Life Patterns

structure from wearable sensors

by

Brian P. Clarkson

Thesis Advisor: Prof. Alex Pentland
Thesis Readers: Prof. Trevor Darrell
Prof. Joseph Paradiso

Prologue

Motivation

The Question: Can we computationally structure a sensor record of an individual's day-to-day life?

sensor record - "dumb" measurements via sensors

structure - similarity measure, perplexity, prediction, classification

The Applications:

- memory prosthesis
- automatic diary
- the frame problem
- "the past explains"
- context-aware agents

Vannevar Bush - memex (1945)

Wearable Sensing

- Steve Mann - wearable cameras (1997)
- Thad Starner - Patrol system (1999)
- Farringdon et. al. - sensory badges & jackets (1999)
- Jennifer Healey - wearable bio-sensing (1998)

Context-Awareness

- Brad Rhodes - Remembrance Agent (2000)
- Lamming & Flynn - “Forget-me-not” (1994)

Robotics & AI

- Grimson et. al. - long-time monitoring of a site (1998)
- Jogan & Leonardis - localization via panoramic views

Prologue

Talk Outline

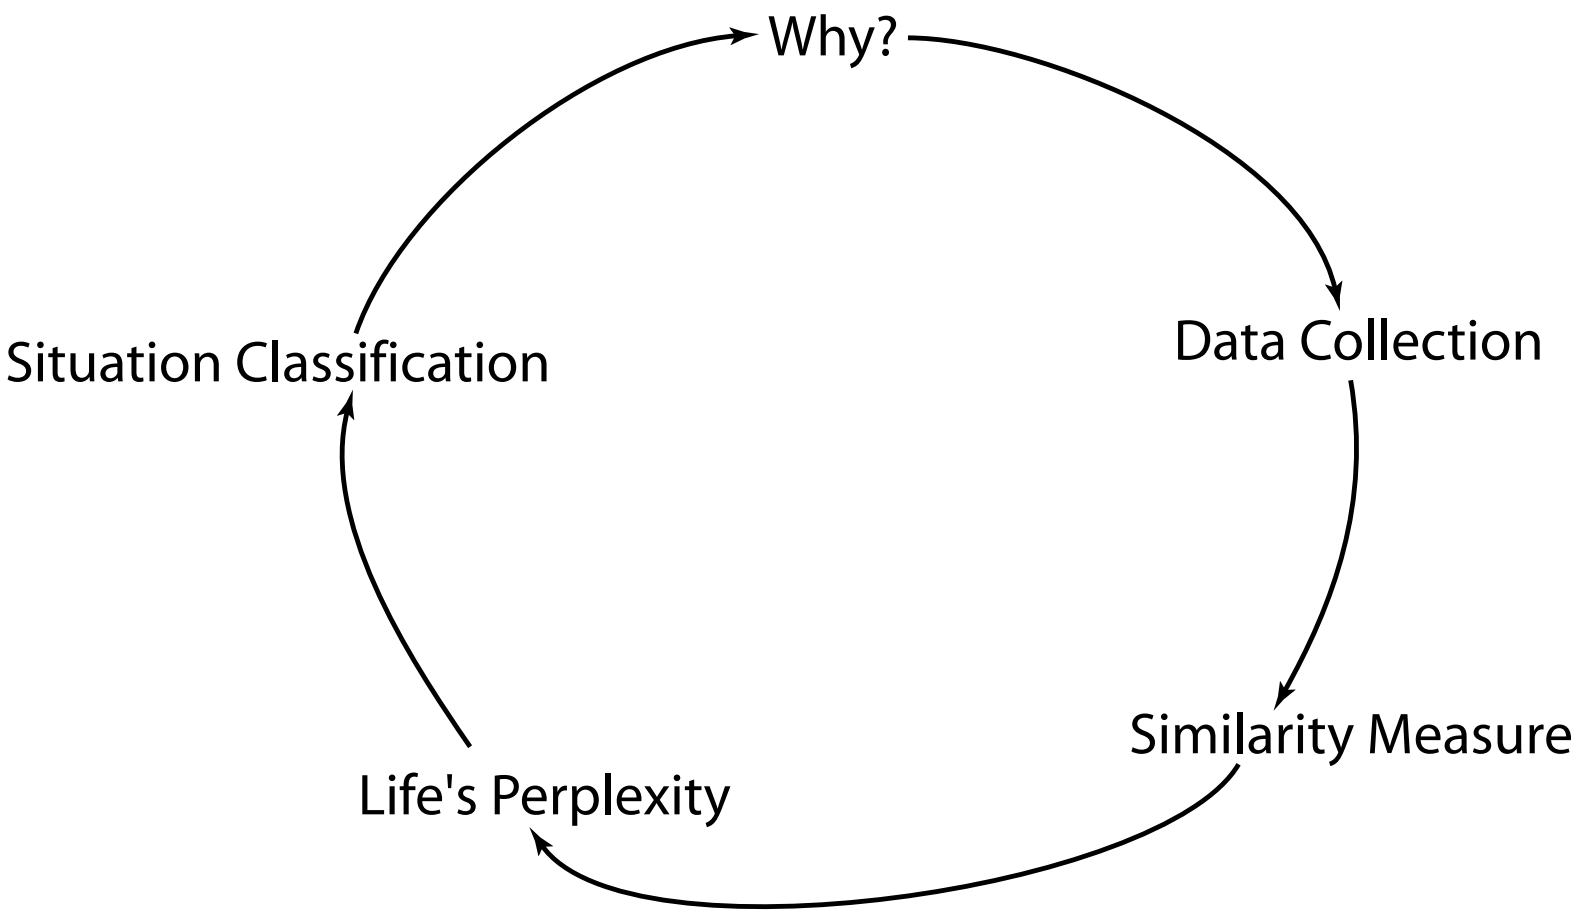

The “I Sensed” Data Set

Mr. Kawara On [1933-]

Date paintings

The “I Sensed” Data Set

Mr. Kawara On [1933-]

The “I Got Up” Series

I Went
I Met
I Am Alive
I Read ...

The “I Sensed” Data Set

Data Collection Wearable

The “I Sensed” Data Set

Data Collection Wearable

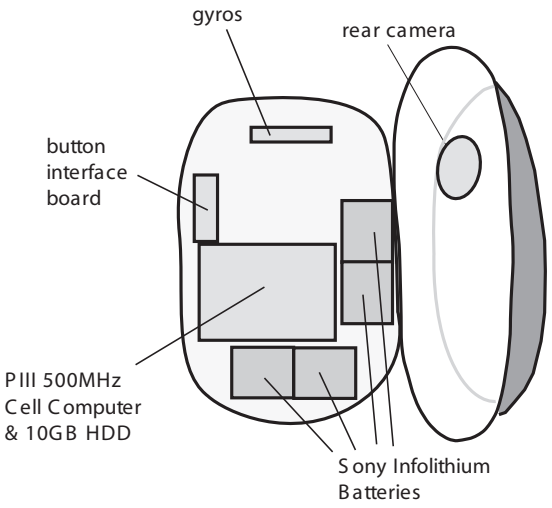

demo available

The "I Sensed" Data Set

Insect Perception

Key Properties: low resolution, wide field-of-view

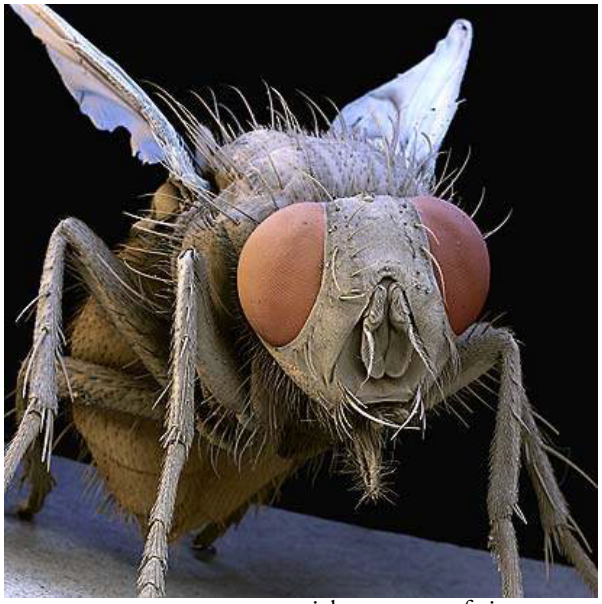

copyright: www.eyeofscience.com

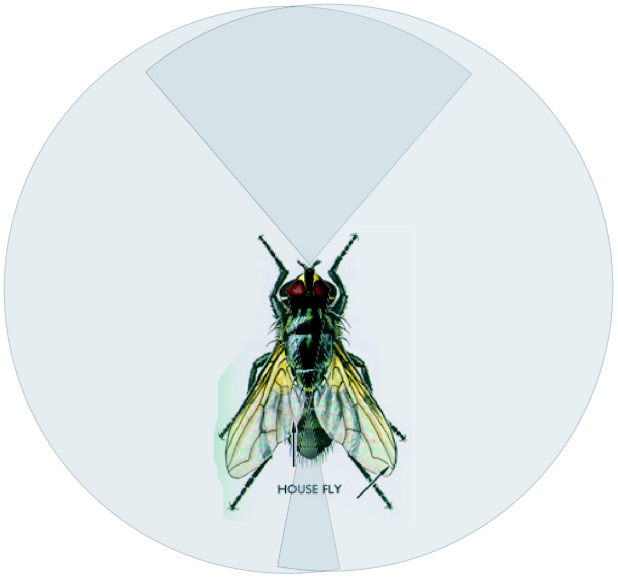

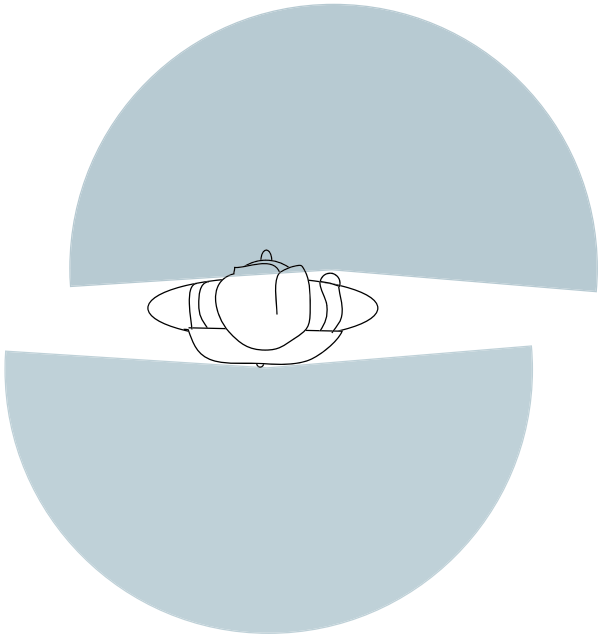

compared to humans:

- 100,000 times fewer photoreceptors
- 360-deg field of view
- 800 eye units, each having 8 photoreceptors

The Similarity Measure

Peripheral Perception

Peripheral vs. Attentive


Walking over a bridge


Shopping at BestBuy

Without target of attention


Renting a video


Working at the desk

With target of attention

Peripheral sensing is robust to
small changes in the environment.


Walking over a bridge


Walking over a bridge

Direct image matching without correspondence is potentially useful!

The Similarity Measure

Principle Components

Front View Eigenvectors

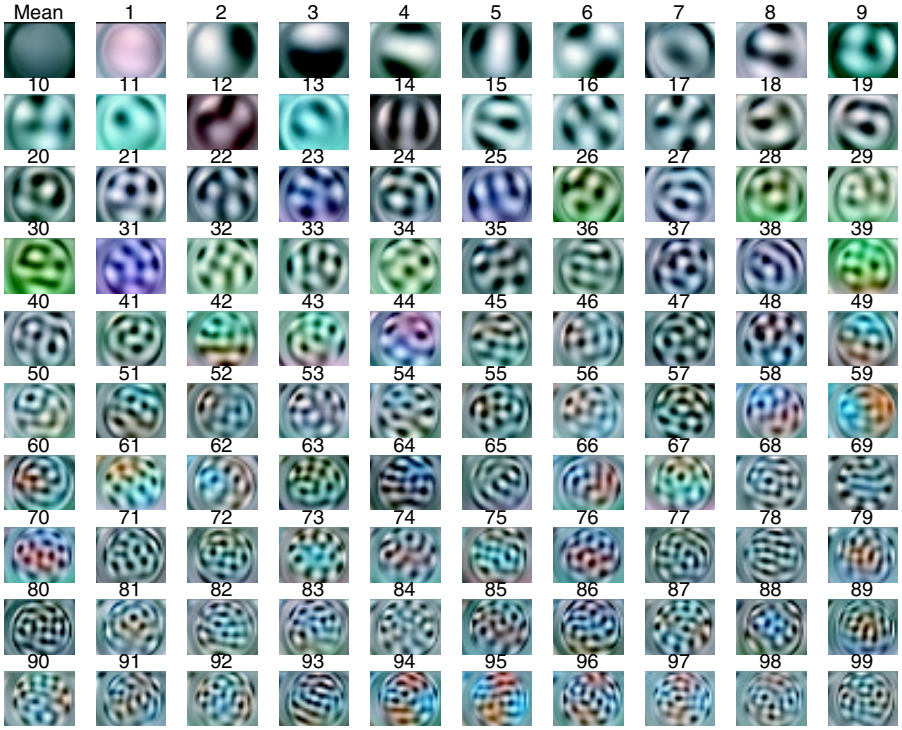

Rear View Eigenvectors

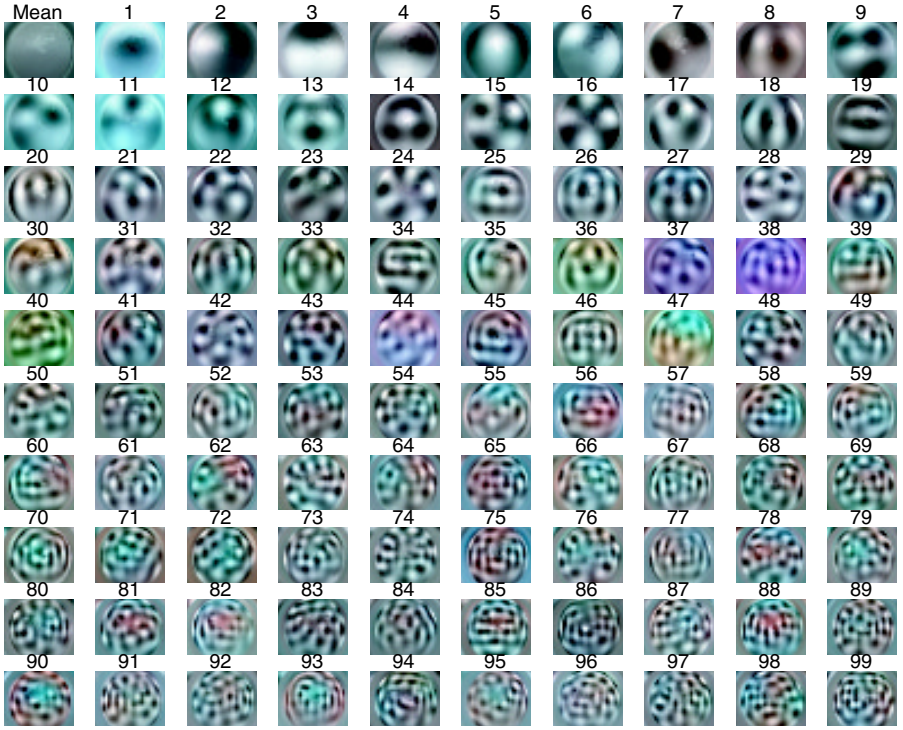

Trained on 32x24 pixel images over 30 days of video.

Bell & Sejnowski - PCA of natural scenes yields Fourier basis
ICA of natural scenes yields localized edge filters

The Similarity Measure

Variance Accounted For

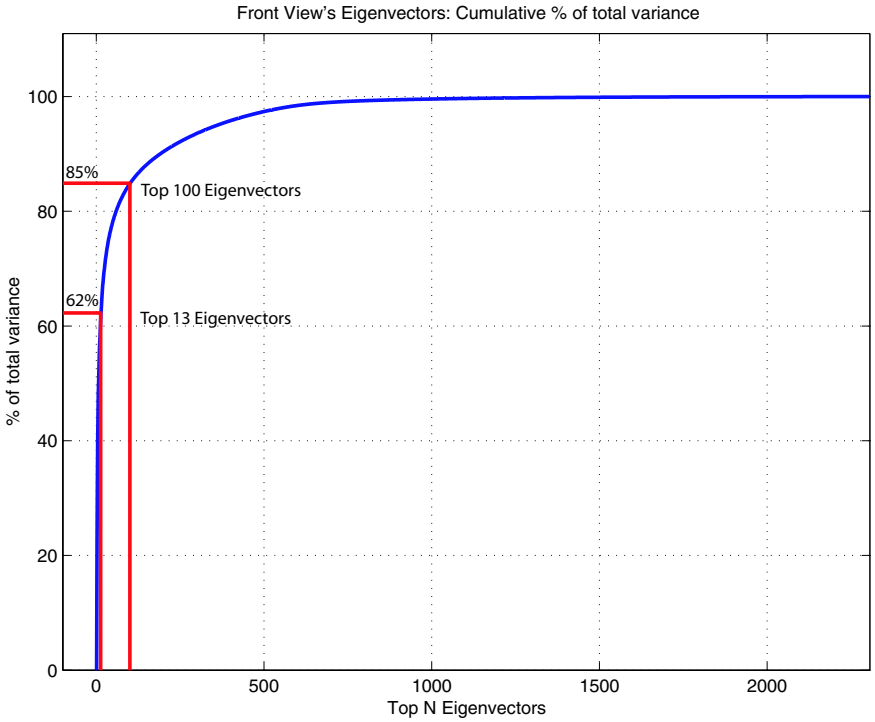

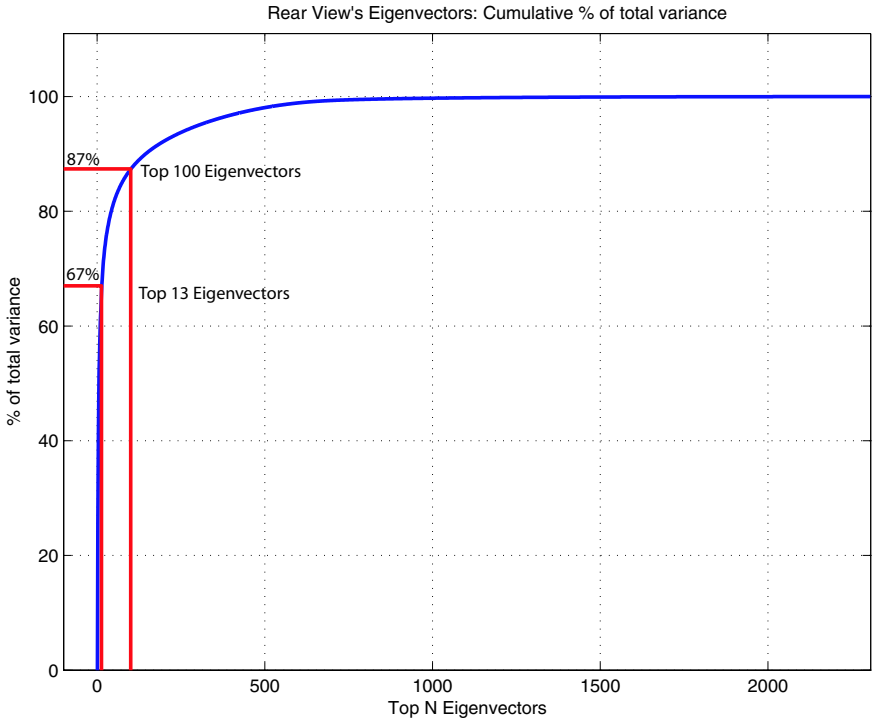

- Compression without loss of detail is difficult.
- We use the top 100 eigenimages for the remaining experiments.

The Similarity Measure

Feature Computation Pipeline

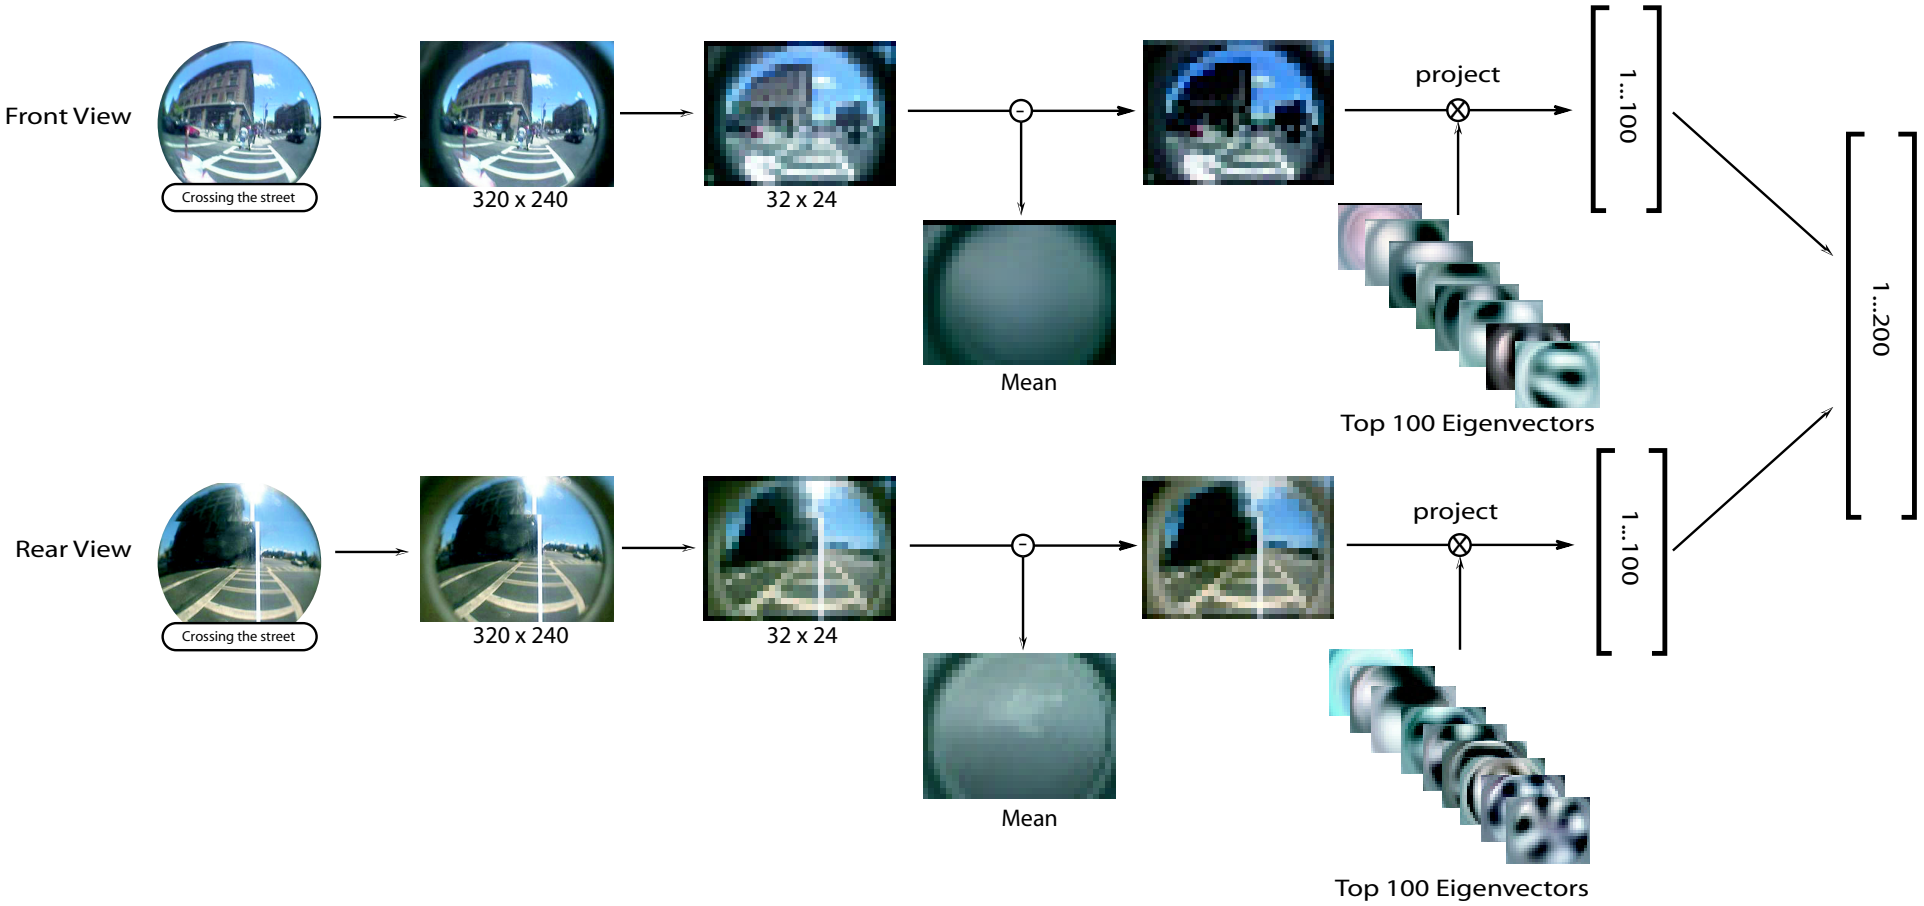

- The result is a 200-dimensional feature vector per frame
- To compare frames, we use the L-1 norm.

The Similarity Measure

Sequence Similarity

How do we compare a pair of frames sequences?

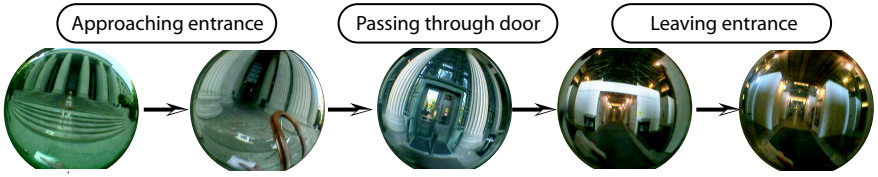

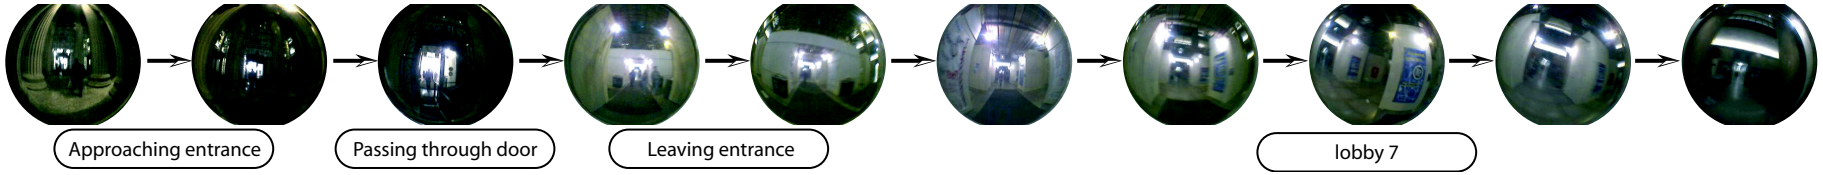

The Similarity Measure

Sequence Similarity

Match one sequence to the other and accumulate frame-by-frame similarities.

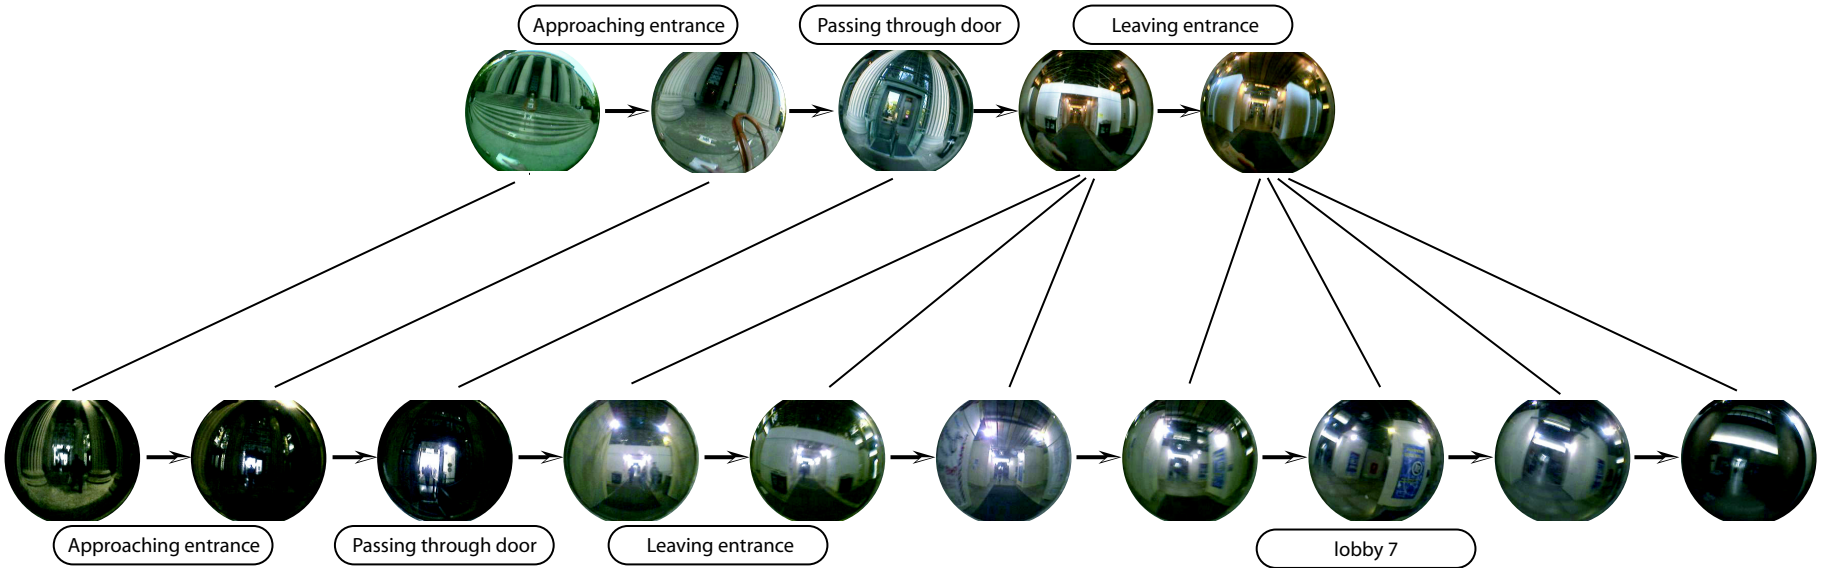

The Similarity Measure

Alignment Model

A sequence of frames becomes a Hidden Markov Model.

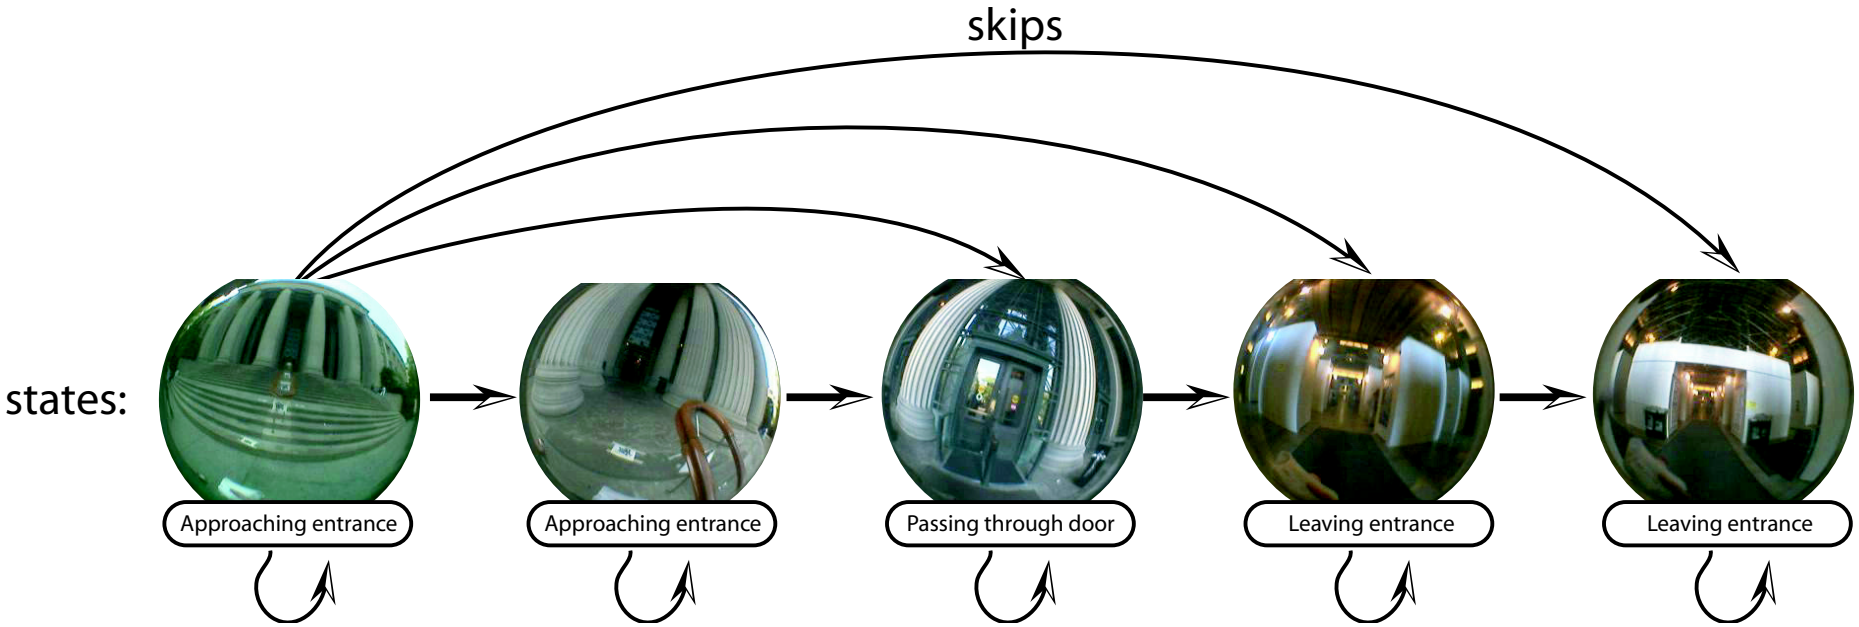

the transitions...

The Similarity Measure

Alignment Model

α-transitions

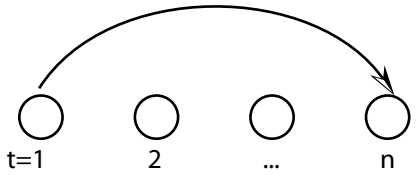

equal probability

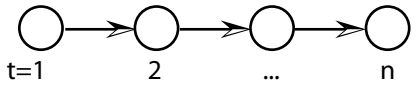

β-transitions

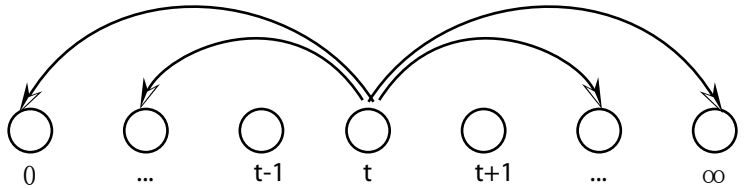

low constant probability

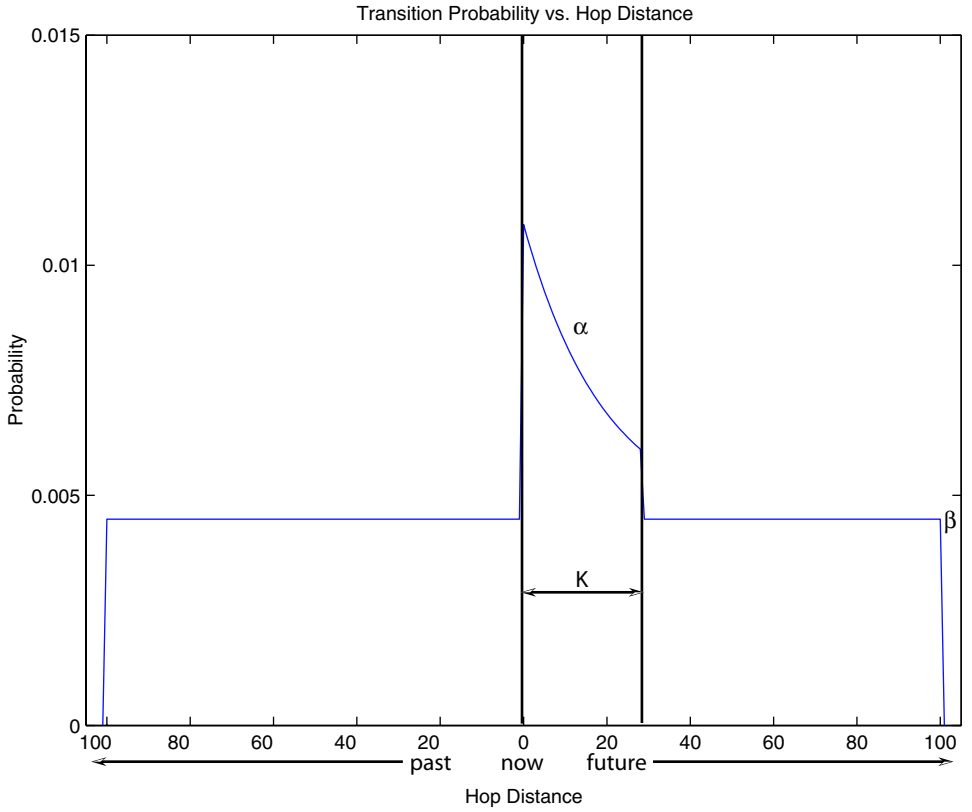

The Similarity Measure

Alignment Path

The Viterbi Algorithm produces the best possible alignment.

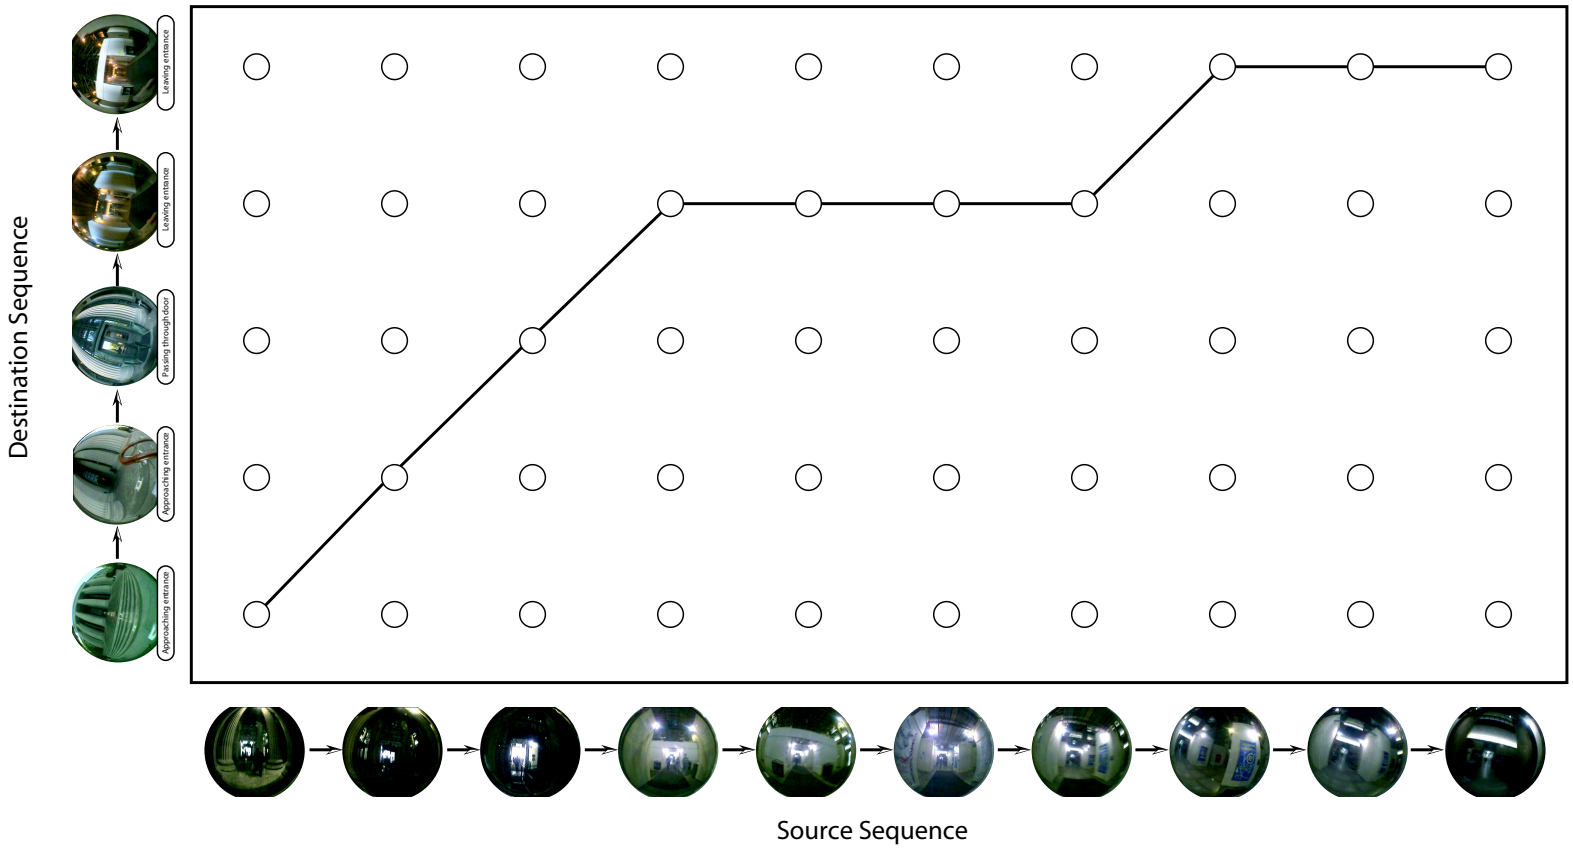

The Similarity Measure

An Example: walking to lab

these are very similar events...

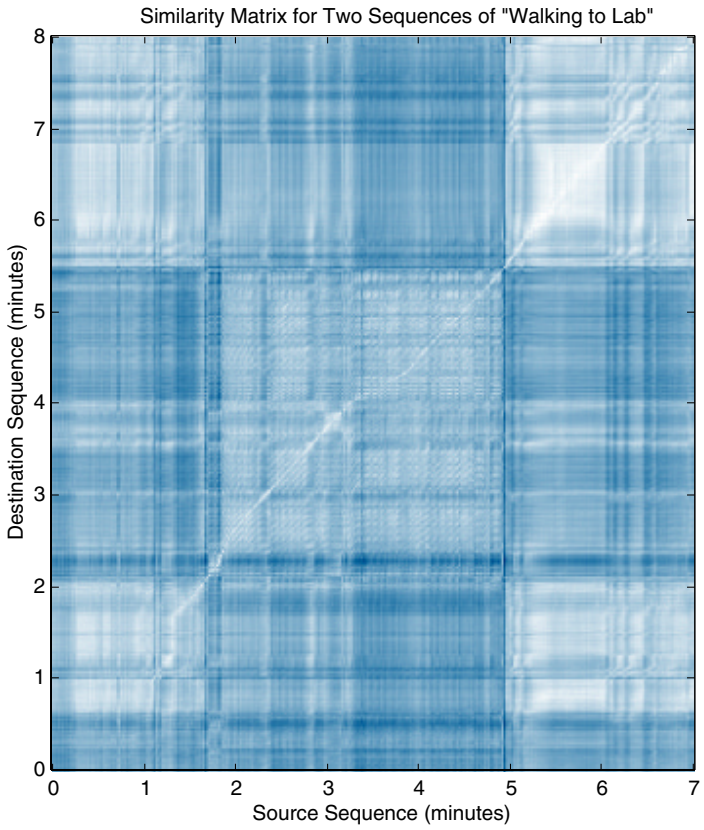

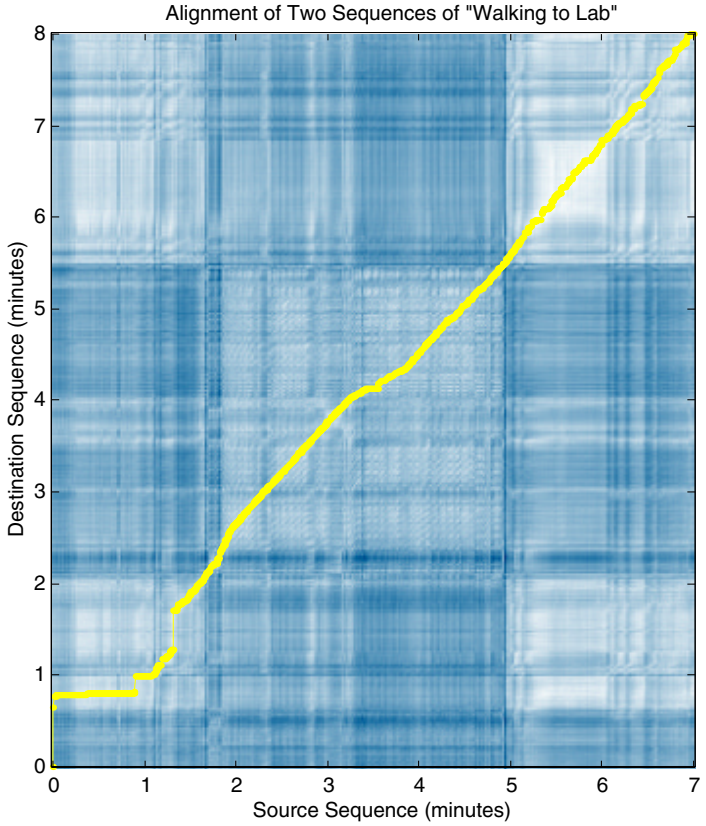

3000 frames thus 3000 states (computationally heavy!)

demo available

The Similarity Measure

Another Example: leaving class

similar at times, dissimilar at others

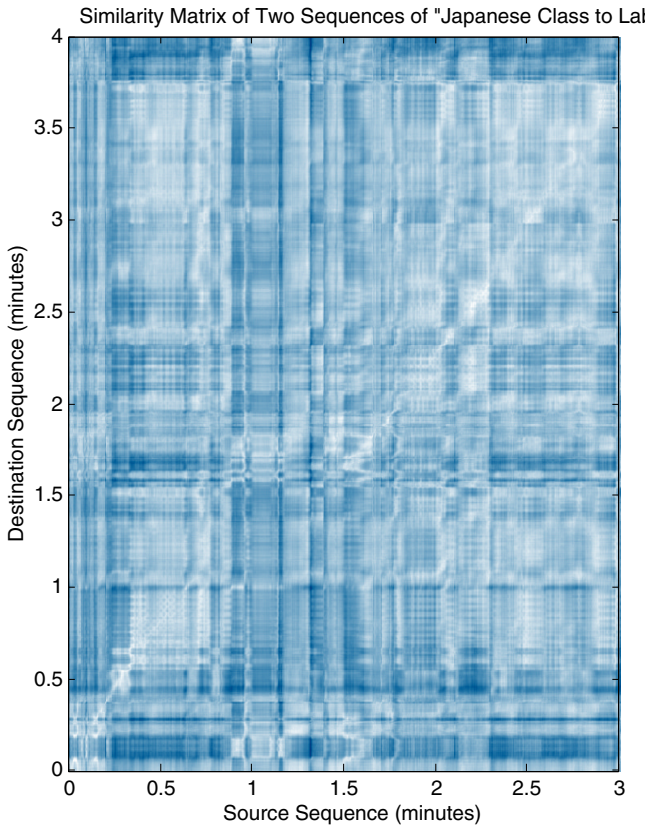

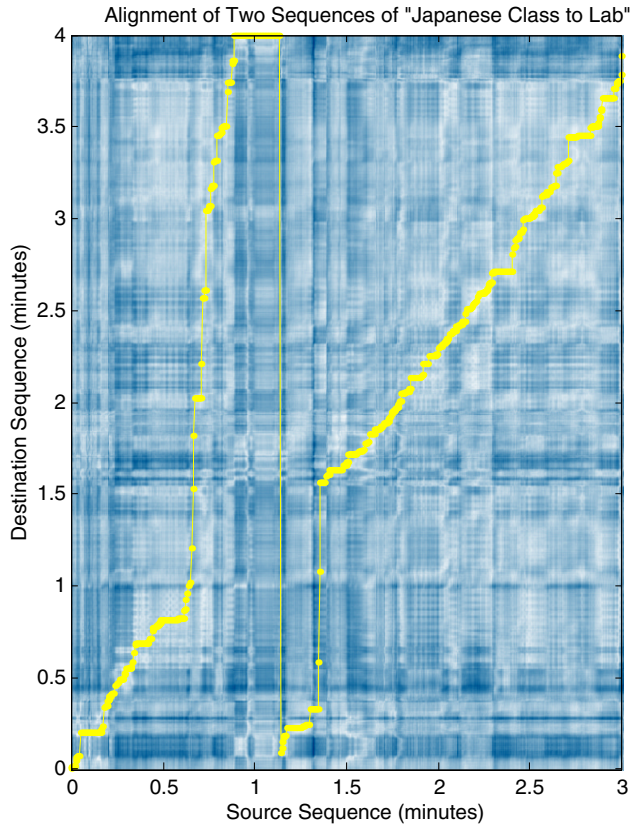

demo available

The Similarity Measure

Run Length Encoding

Remove Redundancy
via Image Similarity

Before
1 day = 200,000 images

After
1 day = 3,000 images

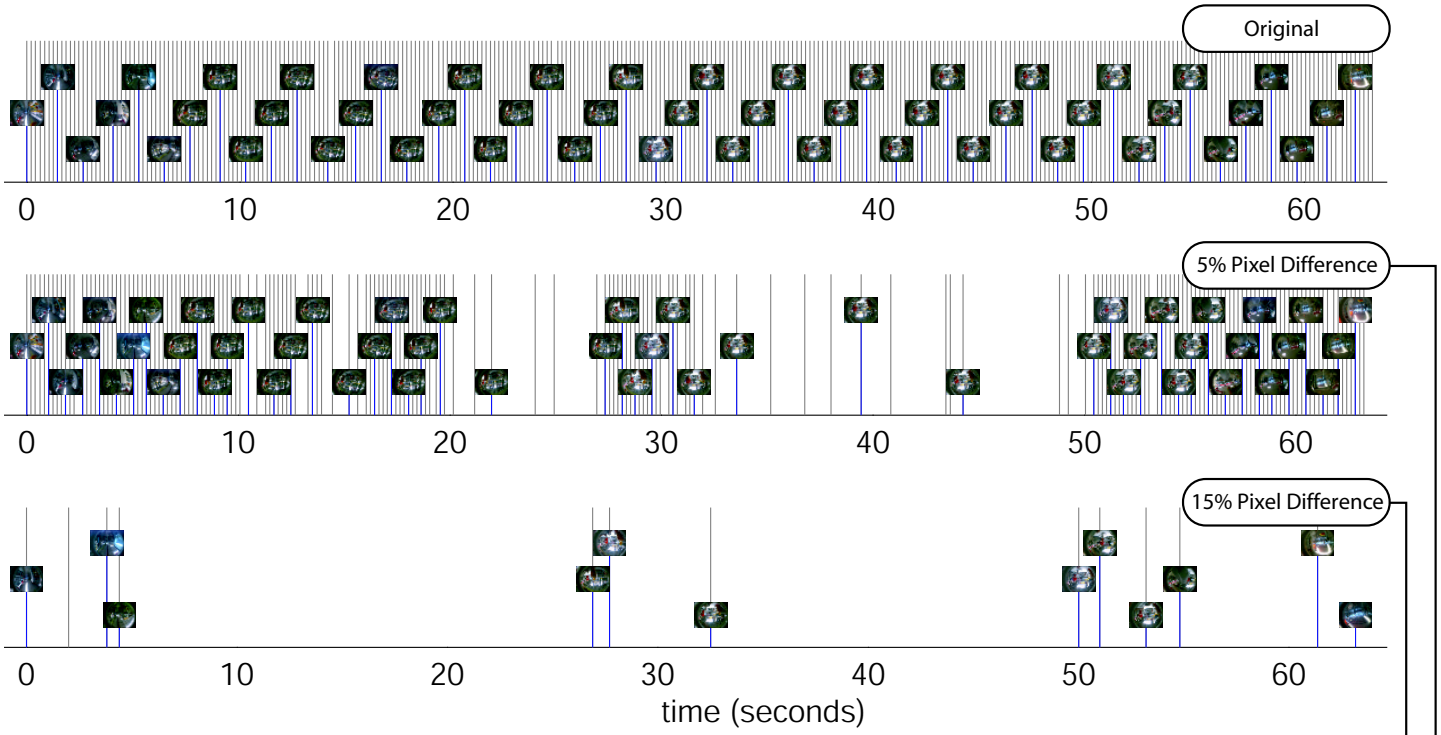

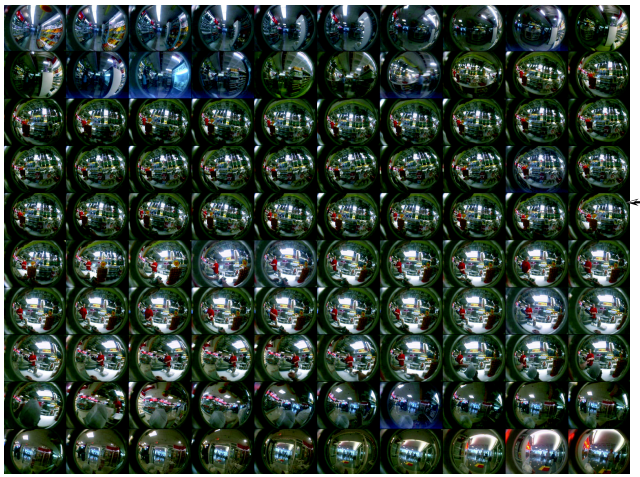

demo available

The Similarity Measure

Alignment of a day

RLE at 15% allows alignment of a pair of days!

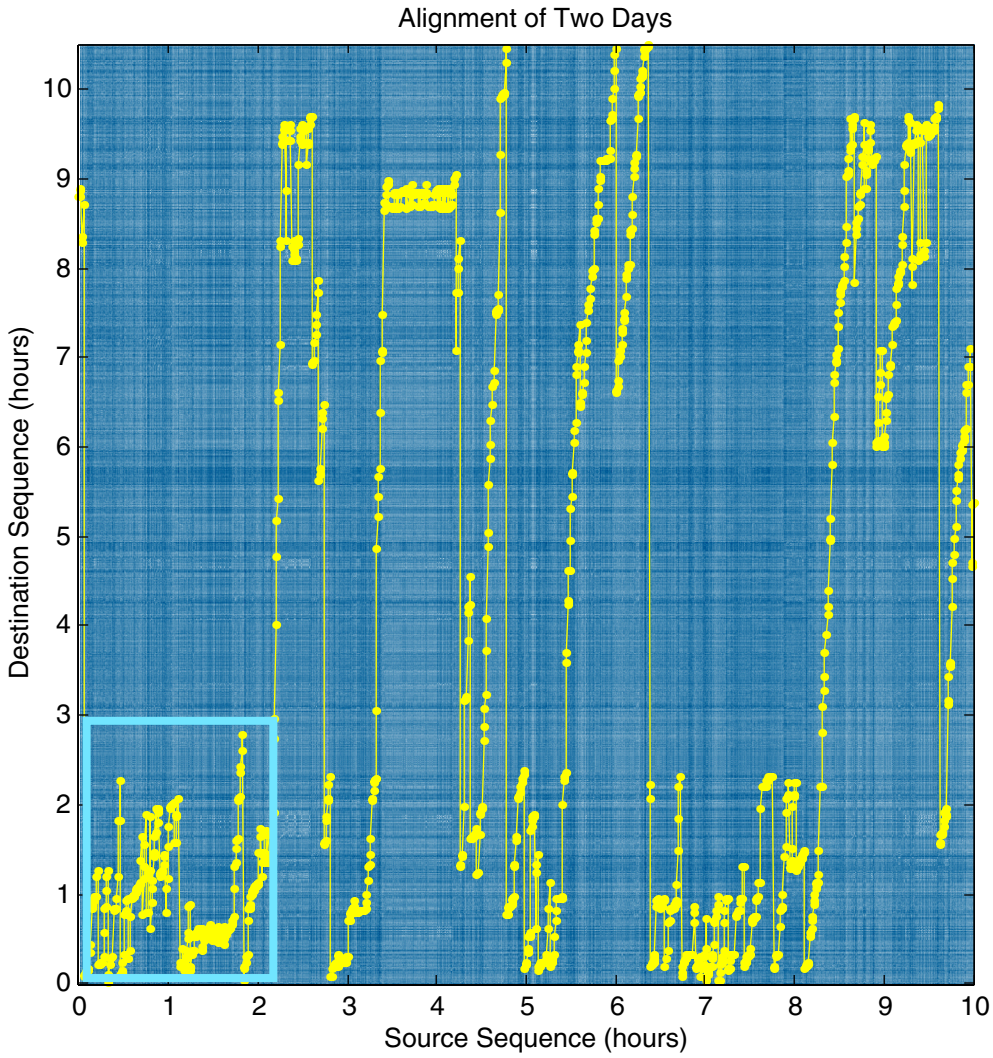

The Similarity Measure

Alignment of a day

The finer detail in the morning...

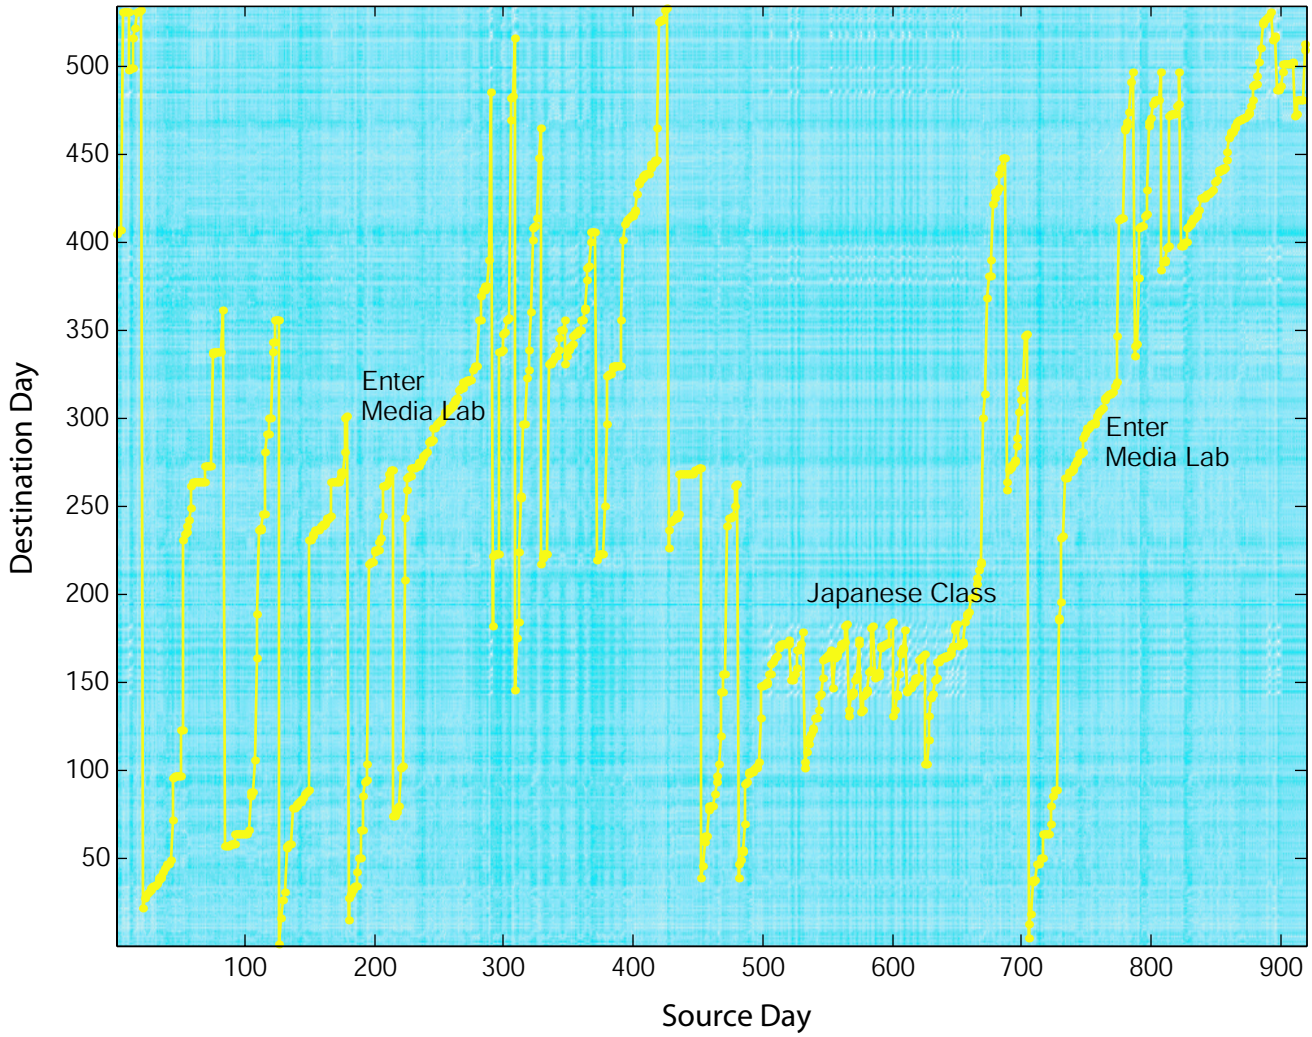

demo available

The Similarity Measure

Alignment of a month

Problem:

- A pair of days typically only align sensibly for a few situations.

Solution:

- Keep adding more sequences to the alignment HMM.

Aligning 1 day to 30 days:

- Build similarity matrix of 30 days at 5 minute resolution.
- Build alignment HMM from each 5 minute chunk in 29 days.
- Align remaining day to the 29 days using the HMM.
- repeat 30 times...

The Similarity Measure

Alignment of a month

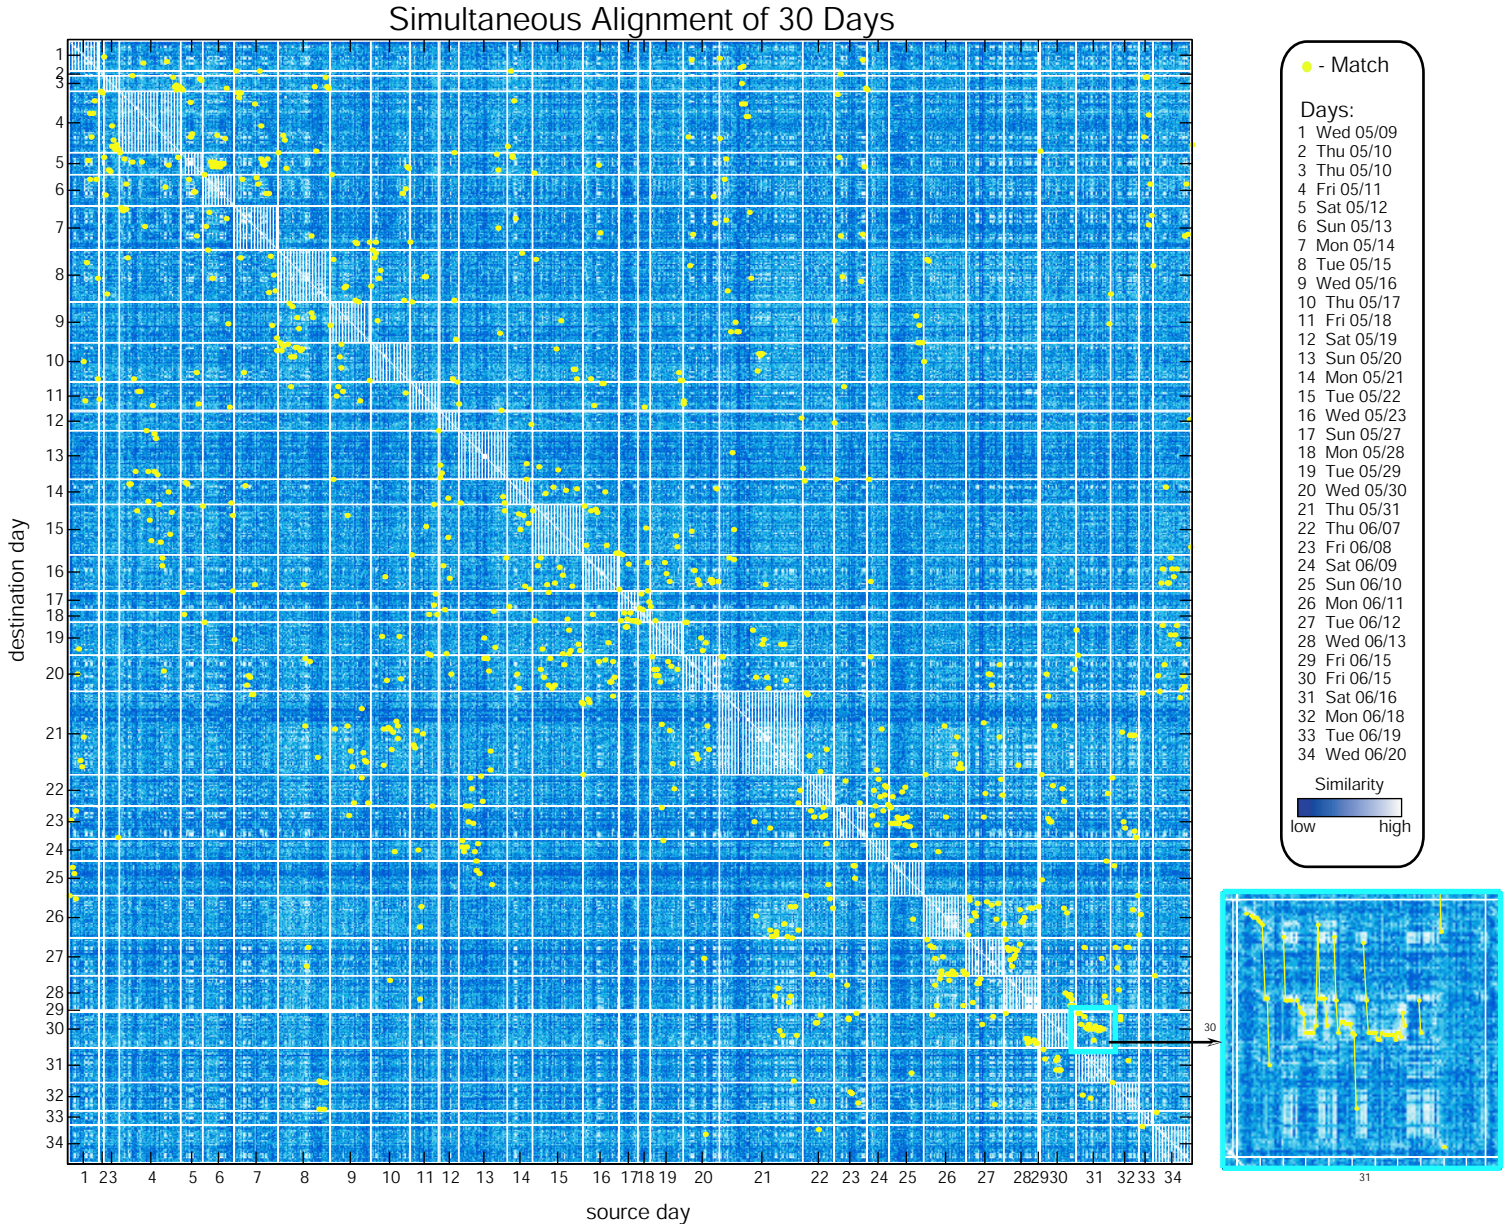

The Similarity Measure

Alignment of a month

Each moment can be aligned to any moment in 29 days.

we will use
this later...

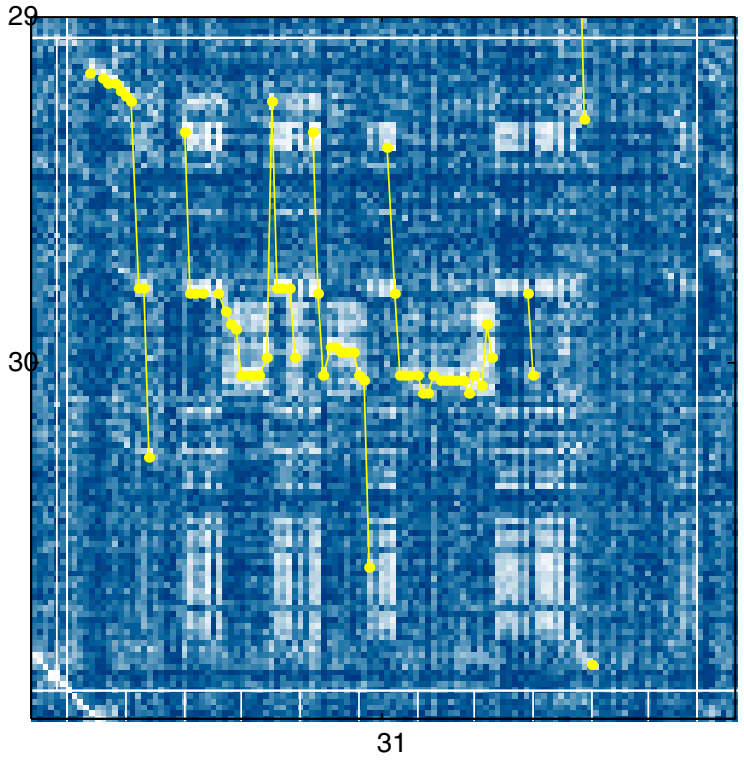

Life's Perplexity

Motivation

"When you come to a fork in the road, take it."

- Yogi Berra

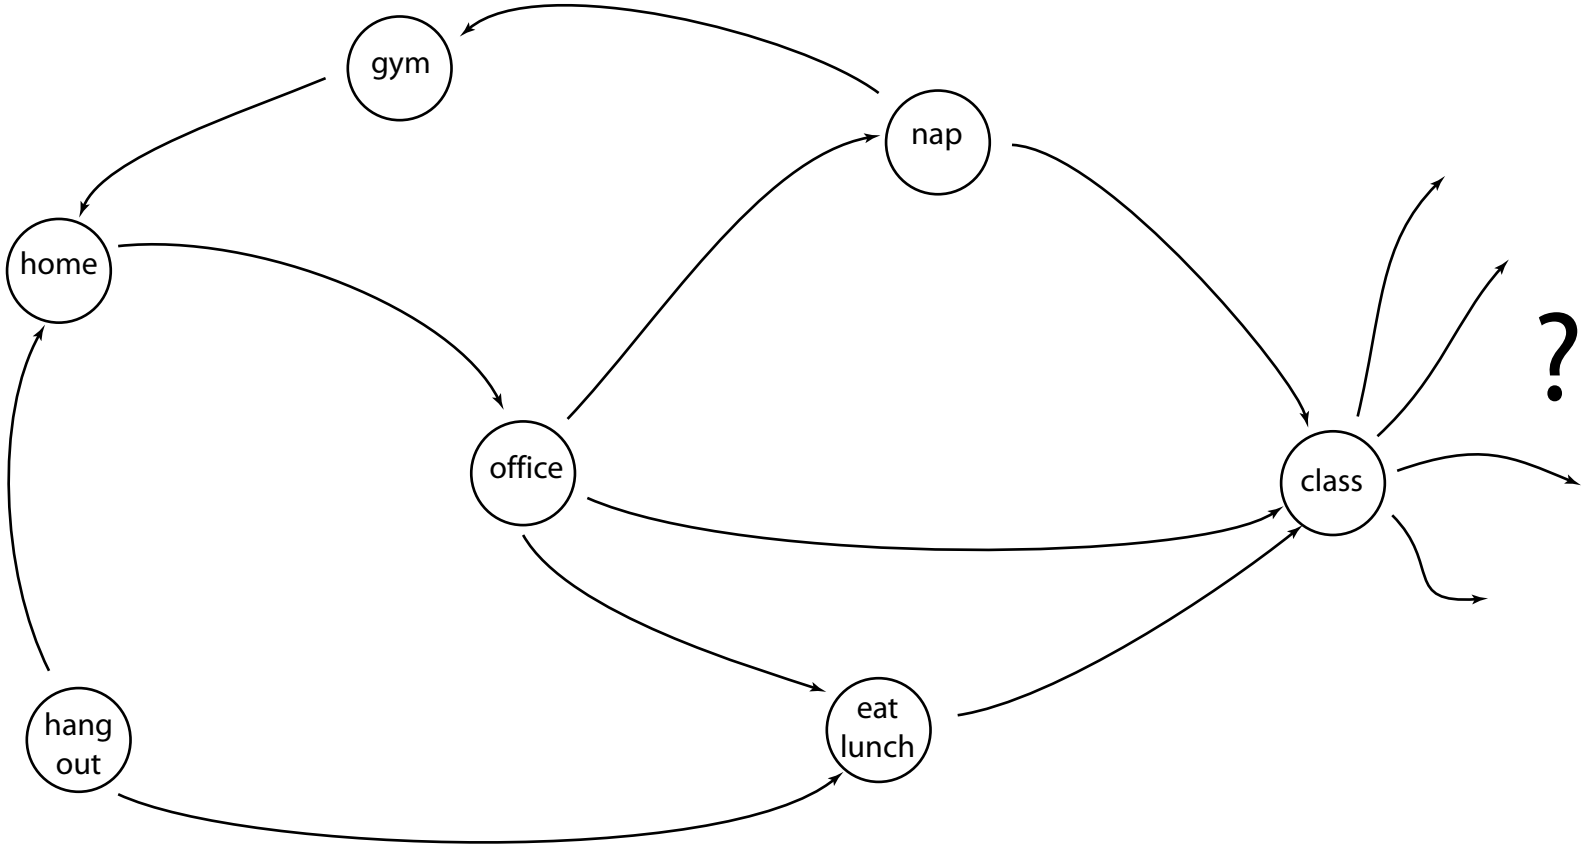

- Where are the decision points? (i.e. what are the nodes)
- What is the perplexity each time a decision is made?
- How consistent is the decision?

Life's Perplexity

Scene Segmentation

β-transitions denote moments of divergence from past behavior.

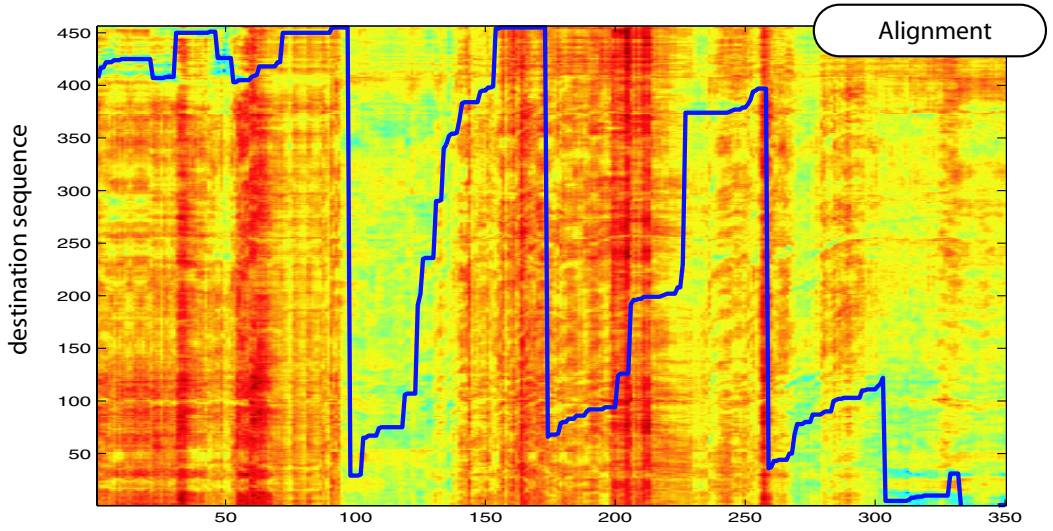

Score each β-transition by its size in time.

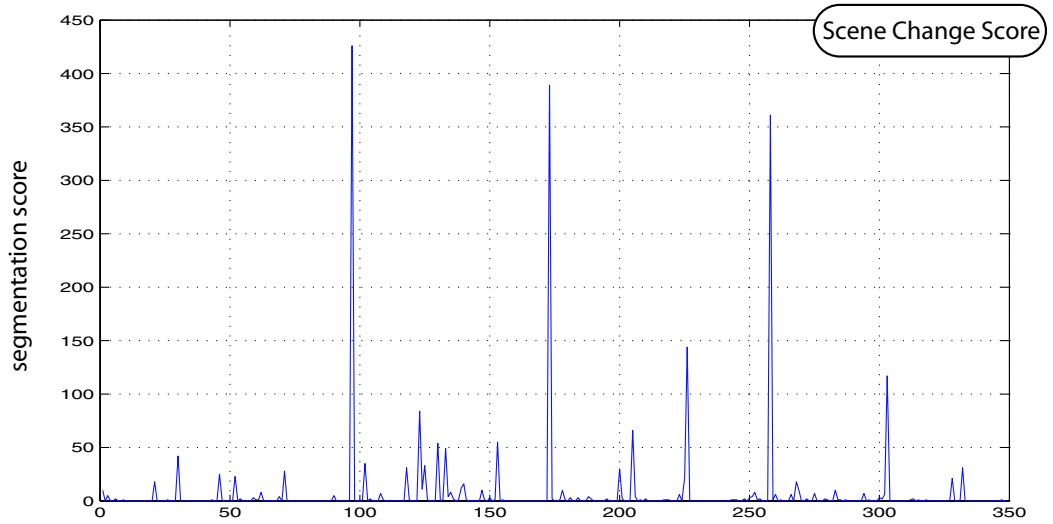

Life's Perplexity

Scene Segmentation

Sweeping a threshold yields a hierarchy of scene segmentations.
30 days = ~1000 scenes

To get scenes cluster using alignment as the similarity measure.

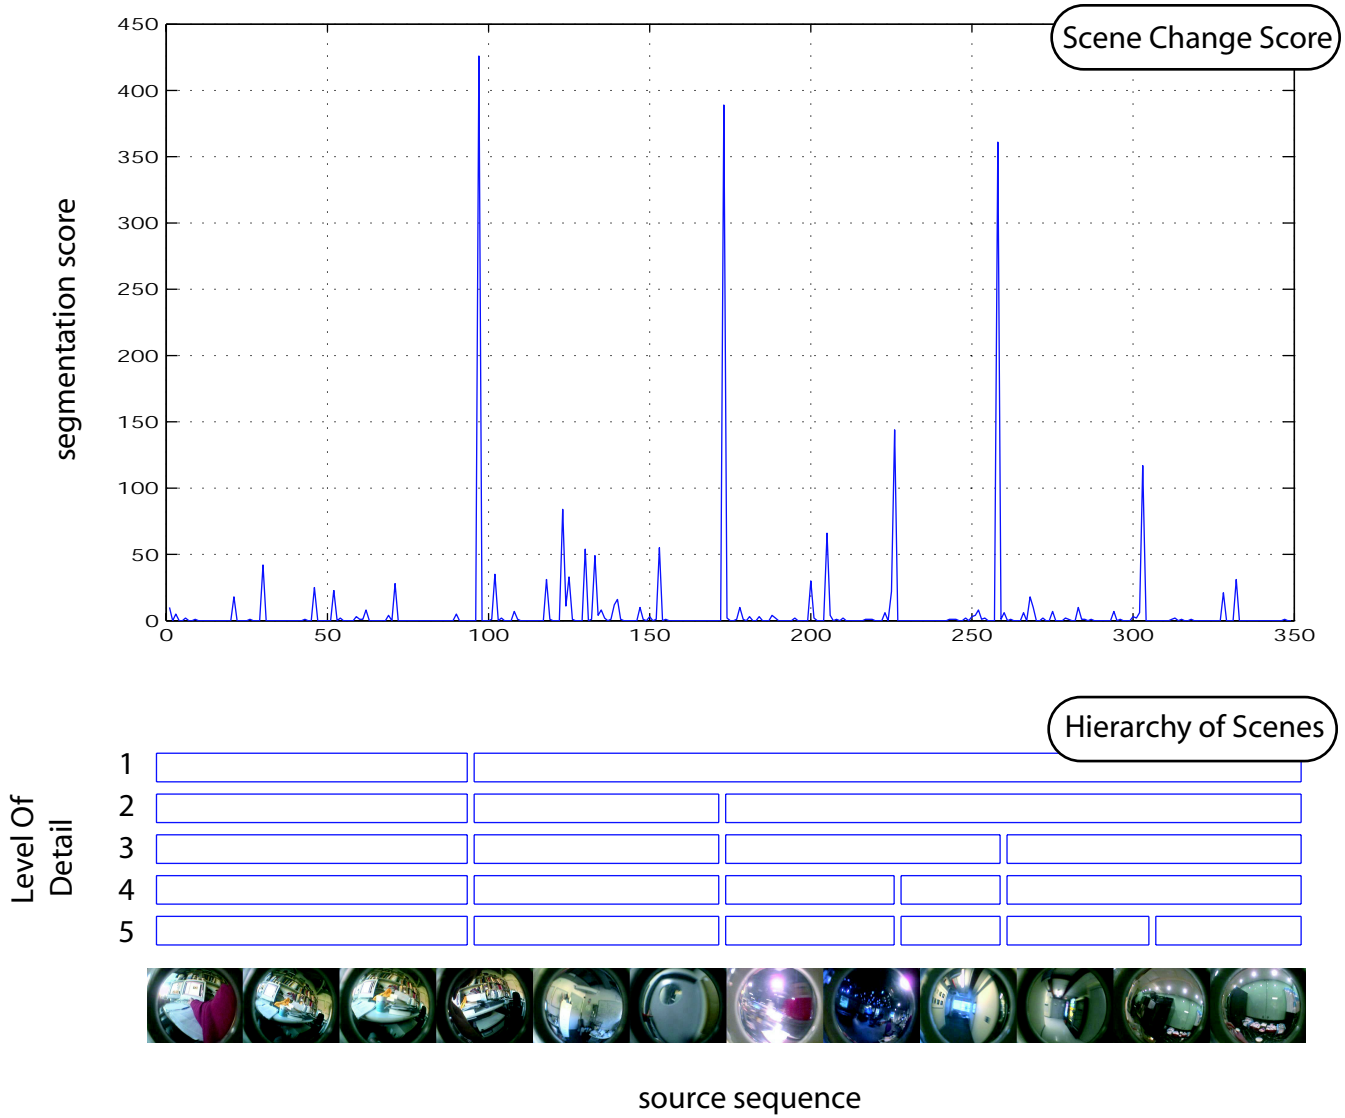

Life's Perplexity

How many nodes?

Degree of redundancy is independent of the # of nodes.

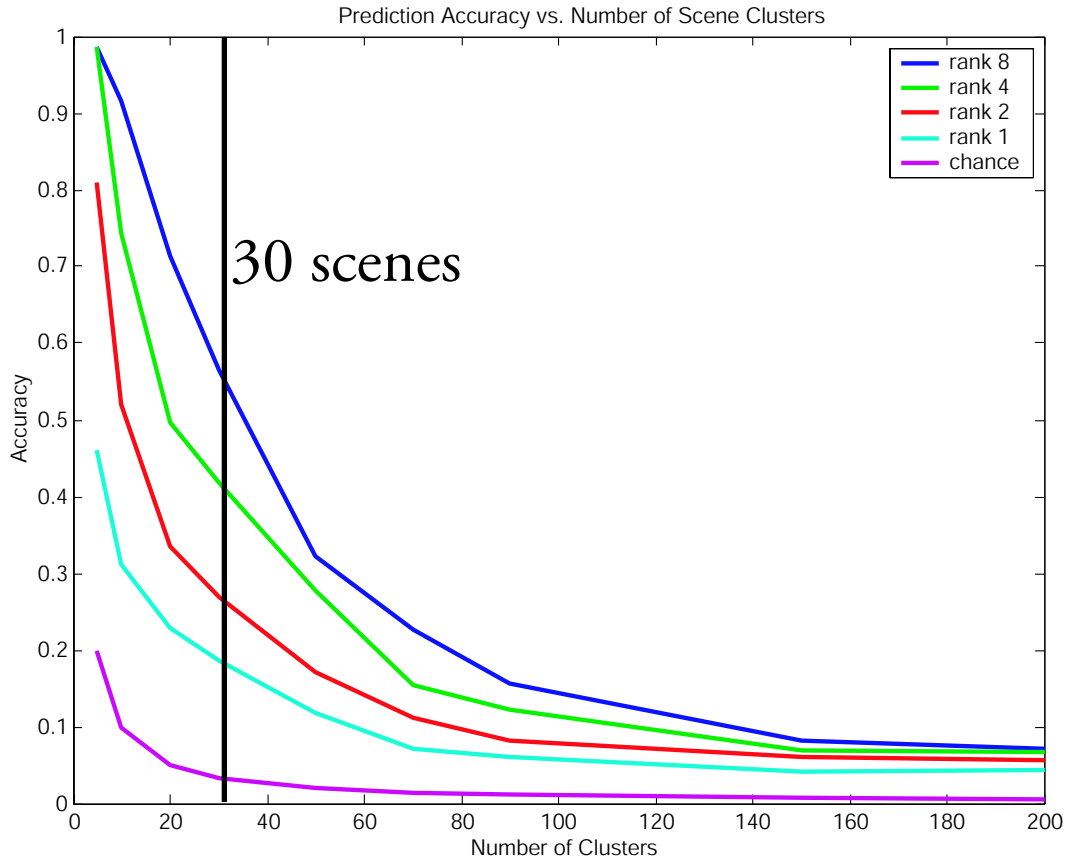

Life's Perplexity

30 Scene Clusters

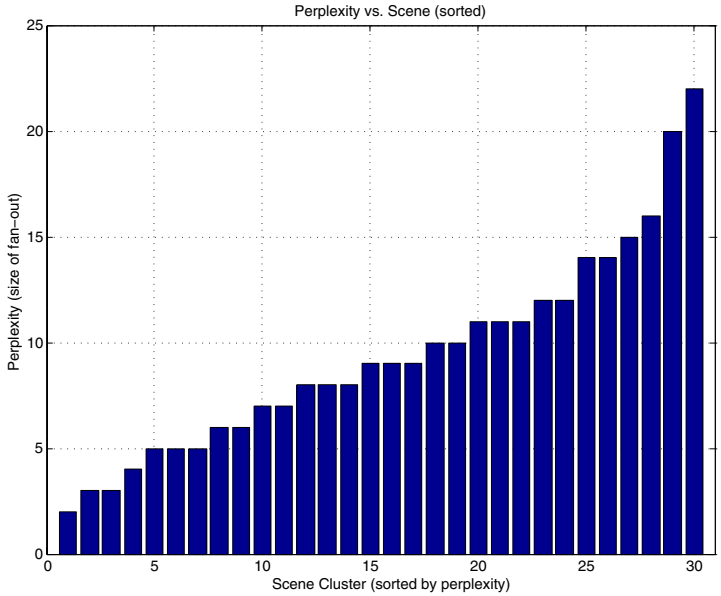

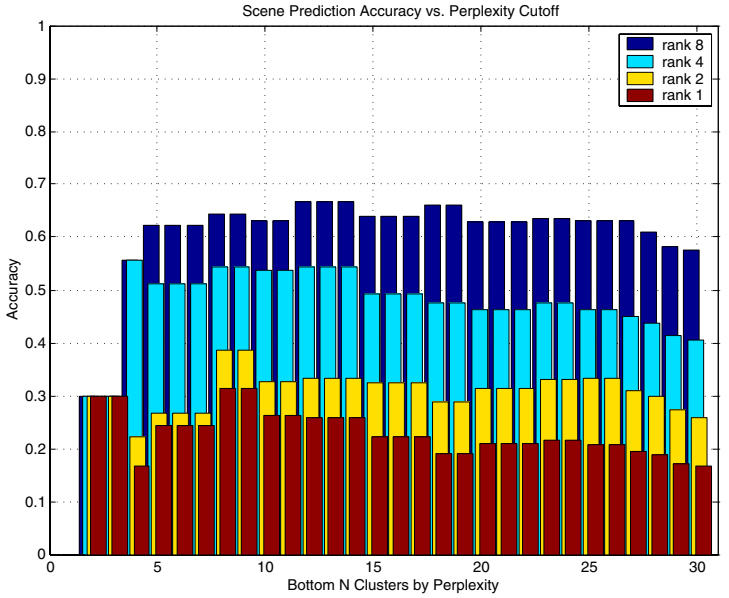

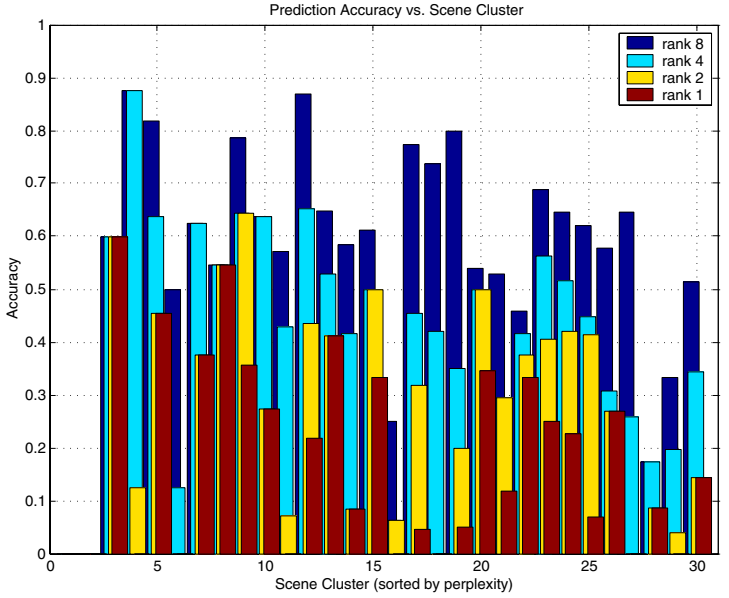

Accuracy is independent of perplexity!

Situation Classification

What is a situation?

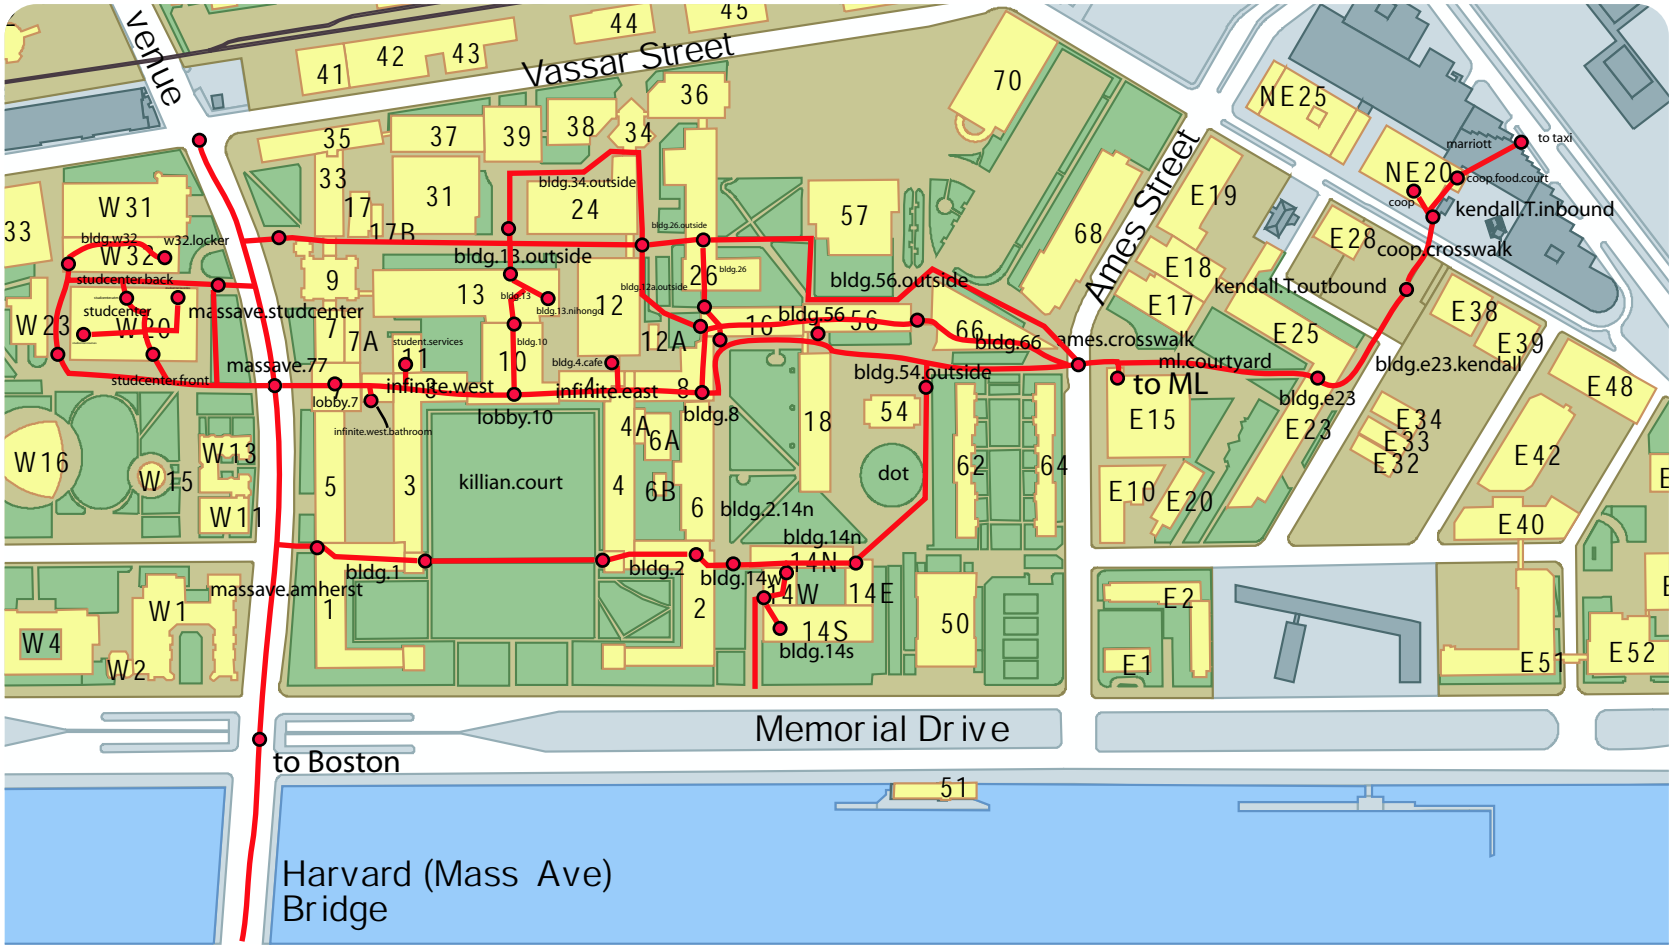

situation = location + activity

Situation Classification

19 Situations

home	apartment
neighborhood	Beacon St., Mass. Ave. (Boston-side)
bridge	Harvard Bridge, Longfellow Bridge
street	Kendall Sq., Boston Downtown, Main St., Memorial Dr., and more
hallway	Infinite Corridor and more
campus	inside & outside of bldg. 56, 66, 7, 10, and more
at work	anything in the Media Lab
elevator	any elevator
stairs	any stairs
office	my office at lab
lab	the area outside of my office
meeting	any meeting
kitchen	kitchen (at home and lab)
bathroom	any bathroom
gym	Dupont
vehicle	taxi, subway, bus
store	any store
restaurant	any restaurant
class	any class

* Every 5 minute interval over 20 days was labeled with its situation(s).

Situation Classification

Context-free

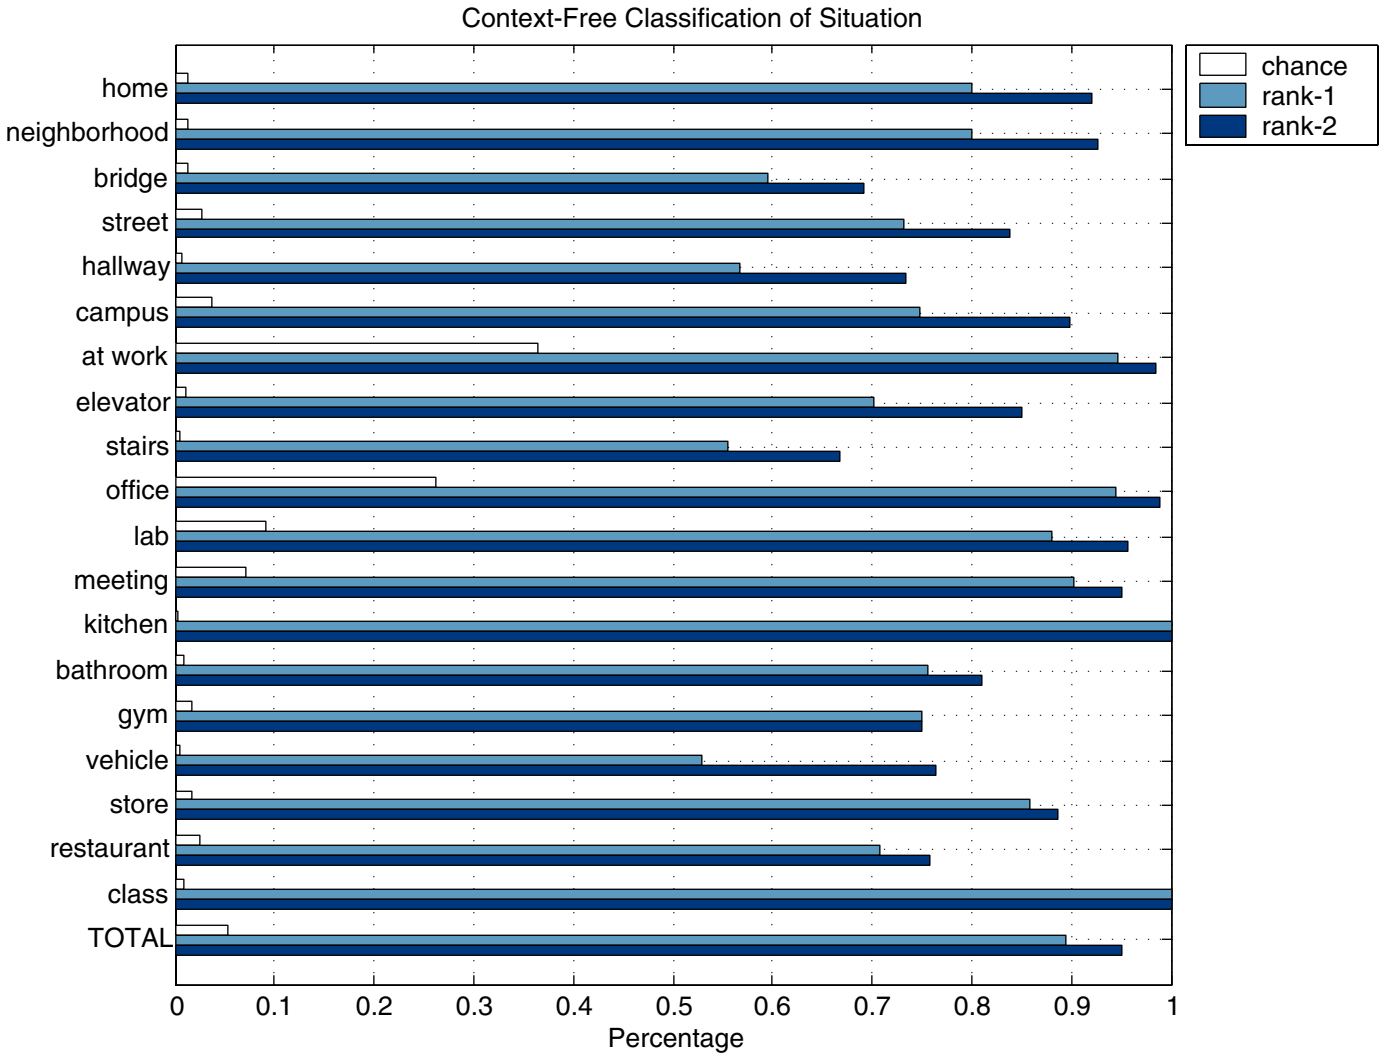

Total:
rank-1 = 89%
rank-2 = 95%

Situation Classification

Far vs. Near in Time

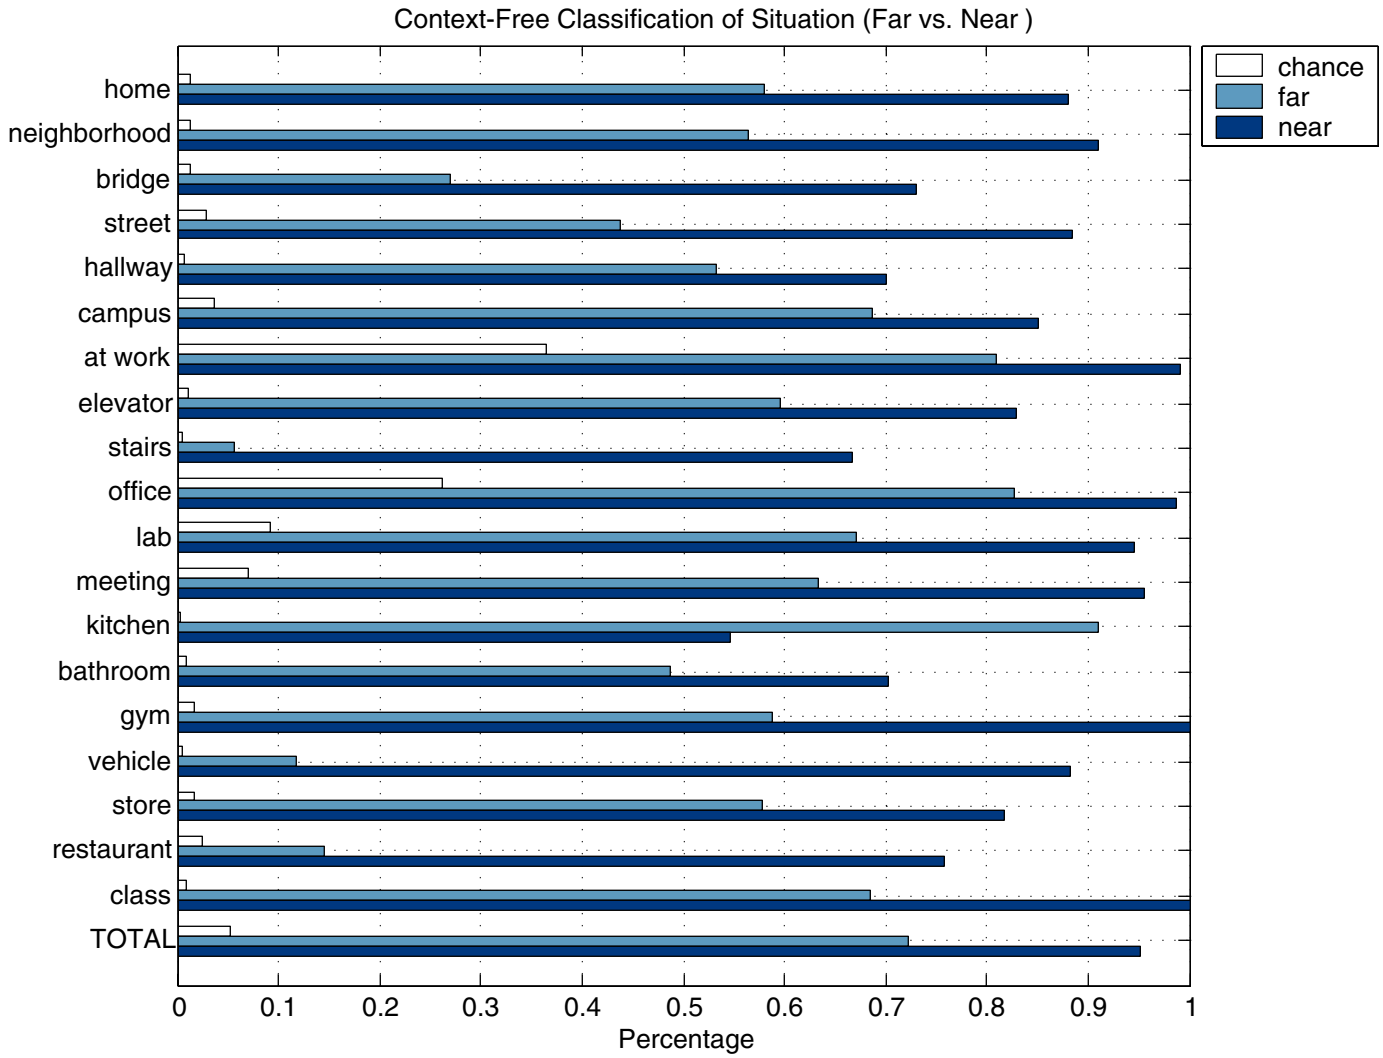

Total:
far = 72%
near = 95%

Situation Classification

with Context

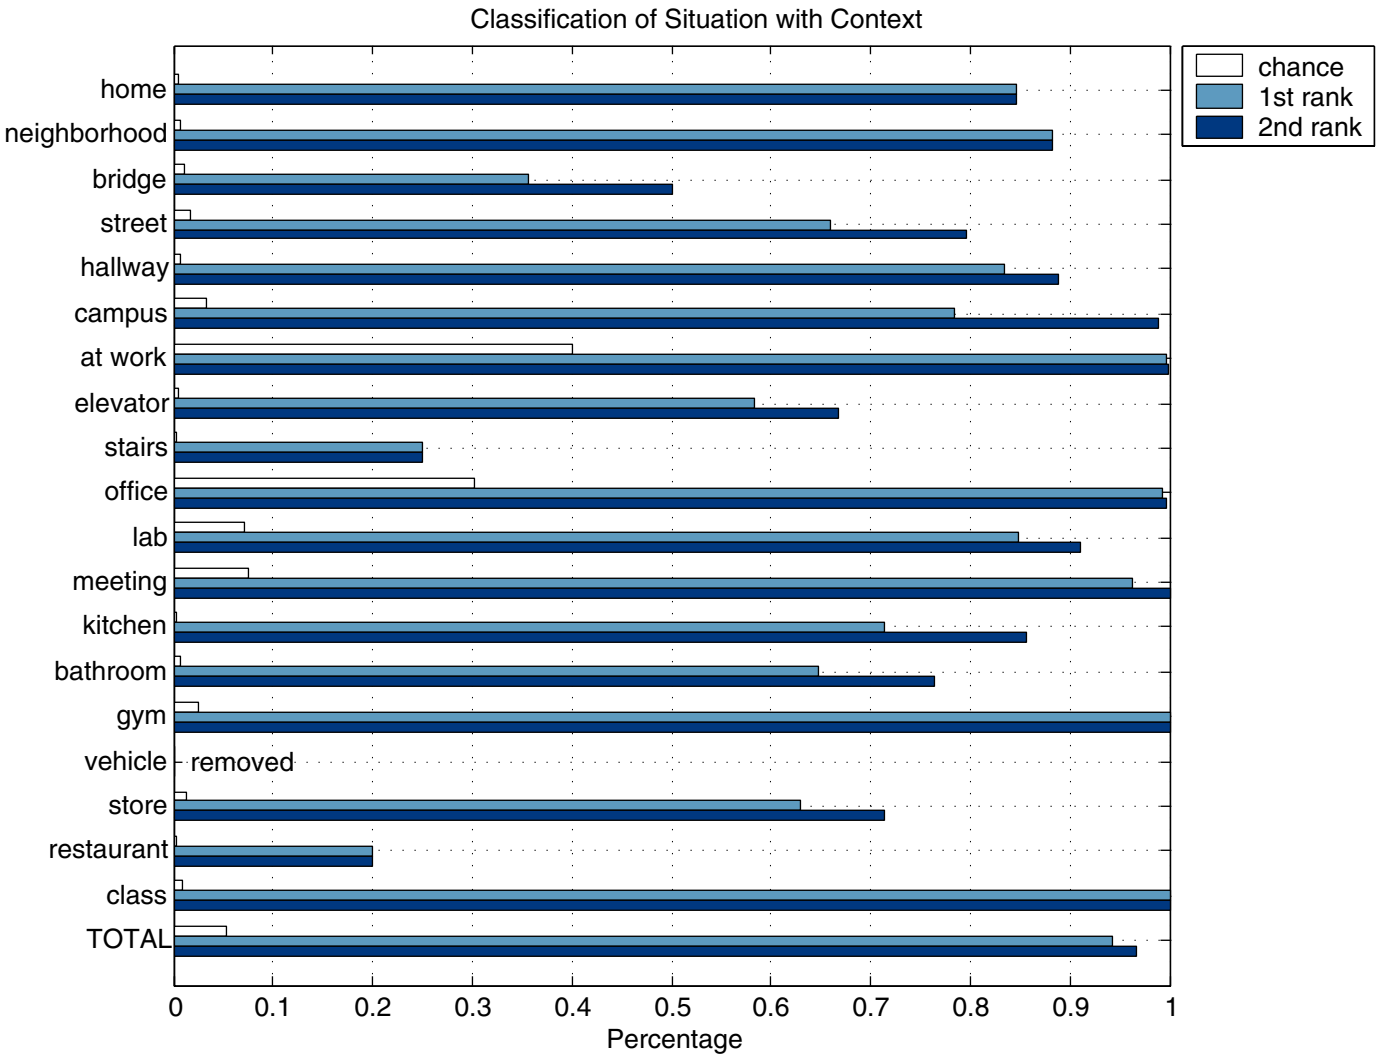

Total:
rank-1 = 94%
rank-2 = 97%

Situation Classification

A Hybrid Classifier

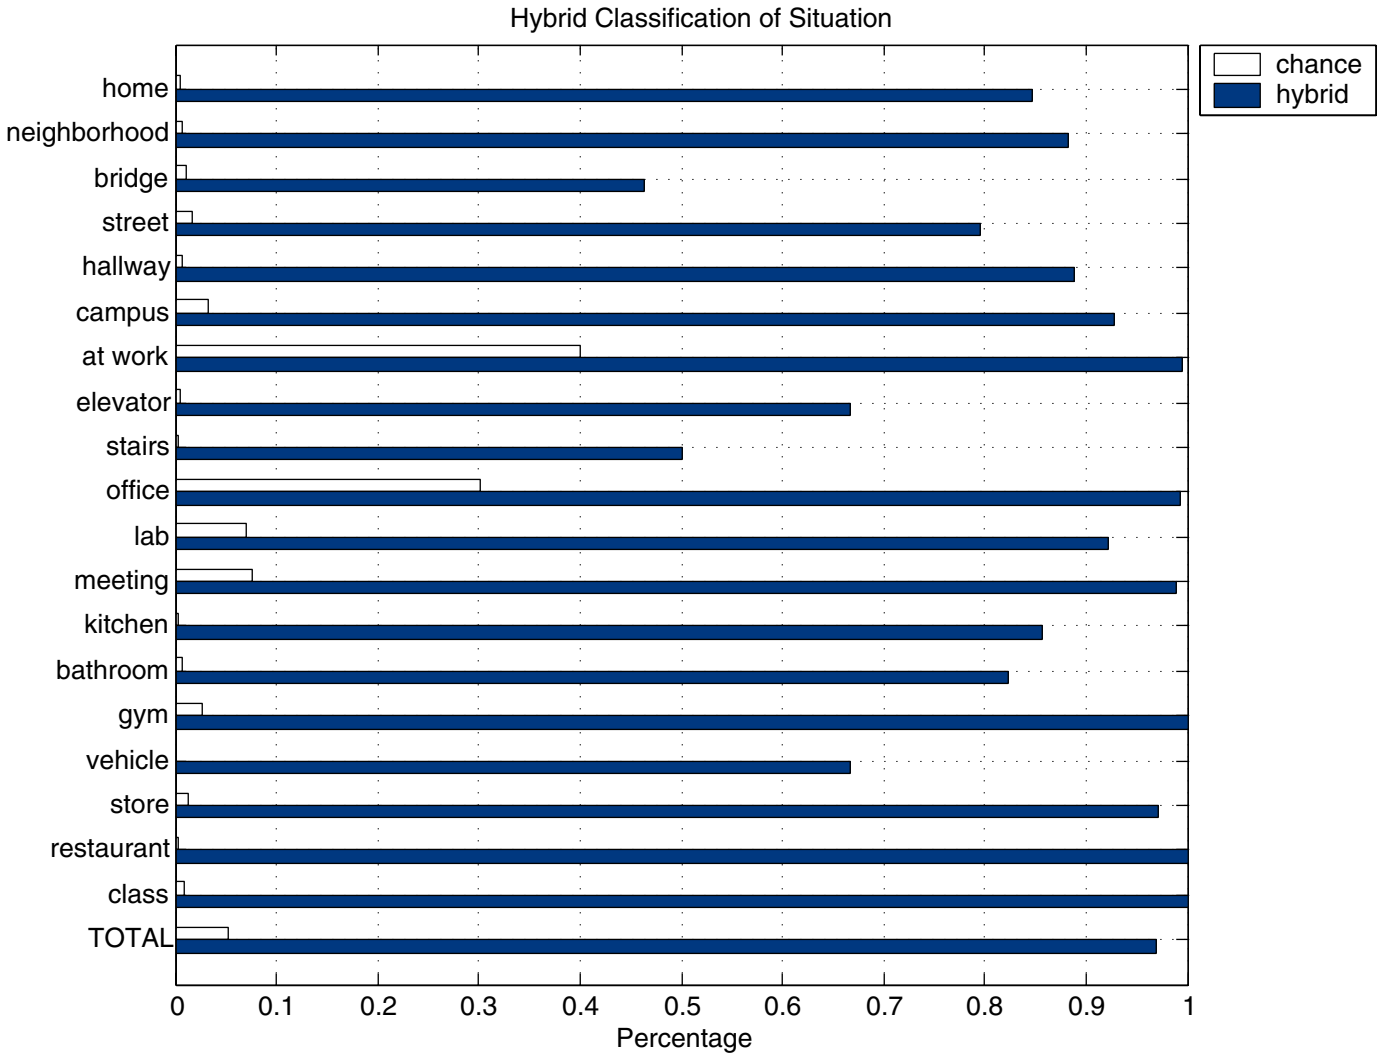

Total:
rank-1 = 97%

Epilogue

Conclusions

What have we shown?

- Capturing complete records of an individual's day-to-day life is practical.
- The capabilities of peripheral and insect-like perception.
- Simple models can capture even the complex structure of human behavior.
- How to compare moments of an individual's life at multiple time-scales.
- That day-to-day behavior is redundant at multiple scales.
- How to classify situations.

Epilogue

Acknowledgements

I am indebited to...

Sandy, my patient advisor
Joe and Trevor, my patient thesis committee
Sumit, Tanzeem and Vikram, my officemates
All the Vismodders: Karen, Rich, Nathan, Steve, Liz, Tony, and many more.

All of my friends in the gymnastics team, the breakin' crew Imobilare, and Japanese class.

I dedicate this work to my little brother, Peter.

Life Patterns

structure from wearable sensors

The End