RoboVis

Alistair Wick

EvoArm

- Small robot arm
- 3 degrees of freedom
- 3D printable
- Controlled with a Python App

- … where to go from here?

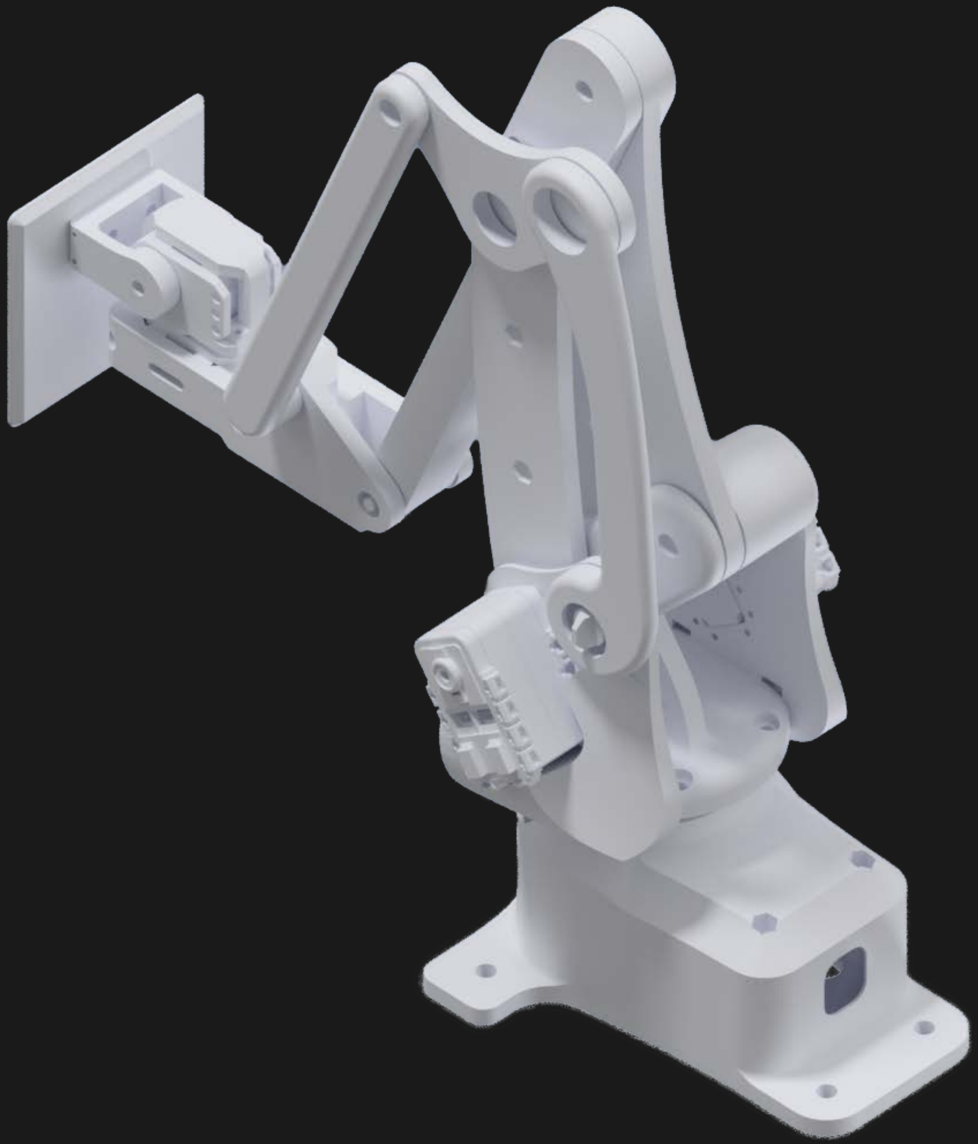

Customization!

- Every 3D-printed arm can be different
- Change mechanics for different purposes

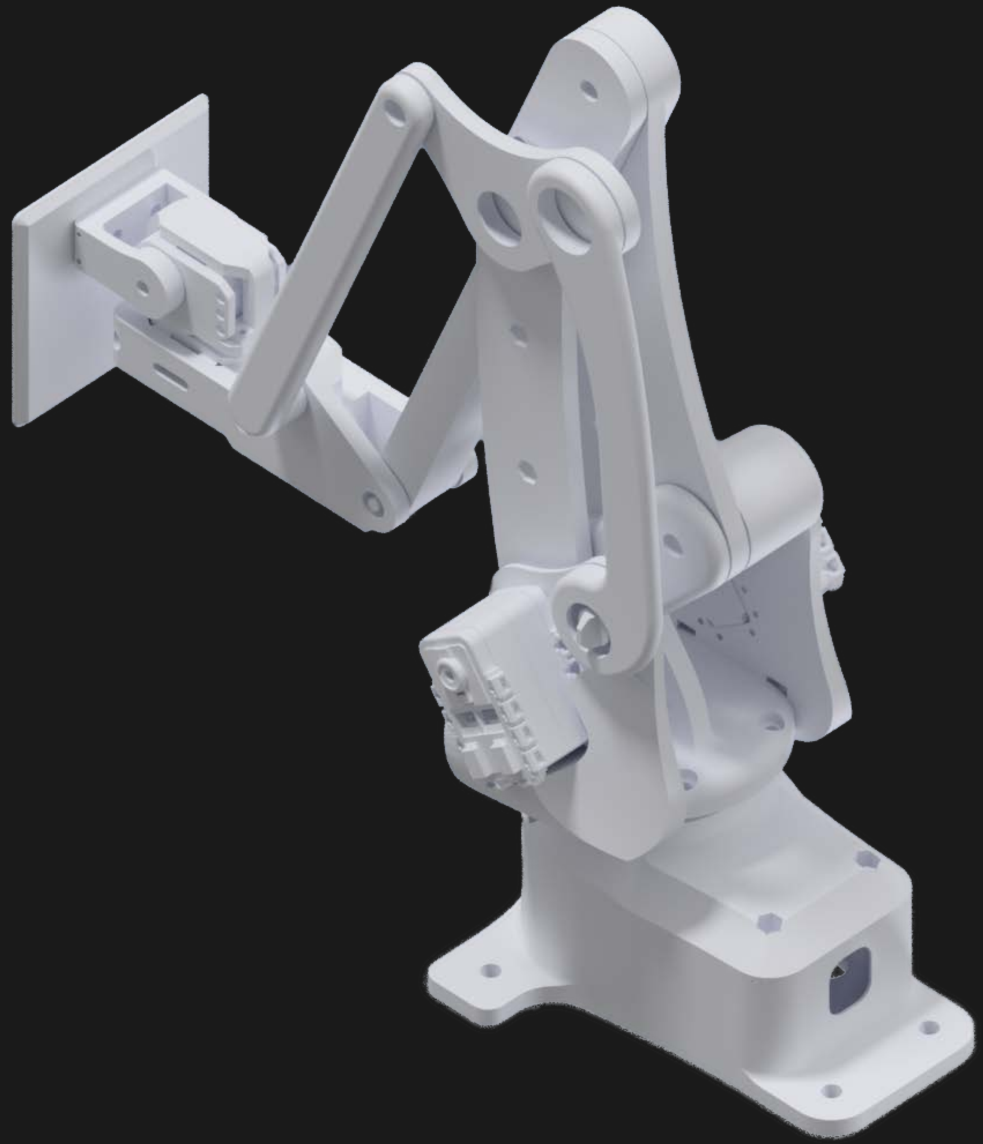

Dimensions
Design constraints
Servo constraints

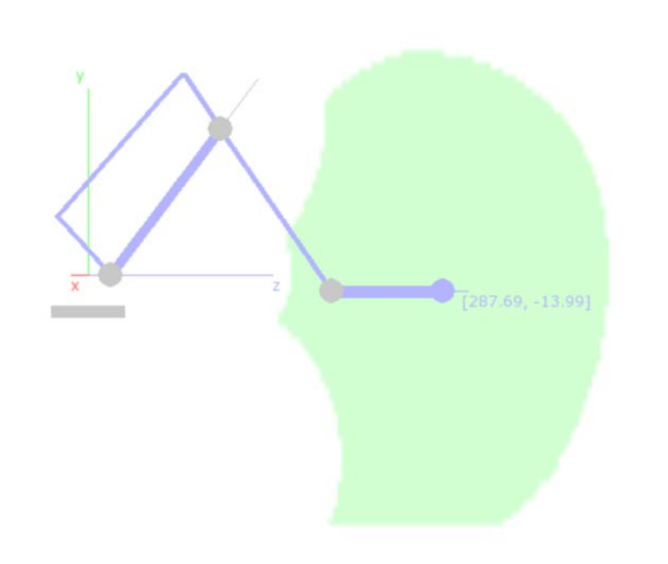

Config

IK

Solution

Dimensions
Design constraints
Servo constraints

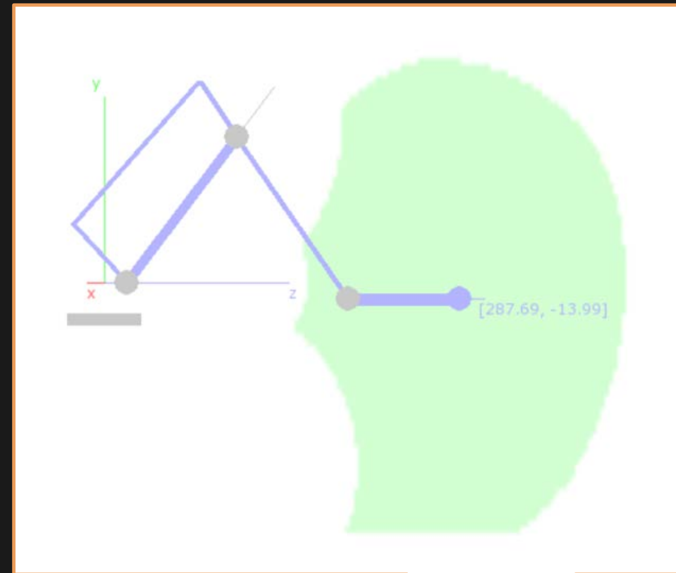

Config

IK

Solution

Exploring Possible Configs

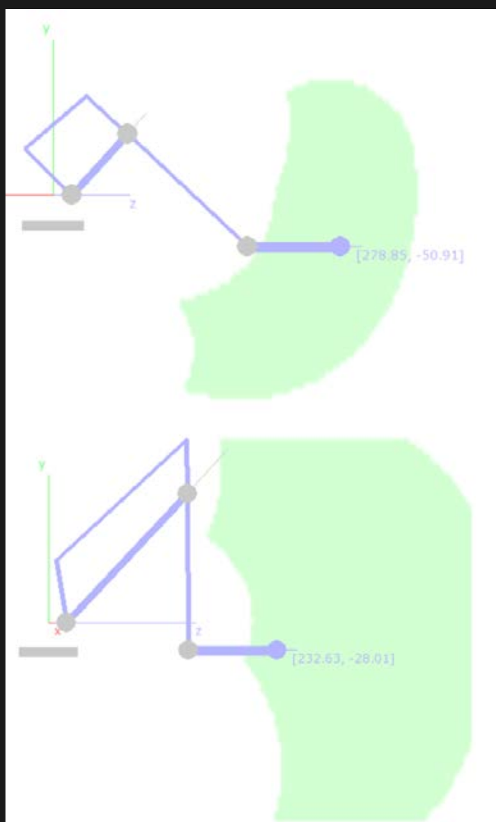

← Changing one of dozens of parameters

- Tedious and impractical to try many designs
- Different people need different capabilities
- Can the exploration process be made accessible?

Dimensions
Design constraints
Servo constraints

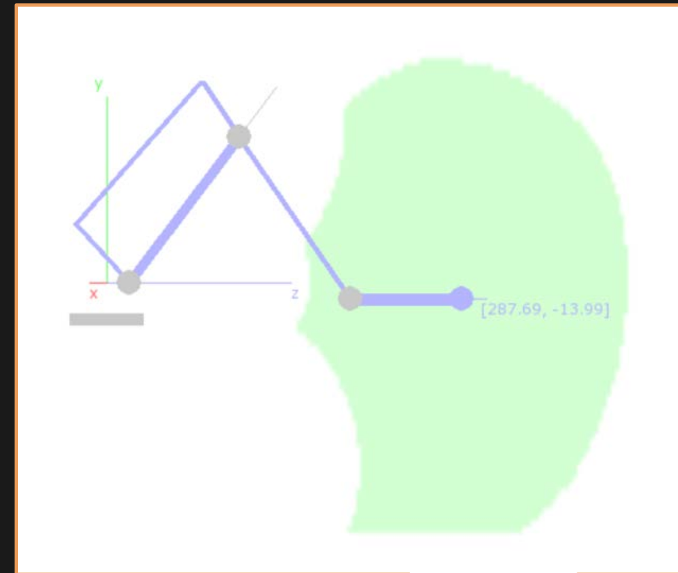

Config

IK

Solution

Want to rapidly iterate – new/different configs

Vis can help!

Vis Idea

- Interactive exploration of design space
- **Data:** Calculated online
  - Reachable points
  - Max load (across reachable space)
  - Max velocity (across reachable space)
- **Design:**
  - Spatial data -> spatial display?
  - Derive attributes?
  - Combine certain parameters?

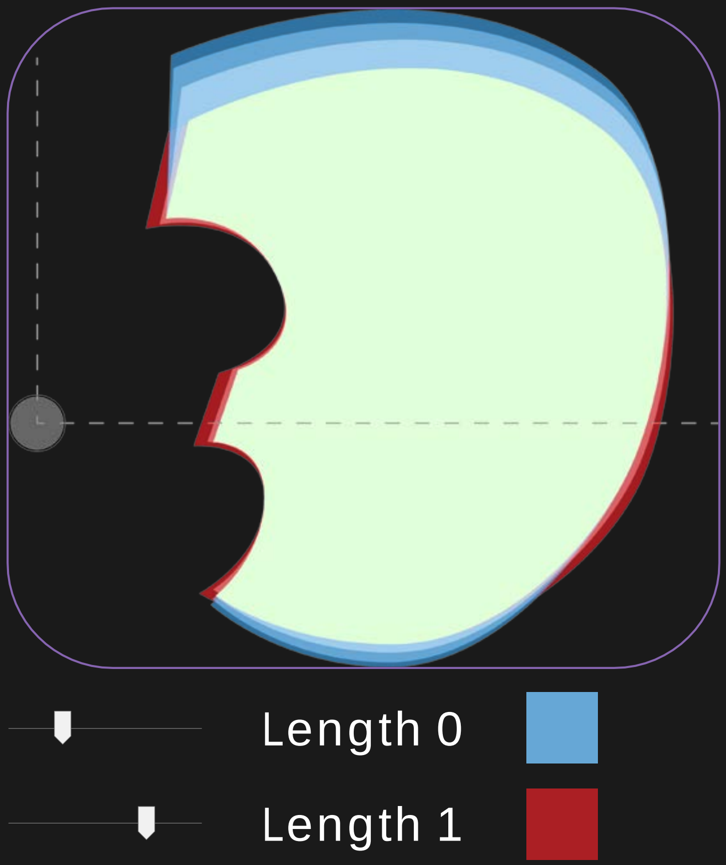

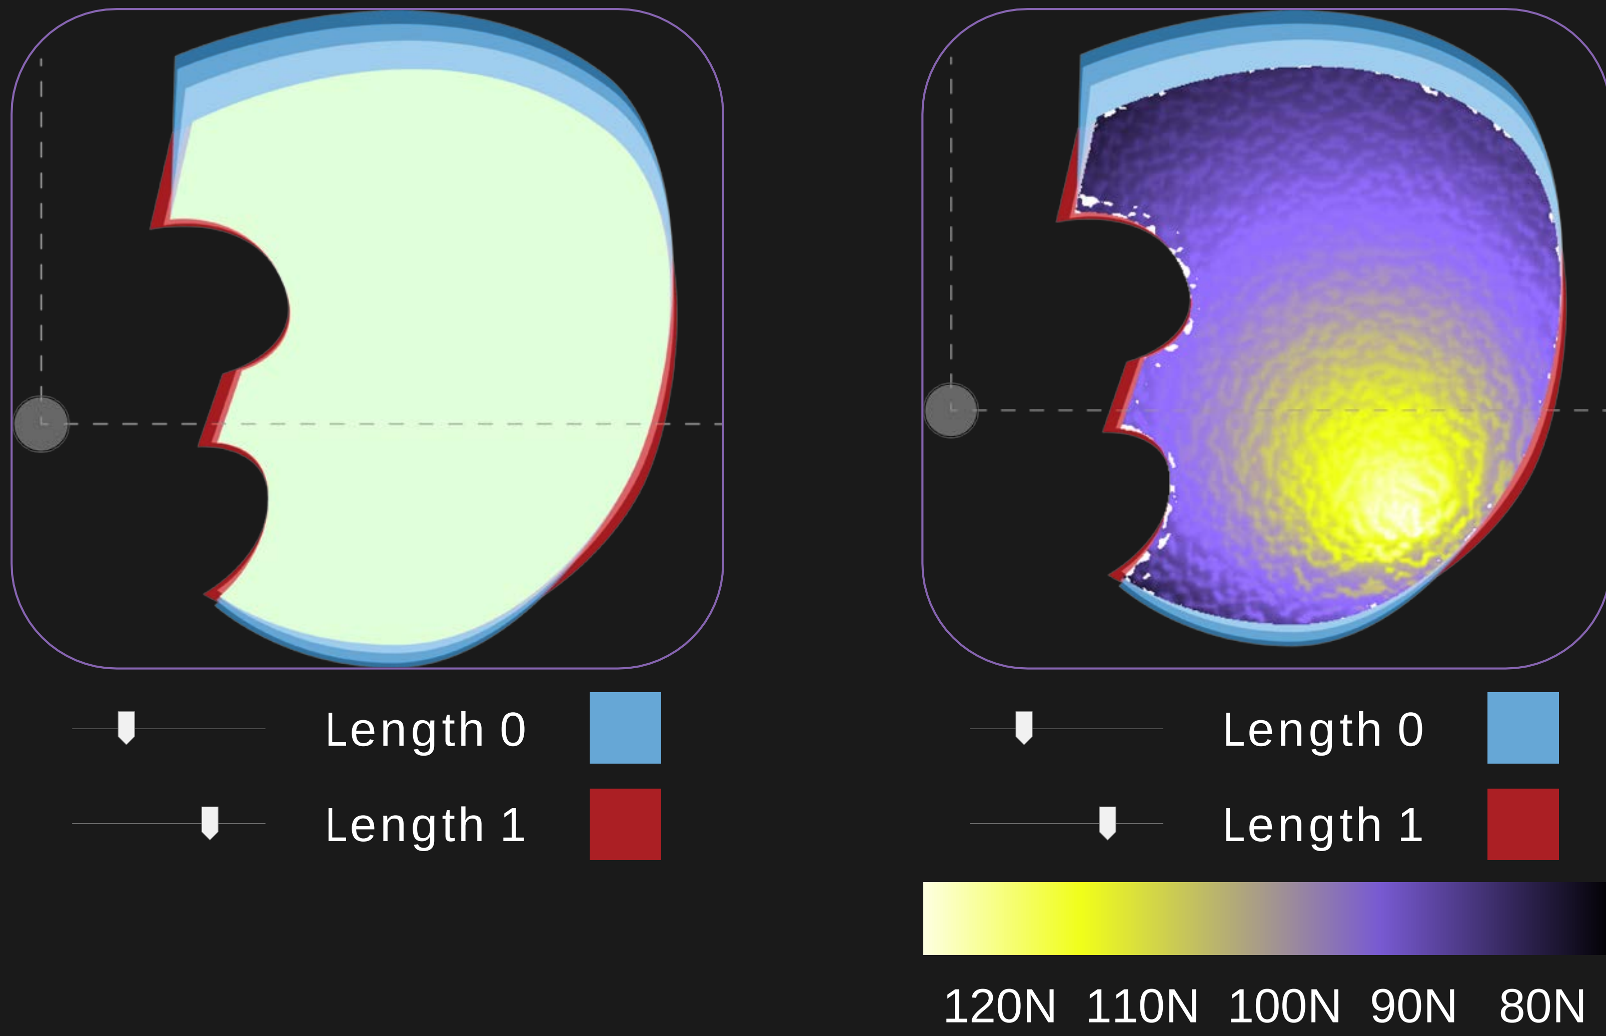
Length 0
Length 1
Length 0
Length 1
120N
110N
100N
90N
80N

PITCH:
VISUALIZING THE ENERGY PERFORMANCE OF A BUILDING

ARASH SHADKAM

WHAT

- ENERGY PERFORMANCE DATA OF A BUILDING (FOR NOW THE BUILDING IS THE CENTER FOR INTERACTIVE RESEARCH ON SUSTAINABILITY/"CIRS")
- TIME-SERIES DATA FROM SENSORS INCLUDING TEMPERATURE AND OCCUPANCY DATA (IF POSSIBLE)
- DERIVED: NORMALIZED ENERGY PERFORMANCE DATA

WHY

- BETTER UNDERSTANDING OF THE BUILDING'S ENERGY PERFORMANCE
- DISCOVERING TRENDS AND CORRELATIONS IN THE ENERGY PERFORMANCE DATA AND IDENTIFY POTENTIAL OPTIMIZATION OPPORTUNITIES IN THE BUILDING'S PERFORMANCE

HOW

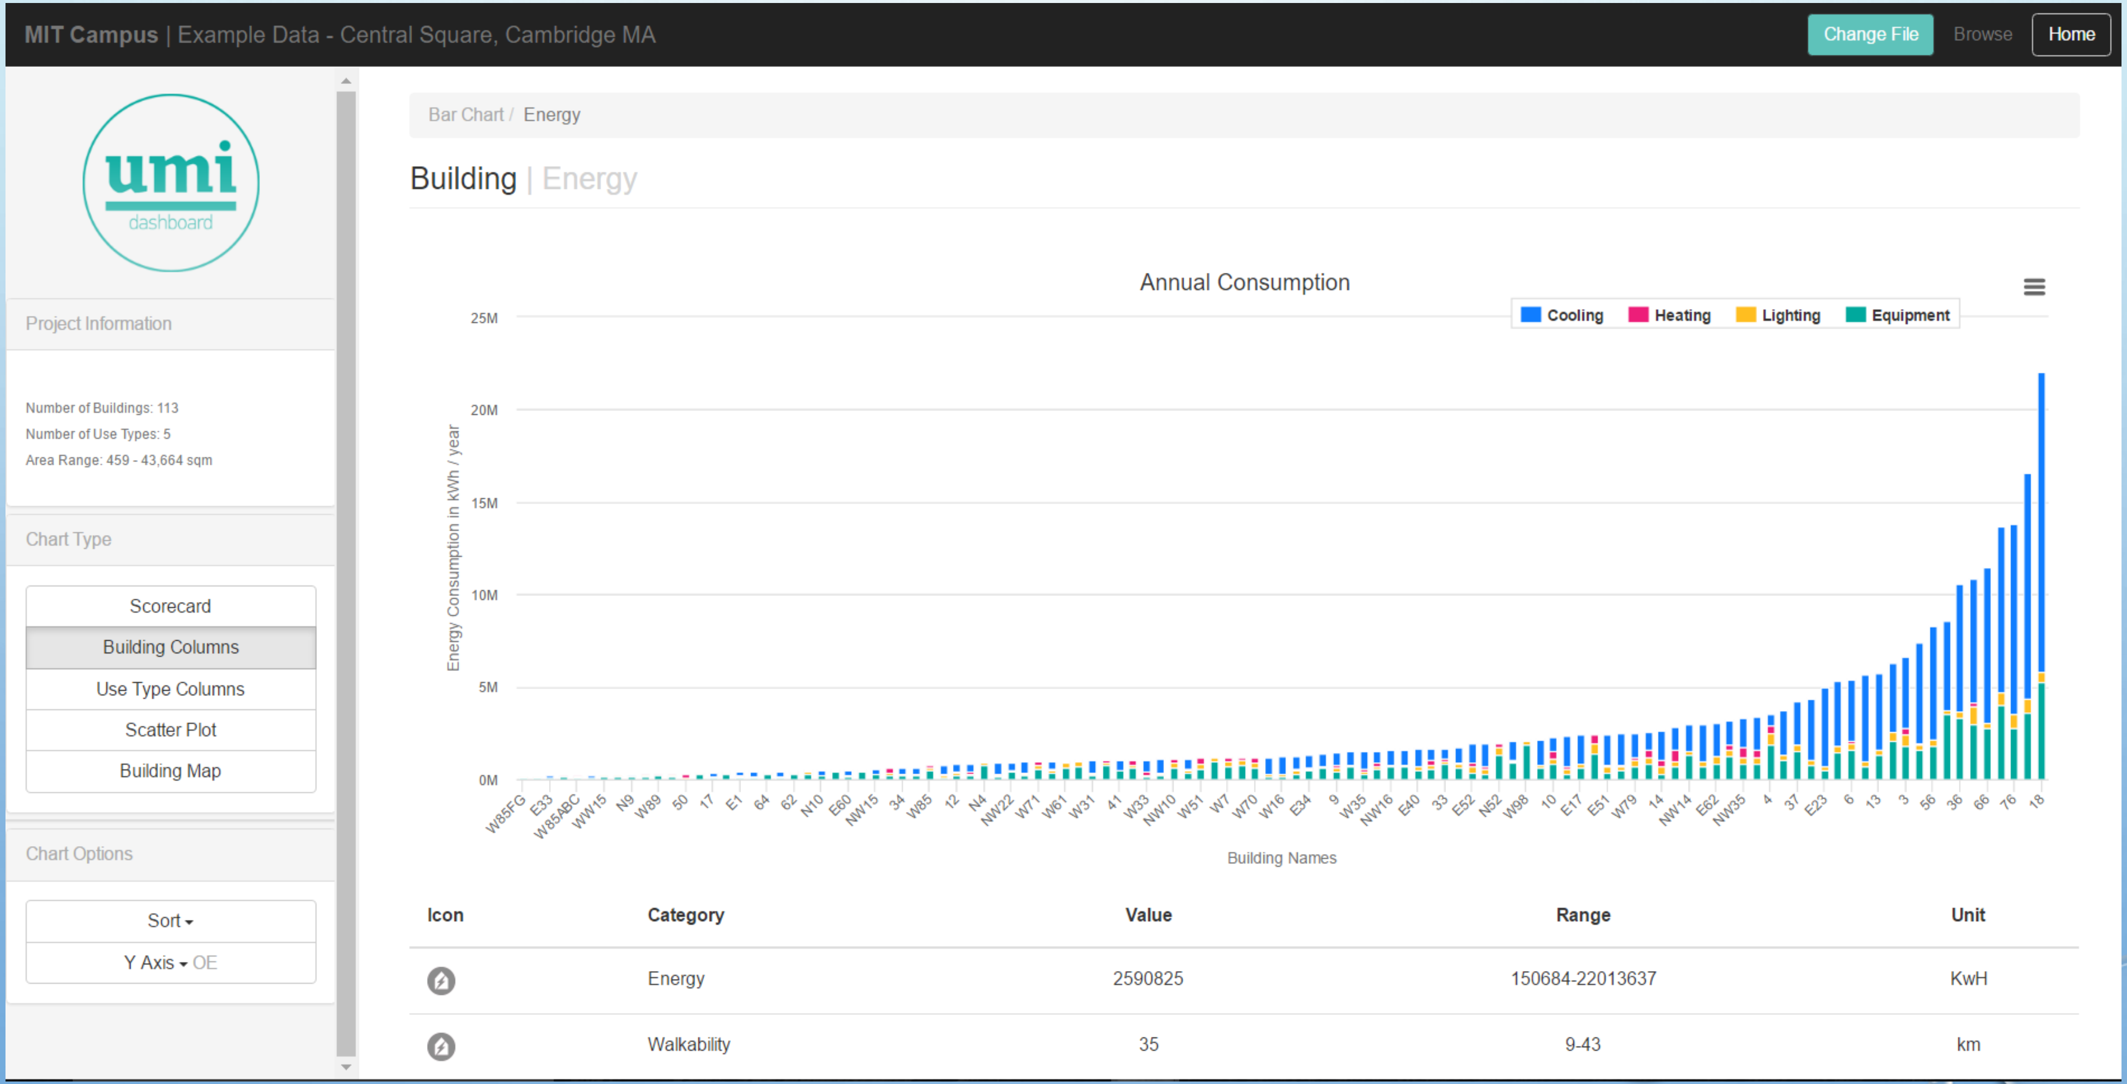

MIT Campus | Example Data - Central Square, Cambridge MA
Change File
Browse
Home
umi
dashboard
Project Information
Number of Buildings: 113
Number of Use Types: 5
Area Range: 459 - 43,664 sqm
Chart Type
Scorecard
Building Columns
Use Type Columns
Scatter Plot
Building Map
Chart Options
Sort
Y Axis OE
Bar Chart / Energy
Building | Energy
Annual Consumption
Cooling
Heating
Lighting
Equipment
25M
20M
15M
10M
5M
0M
Energy Consumption in kWh / year
Building Names
Icon
Category
Value
Range
Unit
Energy
2590825
150684-22013637
KwH
Walkability
35
9-43
km

HOW

- FACET: MULTI-FORM OVERVIEW-DETAIL VIEWS/LINKED HIGHLIGHTING
- MANIPULATE: SELECT
- REDUCE: FILTER/RANGE SLIDERS FOR DIFFERENT TIME SPANS

THANKS!

A VISUALIZATION TOOL FOR COMPUTER PROGRAM PERFORMANCE DEBUGGING

Augustine Wong

WHAT IS COMPUTER PROGRAM PERFORMANCE DEBUGGING?

Diagnosing why a computer program is running slowly

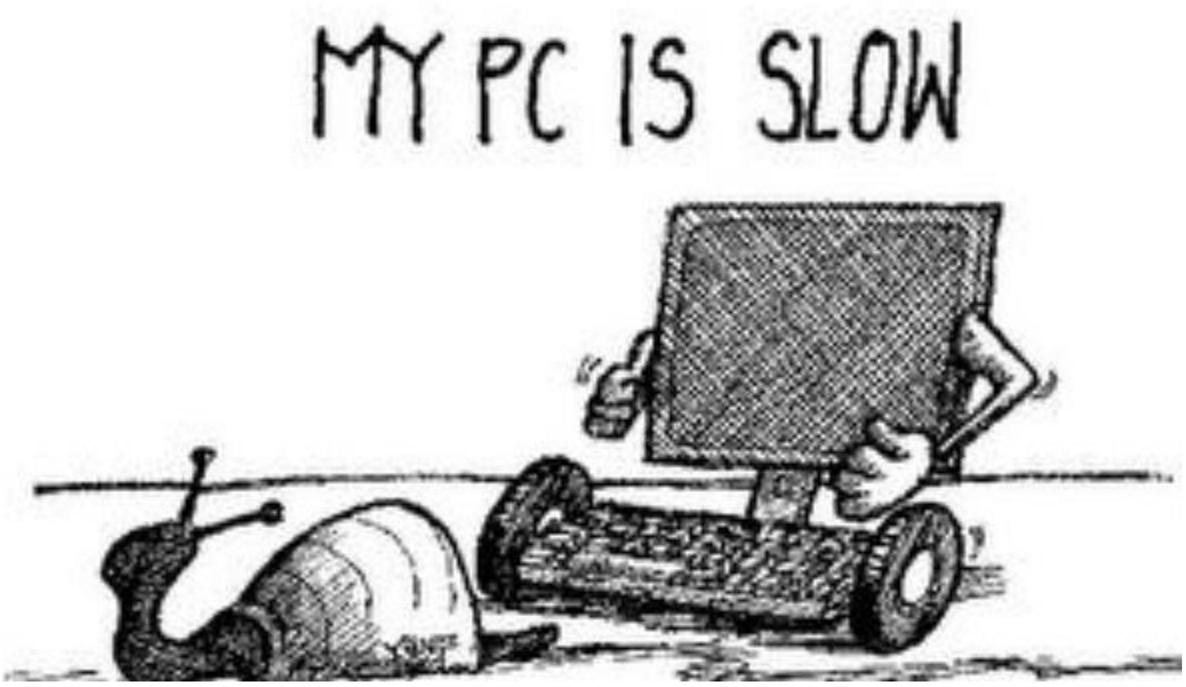

HOW DO VISUALIZATION TOOLS HELP?

Let's look at an existing visualization tool…

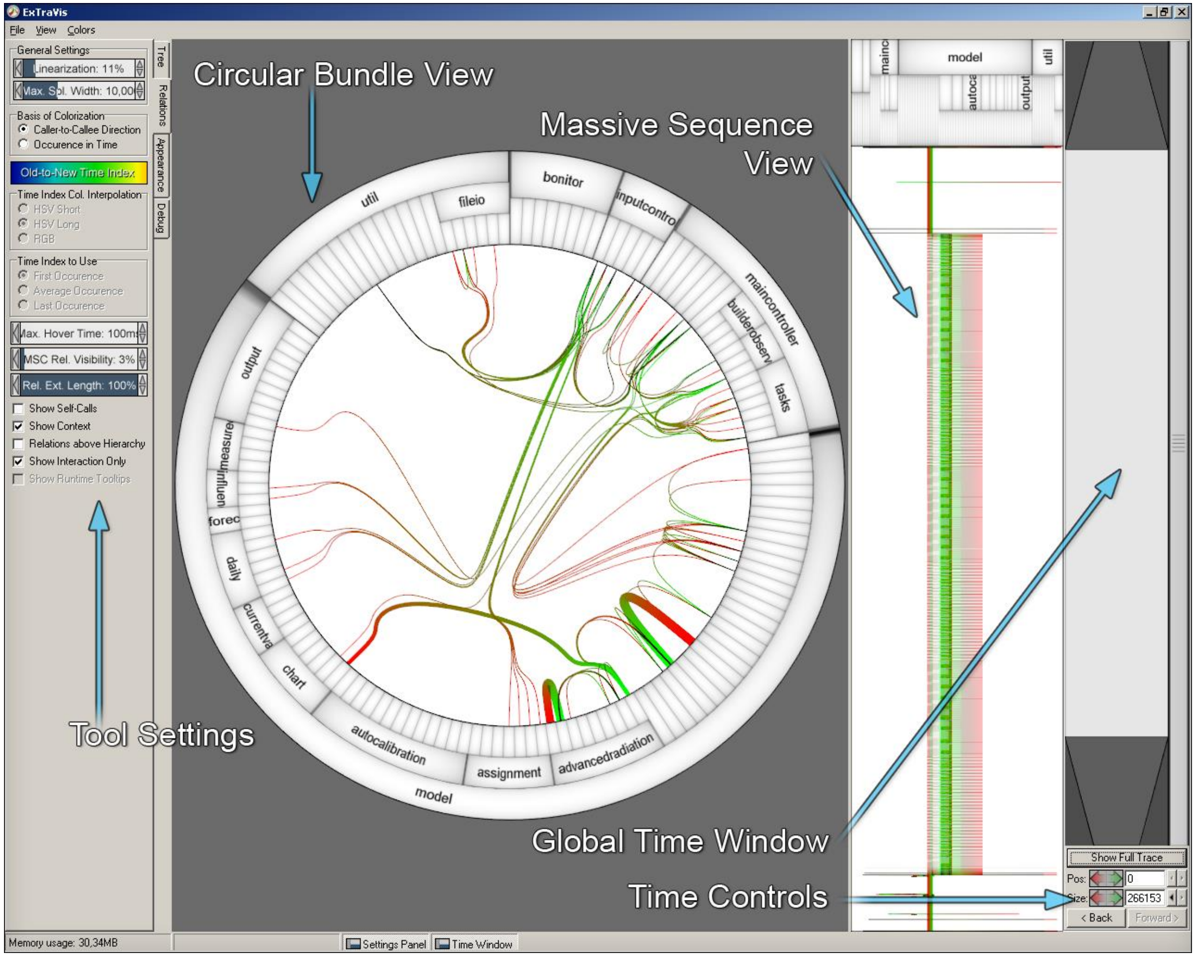

PROJECT OBJECTIVES

Create a visualization tool which:

- Uses the “search, show context, expand on demand” approach
- Visualizes “patterns” of computer program behavior
- Evaluates which patterns are good starting points for initially exploring the computer program

Quantum Annealing Visualization

Austin Wallace
5th year undergraduate student
Integrated Science-Machine Learning

Chimera Graph

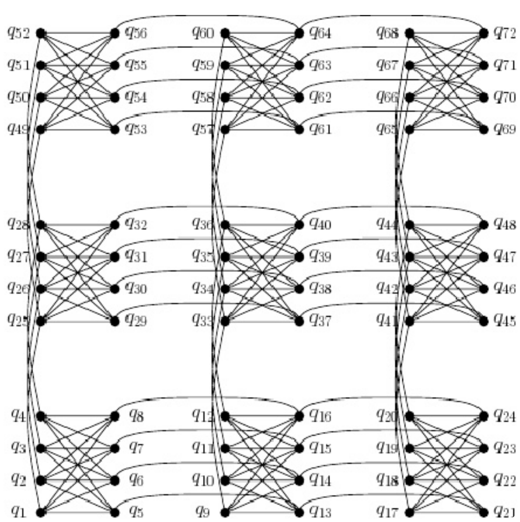

Visualizations For Justifying Machine Learning Predictions

David Johnson

Motivation

- Strengths of ML allowed expansion to diverse fields
- Fields and contexts far removed from traditional ML
- Users not trained in ML

Motivation

- Biggest factor for users is understanding how predictions occur

- Particularly important in[1]:
    - High risk applications like medicine
    - Consumer-facing applications such as Recommender Systems
    - Context-Aware applications

[1] Biran, McKeown. 2014. Justification Narratives For Individual Classification

Justification Visualization

- Visualizations present important evidence for a prediction

- Intensions are to tie in to thesis work

Yelp Visualization Tool

Dilan Ustek
Matthew Chun

Motivation

- Target User: Yelp end-users
- Comparing businesses
- Filtered visualization

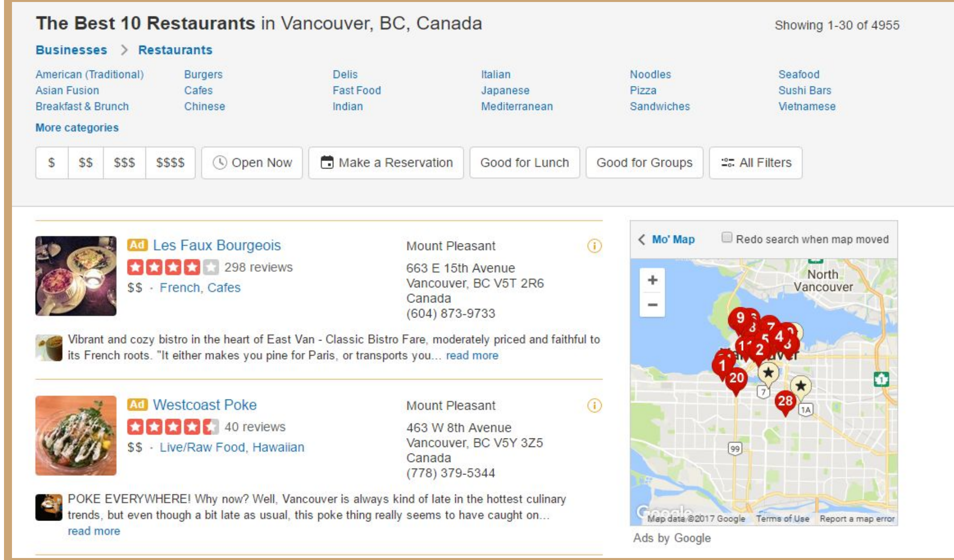

The Dataset

https://www.yelp.com/dataset_challenge/dataset

yelp_academic_dataset_business.json

```
{
    "business_id":"encrypted business id",
    "name":"business name",
    "neighborhood":"hood name",
    "address":"full address",
    "city":"city",
    "state":"state -- if applicable --",
    "postal code":"postal code",
    "latitude":latitude,
    "longitude":longitude,
    "stars":star rating, rounded to half-stars,
    "review_count":number of reviews,
    "is_open":0/1 (closed/open),
    "attributes":["an array of strings: each array element is an attribute"],
    "categories":["an array of strings of business categories"],
    "hours":["an array of strings of business hours"],
    "type": "business"
}
```

yelp_academic_dataset_user.json

```
{
    "user_id":"encrypted user id",
    "name":"first name",
    "review_count":number of reviews,
    "yelping_since": date formatted like "2009-12-19",
    "friends":["an array of encrypted ids of friends"],
    "useful":"number of useful votes sent by the user",
    "funny":"number of funny votes sent by the user",
    "cool":"number of cool votes sent by the user",
    "fans":"number of fans the user has",
    "elite":["an array of years the user was elite"],
    "average_stars":floating point average like 4.31,
    "compliment_hot":number of hot compliments received by the user,
    "compliment_more":number of more compliments received by the user,
    "compliment_profile": number of profile compliments received by the user,
    "compliment_cute": number of cute compliments received by the user,
    "compliment_list": number of list compliments received by the user,
    "compliment_note": number of note compliments received by the user,
    "compliment_plain": number of plain compliments received by the user,
    "compliment_cool": number of cool compliments received by the user,
    "compliment_funny": number of funny compliments received by the user,
    "compliment_writer": number of writer compliments received by the user,
    "compliment_photos": number of photo compliments received by the user,
    "type":"user"
}
```

Scope

- One city but yet to be decided
- Focus on the end users, aka the people who use the Yelp site/app
- Data features to consider … it depends but theme of holistic/detailed comparison
    - Discover the “nuances” behind the existing Yelp data eg. distribution of 5 star restaurants in different price categories
    - More informed decisions for end users

Project Pitch

Information Visualization 2017

Felix Grund

Munich


North Sea
Denmark
Lithuania
Hamburg
Berlin
Poland
Warsaw
Amsterdam
Netherlands
London
Brussels
Cologne
Belgium
Germany
Frankfurt
Prague
Luxembourg
Czechia
Paris
Munich
Vienna
Slovakia
Budapest
Austria
Hungary
France
Switzerland
Slovenia
Zagreb
Milan
Croatia
Belgrade
Bosnia and
Manchester


WINNERS
UEFA CHAMPIONS LEAGUE 2012/13

Your consultants, developers and troubleshooters

CONSULTING

DEVELOPMENT

OPERATIONS

Do you need advice regarding your current IT solution or targeted assistance for a specific project? Our experienced consultants uncover the optimisation potential of your existing strategy and cater to your individual needs. We offer Requirements Engineering, UI/UX Consulting, Project Management and Cloud Strategy from a single source.

EXACTLY WHAT I NEED

Atlassian
Platinum Solution Partner

Who is Scandio?

- 2016:
  - 40 employees
  - 82 clients
  - 176 projects

- Projects:
  - <u>Fixed price</u> ("client pays what's estimated")
  - Time and material ("client pays the hours")

What is a fixed price project at Scandio?

- Efforts range from 5 days - 100 days
- Duration ranges from 3 weeks – 1 year
- Before project starts: effort estimation
- Generally higher risk of “failure”
  - If over estimation in the end, company mostly has to pay (sometimes compromises with client)

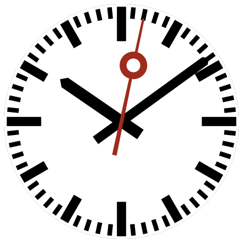

What are the project results?

- Total amount of efforts in the end
  - Exactly as estimated (rare)
  - Less than estimated (sometimes) ☺
  - More than estimated (sometimes) ☹

What are the key attributes?

1. Hours worked
   – Employees track time on project in web app
2. Degree of completion (DOC)
   – Estimated monthly by project lead
3. Hourly rate for project
   – Determined in the beginning dependent on budget and total effort
   – Changes retrospectively depending on 1 and 2

? Questions ?

- When do estimation and degree of completion conflict?
- When are our hourly rates too low?
- How do hourly rates change retrospectively?
- What tendencies can we observe over multiple projects?
- When interfere to maintain project success?
- How can we identify wrong estimations on DOC?
- How do project leads differ in their monthly estimations?
- …

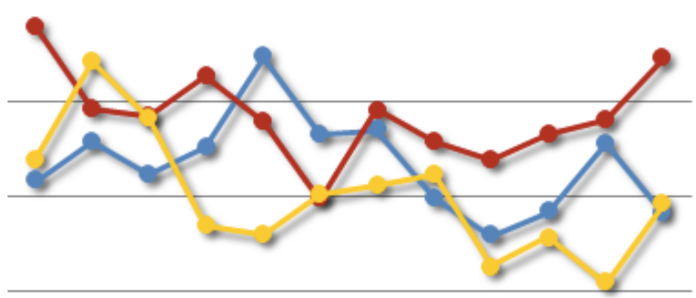

Is there still time?

Time Tracking

Tages-Informationen

Datum	von	bis	Pause	Stunden
13.05.2014	09:45	17:30	0,75	7,00

Buchungs-Informationen

Projektkennung	Tätigkeit	Issue ID	Bemerkung	Stunden
DKMS_OPT1	Dashboard	**DKMSSUPP-95**	Dashboard Anpassungen	1,00
DKMS_OPT1	Dashboard	**DKMSSUPP-96**	Dashboard styling	0,75
DKMS_OPT1	News-Umstrukturierung	**DKMSSUPP-102**	News Styling	1,25
BSH_SUPP2014	PIWIK Support	**BSHP-22**	Test System Sync	0,50
BSH_SUPP2014	Corporate Wiki Support	**BSHSUPP-226**	Test Upgrade	0,25
INF_SEC2014	POC - User Switch	**INFPRJ-18**	DiskussionenTests/Implementierung Tomcat Proxy	1,00
DKMS_OPT1	TPL		Telko mit Guido zu Projektstand	1,00
DKMS_OPT1	Dashboard	**DKMSSUPP-119**	Dashboard Performance	1,25

Project results (good)

Mitarbeiterstunden

Mitarbeiter	Stunden
fgrund	60,00
gschmidl	12,50
Summe	**72,50**

Projektstunden

Tätigkeit	Plan	Ist	Rest
DKMS_OPT1: Blogposts-Plugin	16,00	14,75	1,25
DKMS_OPT1: Dashboard	16,00	15,25	0,75
DKMS_OPT1: IE10 Anpassungen	8,00	2,75	5,25
DKMS_OPT1: News-Migration	8,00	0,75	7,25
DKMS_OPT1: News-Umstrukturieru…	8,00	[illegible]	-2,25
DKMS_OPT1: Patient Faces	24,00	[illegible]	12,75
DKMS_OPT1: TPL	4,00	[illegible]	-13,50
Summe	**84,00**	**72,50**	**11,50**

Project results (bad)

Mitarbeiterstunden

Mitarbeiter	Stunden
fgrund	52,00
jgrabski	10,75
jstadler	9,00
Summe	**71,75**

Projektstunden

Tätigkeit	Plan	Ist	Rest
DKMS_SSO: SSO Link	0,00	14,25	-14,25
DKMS_SSO: SSO Plugin	32,00	57,50	-25,50
Summe	**32,00**	**71,75**	**-39,75**

Thanks.

Visualizing Internal Components of a Convolutional Neural Network

Mahdi Ghodsi - Hooman Shariati

Background:

What is Machine Learning

Machine Learning is taking over.

Applied to many fields:

Bioinformatics, Gaming, Medical diagnosis, Marketing, Machine Vision, ….

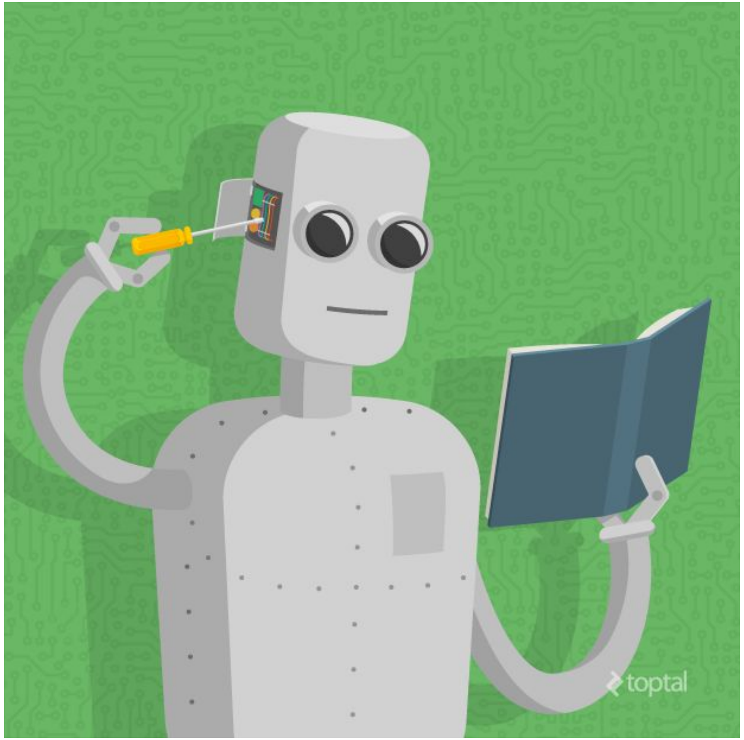

Convolutional Neural Network

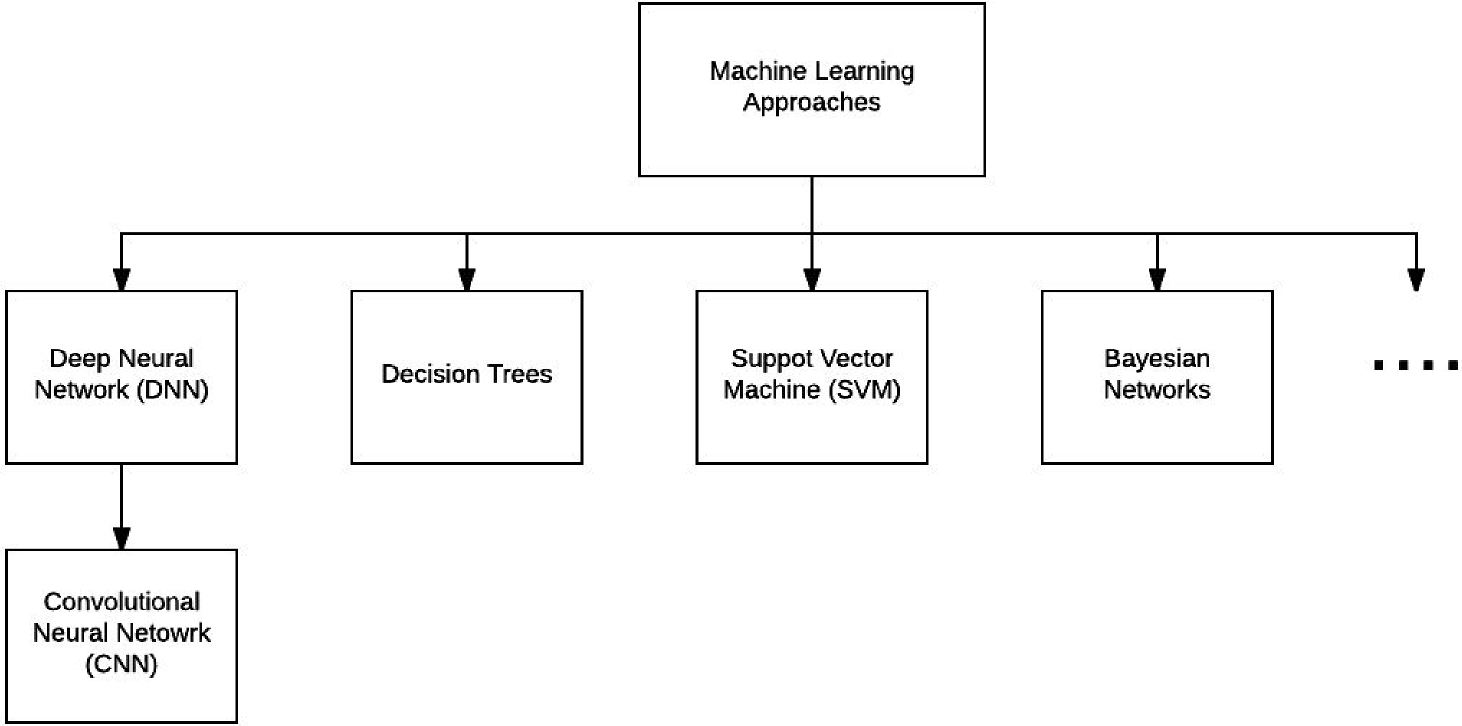

Convolutional Neural Network

The idea has been around since 1980s
But Introduction of GPU computing with 30x speed up gave DNNs a boost

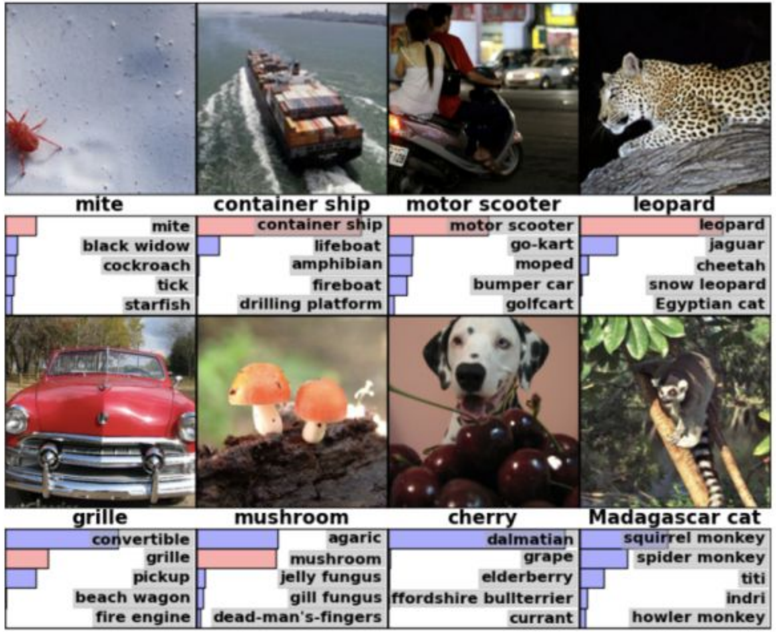

ImageNet Competition

Google Deep Dream

Convolutional Neural Network

Very Popular Research Area

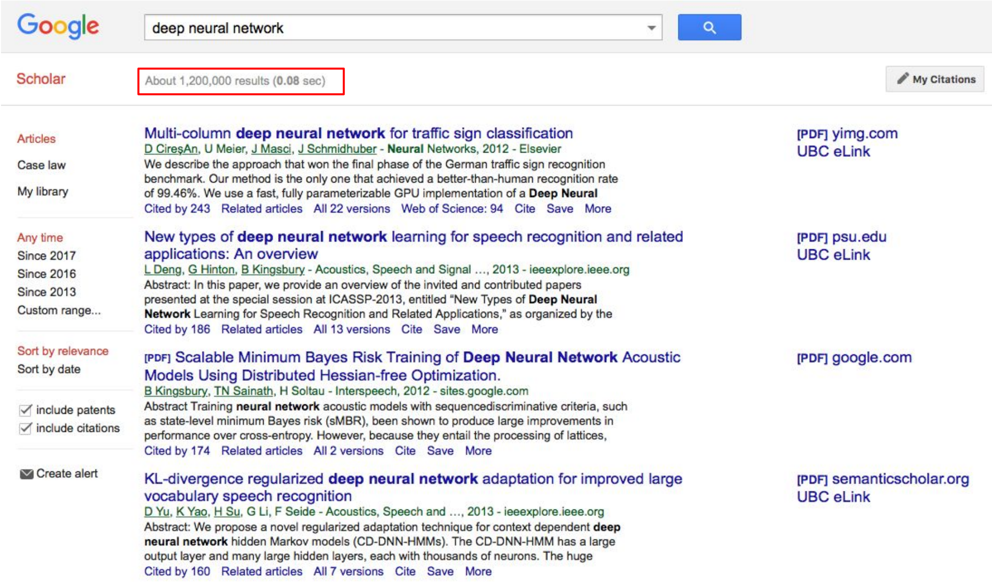

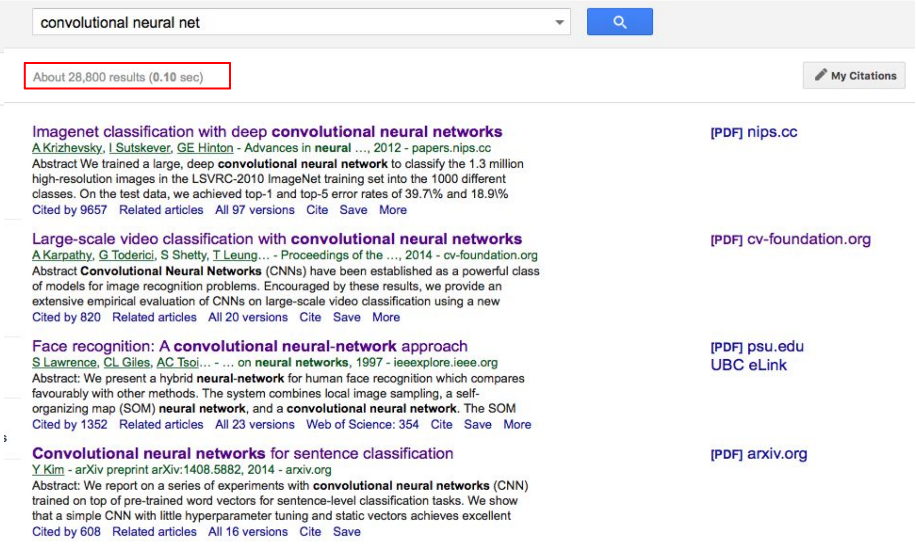

Convolutional Neural Network

However ...

Convolutional Neural Network

How researchers see CNNs

"Neural networks have long been known as "black boxes" be- cause it is difficult to understand exactly how any particu- lar, trained neural network functions due to the large number of interacting, non-linear parts."

Yajin Zhou

Department of Computer Science North Carolina State University

Convolutional Neural Network

How researchers see CNNs

How CNNs looks like

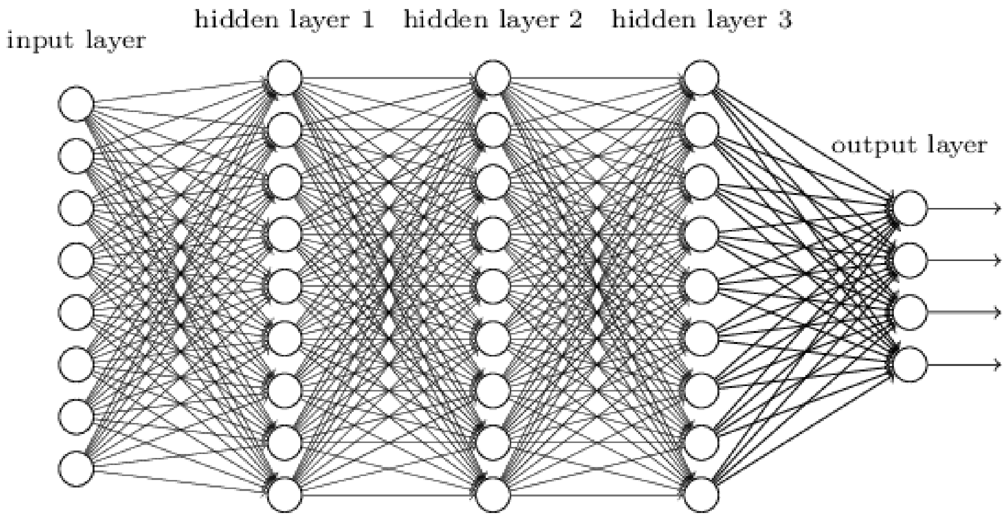

Visualizing and making sense of of CNNs in literature:

Visualizing and Understanding Convolutional Networks By M. Zeiler (NYU)

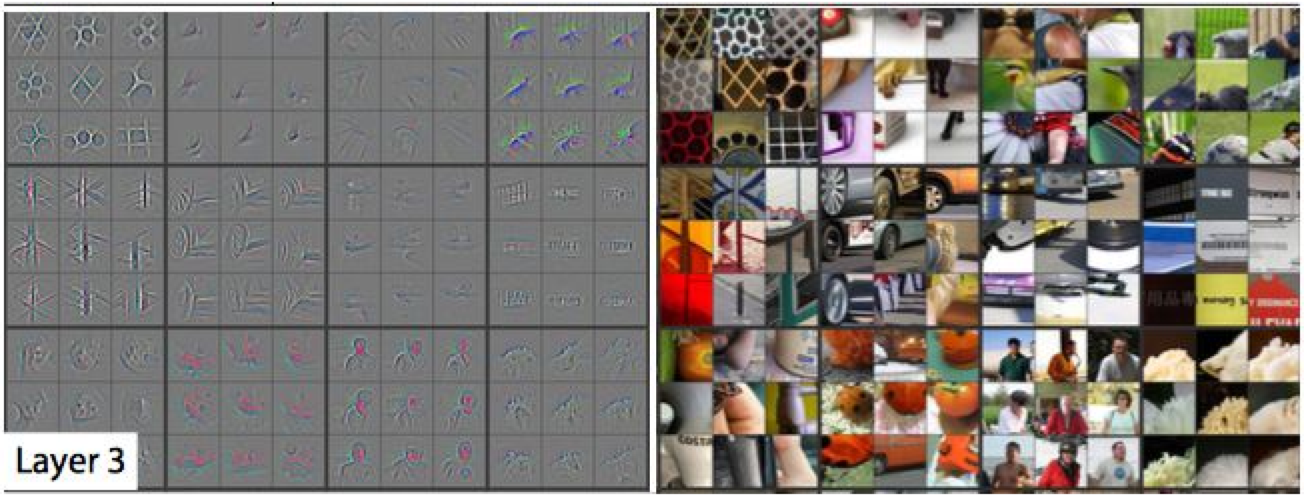

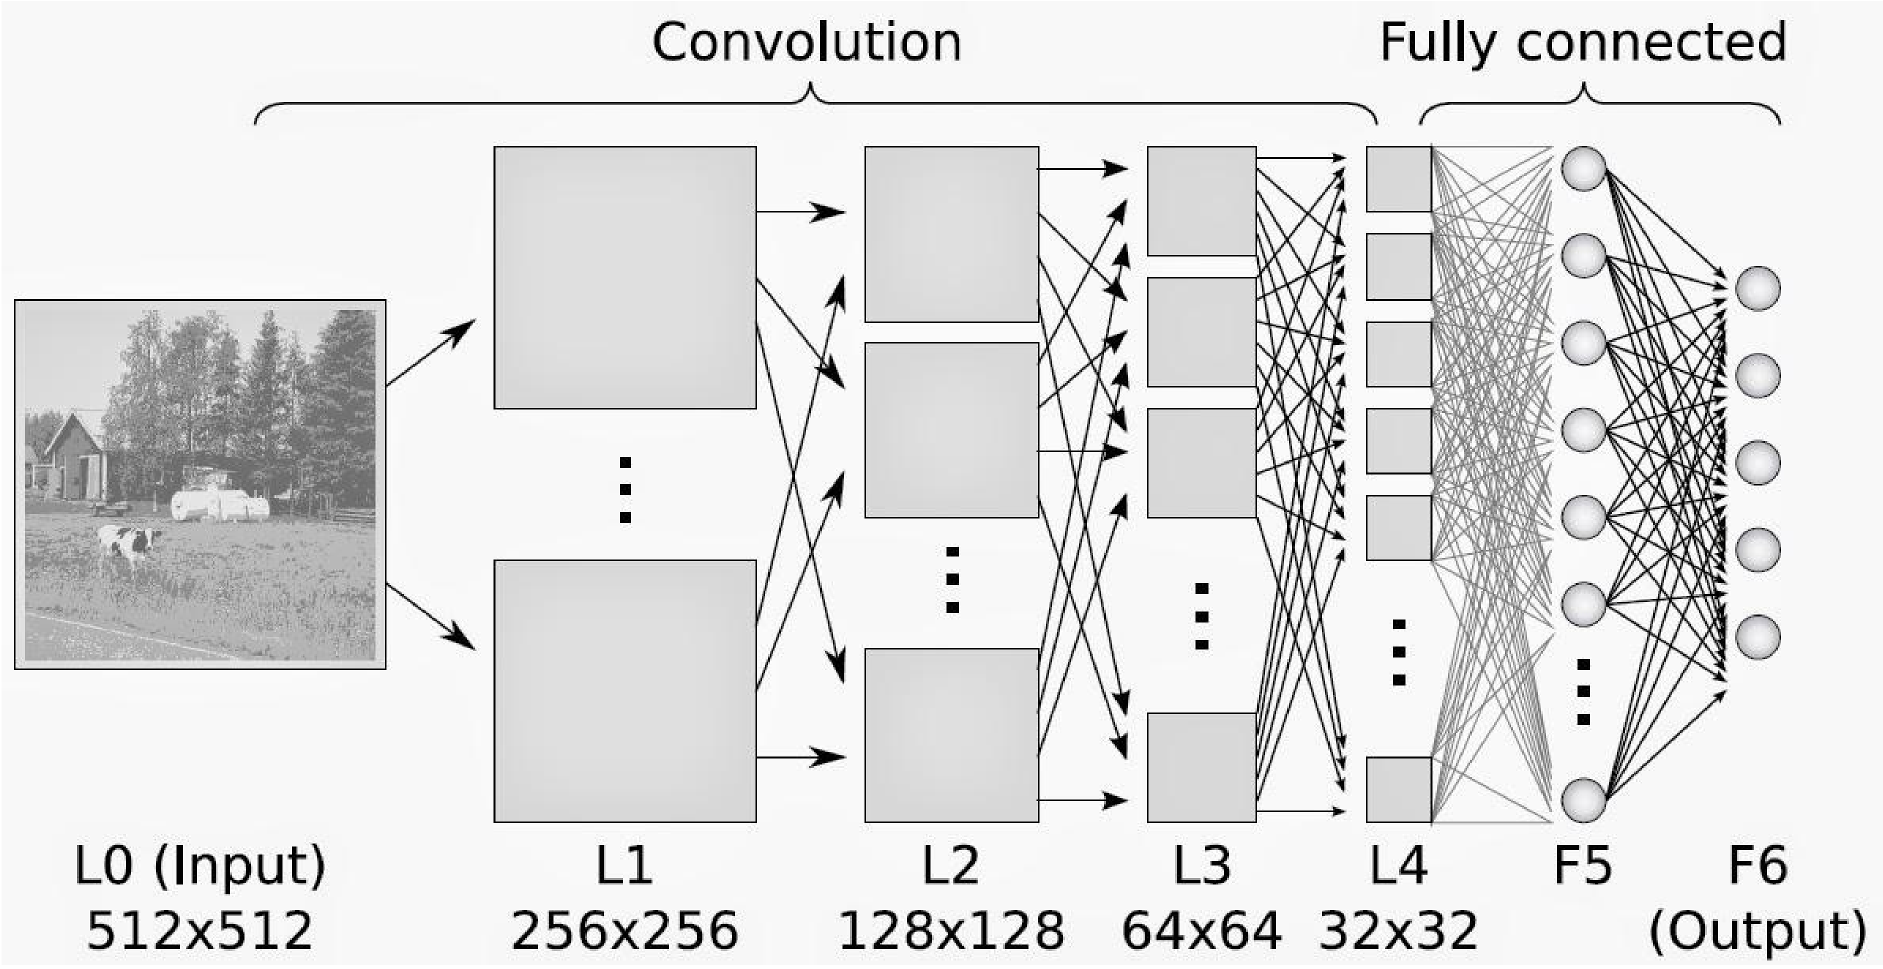
Convolution
Fully connected
L0 (Input)
512x512
L1
256x256
L2
128x128
L3
64x64
L4
32x32
F5
F6
(Output)

Visualizing Ambiguity

James Hicklin

Case Scenario

- Imagine you are Betty
- Just finished chemo for breast cancer
- Typical post-chemo therapy is Tamoxifen for 5 years

Tamoxifen 10-year risk estimates compared to 5-year risk estimates (out of 1000)

Attribute	Change
Breast cancer recurrence	↓ 28
Death from breast cancer	↓ 28
Development of gallstones	↑ 2
Development of endometrial cancer	↑ 16
Stroke	↑ 2

Point estimates…

- Imagine Betty only cared about her chance of dying from breast cancer and her chance of developing endometrial cancer

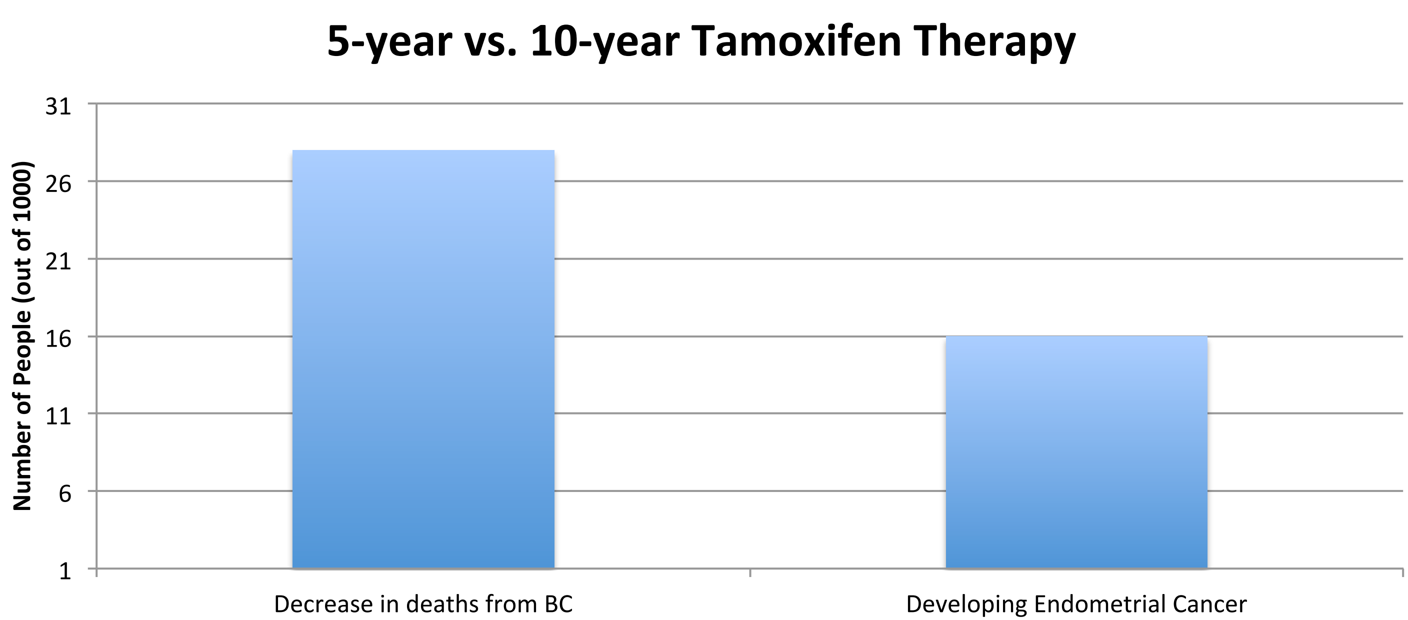

With confidence intervals…

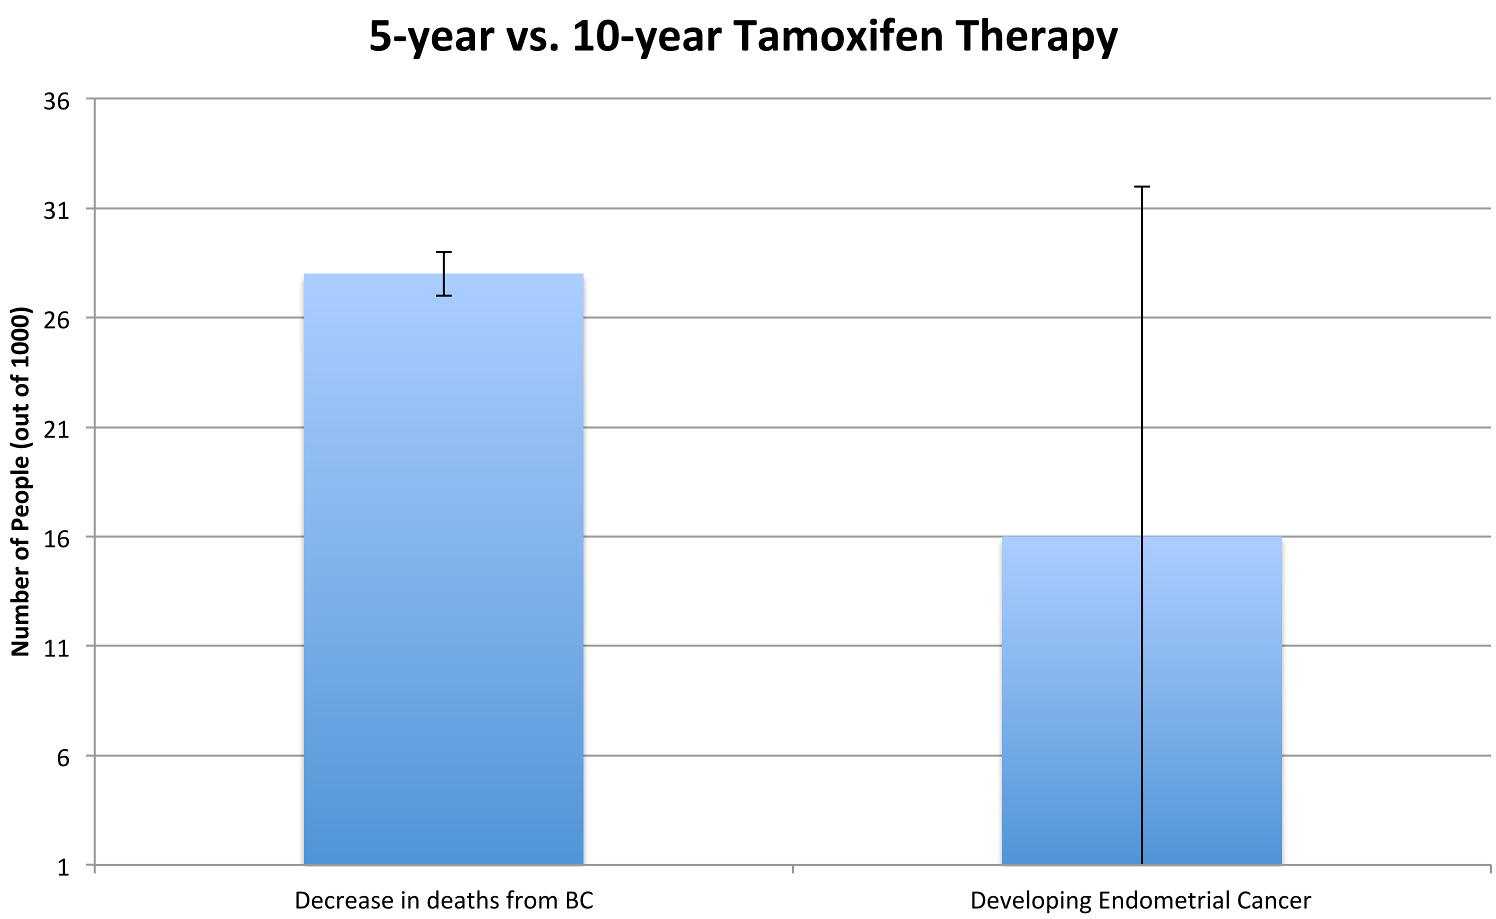

Alternatives to Error Bars

Violin Plots

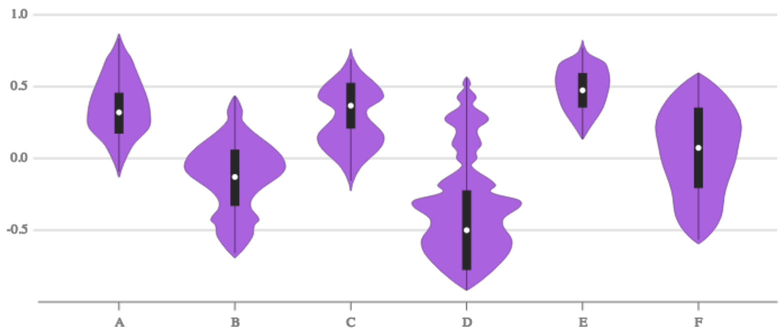

http://www.datavizcatalogue.com/methods/violin_plot.html

Box Plots

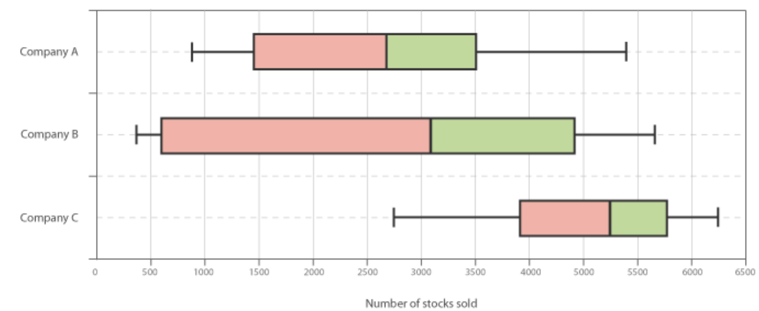

http://www.datavizcatalogue.com/methods/box_plot.html

Dynamic Icon Arrays

Gradient Plots

Project

- Design new visualization to present ambiguity to patients
- Interactivity
  - Adjust bounds of error
  - Show best & worst case scenarios
  - Show how risk estimates might change given different samples

Dviz

Visualizing Distributed Systems
with Stewart Grant and Jodi Spacek

Motivation

- Understanding the behaviour of distributed systems is hard
- Developers need tools for comprehending their systems
- Most distributed systems are designed around FSM
- FSM are often how developers think of their systems
- Can an FSM be generated from an execution so developers can check their mental models?

Concept

- Collect distributed snapshots (state from across the whole system)
- Calculate a distance between each snapshot (xor distance)
- Plot each snapshot at it’s relative distance using clustering
- Connect each snapshot with a time curve

etcd (distributed key value store) puts -> gets

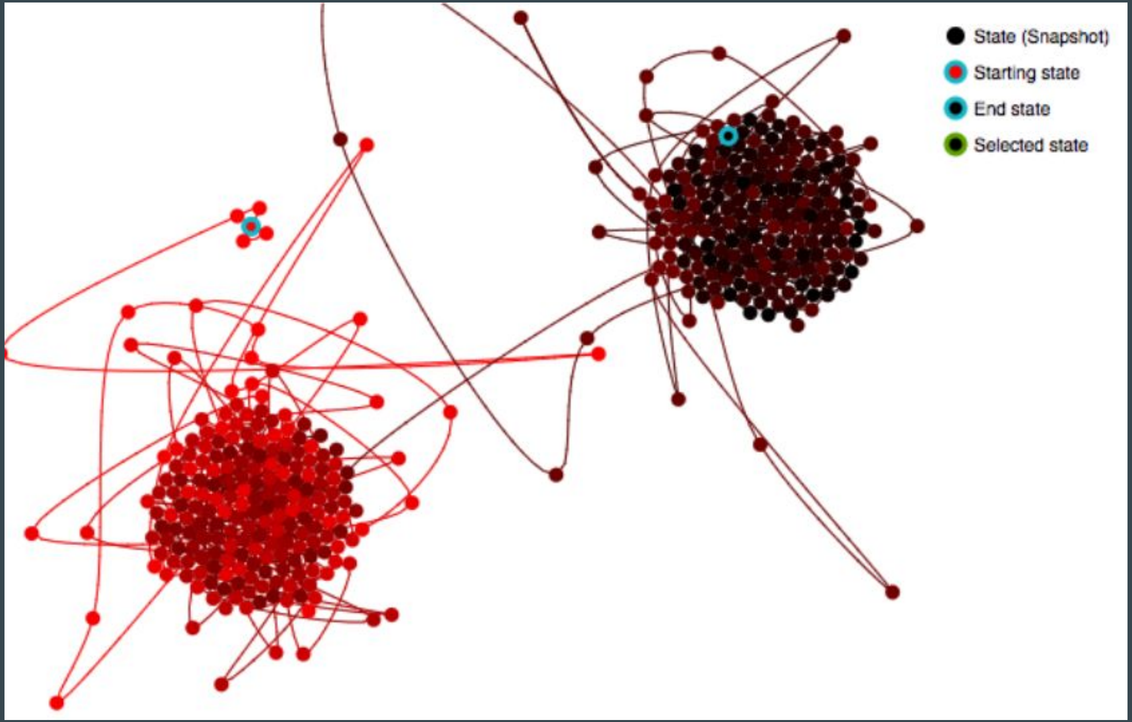

Limitations

- States are not labeled meaningfully
- Semantics of state transitions are not clear
- FSM’s require both

Extensions to Project

Improving Visualization

Interaction Extension

FSM would provide a higher level on which users could zoom in on

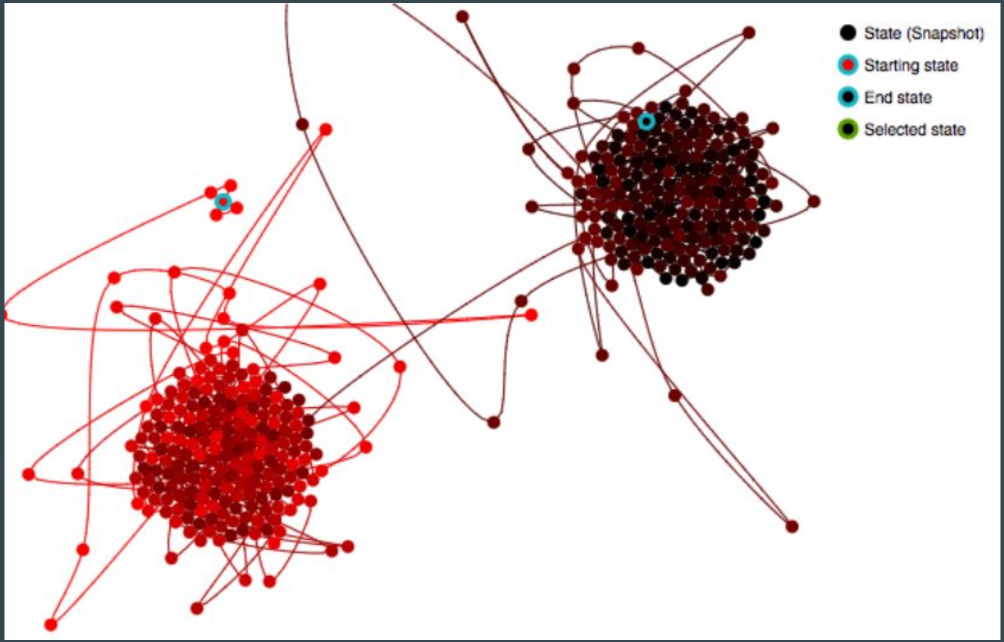

Current

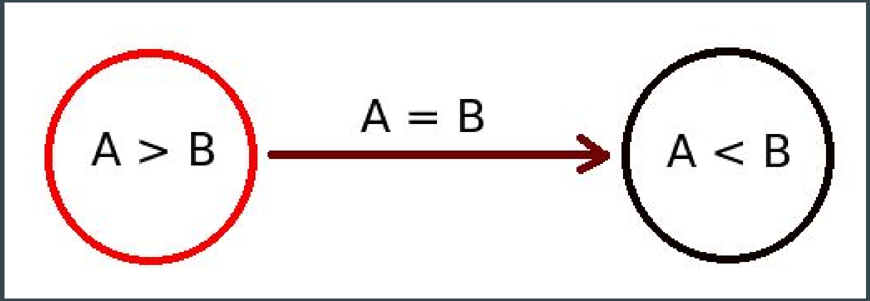

Proposed zoom

Filtering the Clusters

- Partitioning: intrinsic meaning
- Collect data invariants: filter to show aggregate data using existing tool set
- Label: Represent clusters by their invariants
- Visualize transitions: use the diff of cluster invariants

Research Questions

- Scatterplots? Occlusion? Continuous scatterplots?
- Interaction?
- Spatial aggregation? Does it make sense?
- Dimensionality reduction? Too much information?
- Effective color coding?
- Dimensional Ordering, Spacing, and Filtering Approach (DOSFA)? Similarities show patterns?

Why this project is neat

- Stems from an existing body of work
- Has practical applications for debugging distributed systems
- No end of data to represent, can easily be extended after the course

Visualizing patient clusters

Lovedeep Gondara

Problem

Physician researchers are often interested in data exploration before committing to a project.

Generally use descriptive statistics to see if there are any obvious signals.

Is there any specific group of patients that have the worse outcome compared to the rest?

Are there natural groupings in the dataset?

Is there an underlying structure to the data?

Proposed solution

Cluster visualization

Use dimensionality reduction methods such as t-sne.

Plot resulting clusters.

Draw survival plots by cluster membership.

Allow investigation of cluster membership.

Thanks

Spanner, Resurrected.

CPSC 547 Project Pitch

Madison Elliott
February 16, 2017

Background

- Project originated as an MA thesis in the CS department

Background

- Project originated as an MA thesis in the CS department
- New technique that applied low-stretch trees to network visualization

Background

- Project originated as an MA thesis in the CS department
- New technique that applied low-stretch trees to network visualization
- Implemented novel edge-bundling technique

Background

- Project originated as an MA thesis in the CS department
- New technique that applied low-stretch trees to network visualization
- Implemented novel edge-bundling technique
- Does not rely on fixed vertices/fixed layout or explicit hierarchical data structure

Background

- Two iterations submitted for publication:
  1. Graph Drawing (technique focused)
  2. Pacific Vis (more emphasis on motivation and visualization application)

Background

- Two iterations submitted for publication:
    1. Graph Drawing (technique focused)
    2. Pacific Vis (more emphasis on motivation and visualization application)

- Both rejected ☹

Background

- Two iterations submitted for publication:
  1. Graph Drawing (technique focused)
  2. Pacific Vis (more emphasis on motivation and visualization application)

- Both rejected ☹
- Reviewer comments largely yearning for a deeper/more defined motivation

Resurrection Pitch

- Find the motivation!

Resurrection Pitch

- Find the motivation!
- Develop and execute a user study

Resurrection Pitch

- Find the motivation!
- Develop and execute a user study
- Revise and resubmit paper

Why?

- Lots of potential!

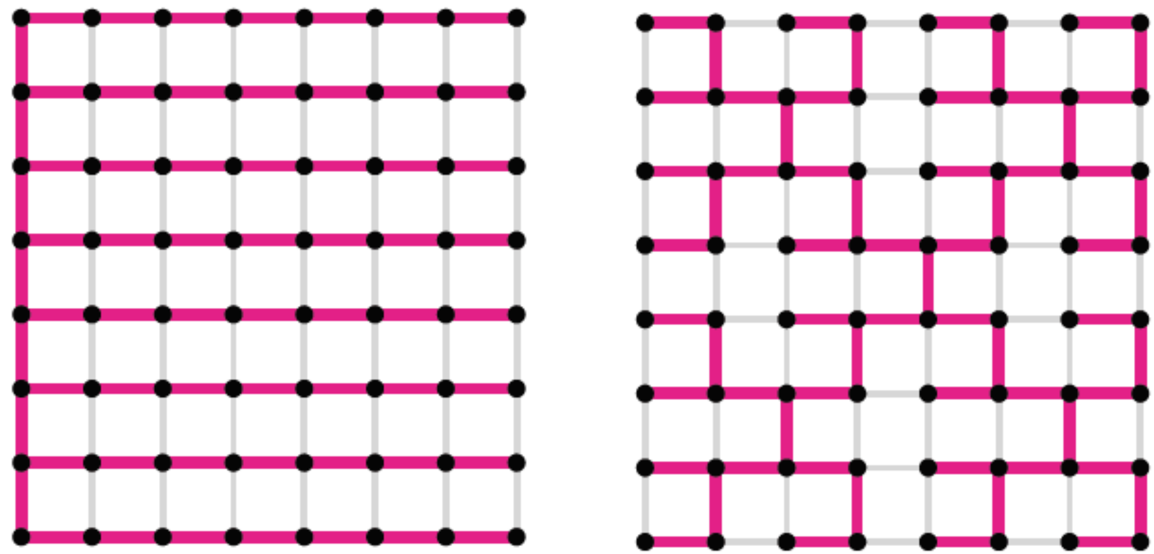

(a) A spanning tree.

(b) Low-stretch spanning tree, as shown in Alon et al. [3].

Figure 3: Comparison between an arbitrary spanning tree and a low-stretch spanning tree for an 8-by-8 grid graph.

Why?

- De-hairball a cluttered network:

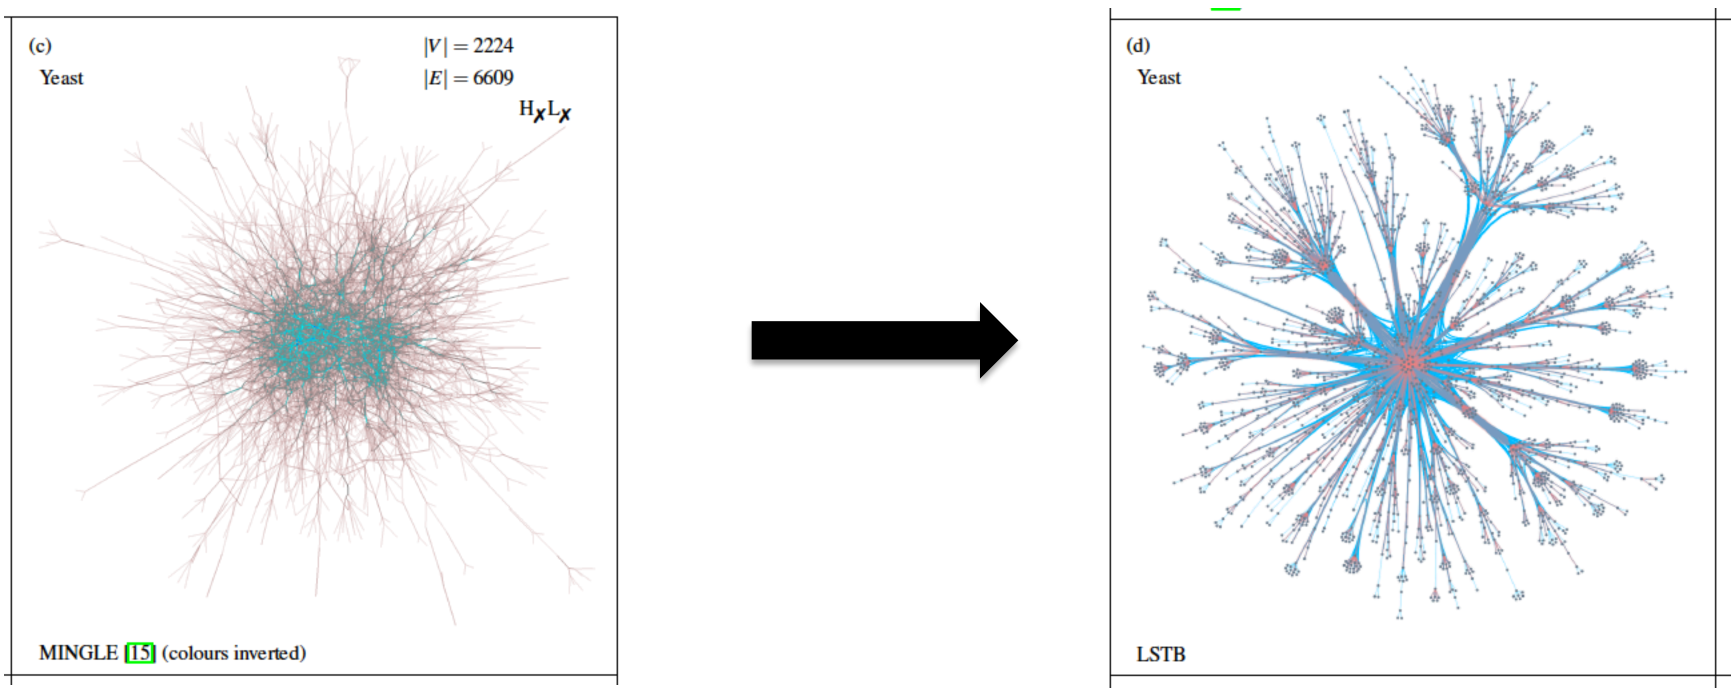

Why?

- Novel, layout free network idioms:

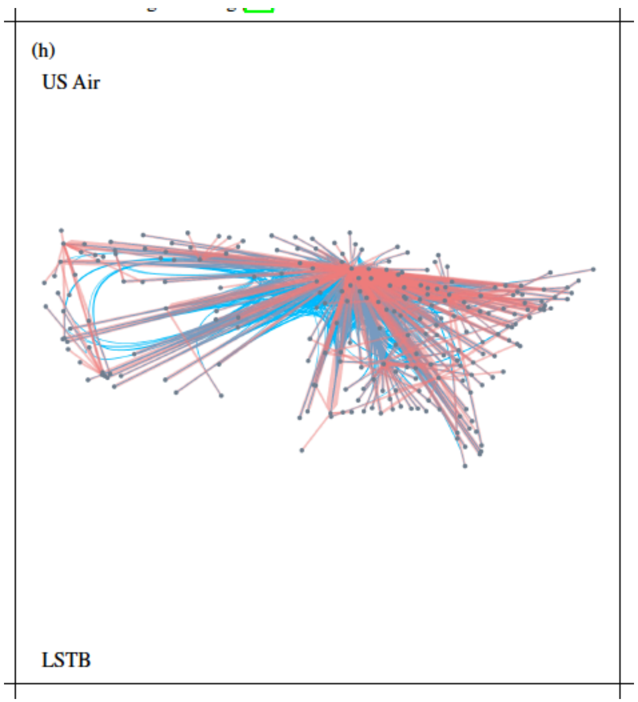

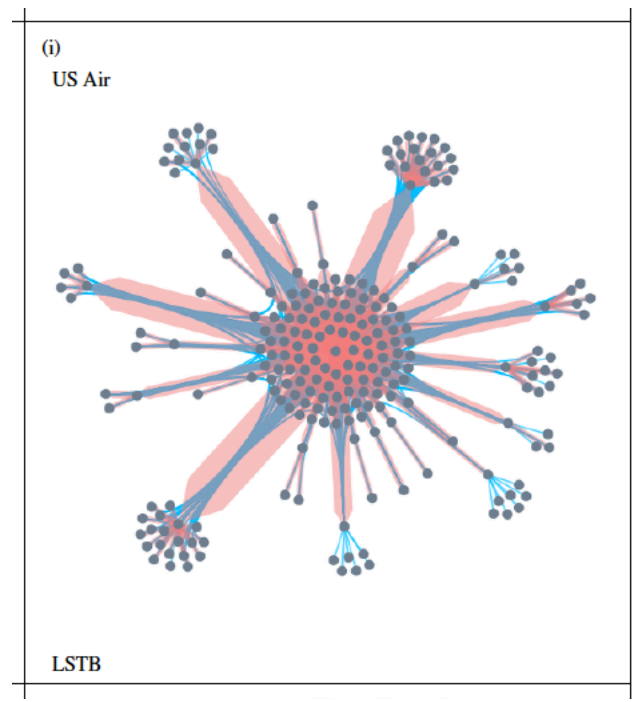

Next Steps

- Complete literature review of network idioms, tasks and taxonomies

Next Steps

- Complete literature review of network idioms, tasks and taxonomies

- Brainstorm new cases where “set” or intuitive network layout is not optimal or necessary for a given task

Questions?

Automatic Grading Service Dataset

NICK BRADLEY

NBRAD11@CS.UBC.CA

Background

Continuous grading service

5.5 GB from **13K** test result records (more coming everyday)

Some data fields *(don't worry if these don't mean anything to you)*

- Grade for every commit each student made
- Test metrics: # tests pass/fail, coverage, duration
- Code metrics: LOC, build failures
- Grade requests: timestamp
- *More data can be pulled from GitHub (diffs, history, branches,…)*

Current Instructor Dashboard

Date	Repo	#Sec	% overall	% pass	% cover	#P	#F	#S	#LOC	Results
02/16 @ 04:40:53	cpsc310project_team40	12.0	72.8	66	95.89	33	17	0	316	
02/16 @ 02:25:33	cpsc310project_team21	16.3	32.09	22	67.43	11	39	0	608	
02/16 @ 01:33:41	cpsc310project_team80	9.2	44	30	95.84	15	35	0	505	
02/16 @ 01:17:14	cpsc310project_team194	10.6	32.28	22	68.4	11	39	0	250	
02/16 @ 01:10:42	cpsc310project_team65	108.6	25.04	12	72.21	6	44	0	439	
02/16 @ 01:09:20	cpsc310project_team78	16.0	77.6	72	95.84	36	14	0	505	
02/16 @ 01:06:58	cpsc310project_team17	18.2	87.2	84	95.73	42	8	0	539	

Current Operational Dashboard

AutoTest Queue:
0

Reference UI:
UP

Geocoder:
UP

Class Portal:
UP

Idea + Impact

Student facing dashboard

- Expanded to CS110, CS210, and CS310 + their corresponding MOOC offerings
- Vis will be used by 1000s of students in production system
- Challenge: make it engaging + promote ‘good’ behaviour
- Feedback: prototype can be made available to current students

Instructor facing dashboard

- Design study with domain expert (current CPSC310 instructor)
- Challenge: needs to scale to 1000s of students

Analysis tool

- Probably only if you are interested in software engineering
- Likely end up as a SE paper

nbrad11@cs.ubc.ca

EMAIL

Visual Methods for Analyzing Motifs in Time-Oriented Data

Soheil Kianzad

PhD student CS

Stock technical analysis

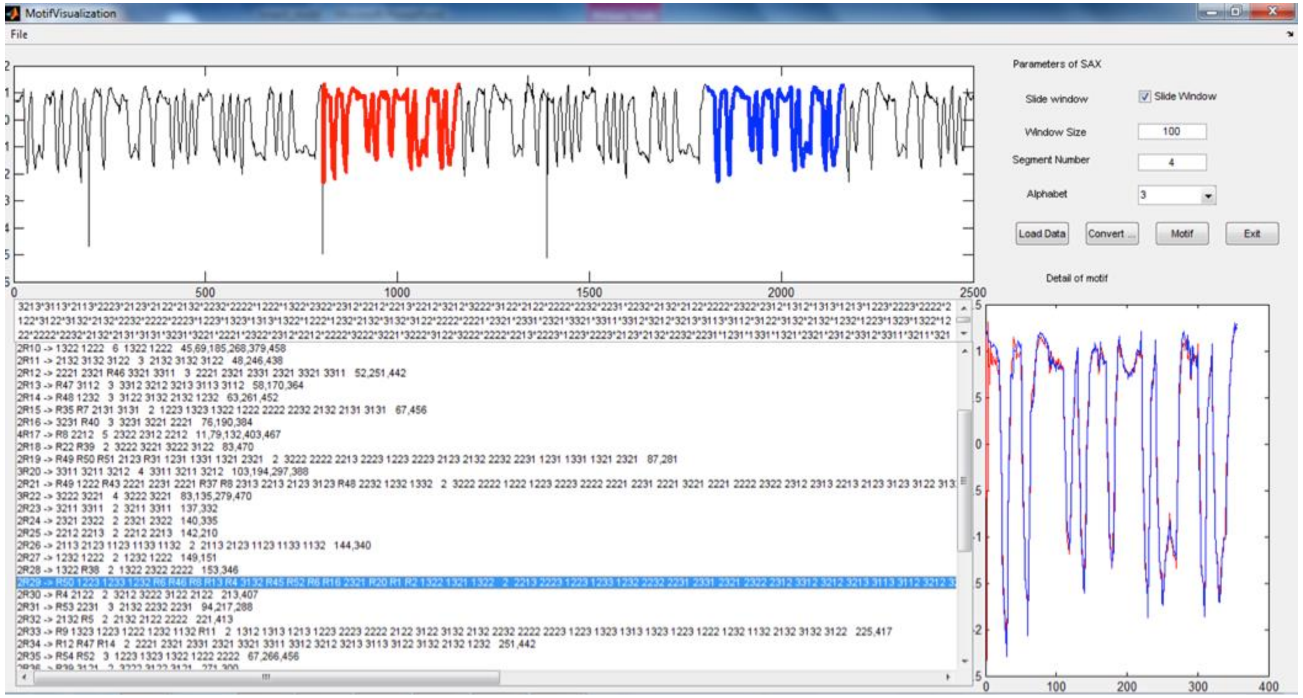

Yuan Li ,GrammarViz, 2012

www.aastocks.com/en/stock/detailchart.aspx?symbol=110000

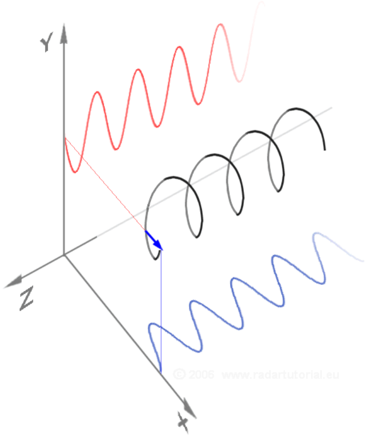

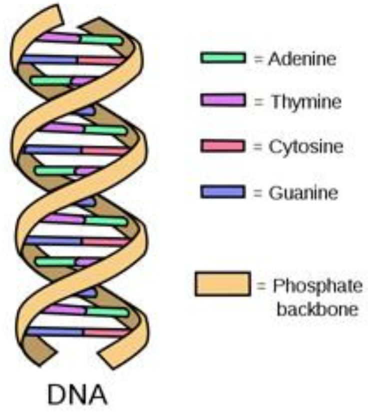

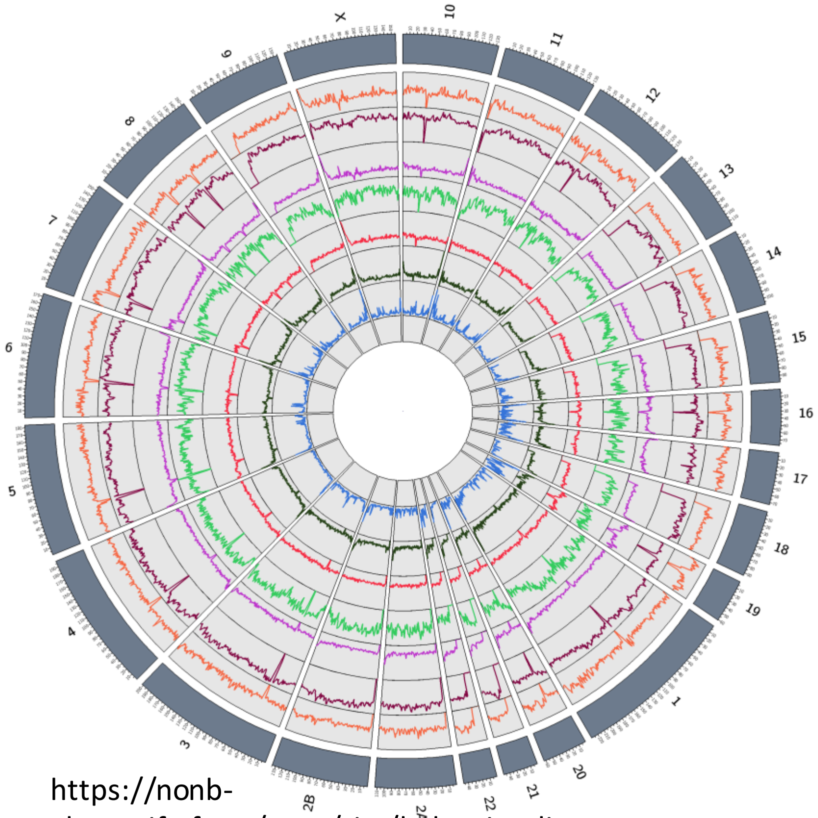

https://nonb-abcc.ncifcrf.gov/apps/site/help_visualize

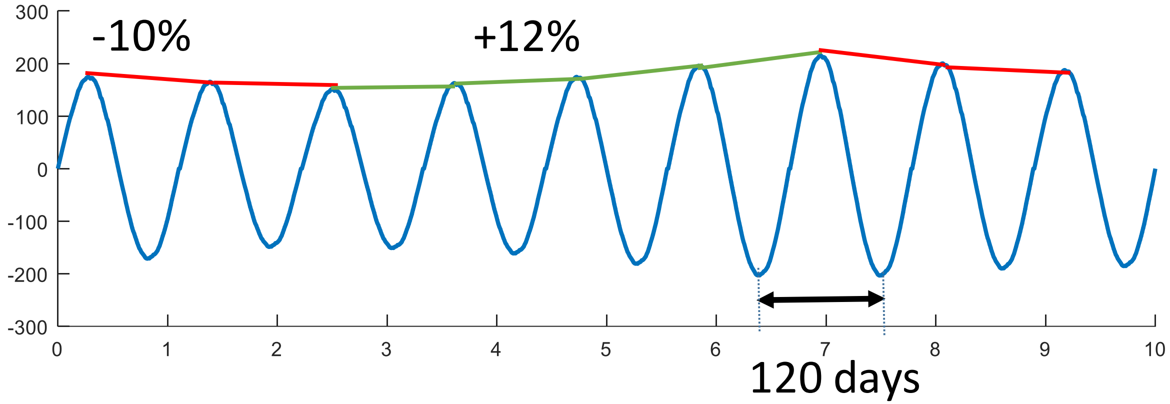

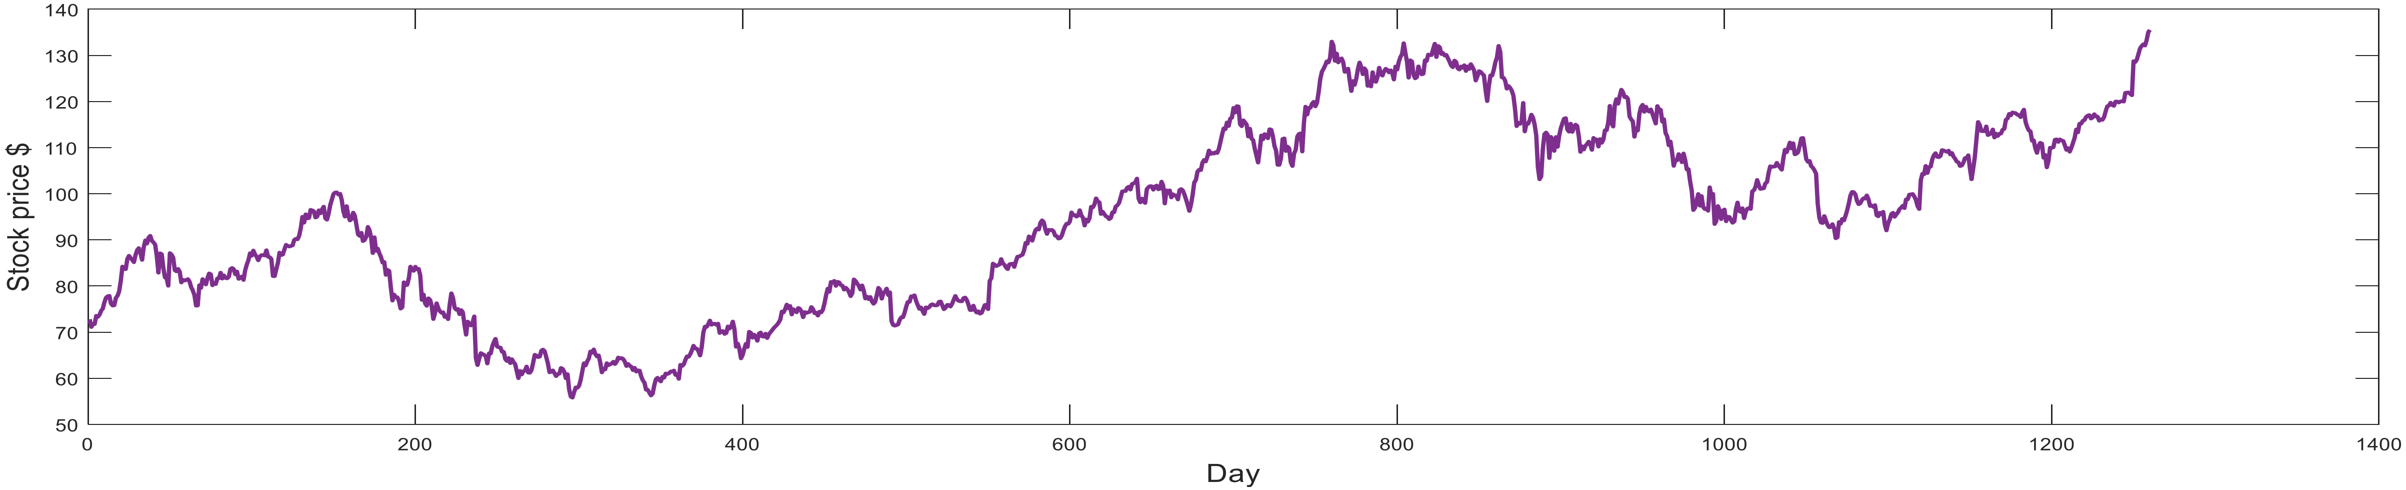

ViSoccer

Visualizing European soccer players

Yann Dubois

Why?

Other sports

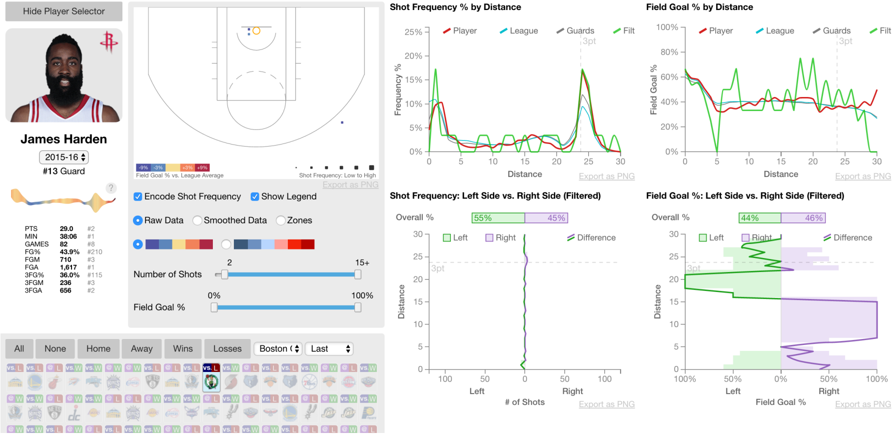

http://buckets.peterbeshai.com/

By region

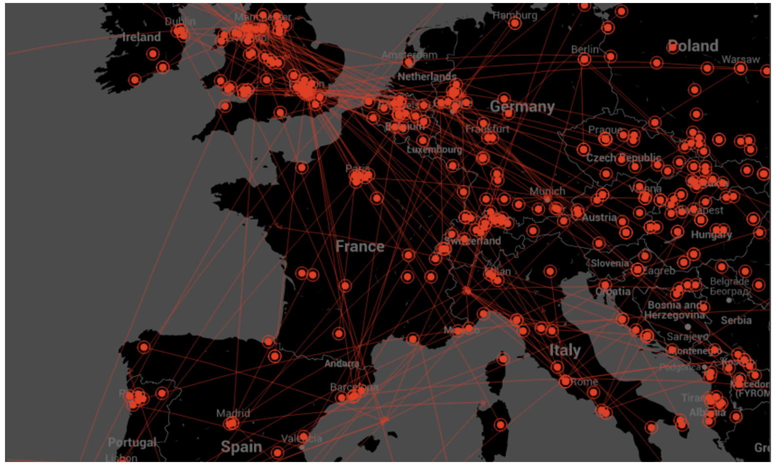

https://www.ibm.com/blogs/bluemix/2016/06/origins-of-soccer-superstars-part2/

World cup

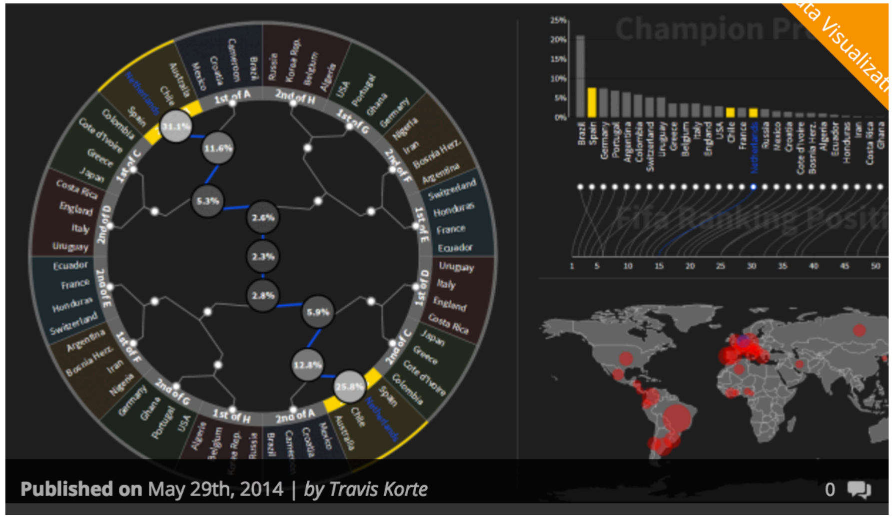

https://www.datainnovation.org/2014/05/predicting-the-world-cup-winner-with-data-visualization/

By game

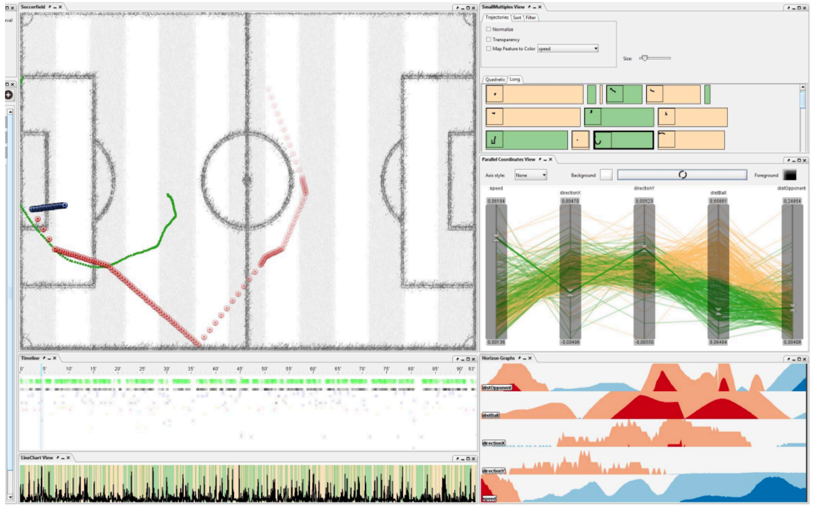

http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=7042477

What?

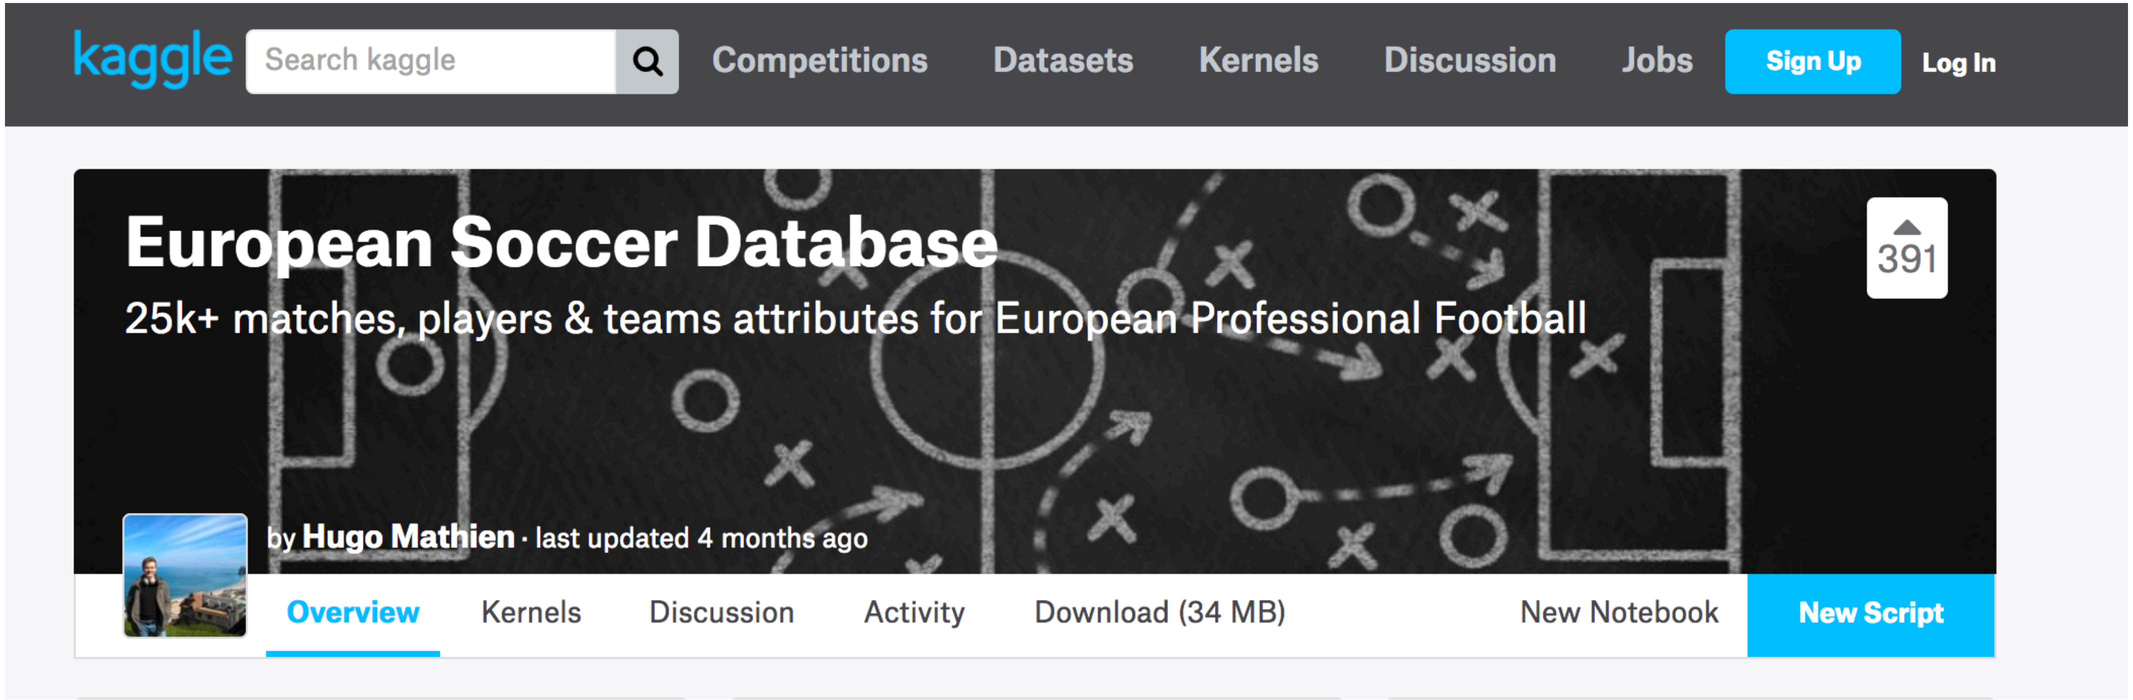

- +25,000 matches
- +10,000 players
- 11 European leagues
- Players and Teams' attributes
- Detailed match events
- Betting odds

+ sports page scrapping

How?

- D3
- P5.js
- Tableau

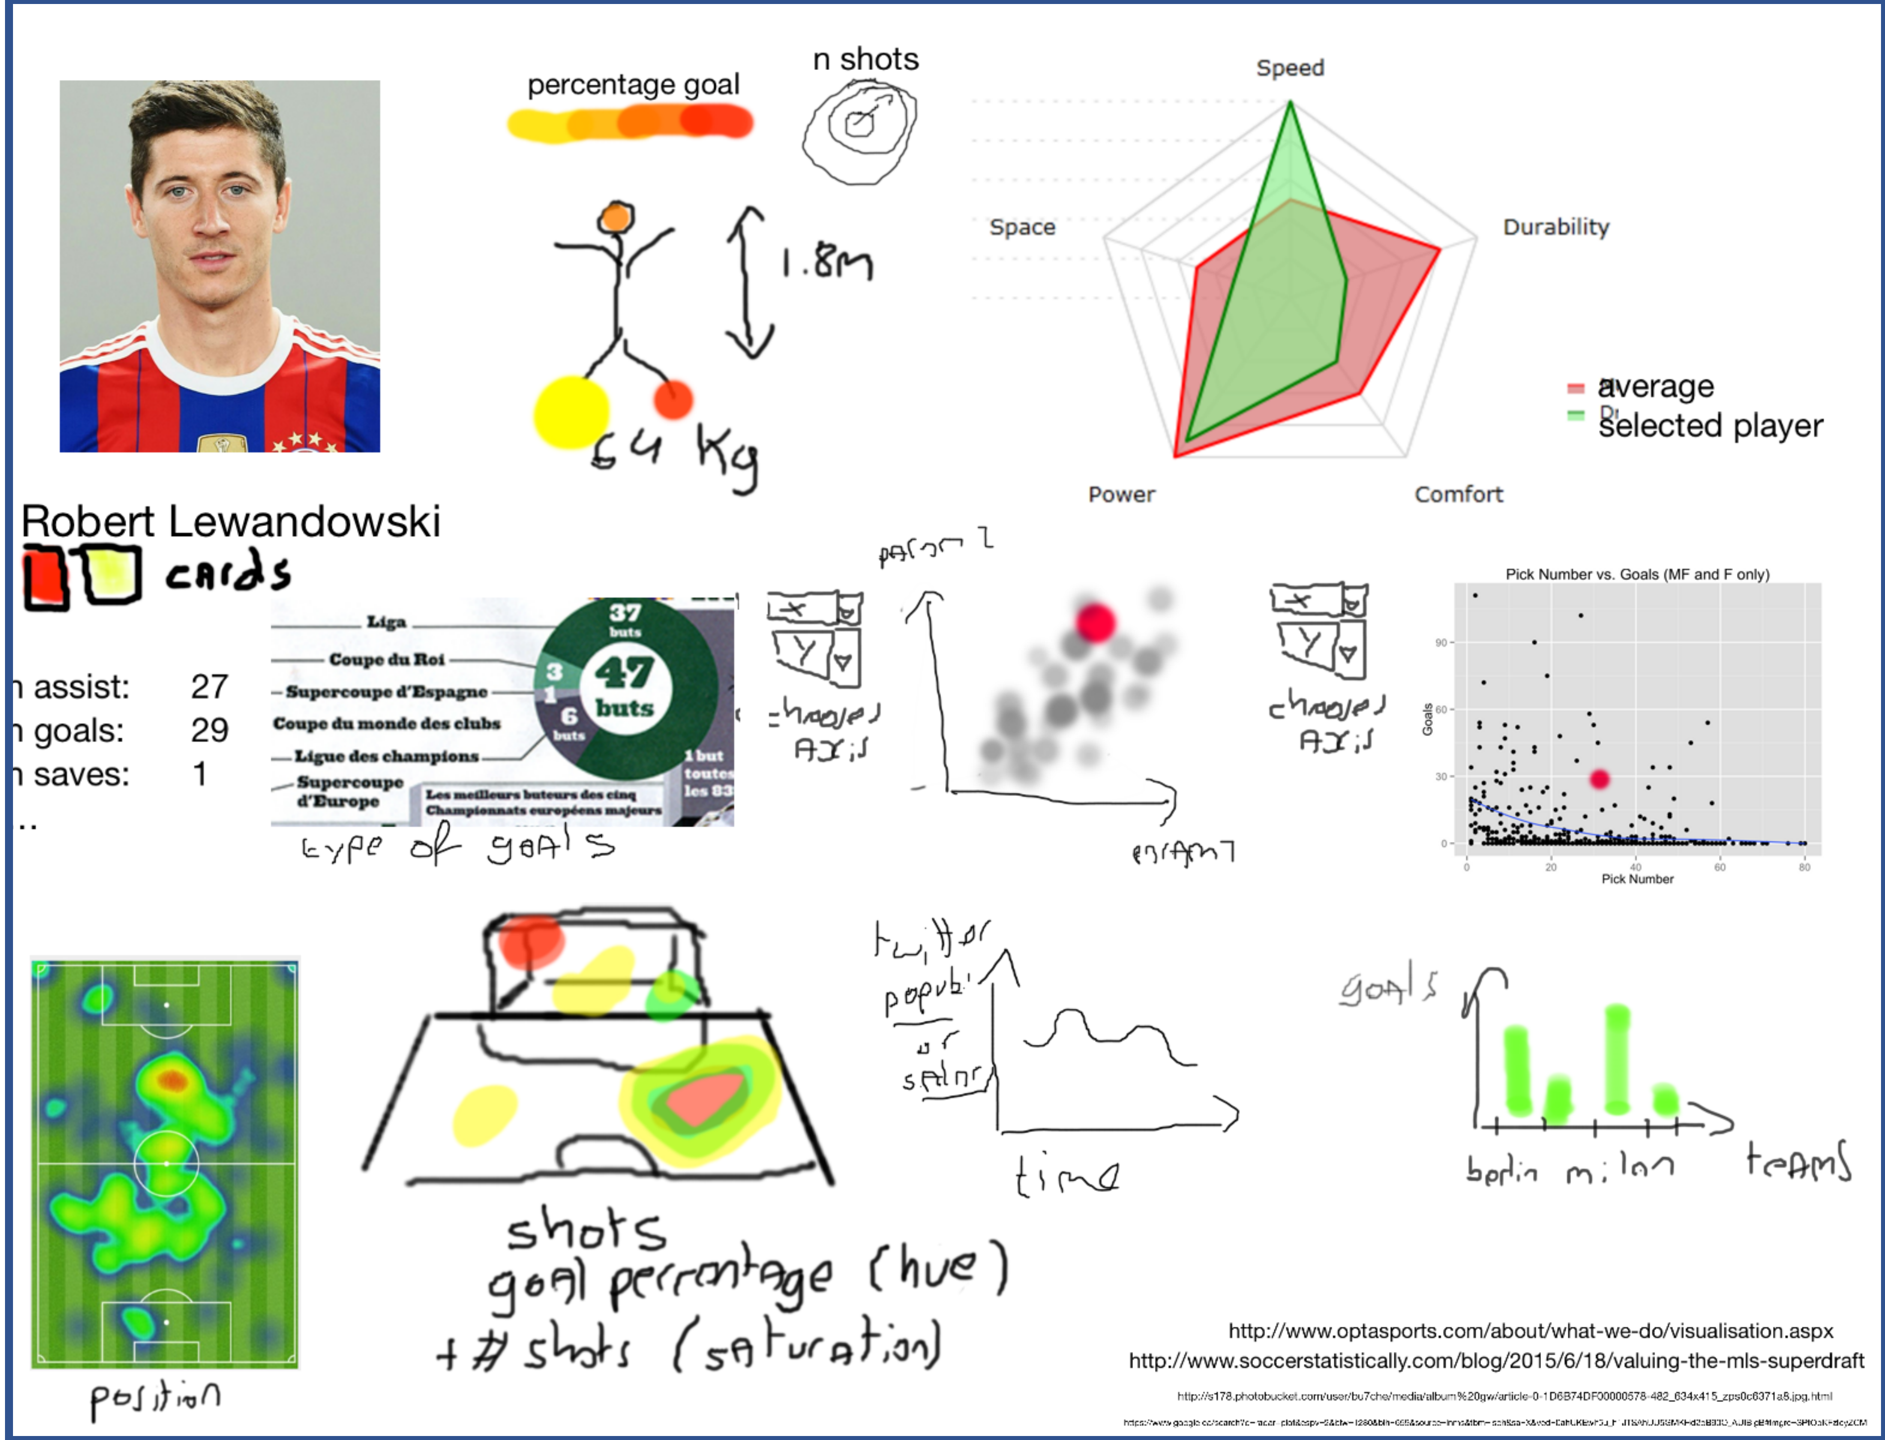

GRADUATE STUDIES – DATA VISUALIZATIONS

JENS LOCHER, ASSISTANT DEAN – STRATEGIC TECHNOLOGIES AND BUSINESS INITIATIVES

WHO ARE WE?

- **Responsible for academic oversight and support for approx. 300 graduate degree programs**
- **Strategic leaders in graduate education at UBC**
- **Support for faculty, programs & students**
- **Central hub for everything related to graduate students**
    - **Communications & Recruitment**
    - **Admission**
    - **Awards**
    - **Thesis & Dissertations**
    - **Doctoral Exams**
    - **Professional Development**
- **Approx. 10,000 graduate students in Vancouver**

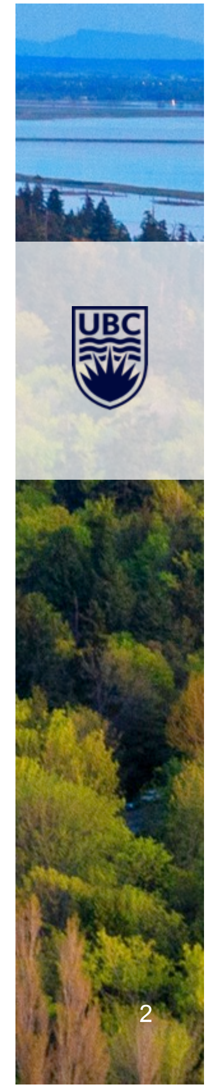

DATA PROJECTS

- Option 1: Canadian Graduate & Professional Student Survey (CGPSS)
  - Satisfaction levels in 13 sections, e.g. general, PD, research experience, financial support, social life
  - Breakdown by discipline, year of study, degree level, gender, etc.
- Option 2: Graduate School data
  - Application data
  - Enrolment statistics
  - Graduation statistics
  - Time in program and completion rates

CGPSS

	2010	2013	2016	Total
Doctoral	13 812	18 377	18 822	51 011
Research Masters	13 593	17 546	18 086	49 225
Other Masters	11 213	15 741	16 834	43 788
Total	38 618	51 664	53 742	144 024
univ	38	48	50	

CGPSS

Desired Outcomes:

1. Visualize key findings from 2016 study
2. Time comparison: 2010 to 2013 to 2016
3. Benchmarking: program vs. UBC vs. Canada

Audiences:

- Students
- Units (access controlled), e.g. program or department dashboard
    - Department Head
    - Program Director
    - Faculty

GRADUATE SCHOOL DATA (CURRENT)

REGISTRATION PERIOD (EXTRACT DATE)	APPLICATIONS	OFFERS	ACCEPTS (CAME)
1995 (March-01-1996)	7,411	2,793	1,732
1996 (March-01-1997)	7,301	2,892	1,821
1997 (March-01-1998)	7,592	2,885	1,786
1998 (March-01-1999)	6,864	2,879	1,826
1999 (March-01-2000)	7,550	3,278	2,074
2000 (March-01-2001)	7,420	3,207	2,066
2001 (March-01-2002)	7,899	3,305	2,156
2002 (March-01-2003)	10,170	3,841	2,548
2003 (March-01-2004)	11,778	3,989	2,728
2004 (March-01-2005)	10,339	4,090	2,802
2005 (March-01-2006)	9,729	3,933	2,671
2006 (March-01-2007)	9,935	4,069	2,690
2007 (March-01-2008)	9,720	4,042	2,672
2008 (March-01-2009)	9,859	4,378	2,907
2009 (March-01-2010)	11,767	4,810	3,363

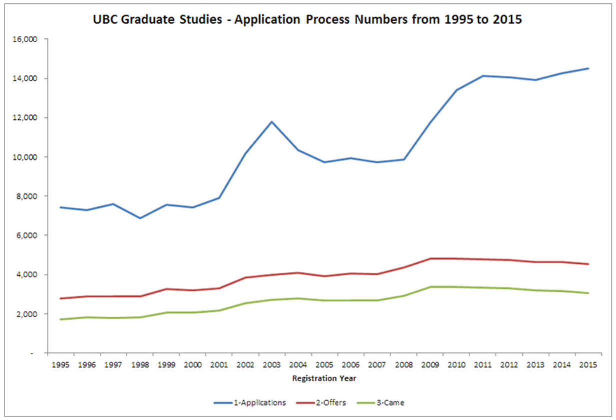

GRADUATE SCHOOL DATA (CURRENT)

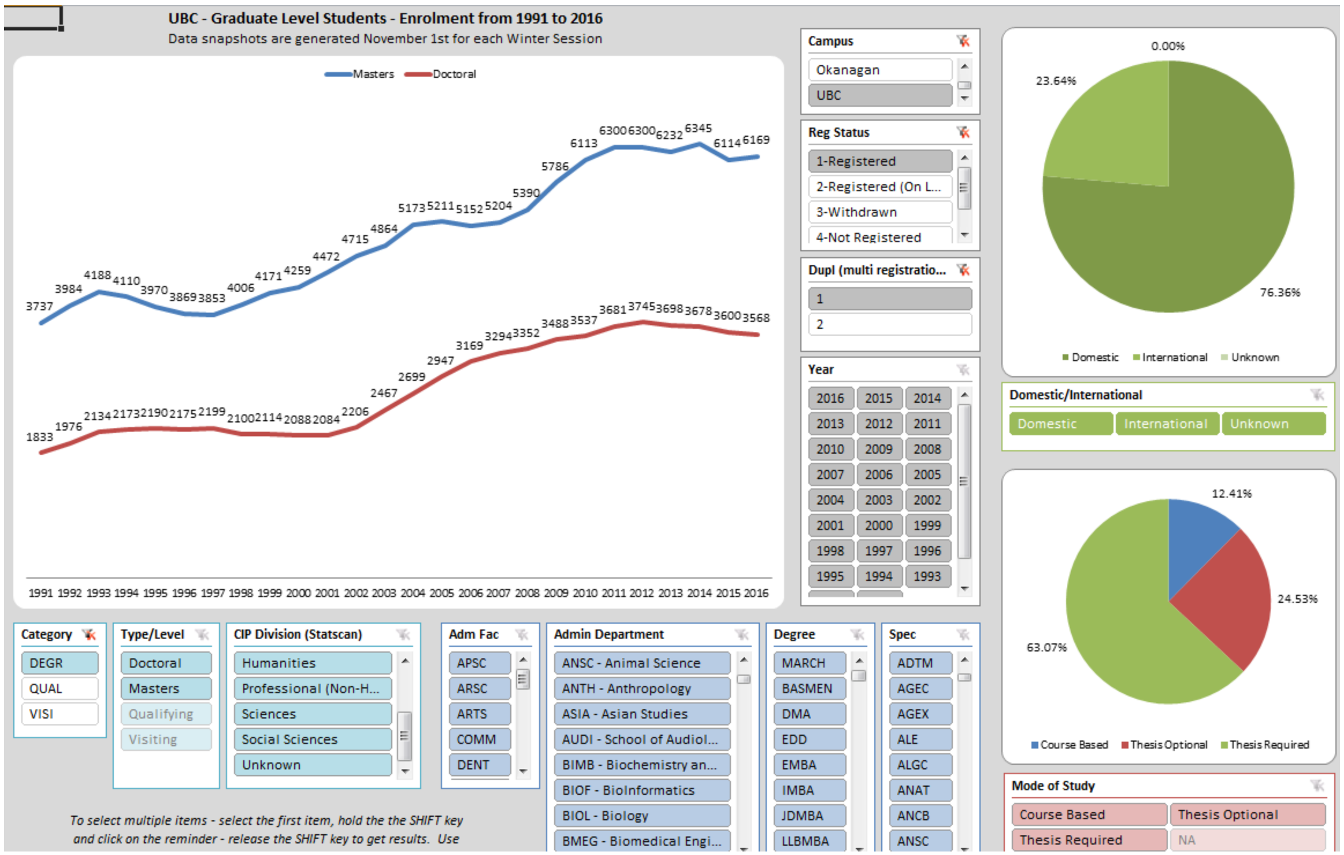

GRADUATE SCHOOL DATA (ALTERNATIVE EXAMPLE)

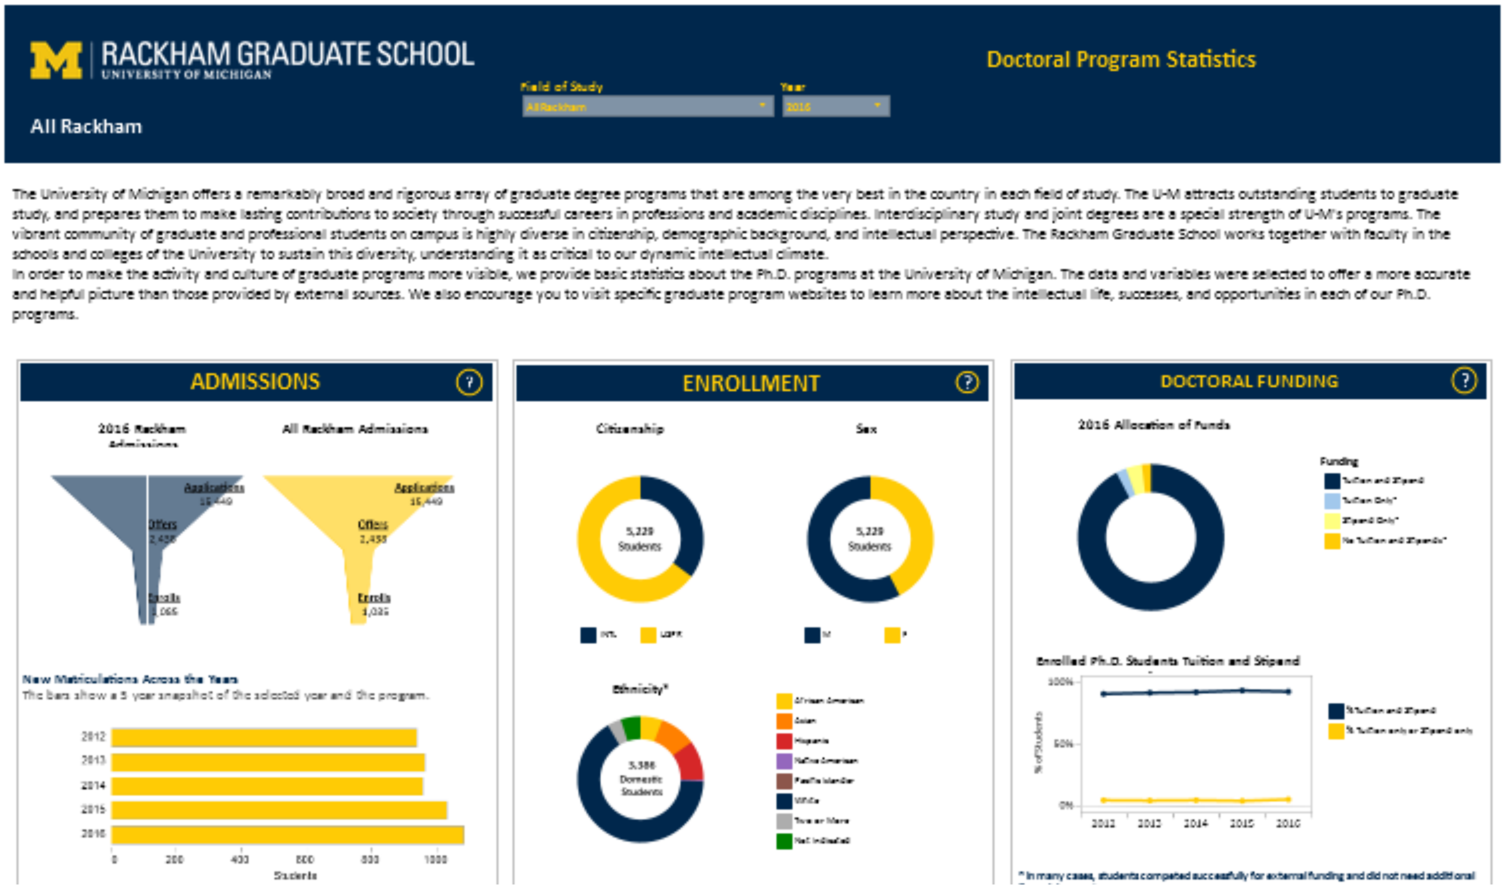

TEAM

Louise Mol

Systems and Data Analysis Manager

Jens Locher

Assistant Dean


UBC
THE UNIVERSITY OF BRITISH COLUMBIA
THE UNIVERSITY OF BRITISH

Visualizing Trends in Product Recommendations

Q.I. Leap Analytics

Who are we?

Q.I. Leap Analytics

- Team of data scientists
- Solutions for retail stores
- 2 products
  - Recommender System
  - Interactive Dashboard

What is a recommender system?

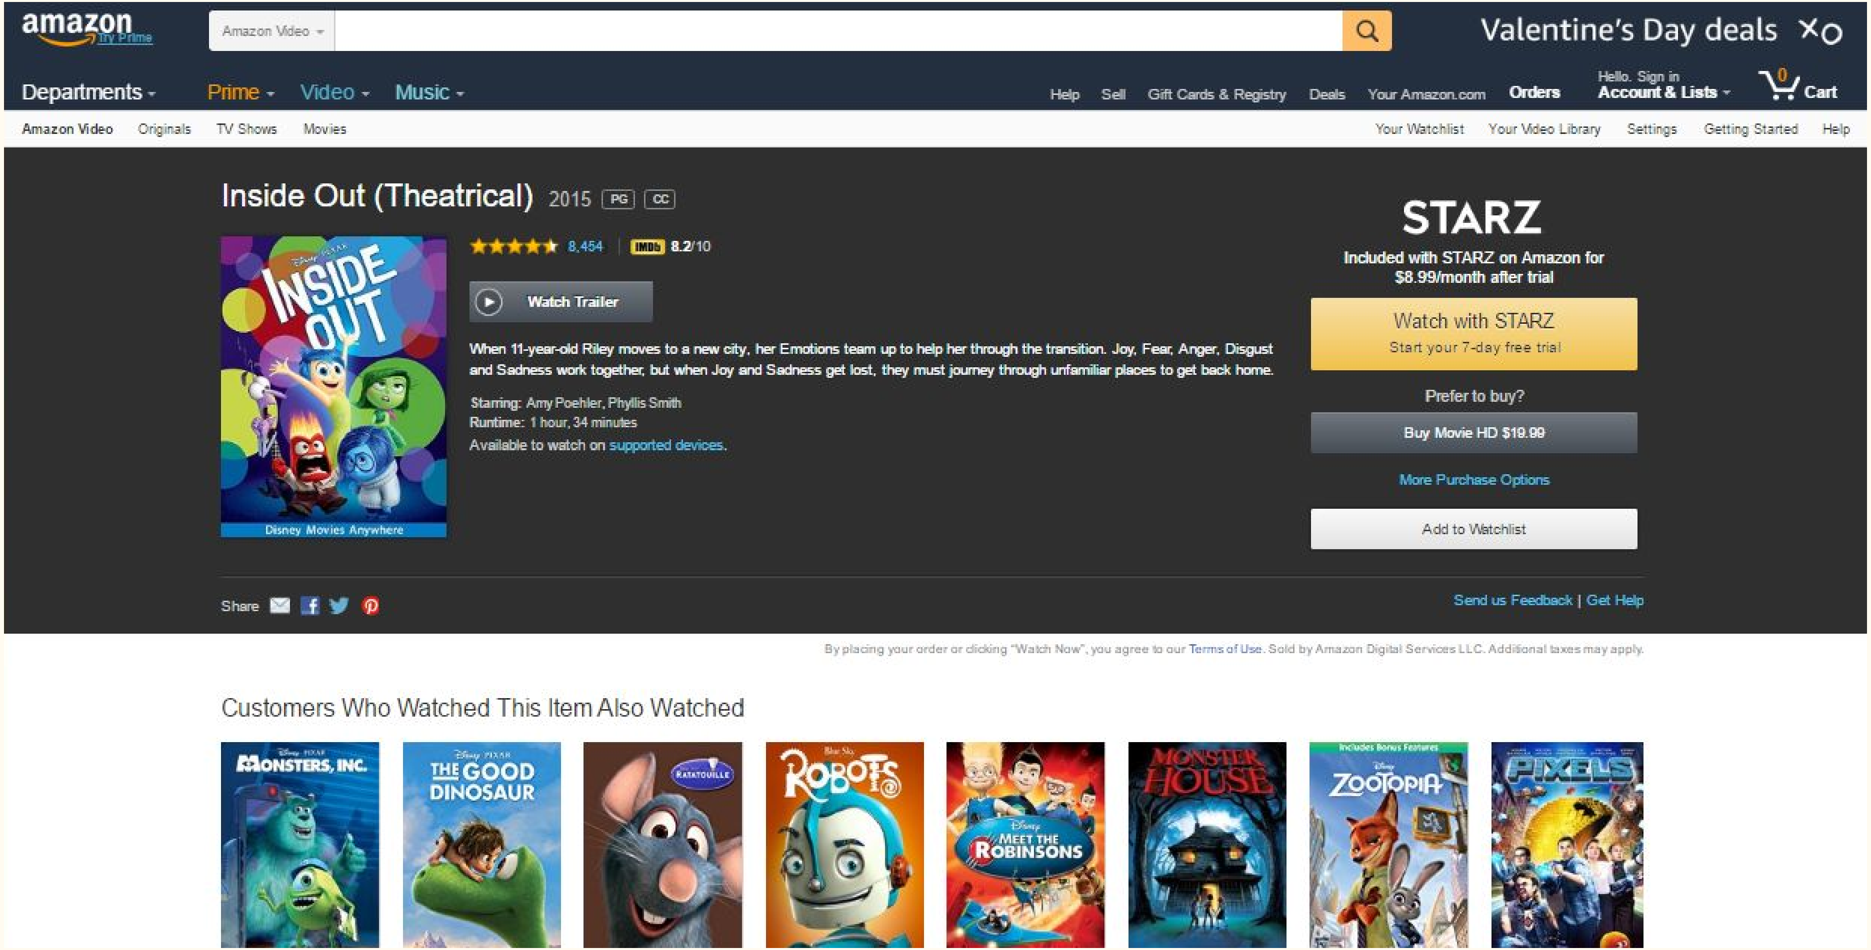

What's the visualization task?

End user: Business that is using the Recommender System

End user desires:

- Which items recommended
- Trends in item recommendations
- Cluster users with similar purchase history
- Cluster items with similar buying history

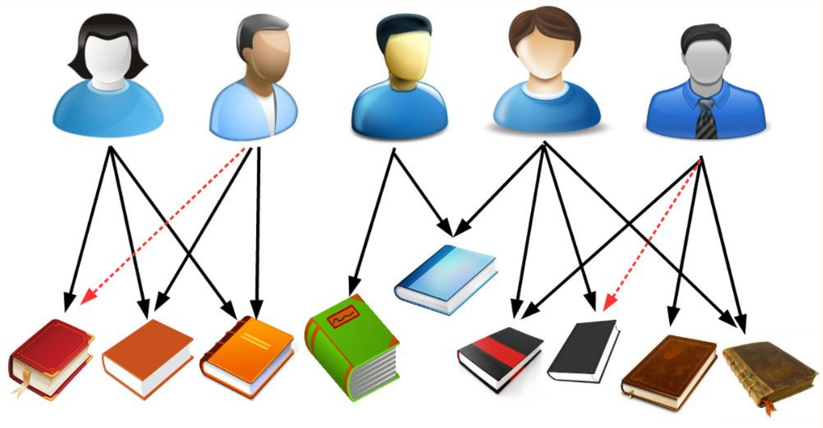

What kind of data would you have to work with?

Transaction data for online store

- 50,000 transactions
- 2,000 unique items
- 13,000 unique customers
- With time, date, city of purchase

Generated recommendation data

- Customer, item viewing history, top 10 recommended items (with scores)

Benefits beyond the classroom

- Implemented in our dashboard product so customers would get to see how their recommender system is being used
- Possibility of internship on completion of project

- Talk to me afterwards if interested in the project!

Q.I. Leap
lauren.fratamico@qileap.com