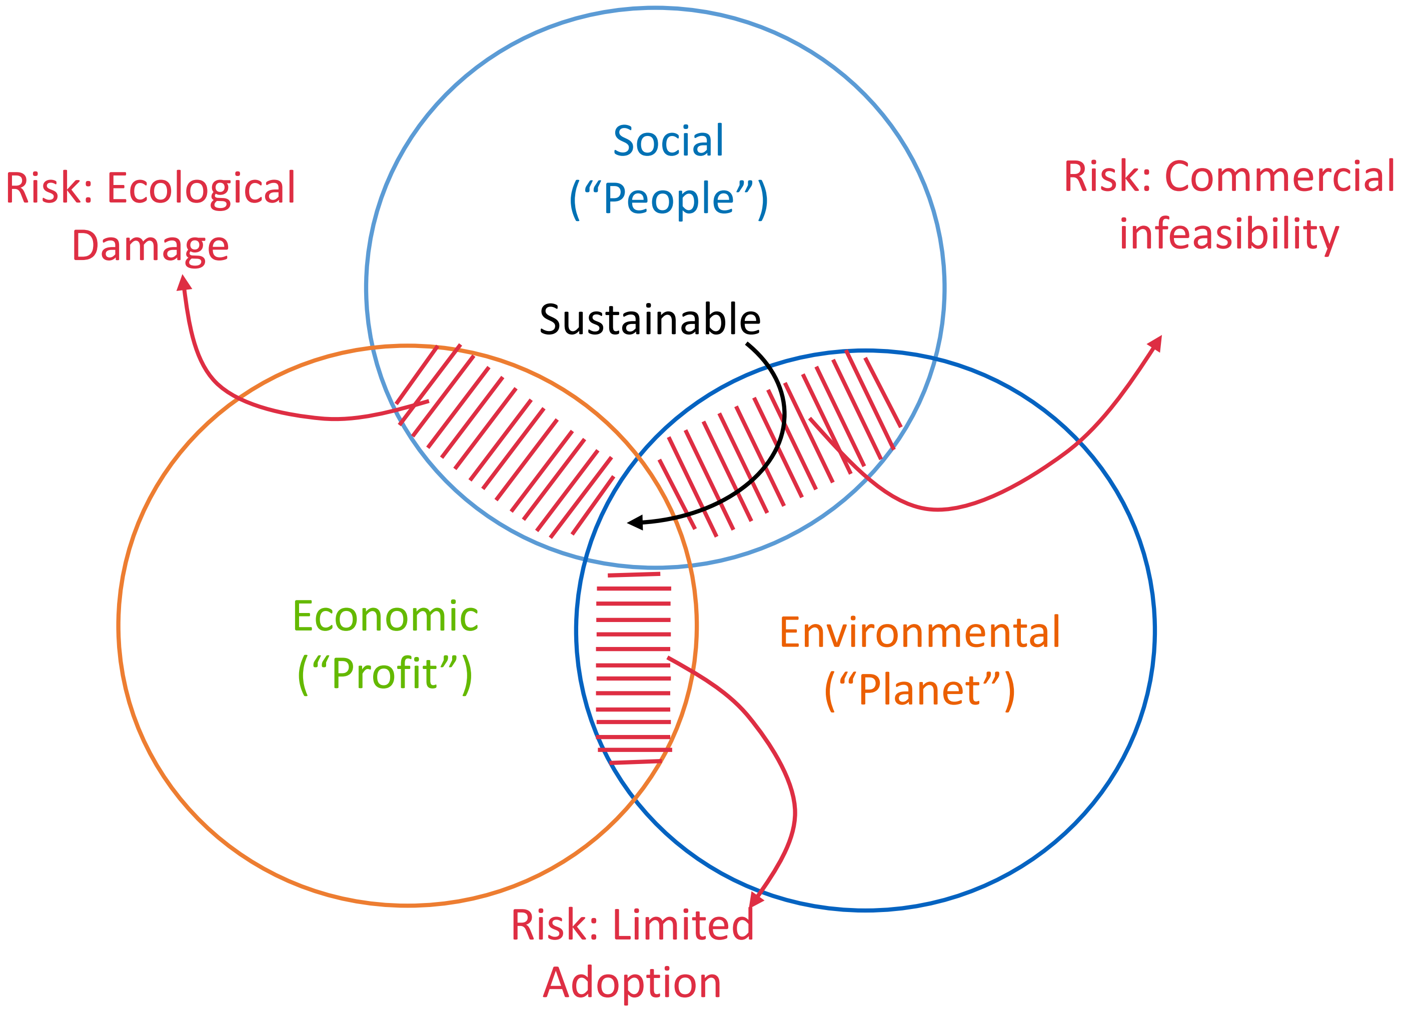

Figure Credit: A. Agogino, UC Berkeley

Environmental Sustainability

- Life Cycle View

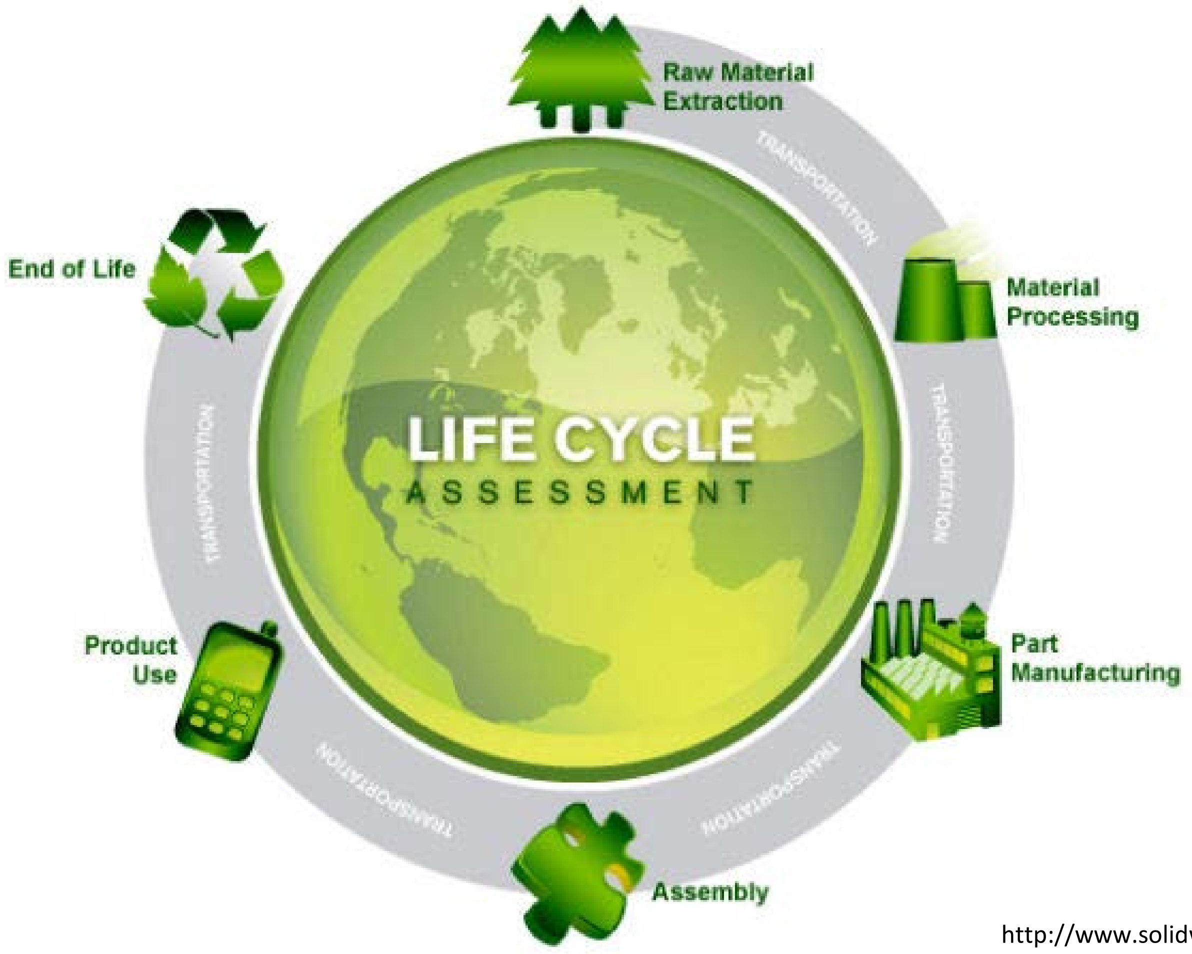

http://www.solidworks.com/sustainability/design/2722_ENU_HTML.htm

- Environmental impact factors (e.g., carbon, water, toxicity, etc.)

Lifecycle View of IT

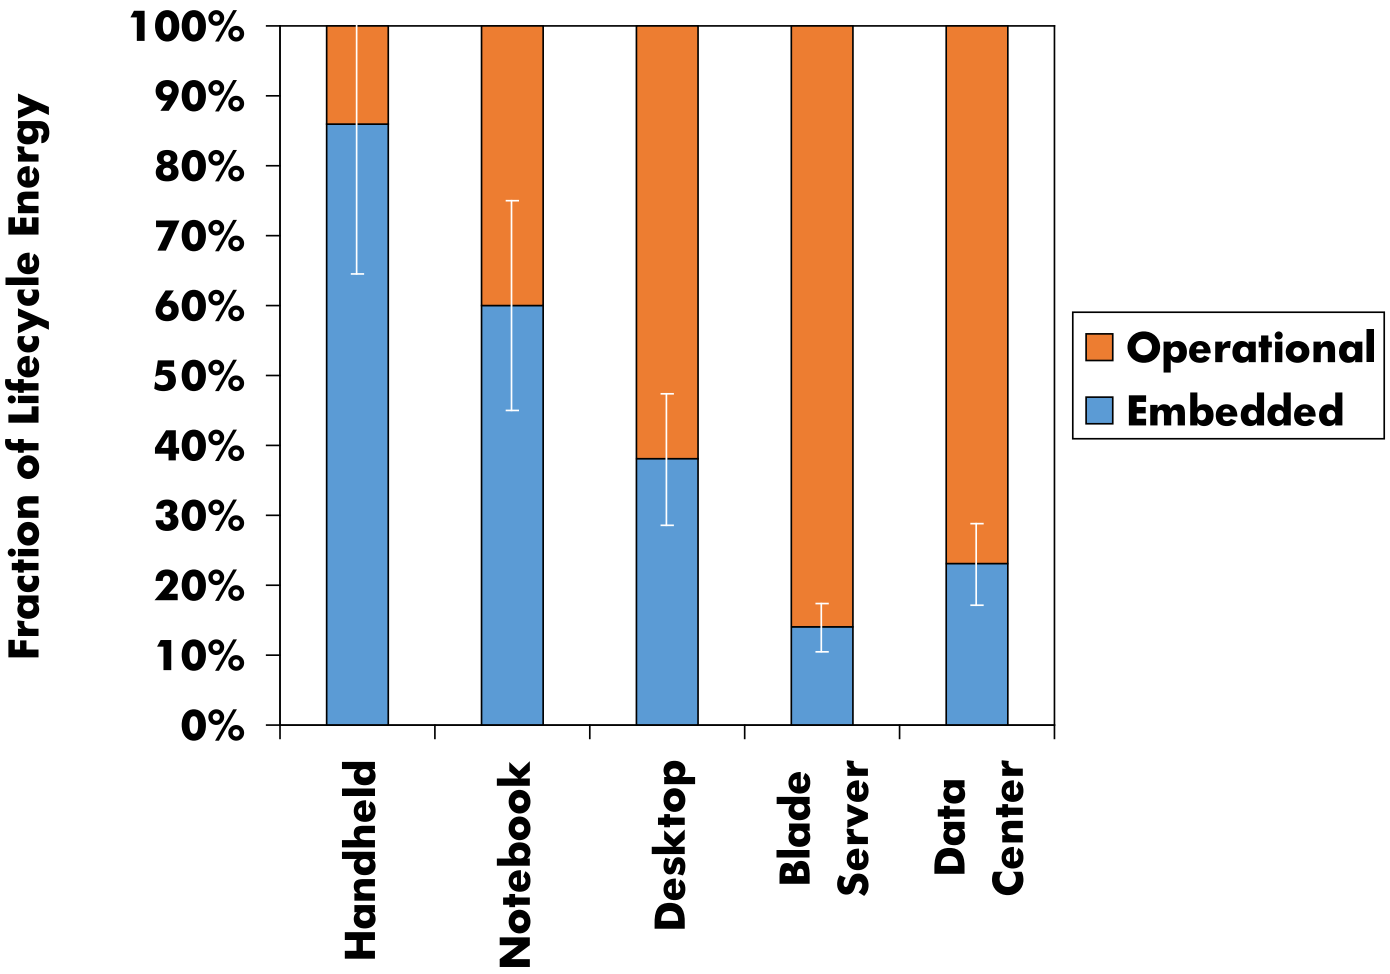

Results are illustrative only.
Actual footprint may differ.

Ref: IEEE Computer 2009

Outline

- Data center from an energy perspective
- Data center chiller operation – temporal data mining [KDD 2009]
- Thermal anomaly detection [SenserKDD 2010]
- PV generation forecasting [AAAI 2012]
- Net-zero Data Center [SIGMETRICS 2012]

Cloud Data Center

Energy Supply and Demand Side

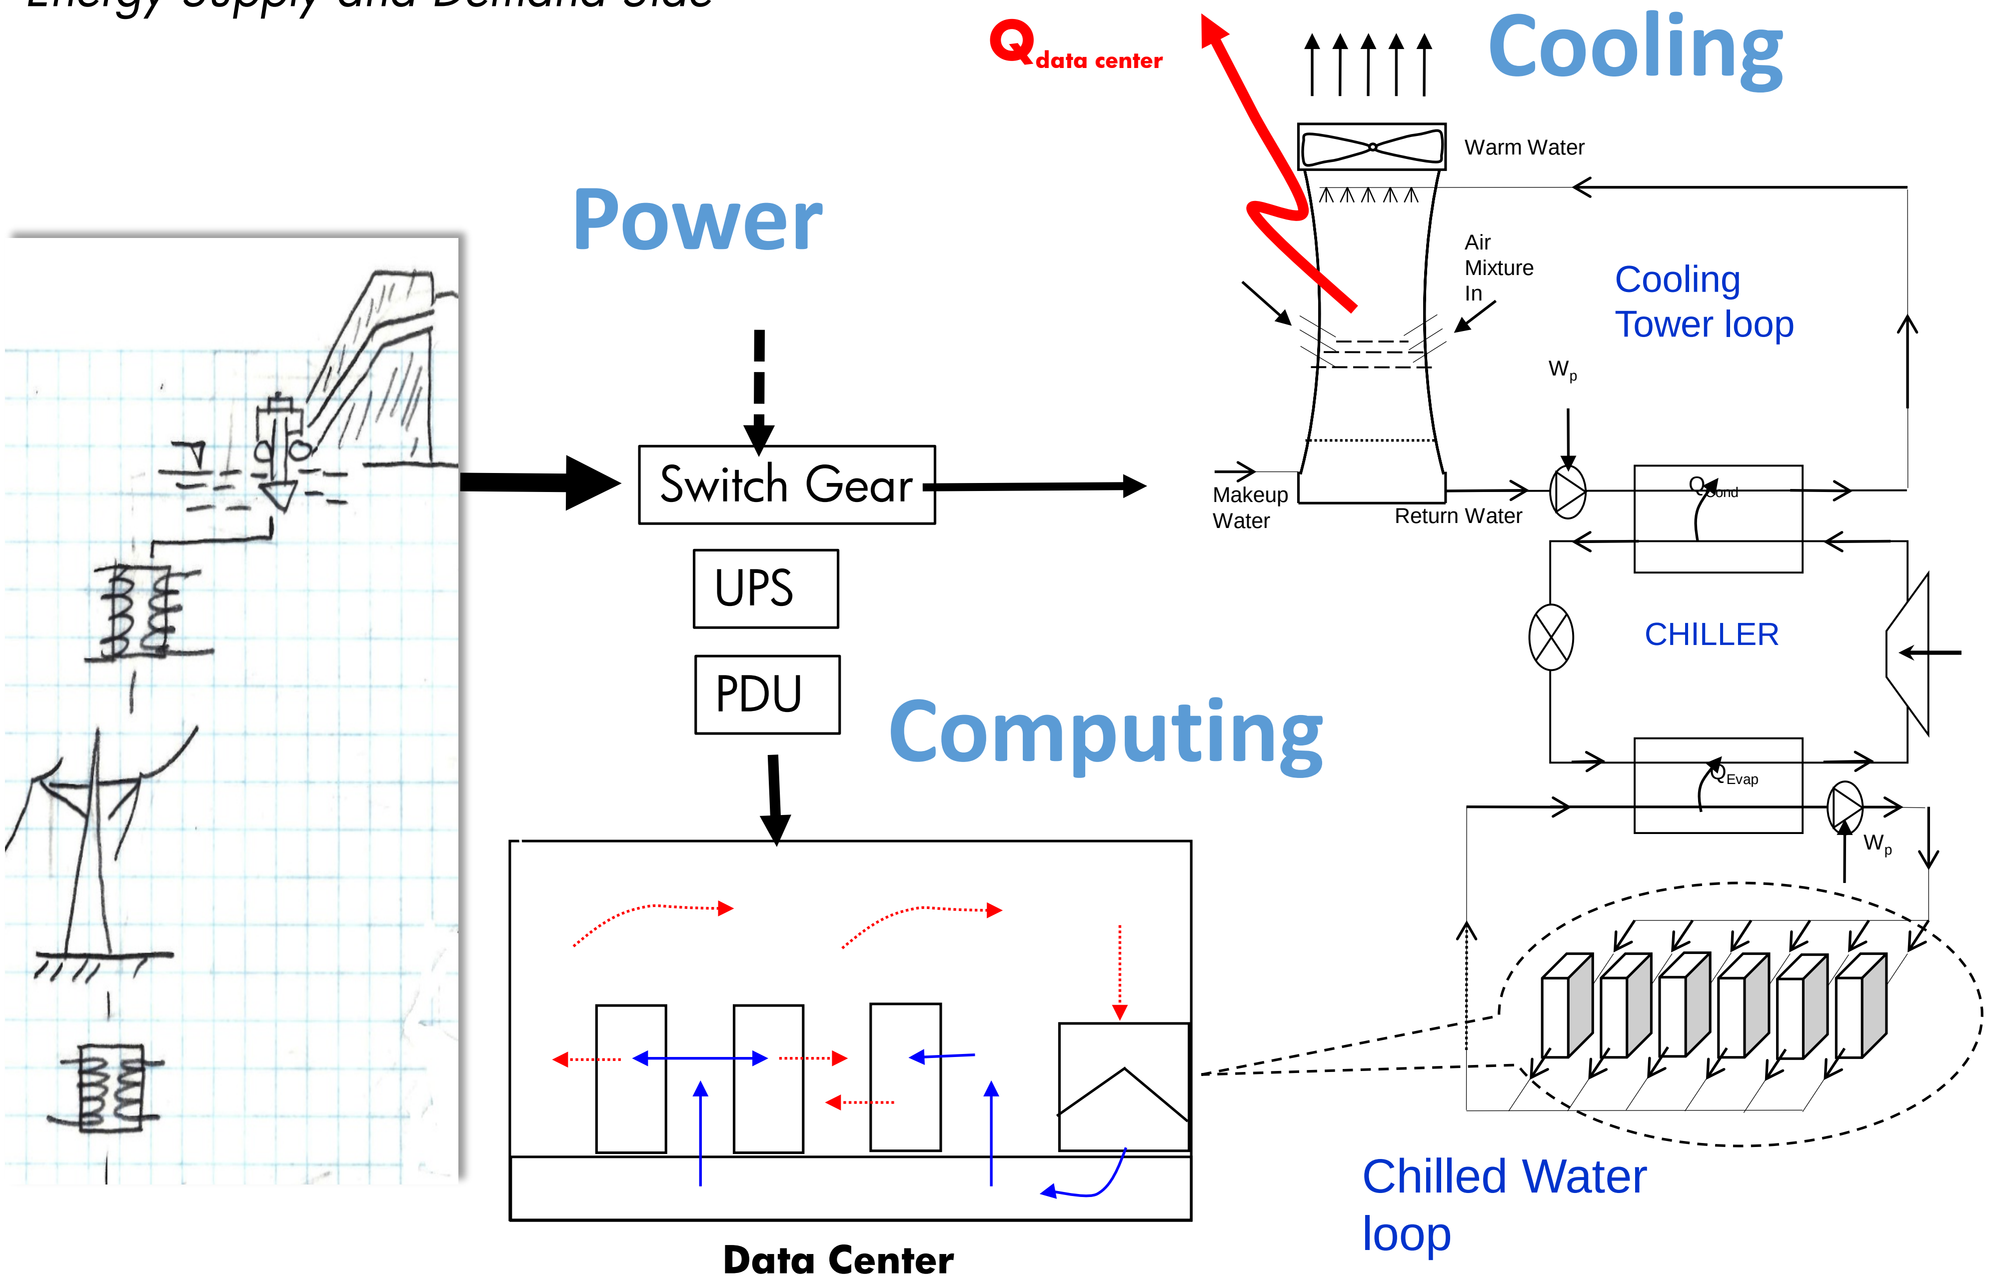

Sustainable Data Center

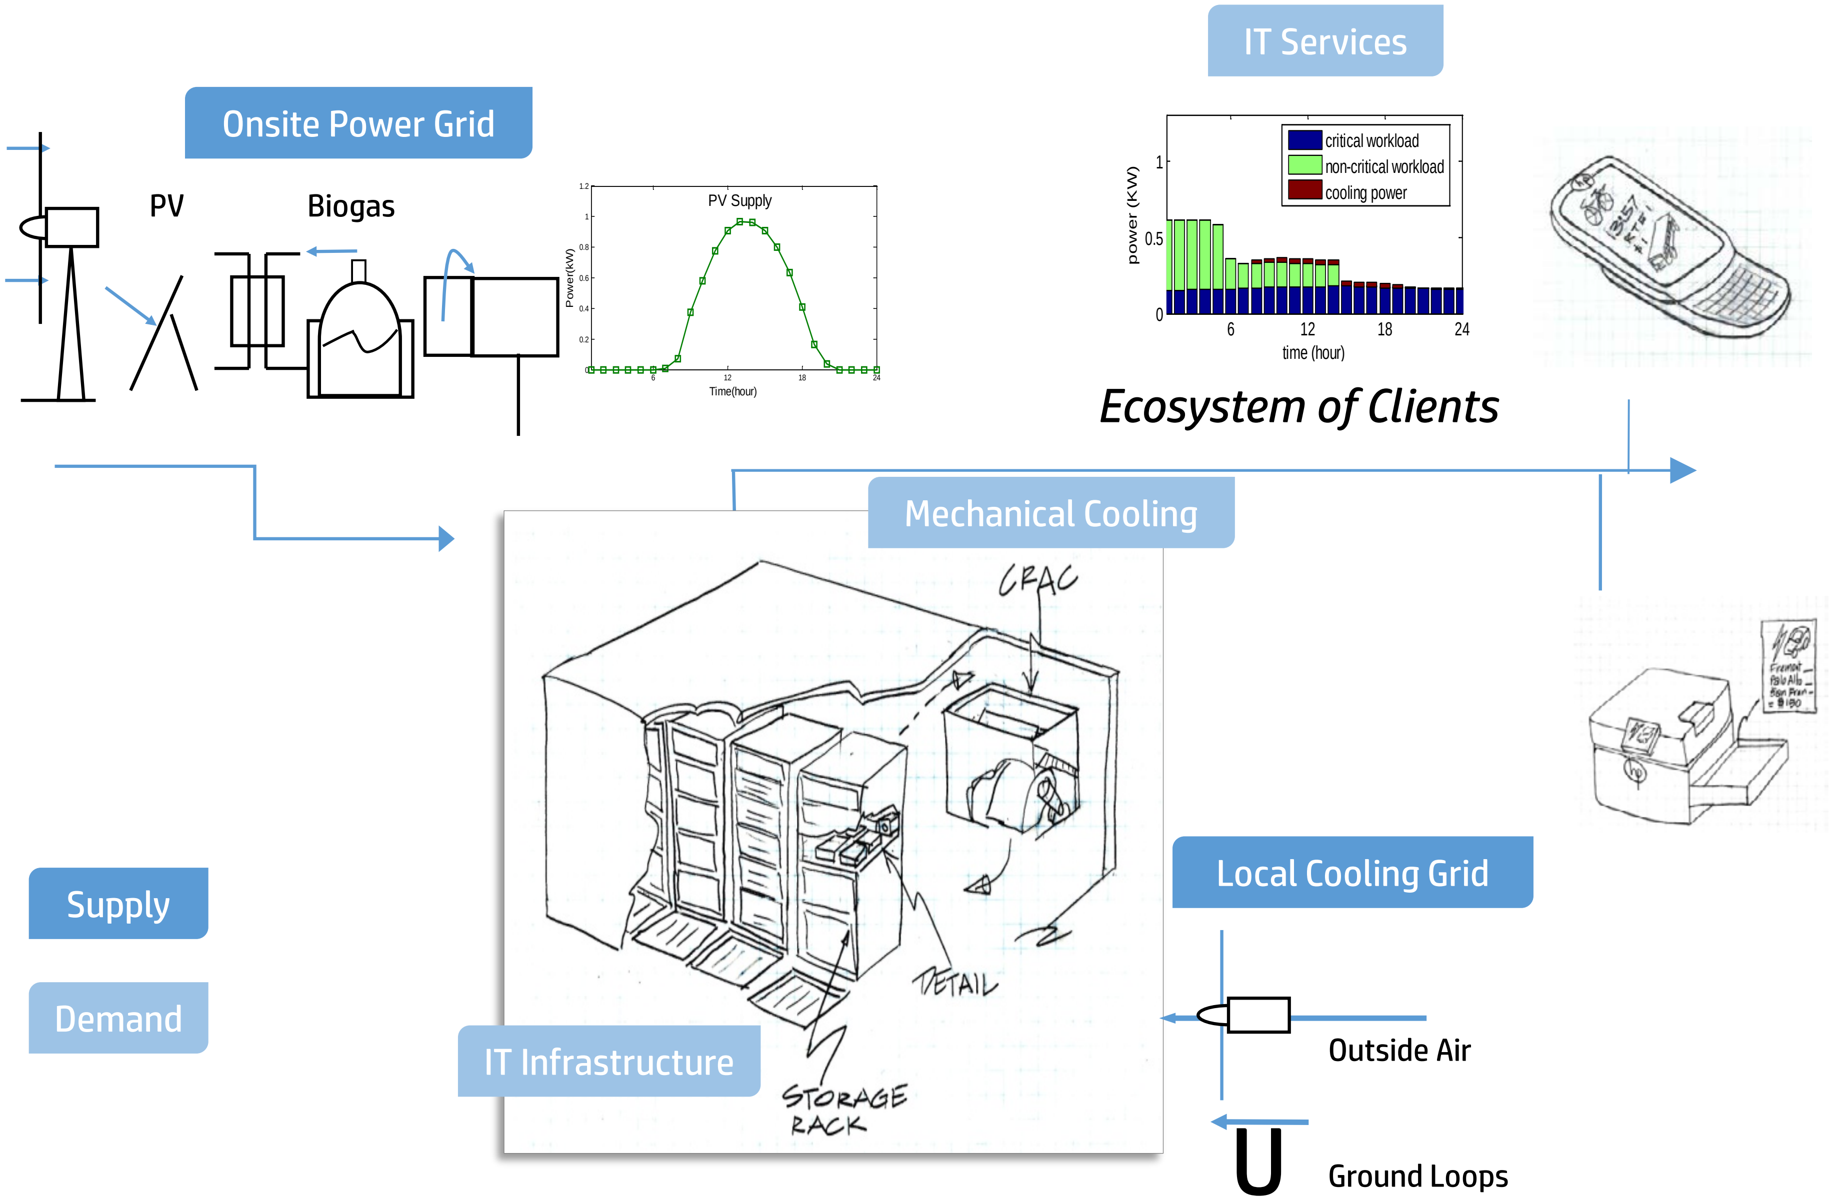

Typical Data Center Energy Use

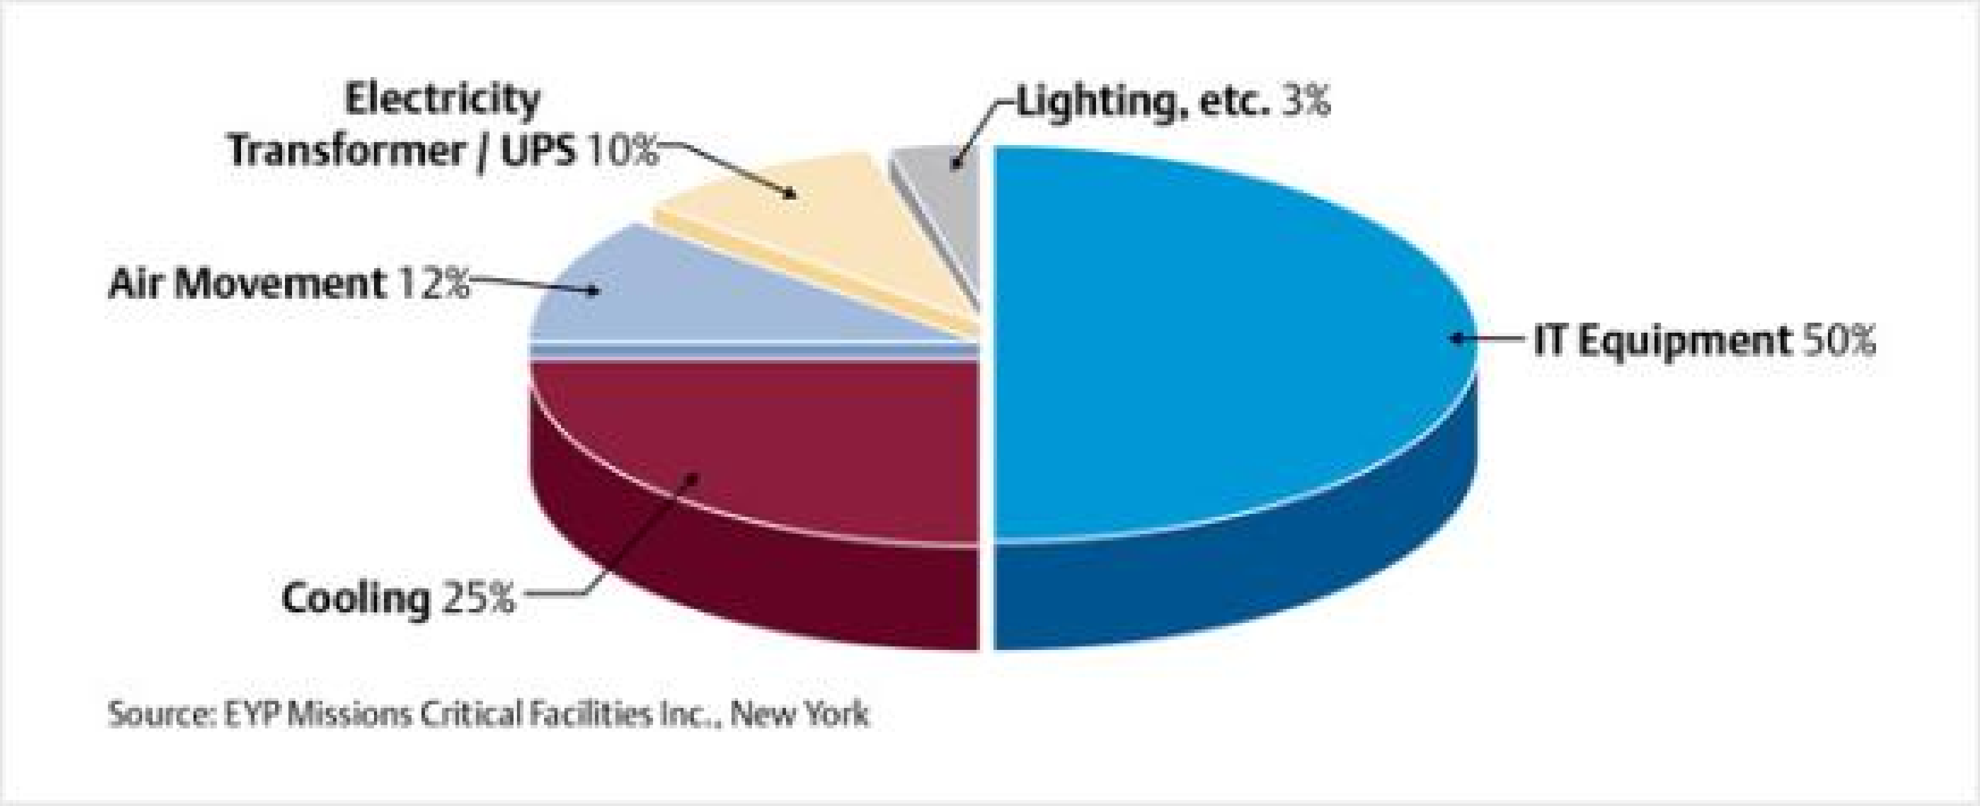

Sustainable Operation of Data Center Chillers – Temporal Data Mining

Data Center Cooling Infrastructure

Consumes from 1/3 up to 1/2 of total power consumption

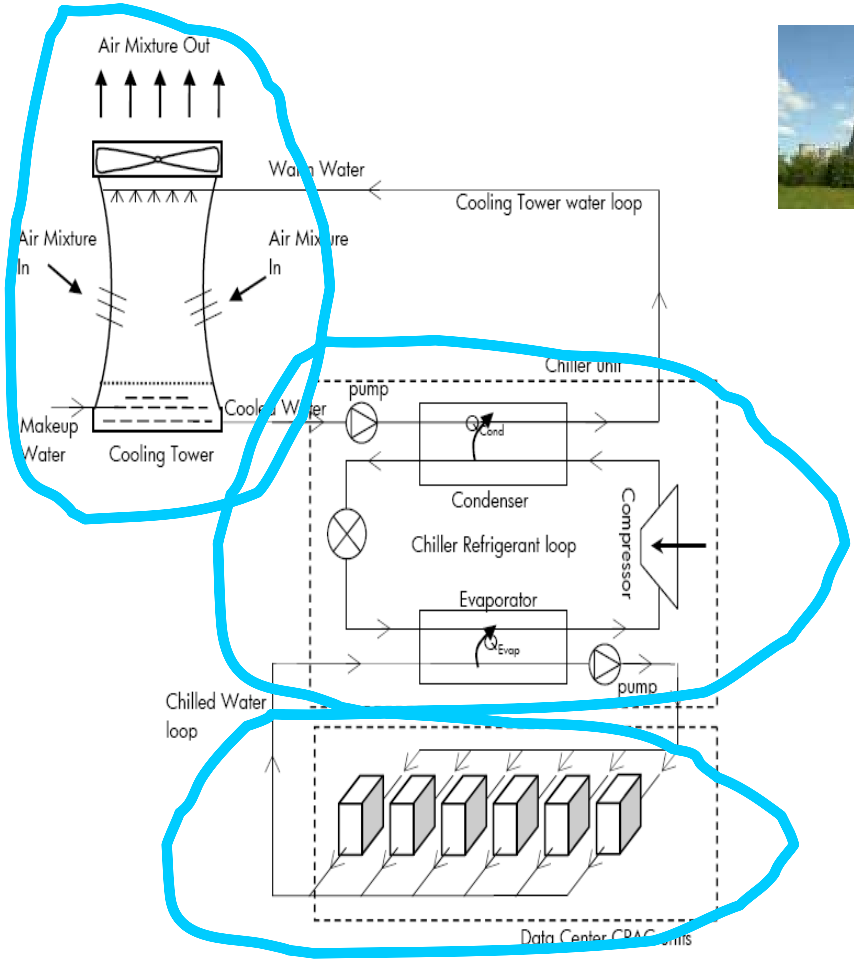

Cooling Towers

Chiller Unit

Water Return (T_{in}) →  → Water Supply (T_{out})

Computer room air-conditioner (CRAC)

Ensemble of Chillers

- Challenging to operate efficiently
  - Complex cyber physical system
    - Dynamic
    - Heterogeneous
    - Inter-dependent
    - Many constraints
  - Accurate models not available
  - Rapid cycles undesirable – reduce lifespan
- Domain experts determine settings based on heuristics
- Can it be automated through a data-driven approach?

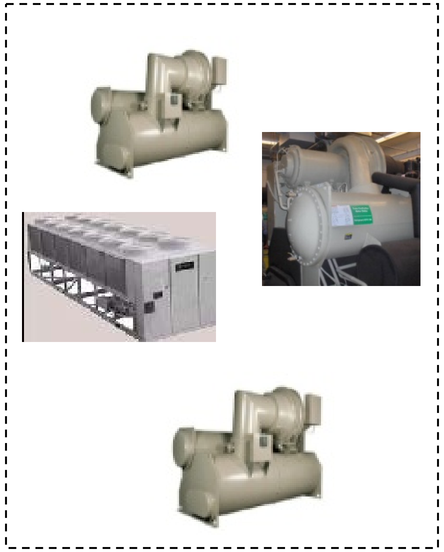

Chiller Ensemble

- Which unit to turn ON/OFF?
- At what utilization?
- How to handle increase/decrease in cooling load?

Problem Statement

- Given the following chiller time series
  - utilization levels
  - power consumption
  - cooling loads
- Is it possible to determine which operational settings are more energy efficient?
- And then use this information to advise data center facility operators

Some Terminology

- IT cooling load
- Chiller utilization
- Chiller power consumption
- Chiller Coefficient of performance (COP)

$$\frac{\text{Cooling Load}}{\text{Power consumption}}$$

Our approach

- Goal: Sustainability characterization of multivariate time series data
  - Chiller utilization data

- Four Main Steps
  - Symbolic representation
  - Event encoding
  - Motif mining
  - Sustainability Characterization

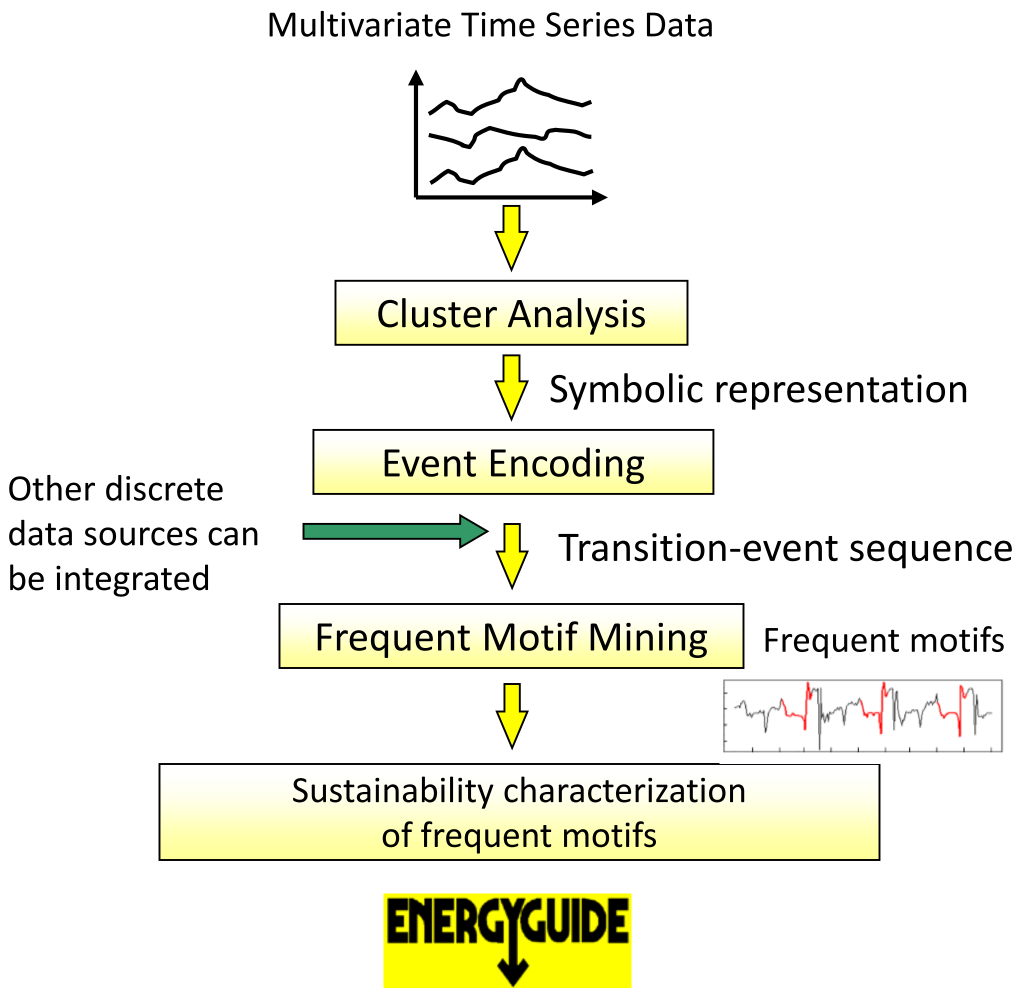

Detour to Clustering (k-means)

[Slides adapted from Sriram Sankararaman]

Supervised vs Unsupervised learning

- Supervised learning: Given $(x_i, y_i), i = 1, \ldots, n$, learn a function $f : X \to Y$.
  - Categorical Y: classification
  - Continuous Y: regression
- Unsupervised learning: Given only $(x_i), i = 1, \ldots, n$, can we infer the underlying structure of X?

Why Unsupervised Learning?

Cluster analysis

Discover groups such that samples within a group are more similar to each other than samples across groups.

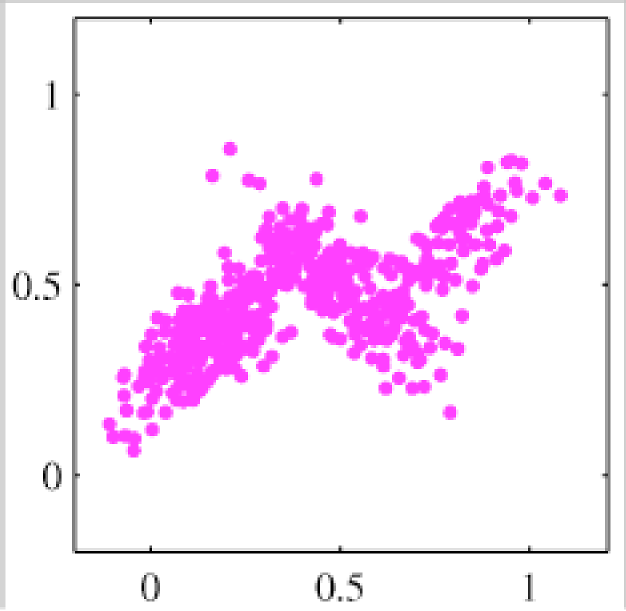

Cluster analysis

Discover groups such that samples within a group are more similar to each other than samples across groups.

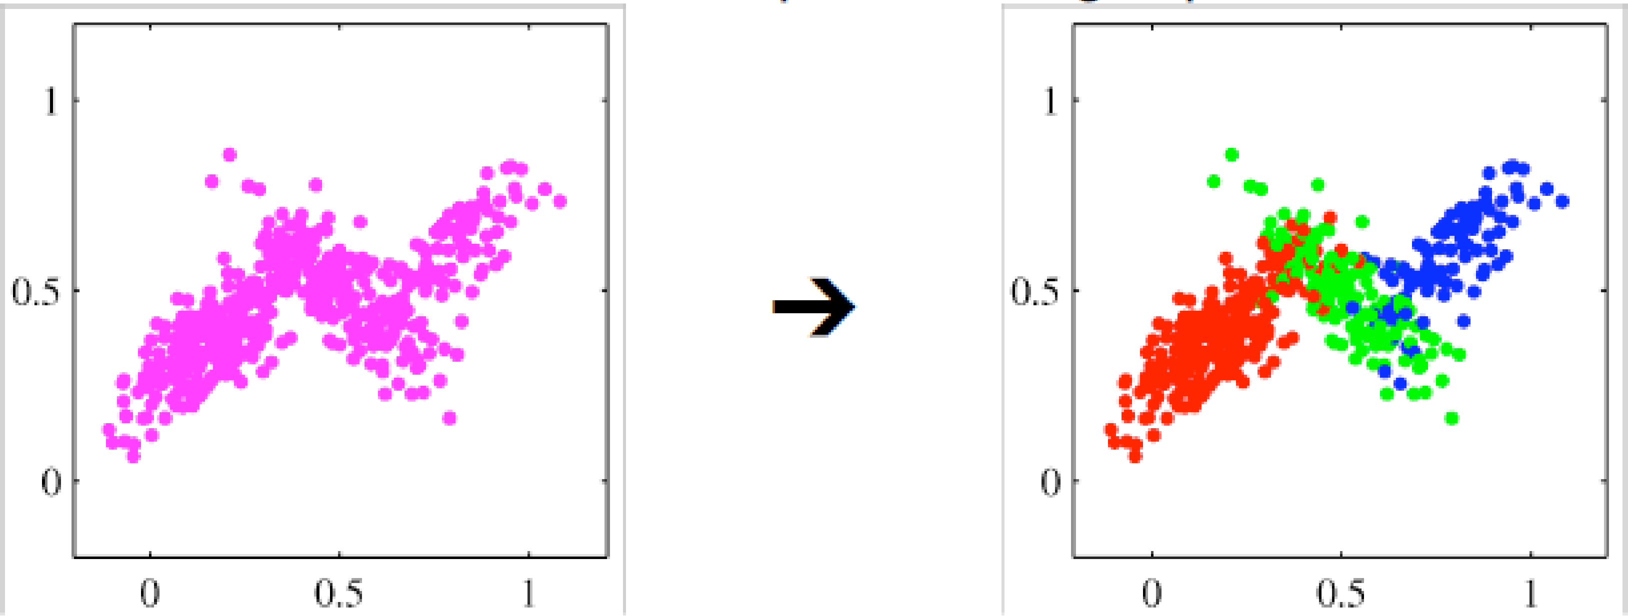

Image Segmentation

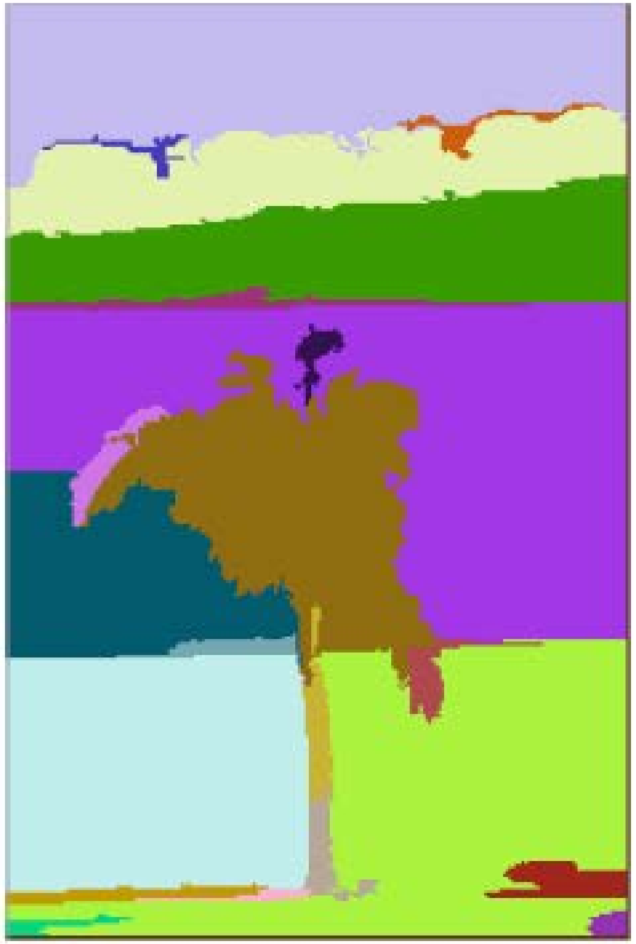

http://people.cs.uchicago.edu/ pff/segment

Ingredients of cluster analysis

- A dissimilarity function between samples.
- A loss function to evaluate clusters.
- Algorithm that optimizes this loss function.

k-means

- Partition data set into k parts
  - k is an input
- Each part is represented / summarized by its mean (centroid)
- Distance function: Euclidean
- Objective: minimize distance of points from their partition mean
  - SSE
- Chicken and egg problem
  - If know partition means → can assign points
  - If know partitions → can find means

K-means

- Iterative procedure
- Pick initial k means randomly
- Two steps in each iteration

- Assignment Step
  - Assign points to means
- Update Step
  - Compute new means

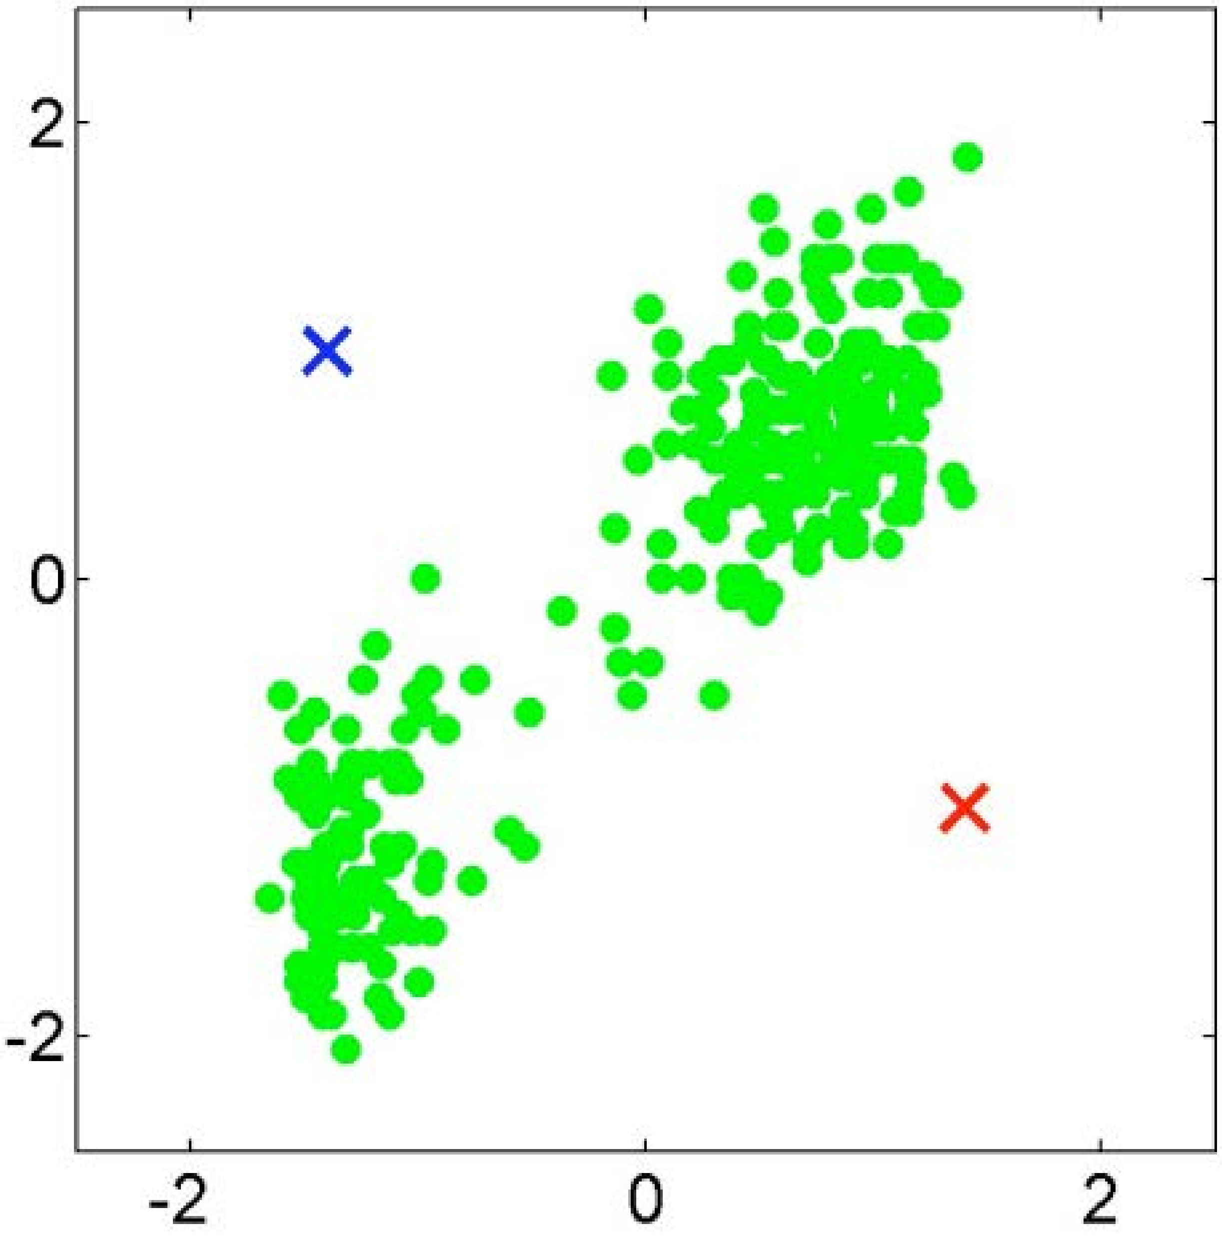

2
0
-2
-2
0
2

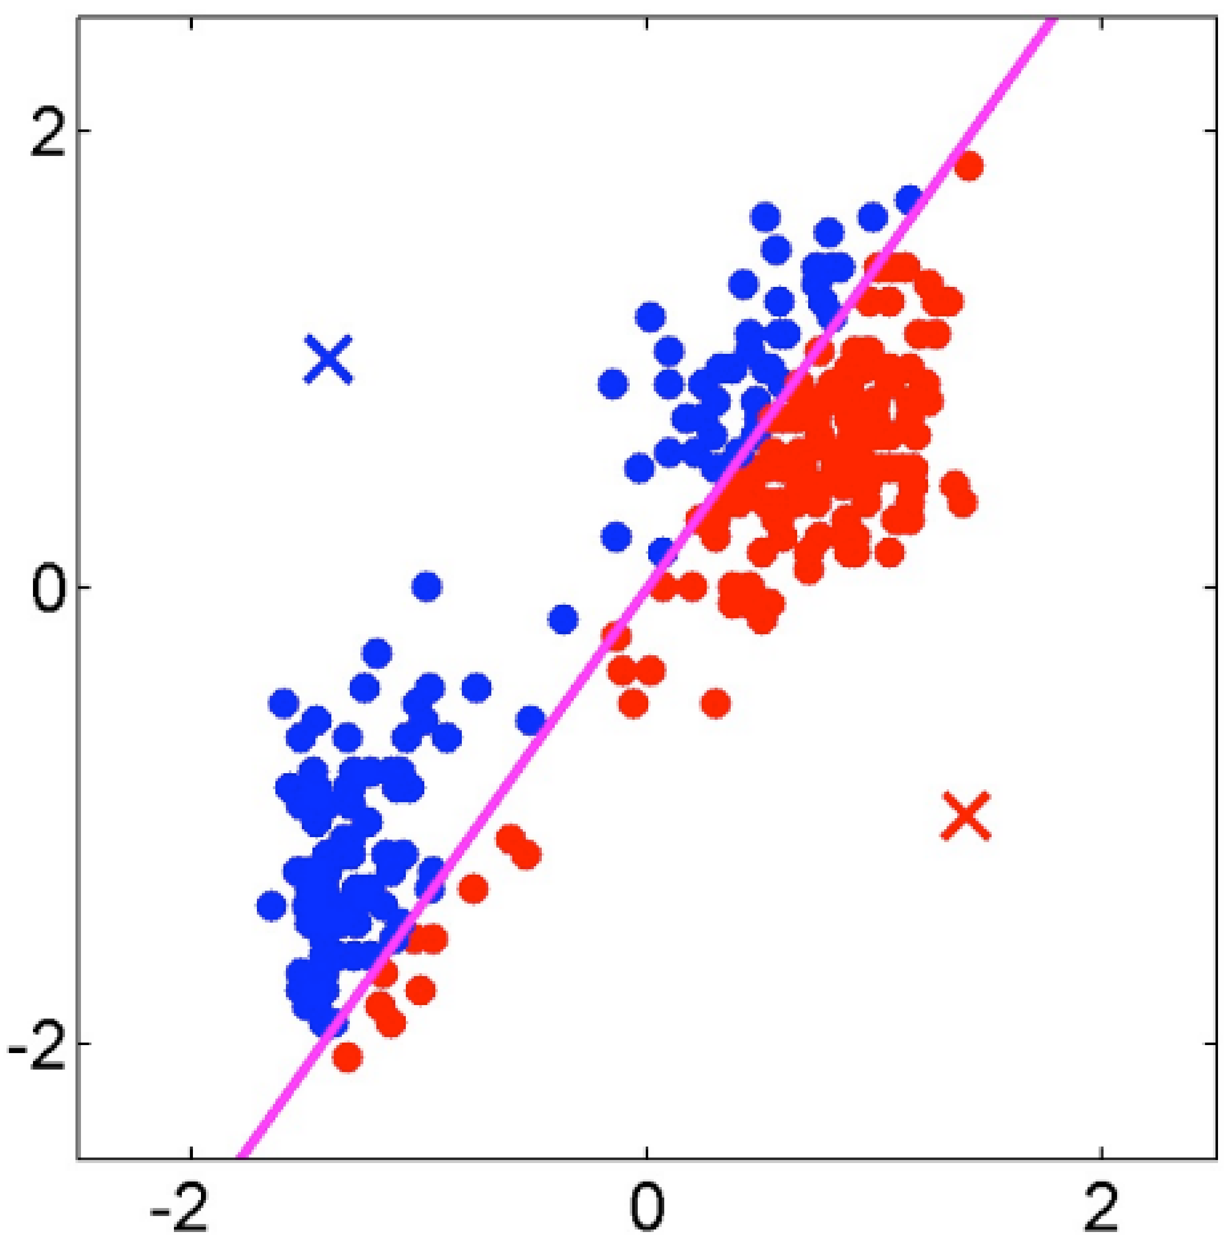

2
0
-2
-2
0
2

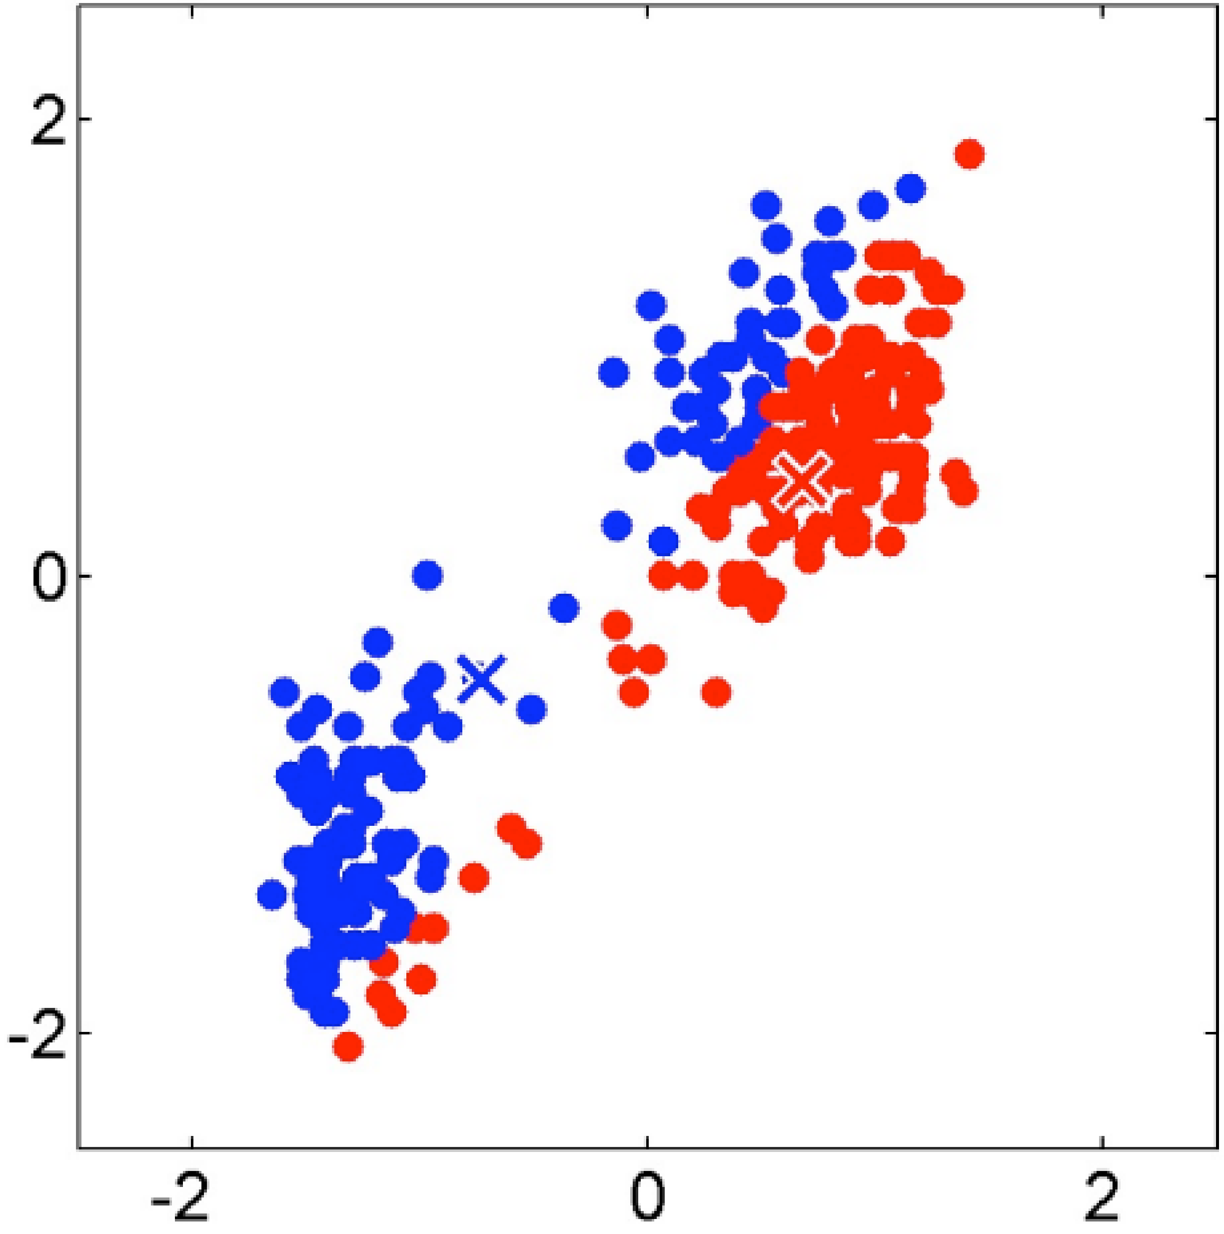

2
0
-2
-2
0
2

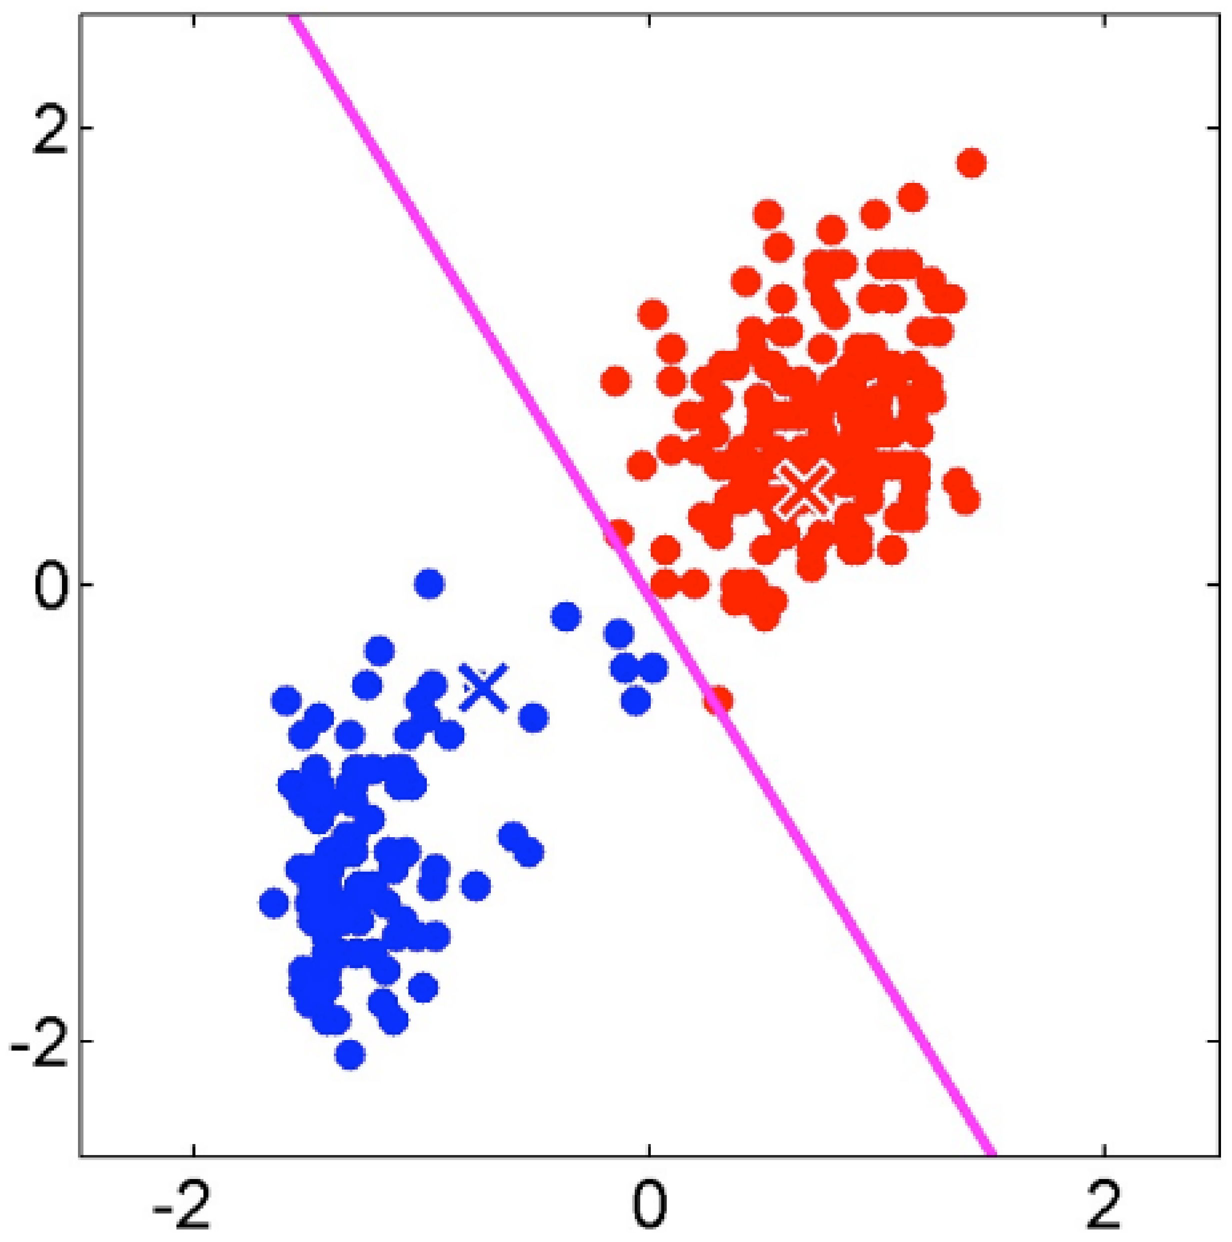

2
0
-2
-2
0
2

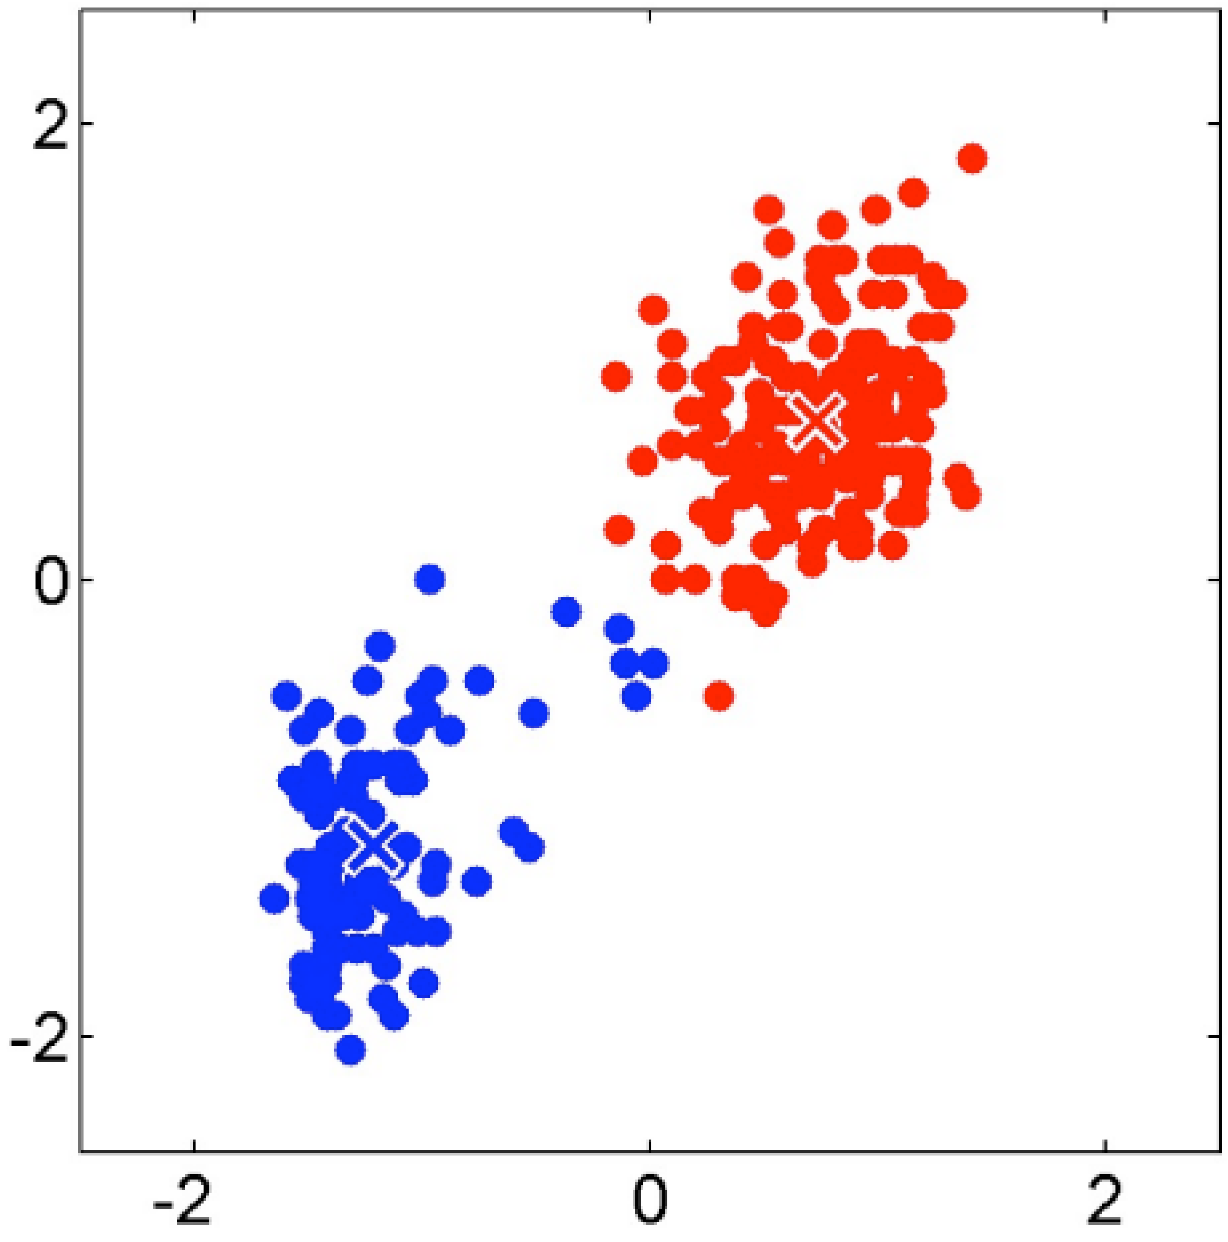

2
0
-2
-2
0
2

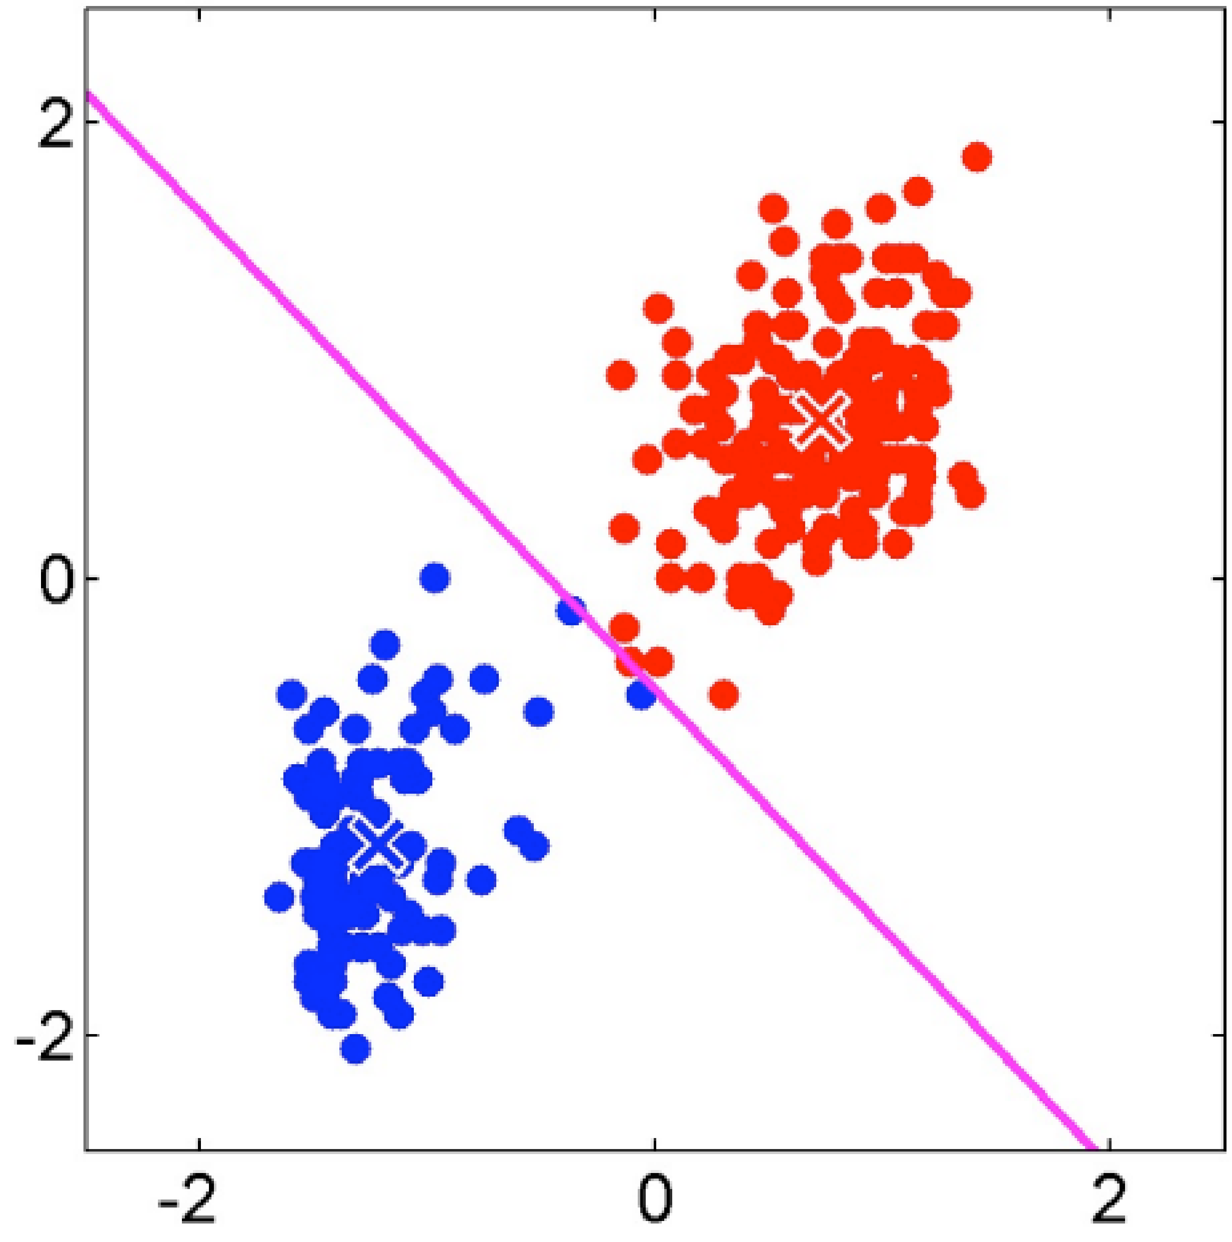
2
0
-2
-2
0
2

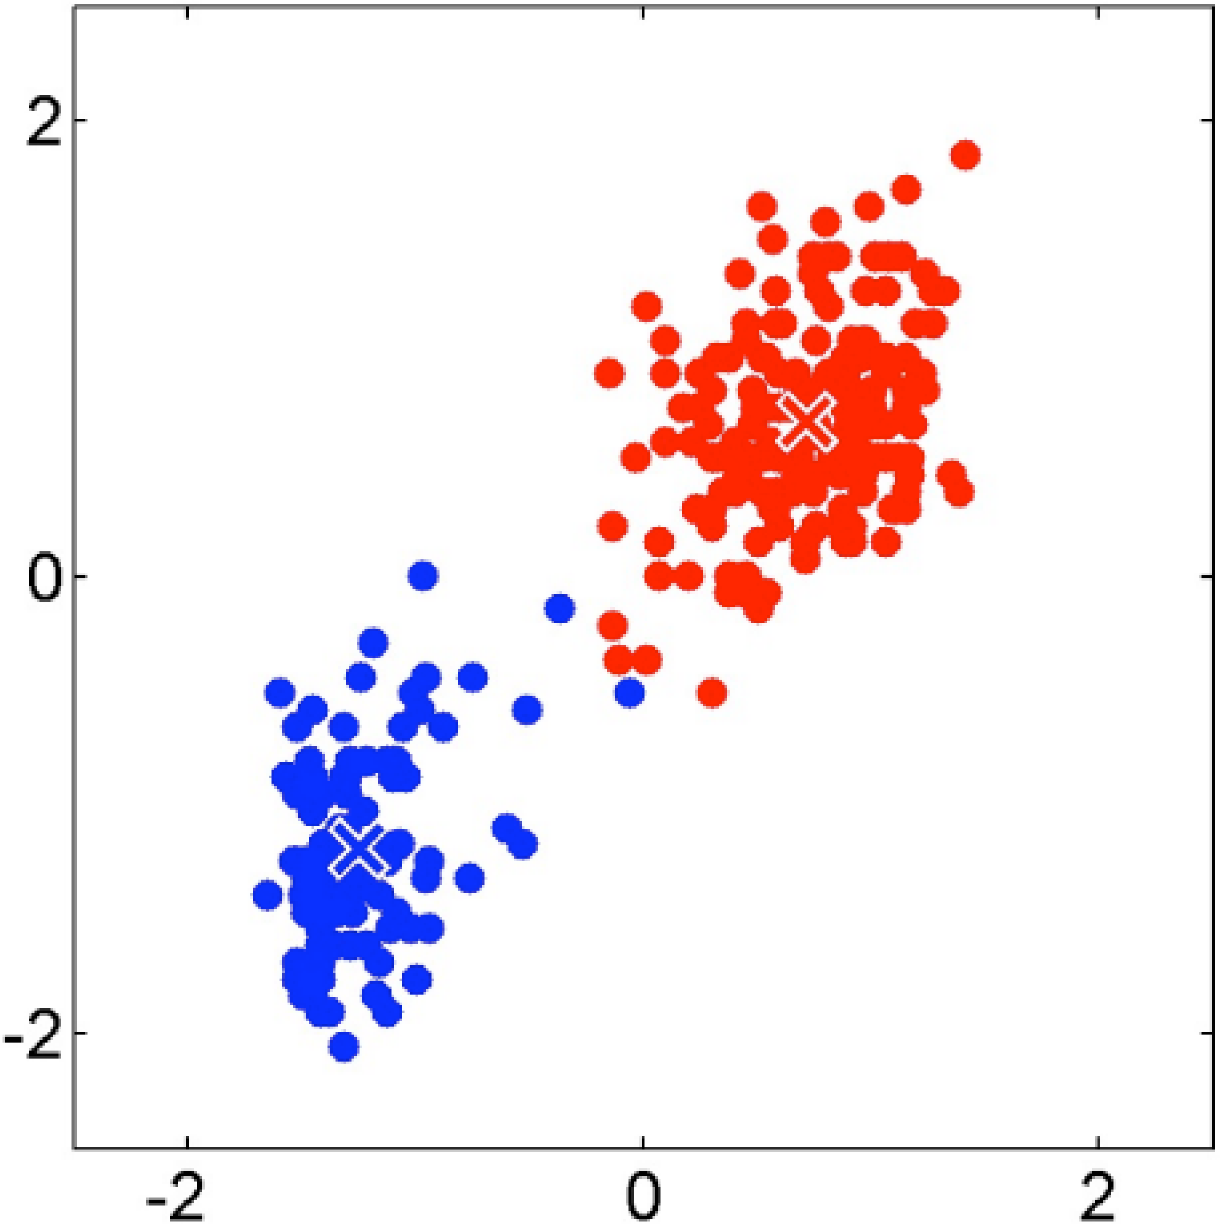
2
0
-2
-2
0
2

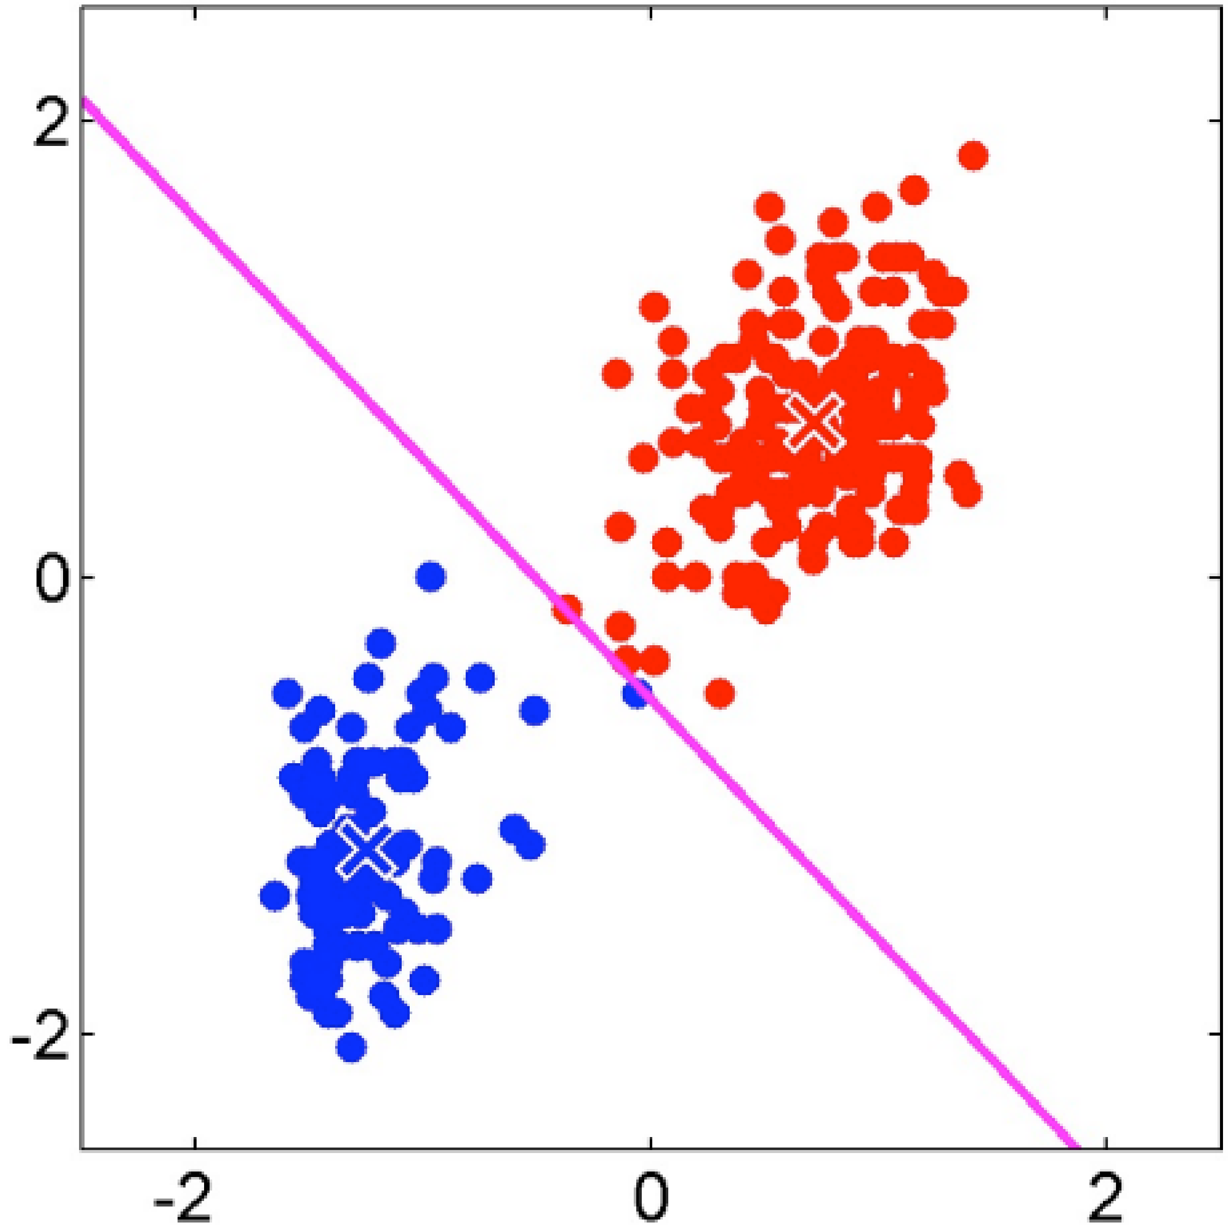

2
0
-2
-2
0
2

Our approach

- Goal: Sustainability characterization of multivariate time series data
  - Chiller utilization data
- Four Main Steps
  - Symbolic representation
  - Event encoding
  - Motif mining
  - Sustainability Characterization

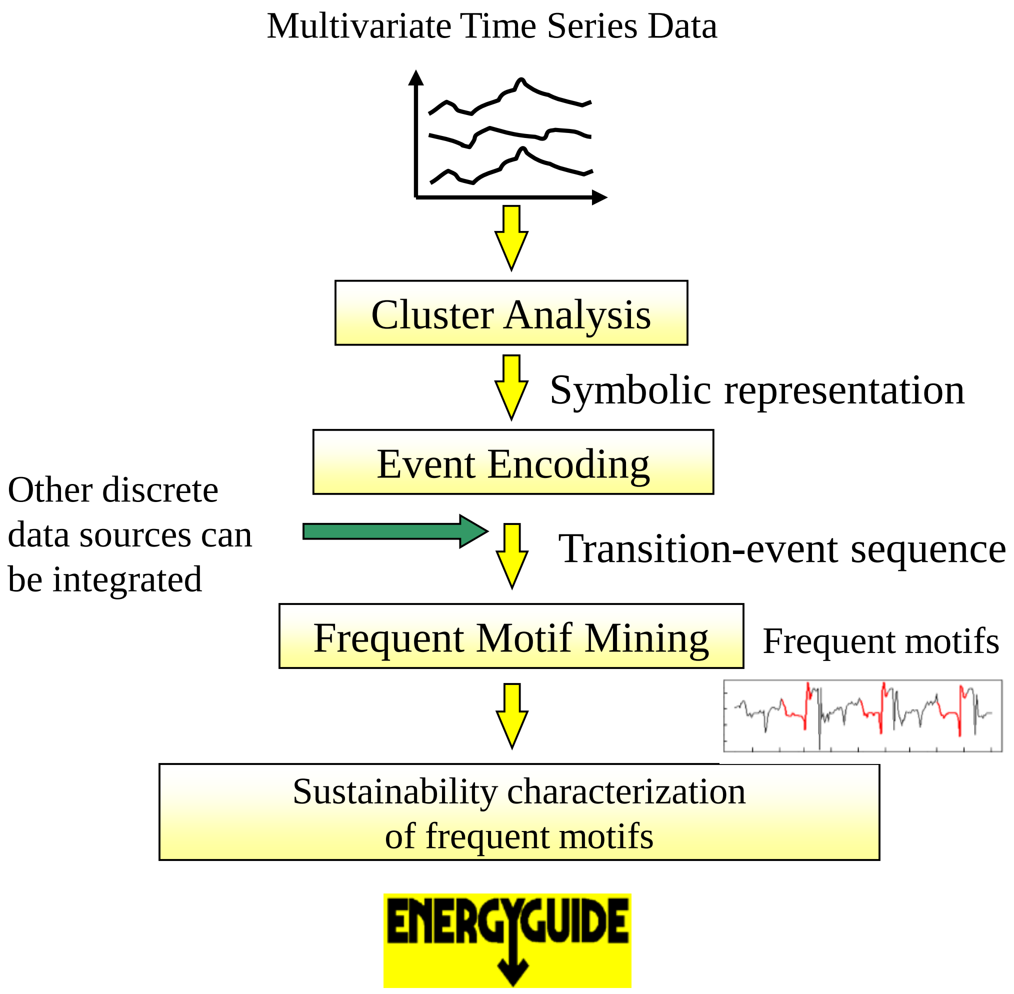

Clustering

- Individual vector: Utilization across all chiller units
- Raw Data: Sequence of such vectors
- Perform k-means clustering
- Use cluster labels to encode multi-variate time series

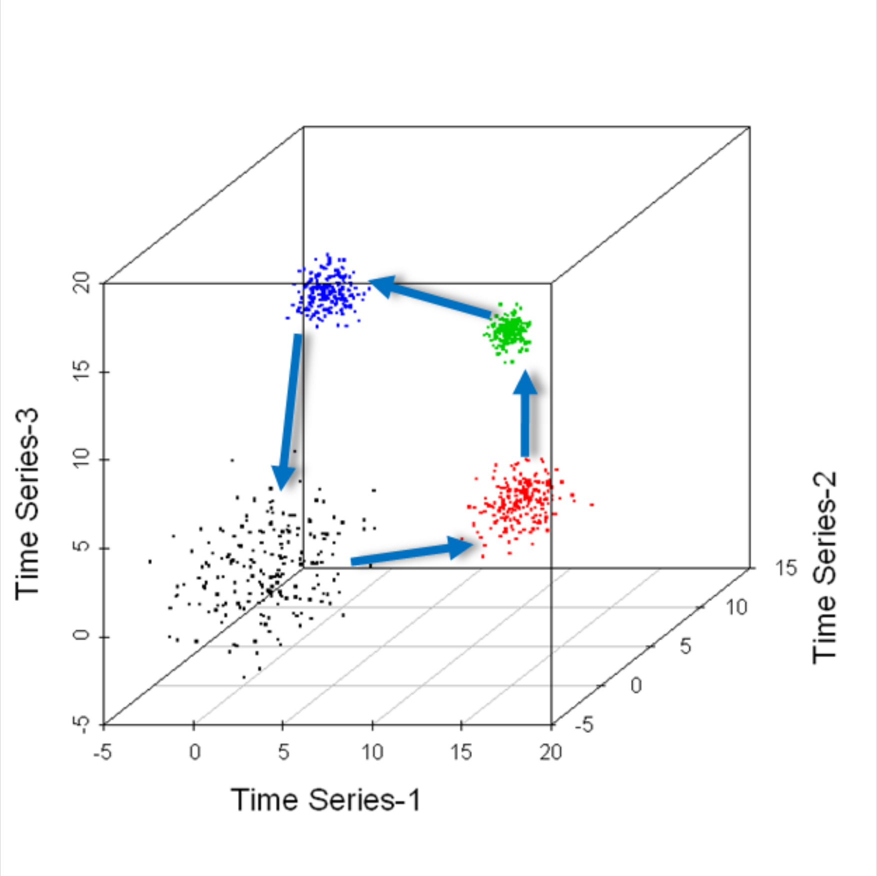

Redescribing time series data

- Perform run-length encoding:
  - Note transitions from one symbol to another
- Higher level of abstraction
  - Transition events

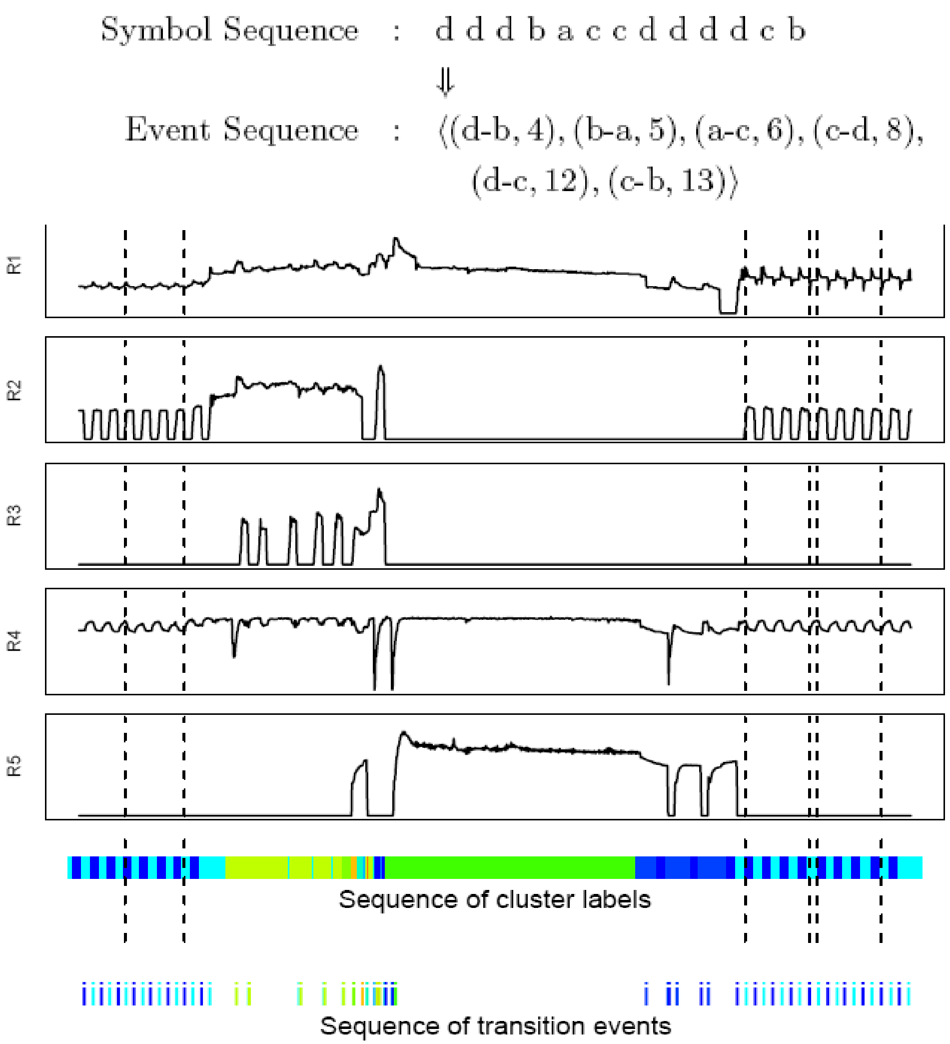

Methodology Summary

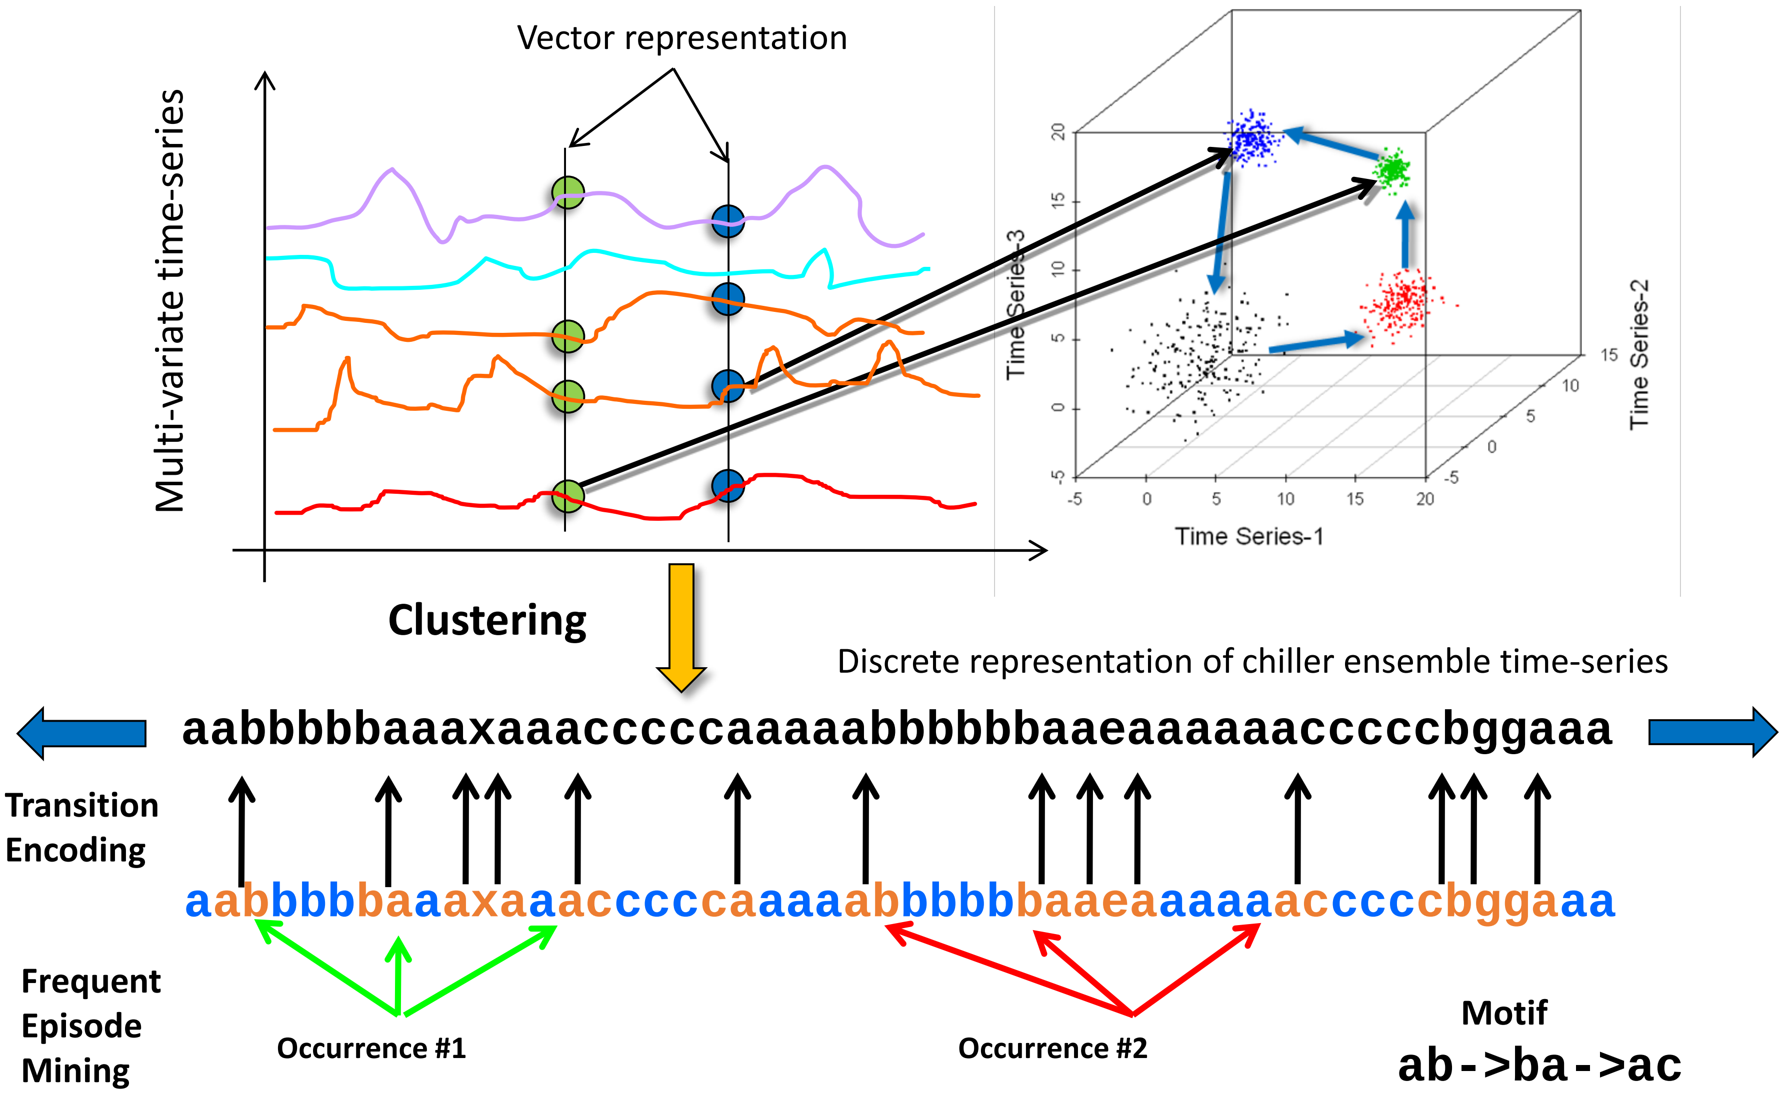

Sustainability characterization of Motifs

- Average motif COP (coefficient of performance)
  - Indicates cooling efficiency of a chiller unit
    - $\text{COP} = \dfrac{\text{IT Cooling Load}}{\text{Power consumed}}$

Experimental Results

- Data
  - From HP R&D data center in Bangalore
    - 70,000 sq ft
    - 2000 racks of IT equipments
  - Ensemble of five chiller units
    - 3 air cooled chillers
    - 2 water cooled chillers
  - 480 hours of data
- 22 motifs found in the data

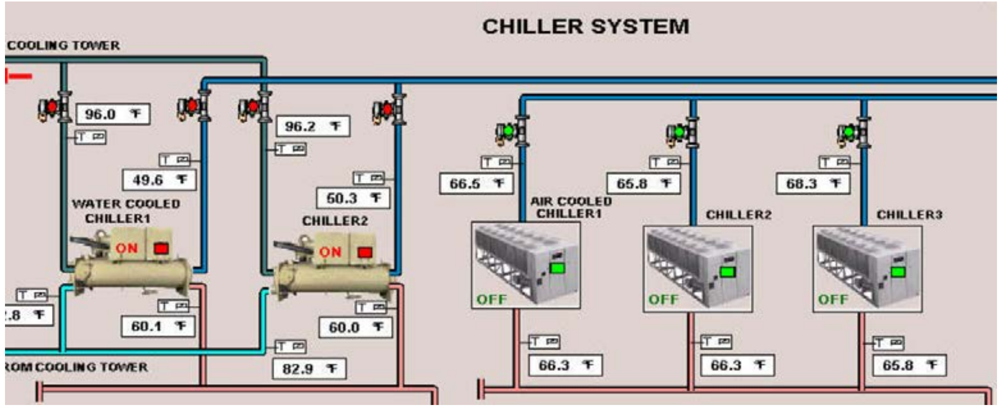

Two Interesting Motifs

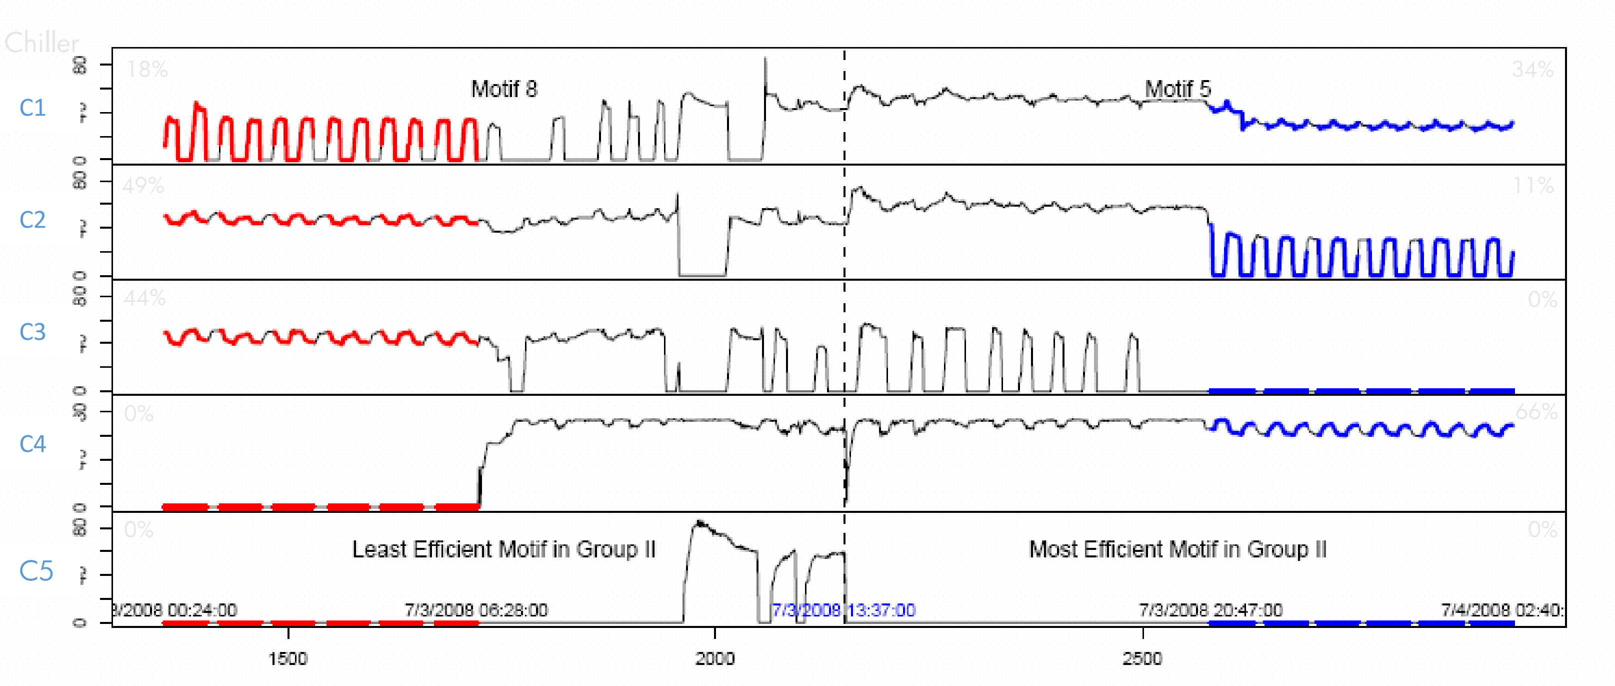

C1, C2, C3 → Air cooled

C4, C5 → Water cooled

Motif 8

Motif 5

	Motif 8	Motif 5
COP	4.87	5.40
Units operating	3 air-cooled	2 air-cooled, 1 water cooled

38

Potential Savings

	Load (KW)		Most Efficient Motif	Least Efficient Motif	Potential Power Savings	
	Ave.	Std			KW	%
Group II	2089	35	5	8	41	9.83%

- Annual saving from operating in Motif 5 instead of Motif 8
  - Cost savings = \$40,000 (~10%)
  - Carbon footprint savings = 287,328 kg of CO_2

Thermal Anomaly Detection – An Example

Thermal Anomalies

- Temperature in a DC needs to be monitored
  - Too cold → waste of energy
  - Too hot → equipment failures
- Racks instrumented with temperature sensors
- How to detect anomalies?
  - Threshold
  - Models
  - …

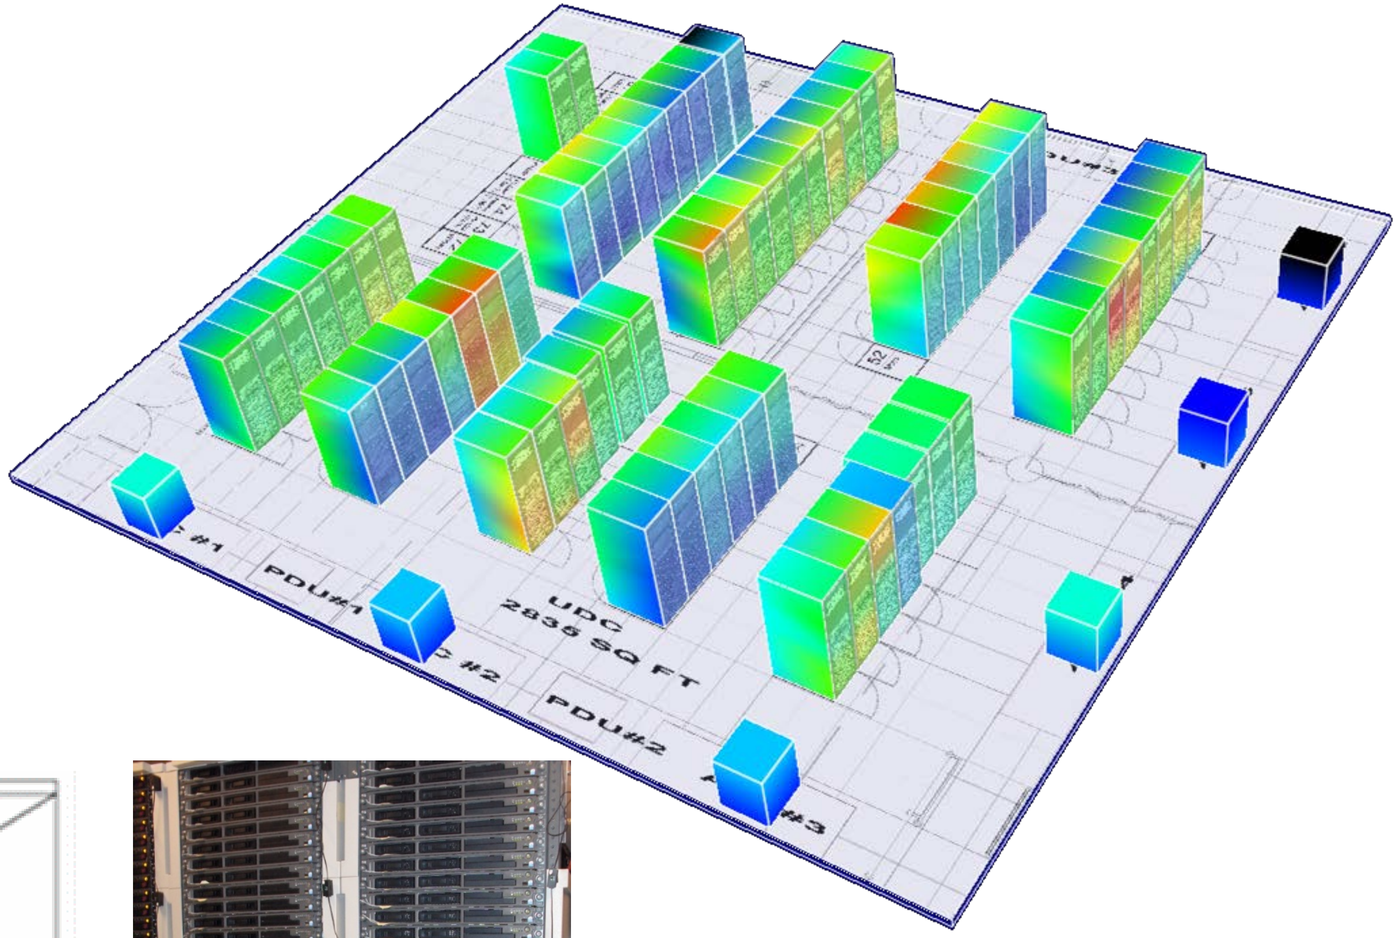

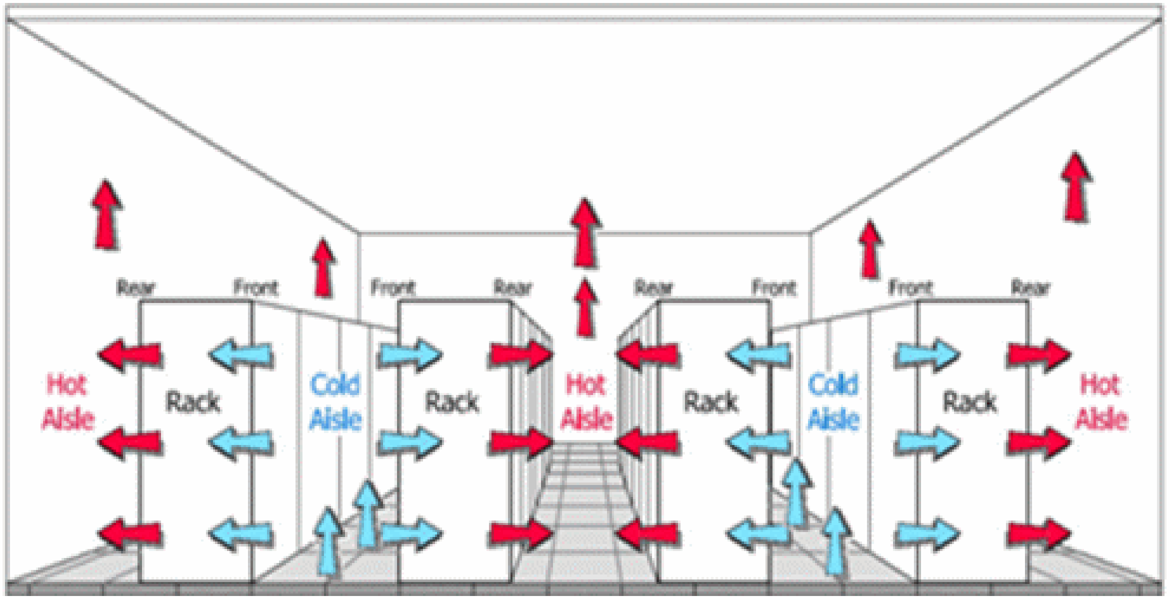

Detour – Principal Component Analysis (PCA)

Some slides from Melanie Mitchell

Linear Dimensionality Reduction

High dimensional point $(X_1, X_2, \ldots X_{400})$

Linear Transformation or Projection

Low dimensional point $(Z_1, Z_2, \ldots Z_{20})$

Data

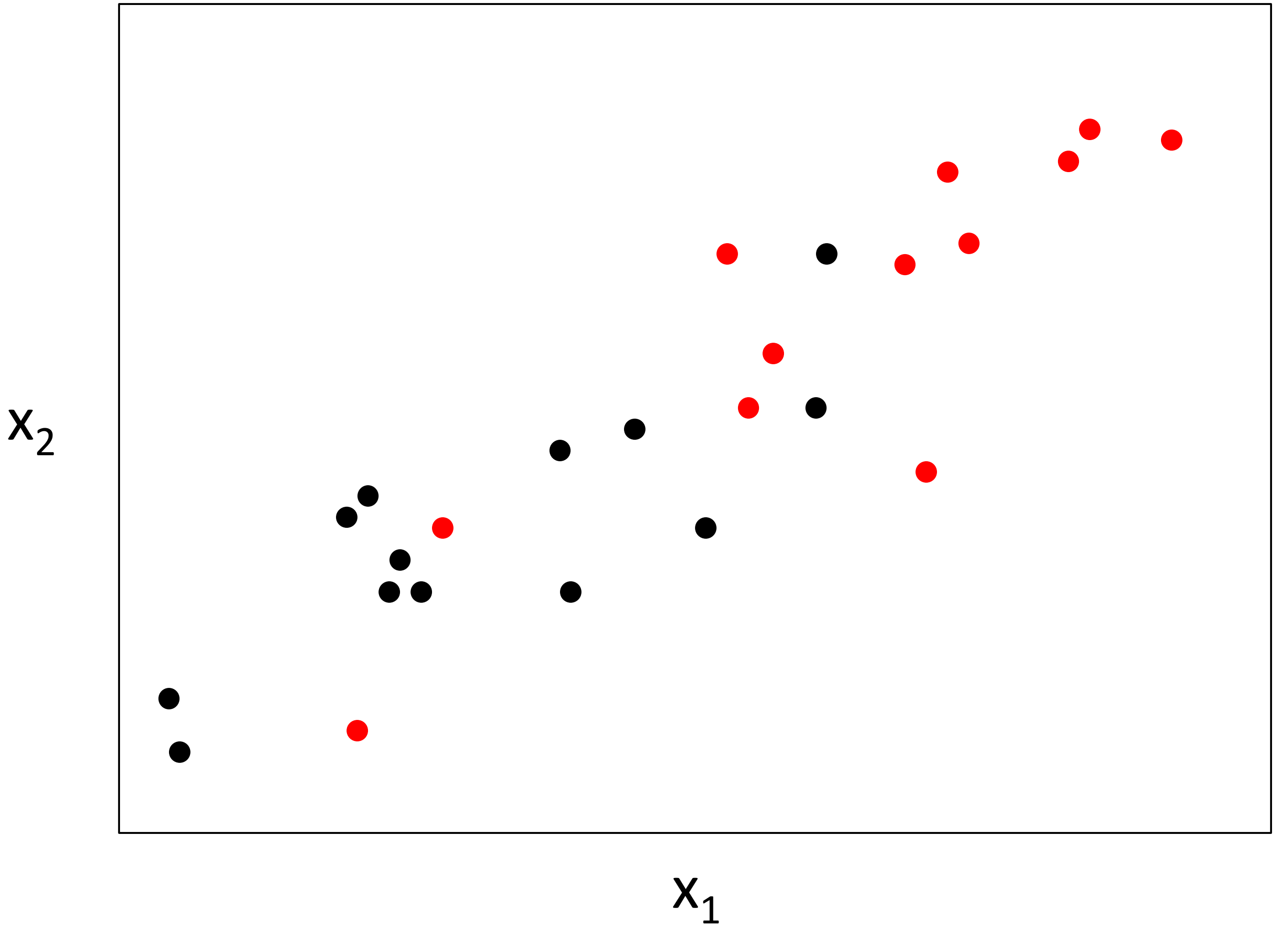

Data

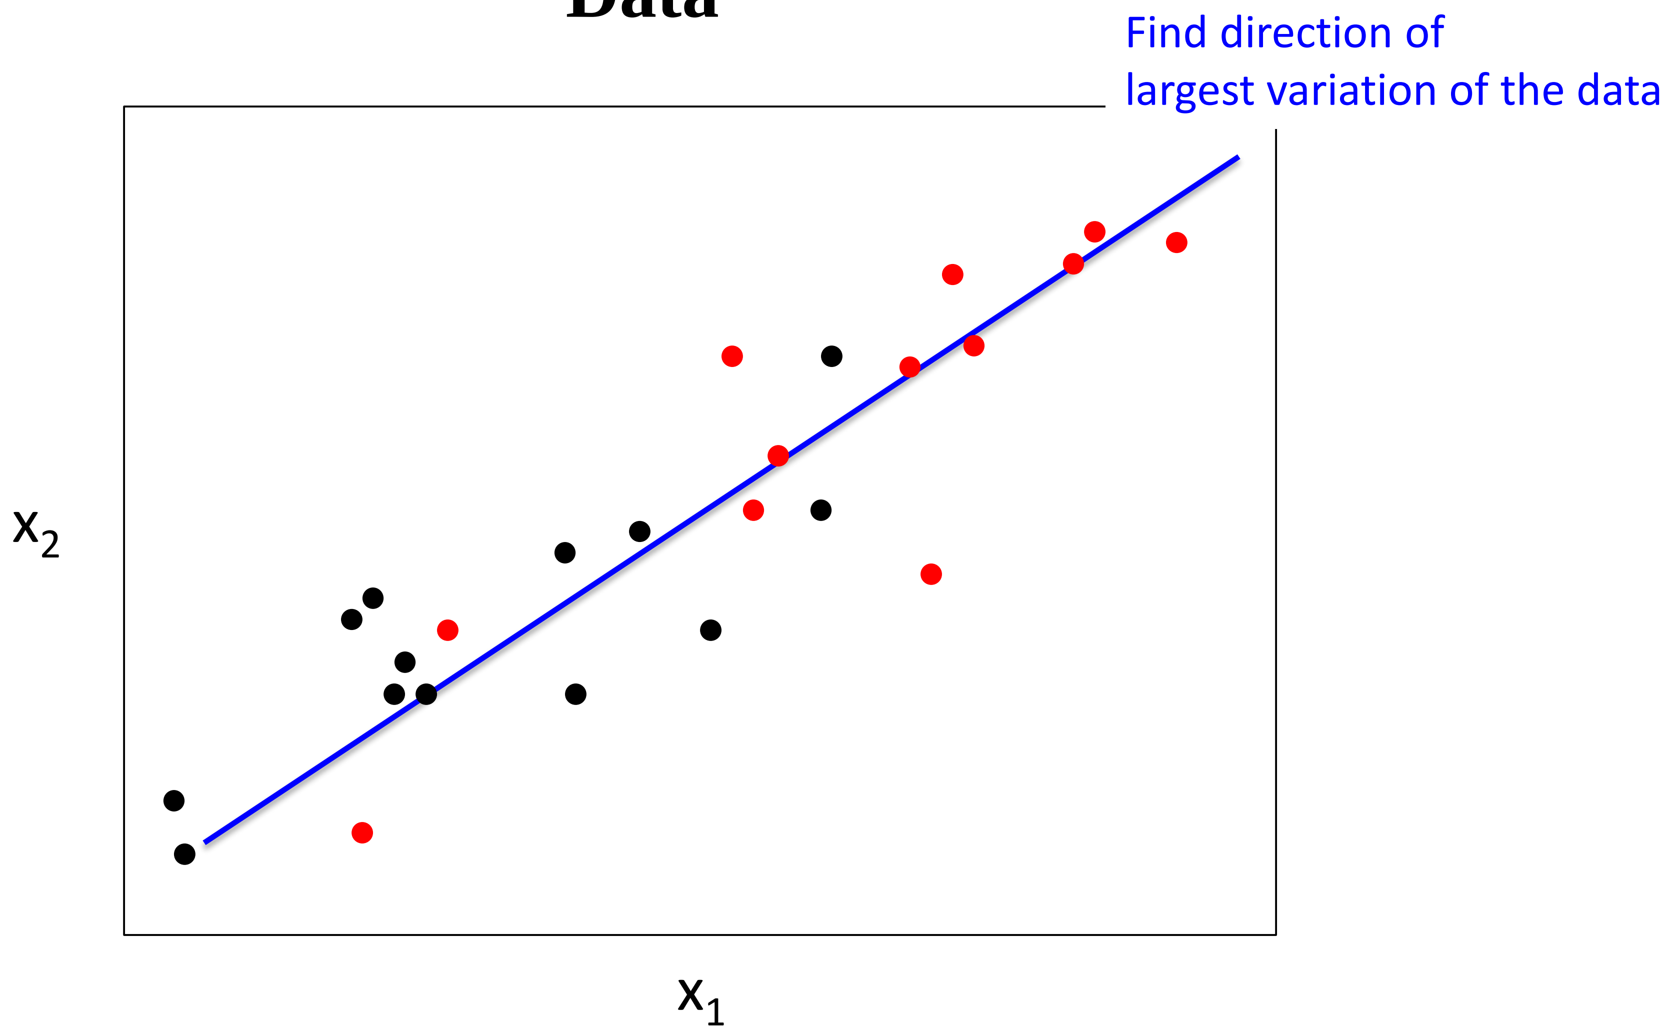

Data

First principal component: direction of largest variation of the data

Second principal component

Gives direction of second largest variation

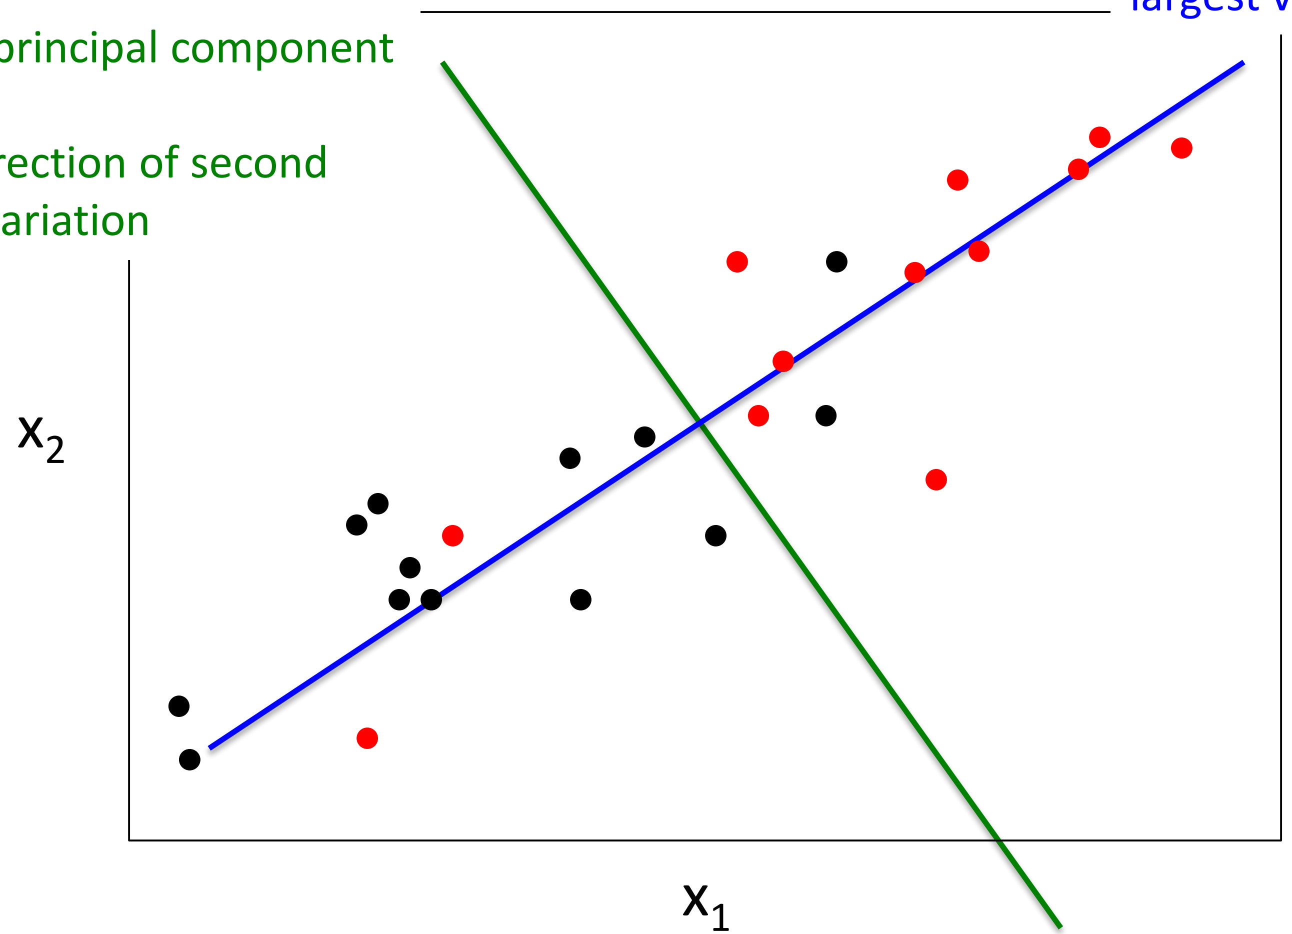

Rotation of Axes

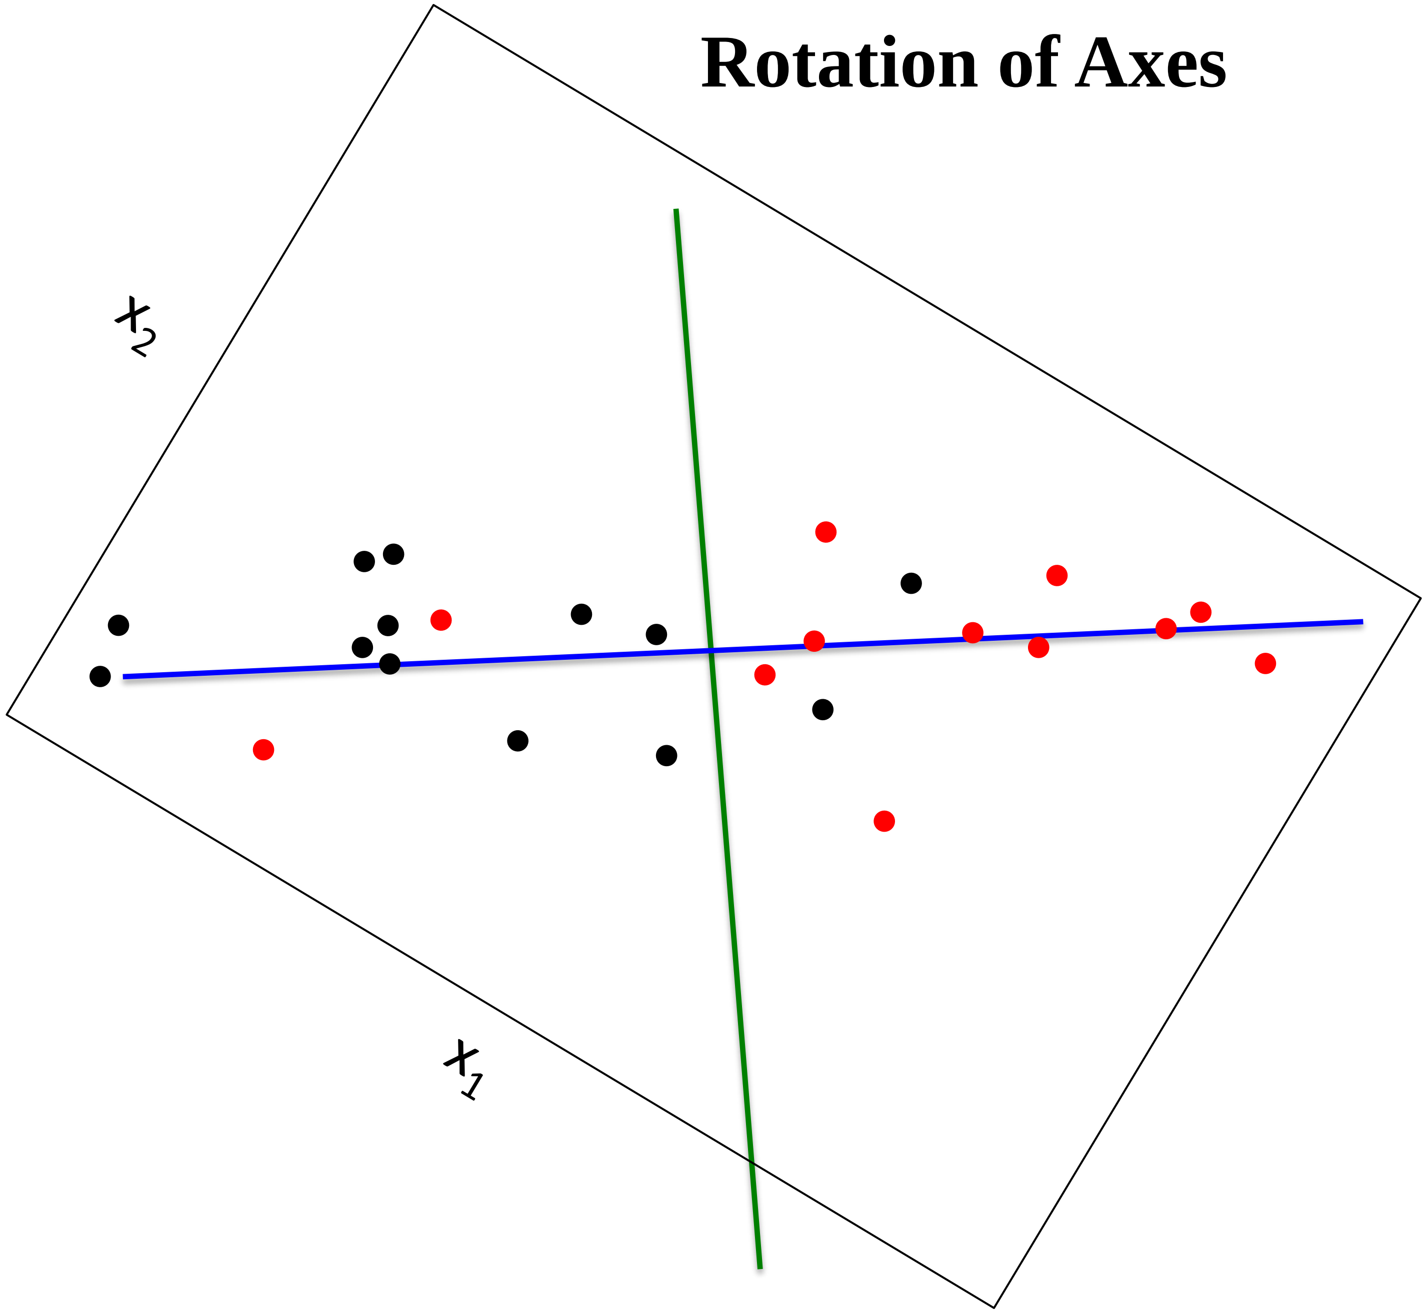

Dimensionality reduction

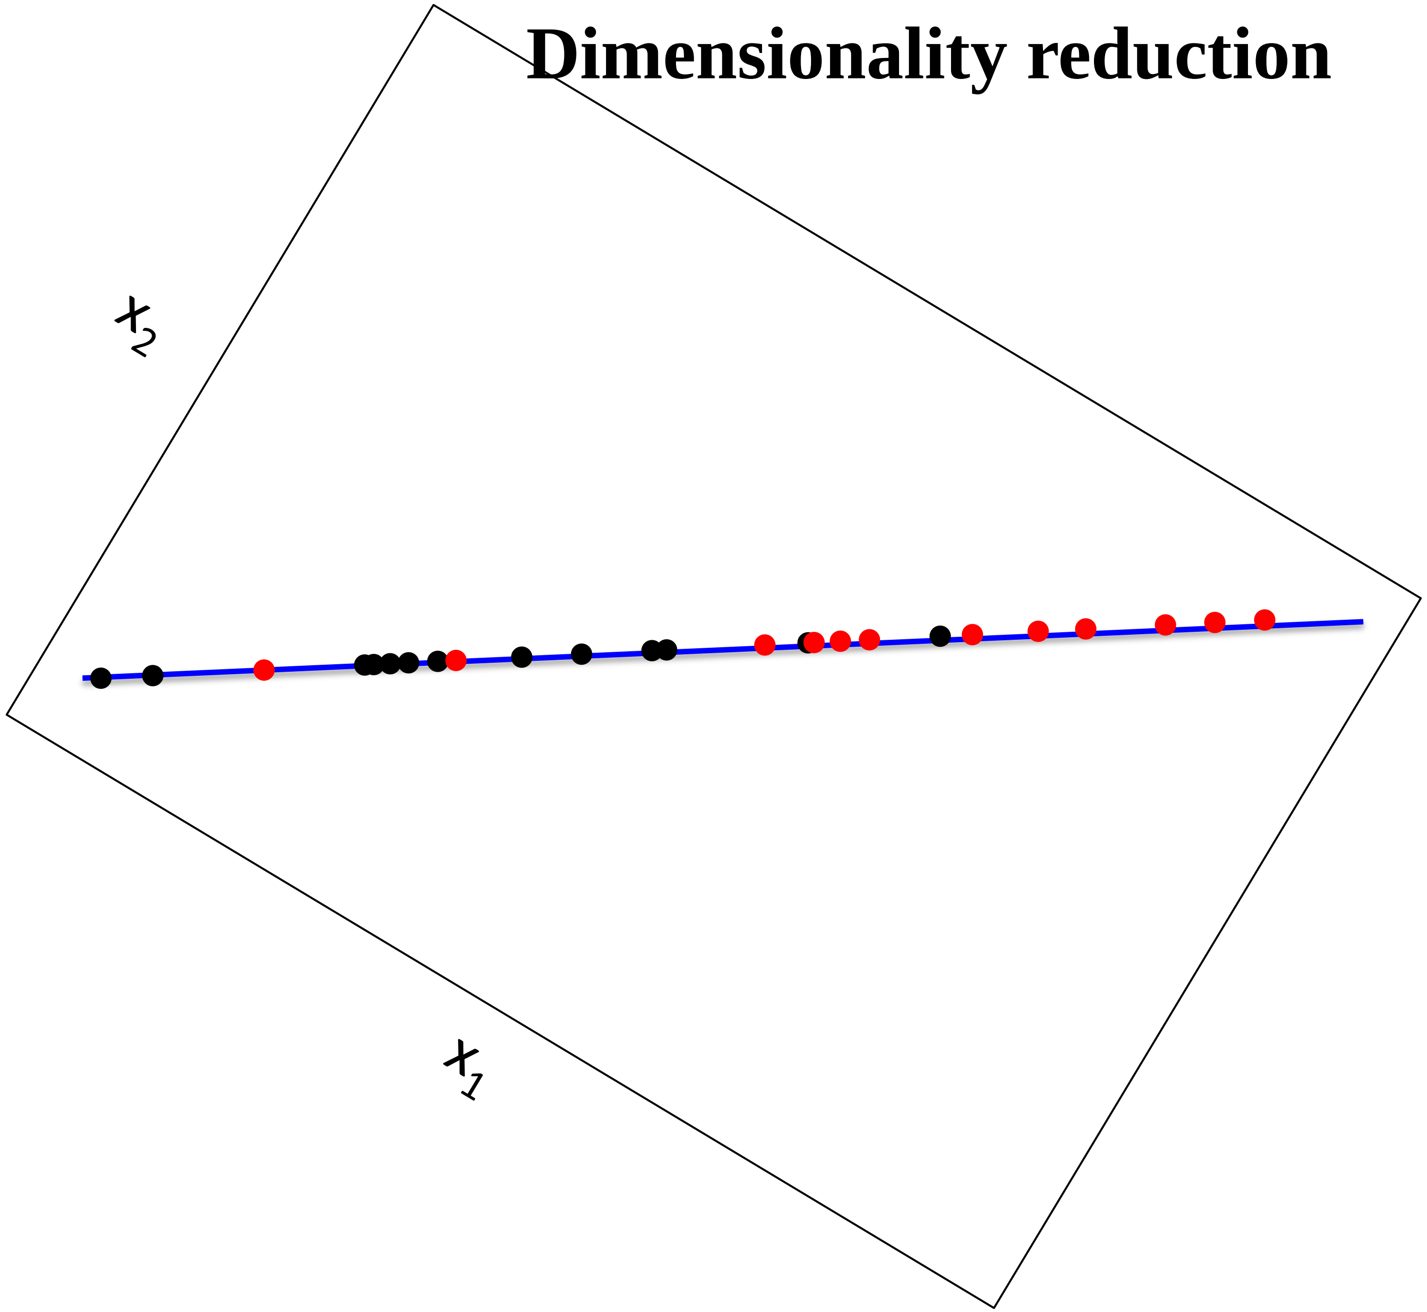

Anomalous Thermal Behavior Detection using PCA

- Example: Event (Anomaly) Detection

Period of increased energy consumption (17 % increase)

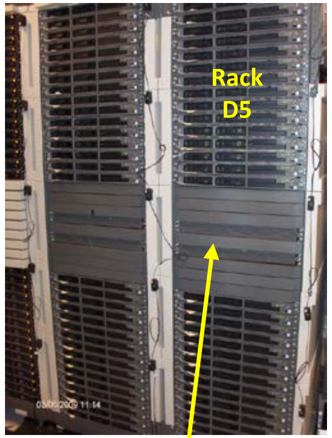

Network Switch

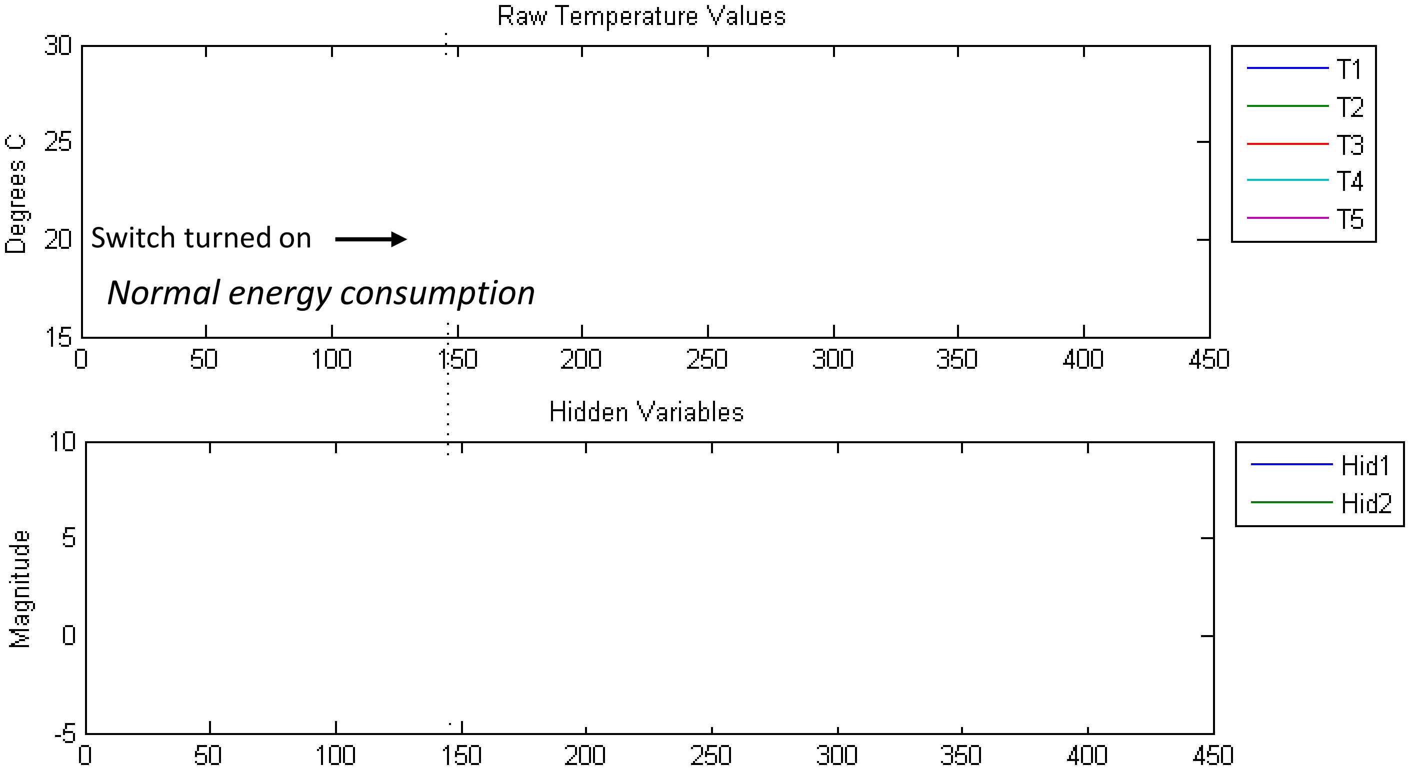

Start: 2009-09-28 16:44:34

End: 2009-09-28 23:58:34

[SensorKDD 2010]

PV Generation Forecasting

Fine grained PV Prediction using Bayesian Ensemble

- **Motivation**
  - Integration of renewable sources is an important goal of the smart grid effort
  - PV output is variable and intermittent
  - Knowledge of future PV output enables demand-side management and “shaping” in data centers

- **Problem addressed**
  - Predict PV output for the next day

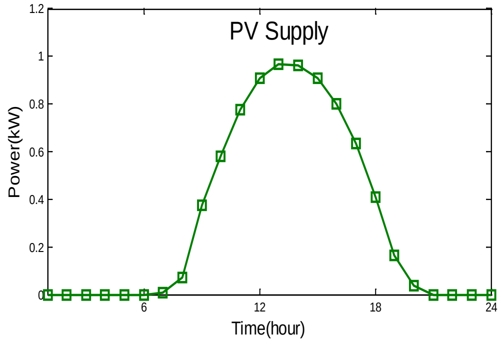

- **Data**
  - Historical PV output data for about 9 months from the HPL Palo Alto site
  - Weather data

[AAAI 2012]

Fine grained PV Prediction using Bayesian Ensemble

- **Approach**
  - Extract daily profiles from training data
  - Use ensemble of predictors
    - Naïve Bayes
    - K-NN
    - Motif based
  - Perform Bayesian model averaging

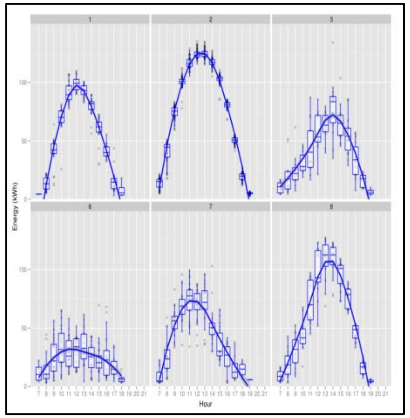

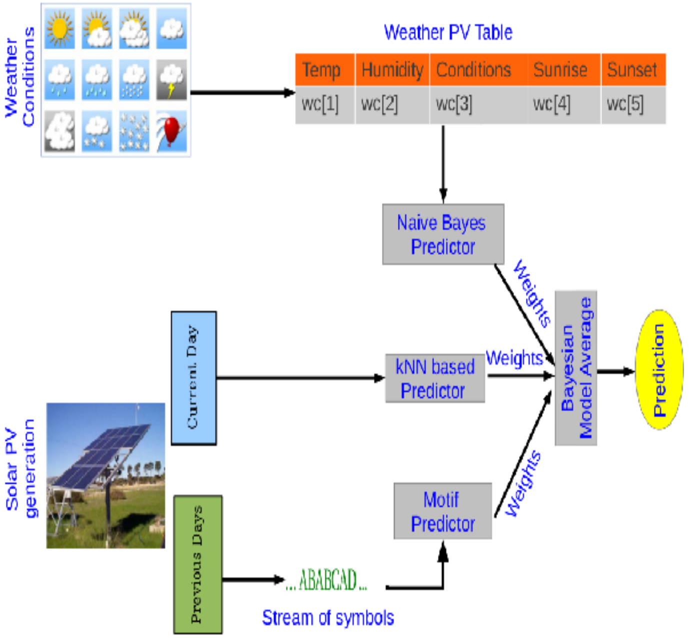

- **Results**

Method	Testing Error		
	Per. Abs. Error	Per. RMS Error	Rel. Abs. Error
PreviousDay	20.54	20.65	20.81
ARWeather	18.54	18.31	19.73
Stagewise	12.77	12.68	15.66
Ensemble2	10.04	10.01	10.01
Ensemble3	8.13	8.21	8.34

[AAAI 2012]

Fine grained PV Prediction using Bayesian Ensemble

- **Results**

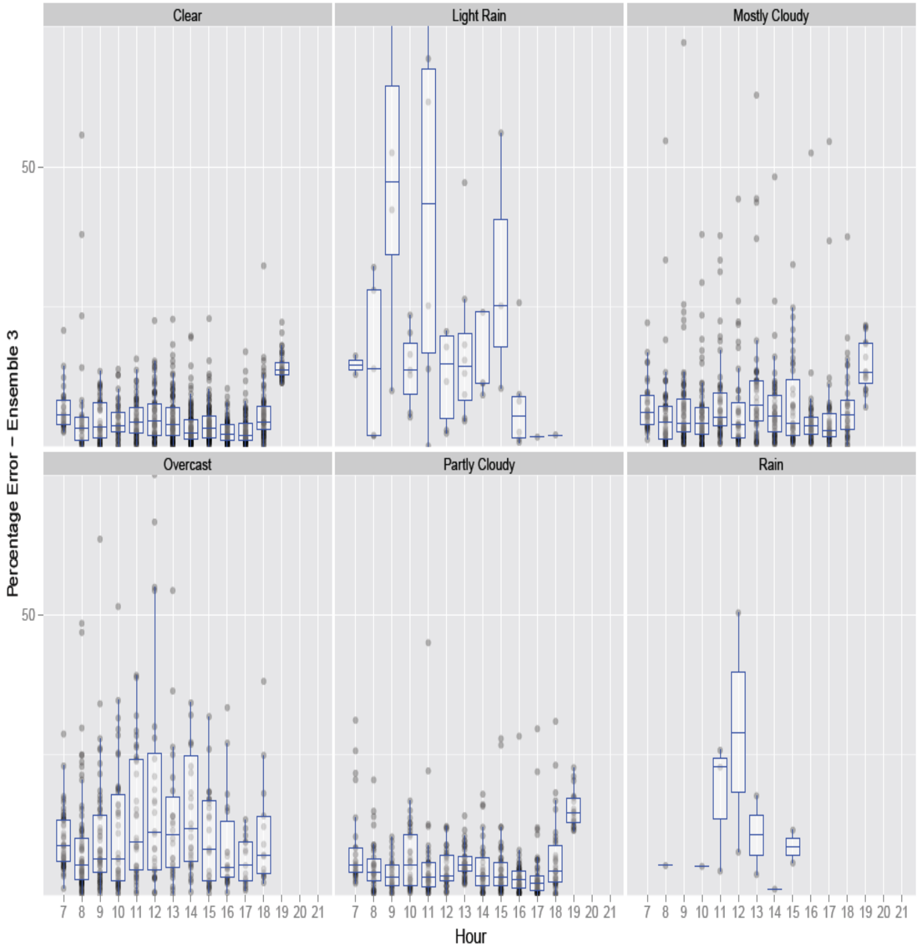

Error by weather condition

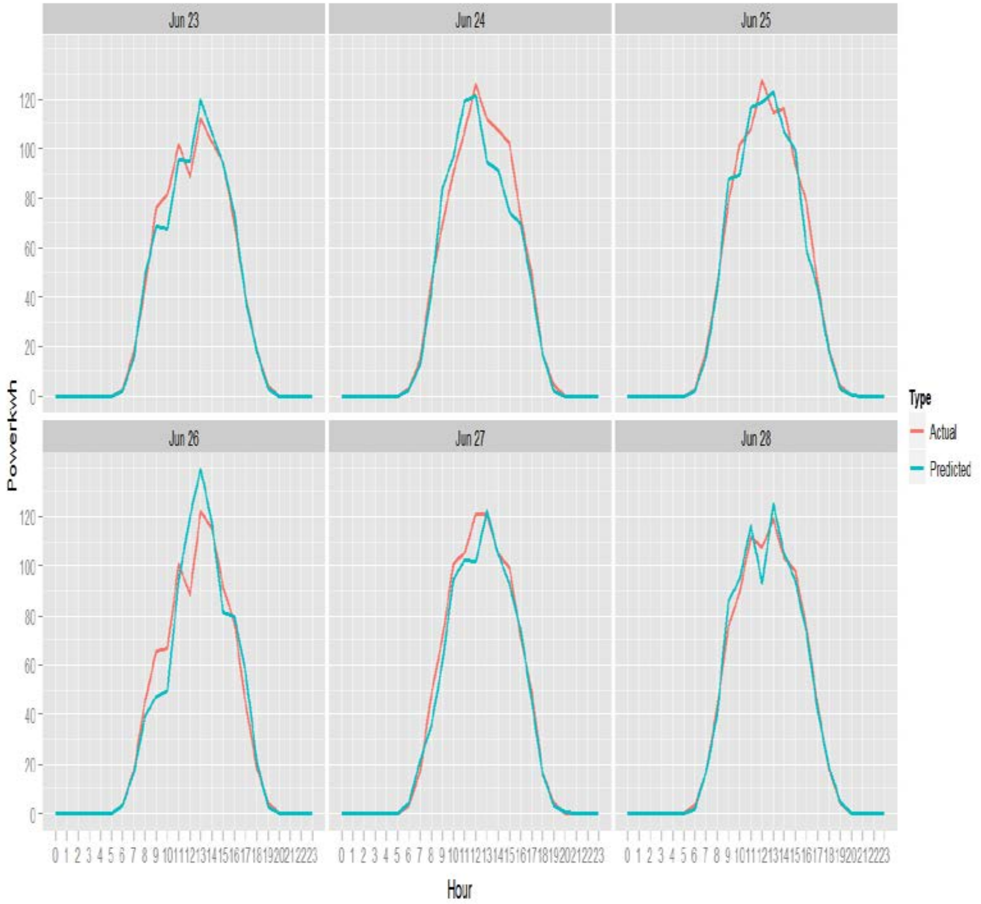

Actual versus predicted

The Net-Zero Energy Data Center

Implementation in Palo Alto

PV micro grid

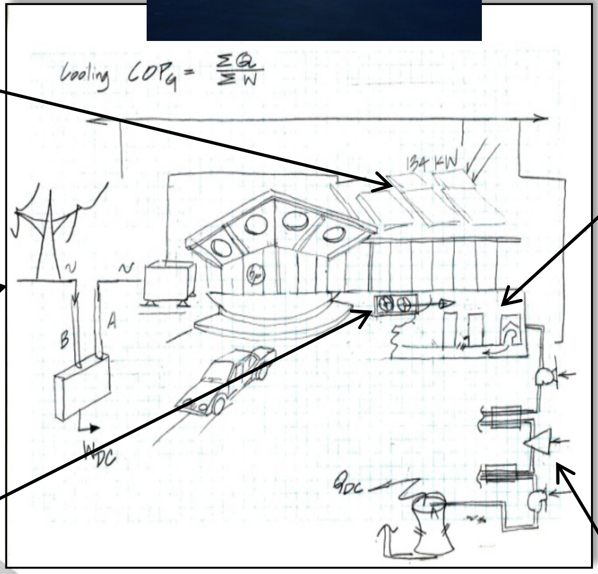

Data center

Outside air

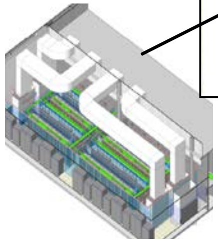

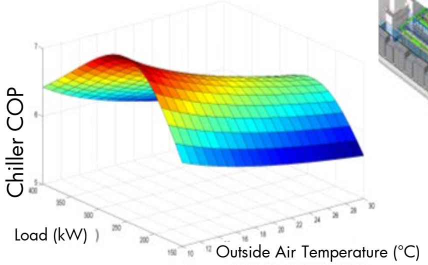

Data center supply side

Cooling
infrastructure
power demand

Data center demand side

Net-Zero Energy Methodology and Integration

Prediction

Supply-side Prediction
- Renewable power prediction
- Cooling capacity prediction

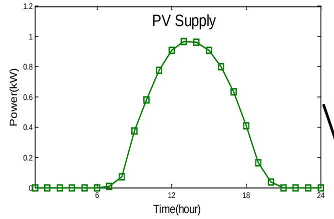

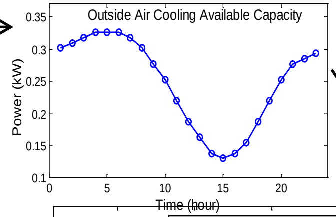

IT Demand Prediction

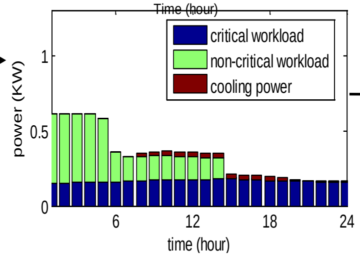

Planning

DC Operation Objectives:
- Net-zero energy operation
- Maximize use of renewable energy
- Minimize dependability on Grid

IT Workload Planning
- Integrate Supply and Demand Side
- IT Demand Shaping
- Power capping

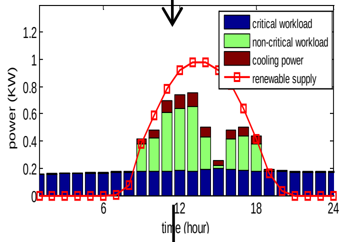

Execution

Dynamic IT Provisioning

Dynamic Cooling Provisioning

Verification and Reporting

Measurement Verification

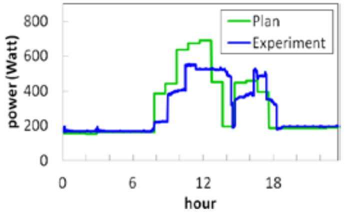

Ref: Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, C. Hyser, "Renewable and Cooling Aware Workload Management for Sustainable Data Centers", ACM SIGMETRICS/Performance, June 11-15 2012, London, UK.

Prediction: Summary

PV Supply Prediction

- Search for most “similar” days in the recent past
- Hourly generation estimated from corresponding hours of “similar” days

Ref: P. Chakraborty, M. Marwah, M. Arlitt, N. Ramakrishnan, Fine-grained Photovoltaic Output Prediction Using a Bayesian Ensemble, in Proceedings of the 26th Conference on Artificial Intelligence (AAAI'12), Toronto, Canada, July 2012

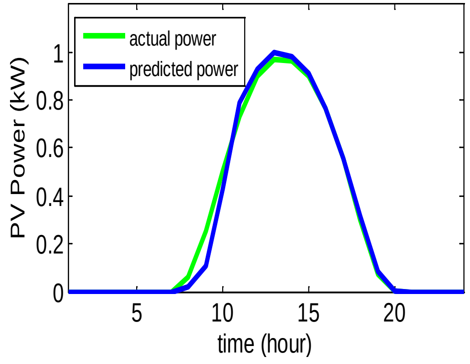

Cooling Capacity Prediction

- End-to-End Energy Modeling

Ref: Breen, T.J. et. al. “From Chip to Cooling Tower Data Center Modeling: Validation of Multi-Scale Energy Management Model”, Proceedings of Itherm, June 2012

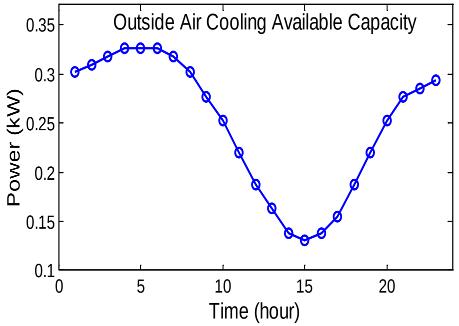

IT Workload Prediction

- Perform a periodicity analysis (e.g., Fast Fourier Transform)
- Use an auto-regressive model to predict workload from historical data

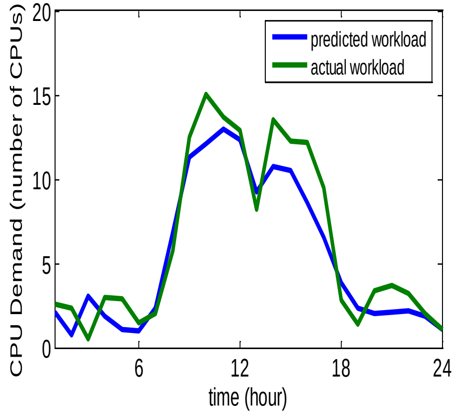

Planning: Supply-Side Aware IT Workload Planning

Planning Flow

electricity price | renewable supply | cooling capacity efficiency | energy storage | IT workload & SLAs | goals

↓

Workload Planning

↓

A detailed workload scheduling and capacity allocation plan
- Workload scheduling plan
- IT resource and power capacity allocation
- Cooling micro-grid capacity allocation

Demand Shaping

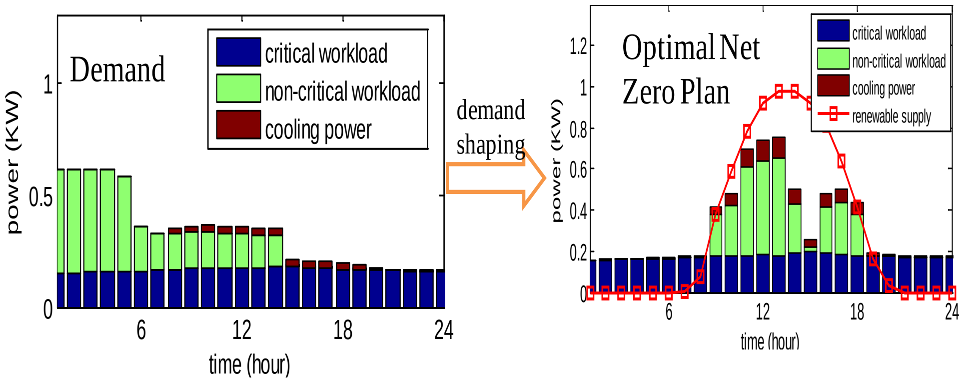

Overall demand is "shaped" according to input constraints and operation objectives

Satisfy critical workload resource requirements

Ref: Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, C. Hyser, "Renewable and Cooling Aware Workload Management for Sustainable Data Centers", ACM SIGMETRICS/Performance, June 11-15 2012, London, UK.

Power and Workload Visualization

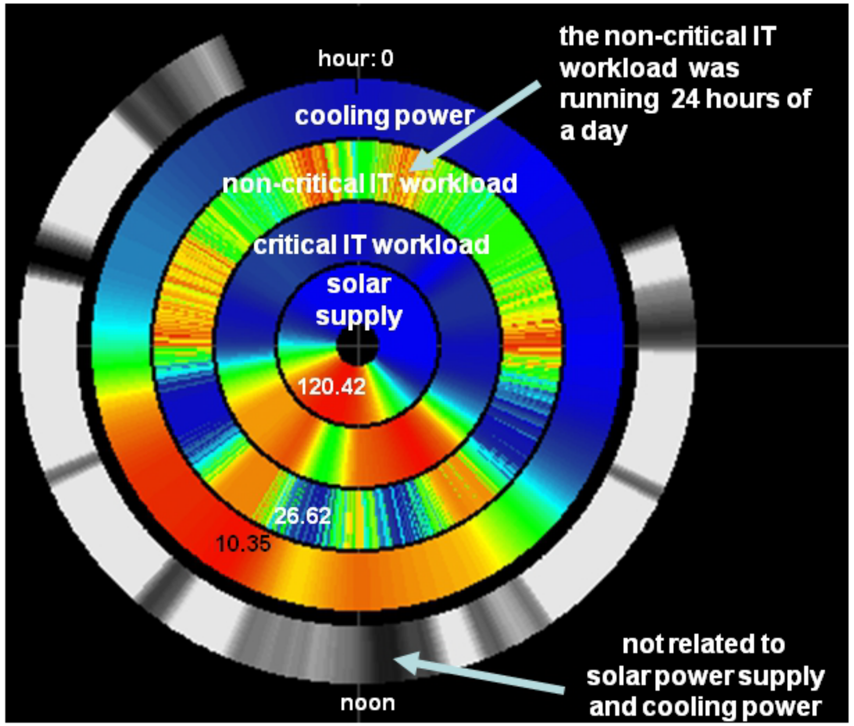

Before optimization

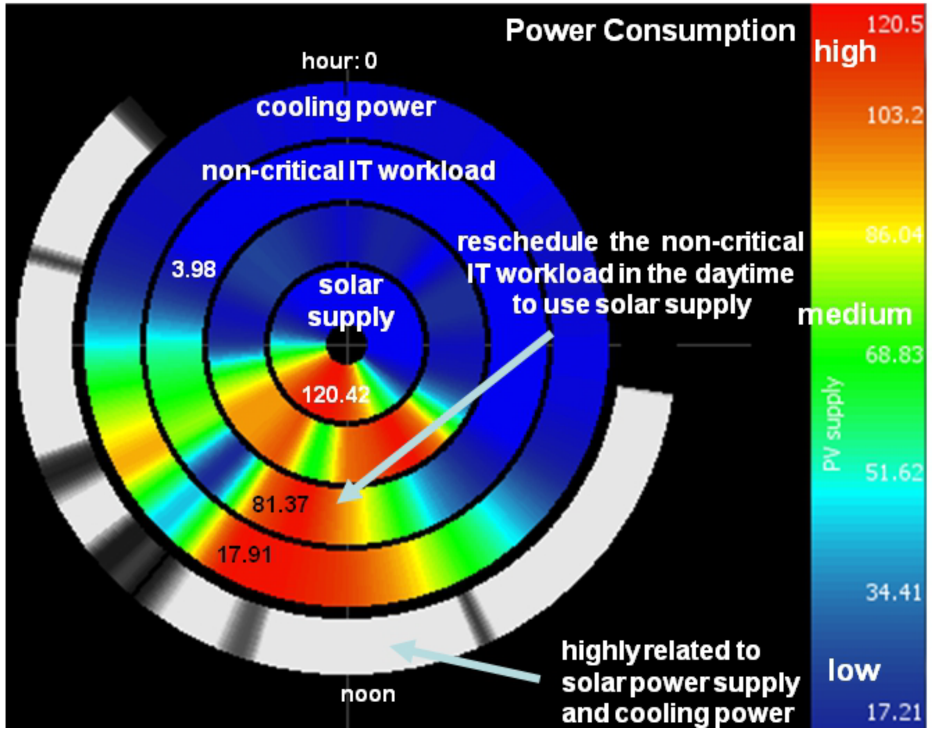

After optimization

[VDA 2013]

Acknowledgements

Martin Arlitt, Cullen Bash, Prithwish Chakraborty, Yuan Chen, Daniel Gmach, Ming Hao, Zhenhua Liu, Freddy Lugo, Chandrakant Patel, Deb Patnaik, Naren Ramakrishnan, Amip Shah, Ratnesh Sharma

References

- **[AAAI 2012]** P. Chakraborty, M. Marwah, M. Arlitt, and N. Ramakrishnan. Fine-grained Photovoltaic Output Prediction using a Bayesian Ensemble, in *Proceedings of the 26th Conference on Artificial Intelligence (AAAI'12)*, Toronto, Canada, 7 pages, July 2012, To appear.
- **[SIGMETRICS 2012]** Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, C. Hyser, "Renewable and Cooling Aware Workload Management for Sustainable Data Centers", ACM SIGMETRICS/Performance, June 11-15 2012, London, UK.
- **[SensorKDD 2010]** Manish Marwah, Ratnesh Sharma, Wilfredo Lugo, Lola Bautista, "Anomalous Thermal Behavior Detection in Data Centers using Hierarchical PCA," in SensorKDD in conjunction with KDD 2010.
- **[KDD 2009]** D. Patnaik, M. Marwah, R. Sharma, N. Ramakrishnan, "Sustainable Operation and Management of Data Center Chillers using Temporal Data Mining," In ACM KDD, June 27 - July 1, 2009, Paris, France.
- **[IEEE Computer 2009]** Amip Shah, Tom Christian, Chandrakant D. Patel, Cullen Bash, Ratnesh K. Sharma: Assessing ICT's Environmental Impact. IEEE Computer 42(7): 91-93, July 2009.
- **[VDA 2013]** Ming Hao, Manish Marwah, Sebastian Mittelstadt, Halldor Janetzko, Daniel Keim, Umeshwar Dayal, Cullen Bash, Carlos Felix, Chandrakant Patel, Meichun Hsu, Yuan Chen, "Visual Analytics for Cyber Physical Data Streams Using Spatio-Temporal Radial Pixel Visualization," submitted to IS&T/SPIE Conference on Visualization and Data Analysis (VDA), 2013.