

Modeling Place

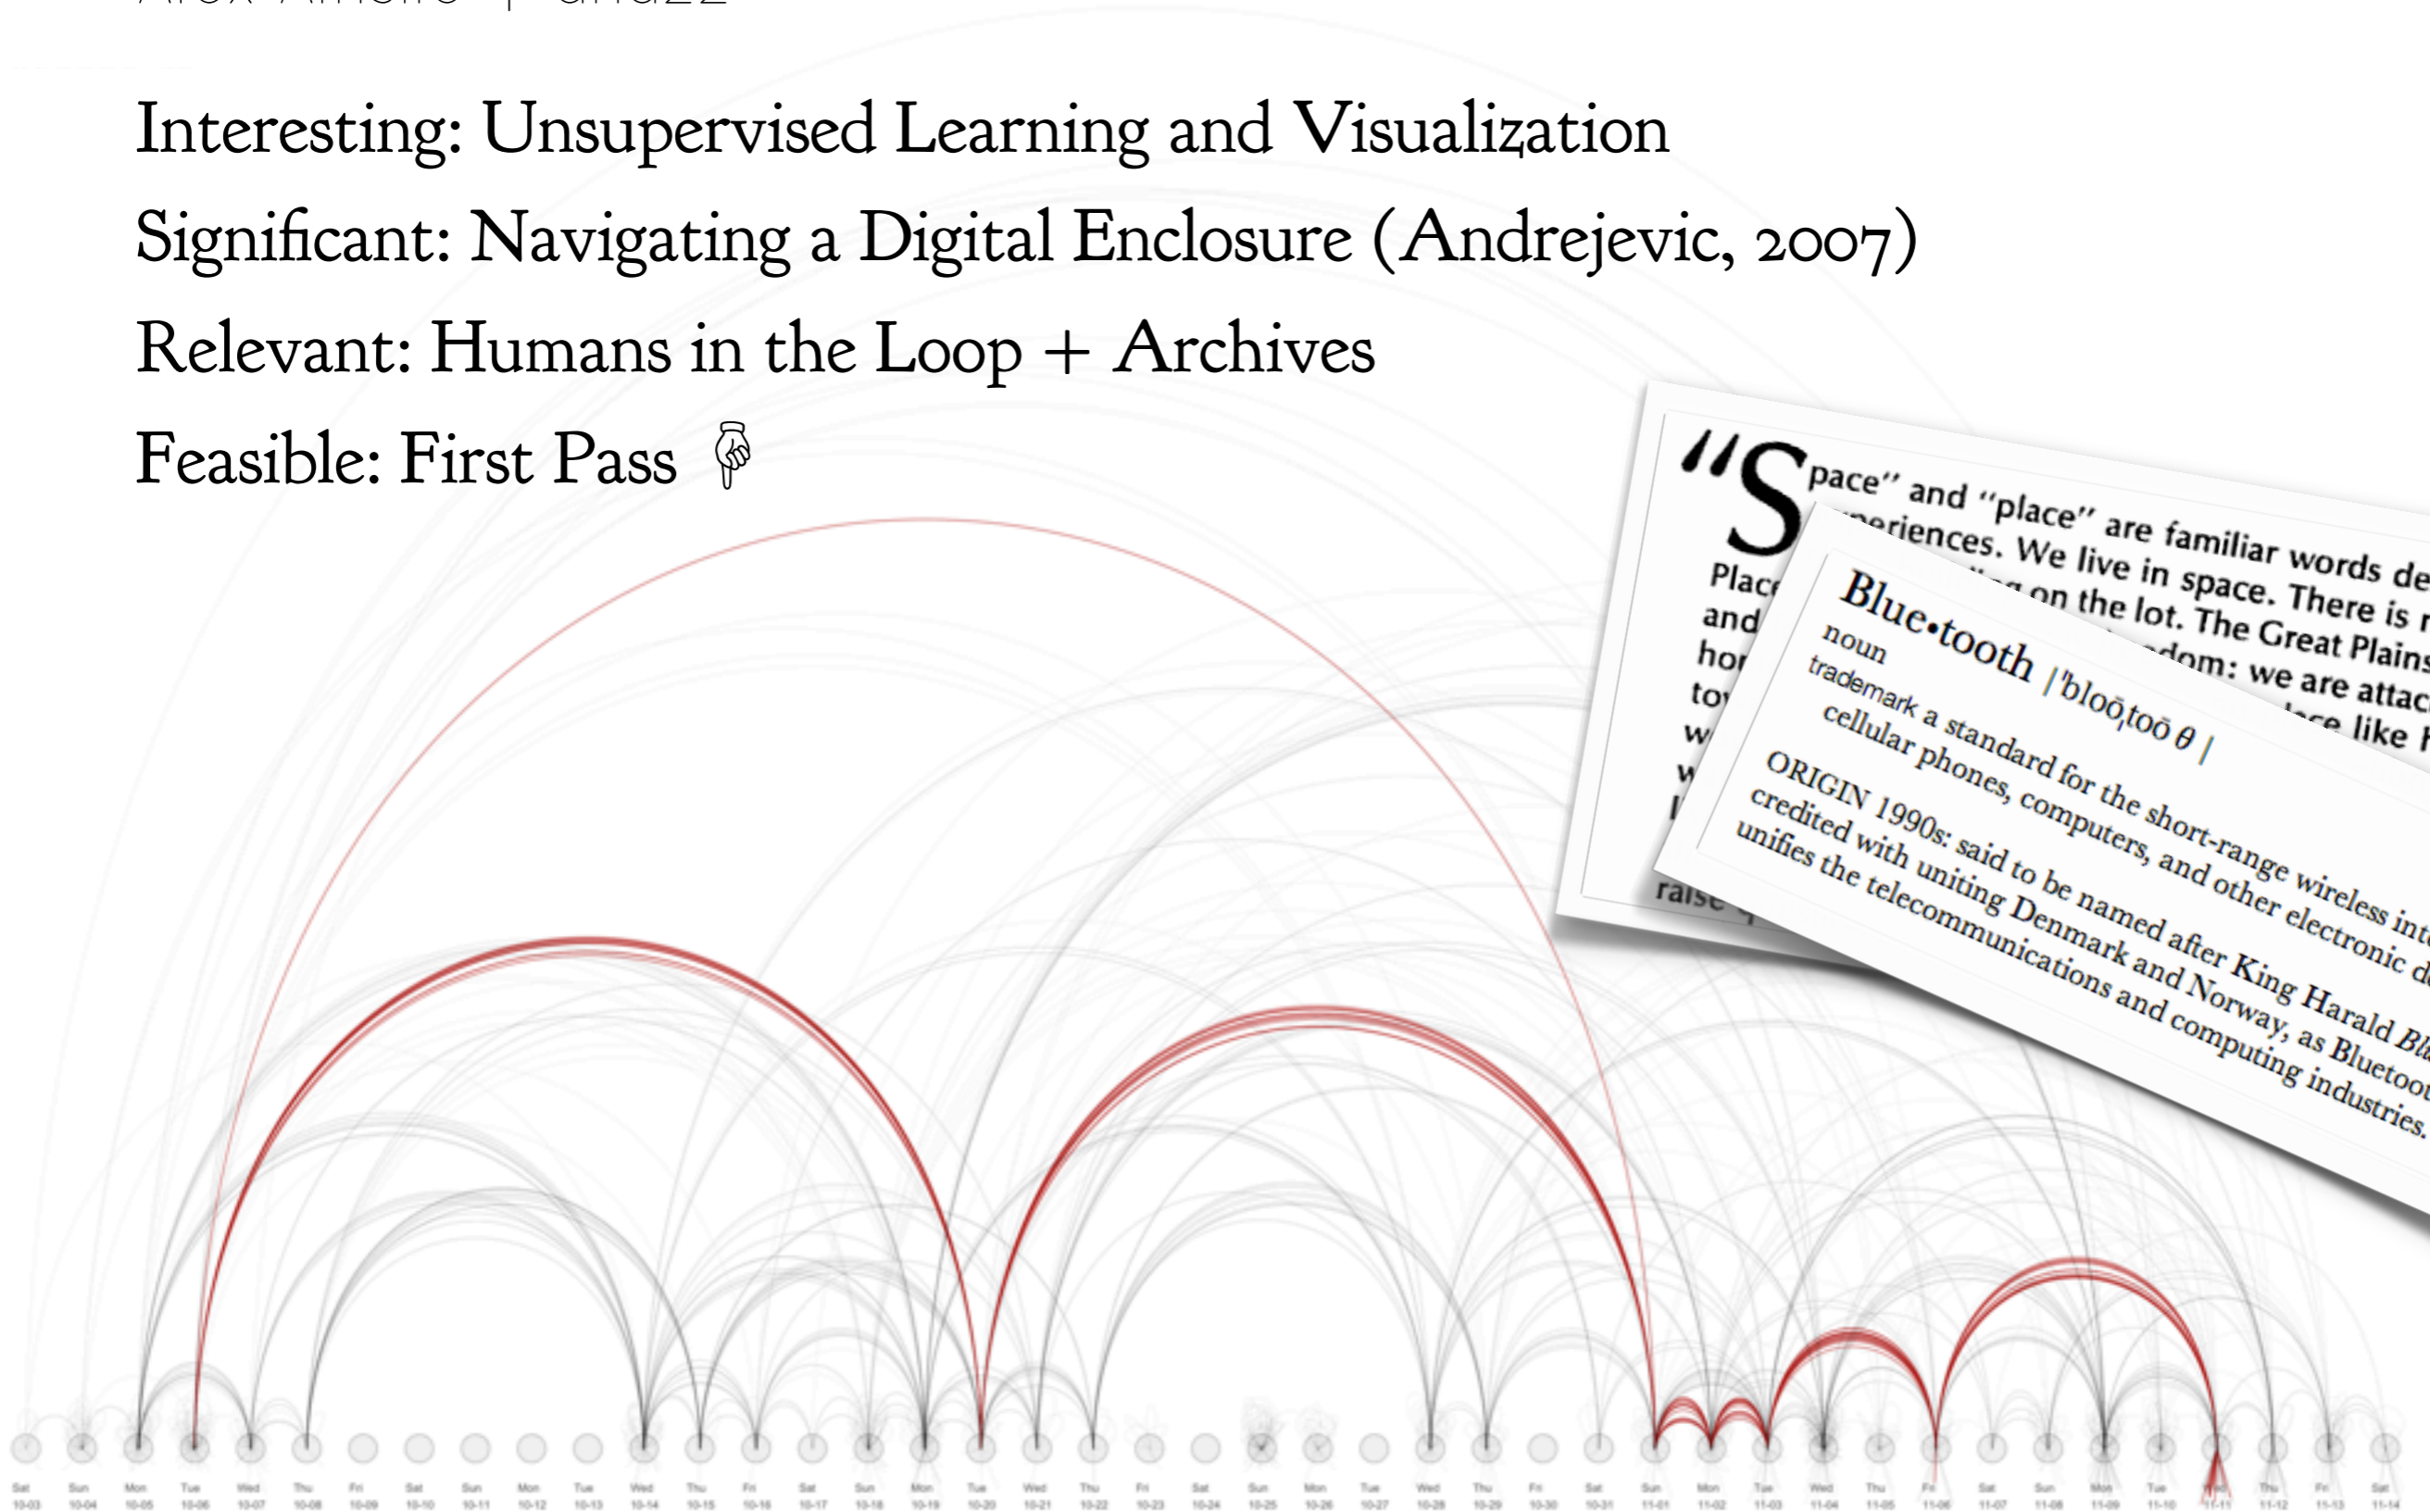
Alex Ainslie | ana22

Interesting: Unsupervised Learning and Visualization

Significant: Navigating a Digital Enclosure (Andrejevic, 2007)

Relevant: Humans in the Loop + Archives

Feasible: First Pass 



Predicting Procedures

Proposer: Devin Kennedy <drk35@cornell.edu>

Structural output prediction: can we predict small programs?

- **Sequences of actions / conditional statements**
- **Motivating example: physicians ordering imaging exams from radiologists**

Given set of symptoms from a patient with abdominal pain → “Run a renal test. If the results indicate a kidney stone may be present, do an abdominal CT with contrast; otherwise, run an abdominal CT without contrast ...”

- **Historical data will be available from Cornell Radiology by the end of this month**
- **Other possible applications?**

Parallel Support Vector Machines

Proposer: Guozhang Wang

guoz@cs.cornell.edu

- Training SVM is not scalable
 - QP solving needs $O(n^2)$ time and $O(n^2)$ memory
- Advances in distributed computing could make SVM at large scale
 - MapReduce for batch computation
 - IPM factorization enable data parallelism (KKT matrix is sparse)
- Other ways?
 - Bagging (SV are few compared to n)
 - But needs to minimize communication

Learning proofs for NuPRL

Proposer: Jean-Baptiste Jeannin, jeannin@cs.cornell.edu

- **NuPRL is a theorem prover, it allows to prove theorems and to check proofs of theorems**
- **A proof consists of a sequence of proof rules**
- **Idea: use machine learning to try to learn this sequence of proof rules from a theorem to be proved**
- **Dataset: theorems proved in the last 15 years by the NuPRL team**
- **Project started in CS6780, where we learnt the first proof rule of a theorem. We would like to go further to prove sequences of rules!**


A human can learn to prove theorems: why not a machine?

Static Analysis of Binary Executables

Proposer: Nikos Karampatziakis

nk@cs.cornell.edu

- High level goal: Predict whether a program is malicious or not
- Extracting CPU instructions as **structured prediction**
 - ▶ Segment input into blocks of code or data using SVM^{struct}
 - ▶ Segmentation can be used for many other tasks
- Novel application domain
- Many interesting algorithmic and engineering problems
- I have already done some of the dirty work




Scheduling Tasks using Natural Language

- Kent Sutherland (khs55)
Rohit Swarnkar (rs634)
- Relevance: Simplify event entry
- Current methods: Simplistic heuristics and text searches (such as in Google Calendar)
- Instead learn from user feedback how to parse complex sentences.
- Train to users' styles instead of forcing a specific syntax

Material Discovery

Ronan Le Bras – lebras@cs.cornell.edu

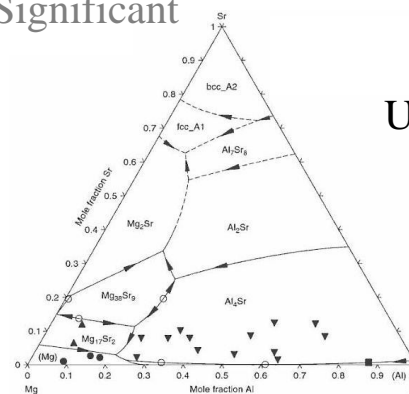
Interesting



Discover new products
in a new area using
newly generated data

[Source: J. M. Gregoire, Department of Physics, Cornell University]


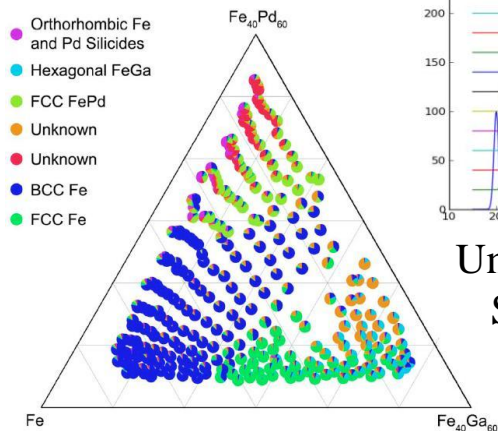
Significant



Material Property
Understanding, Product
Substitute, Resource
Management...

[Source: Methods for phase diagram determination, Ji-Cheng Zhao, 07]

Relevant



Unsupervised Learning,
Structured Prediction,
ICA, Clustering with
Constraints...

[Source: Rapid identification of structural phases in combinatorial thin-film lib. using x-ray diff. and non-neg. matrix factorization, Long C.J. et al., 09]

Feasible



Predicting Wikipedia updates

Ruben Sipos, rs@cs.cornell.edu



TITANIC SINKS FOUR HOURS AFTER HITTING ICEBERG; 866 RESCUED BY CARPATHIA, PROBABLY 1250 PERISH; ISMAY SAFE, MRS. ASTOR MAYBE, NOTED NAMES MISSING

★ Tax credits, highway funds in Senate jobs bill
Reuters - [Andy Sullivan](#), [Doña Chiacu](#) - 4 hours ago
WASHINGTON (Reuters) - With a nervous eye on the November congressional elections, Senate Democrats will unveil tax credits and other proposals on Thursday that aim to bring down the nation's double-digit unemployment rate.
☑ [Video: President Obama Q&A with Senate Democrats](#) 

House Faces Tough Vote on \$1.9T More Debt 
[Los Angeles Times](#) - [USA Today](#) - [New York Times \(blog\)](#) - [ModernHealthcare.com](#)
[all 3,140 news articles »](#) ☑ [Email this story](#)

★ China Renews Opposition to Iran Sanctions
[New York Times](#) - [Alan Cowell](#), [Steven Erlanger](#) - 2 hours ago
PARIS - Adding to a growing catalogue of disputes between Washington and Beijing, a senior Chinese official said on Thursday that pressure for tighter sanctions against Iran over its nuclear program could block chances of a diplomatic ...
[Iran sanctions will derail diplomacy, China warns](#) AFP
[Tehran sends subtle signal on nuclear issue](#) Xinhua
[BBC News](#) - [ABC Online](#) - [Wall Street Journal](#) - [CNN](#)
[all 2,379 news articles »](#) ☑ [Email this story](#)

★ Mr. Brown comes to Washington
[USA Today](#) - 1 hour ago
Good morning. By day's end, Massachusetts will have its first Republican senator in more than three decades. That would be Scott Brown, the up-until-recently little known state lawmaker whose out-of-nowhere victory in last month's special election ...
☑ [Video: Brown to be sworn in immediately](#) WWLP.com
[Brown taking over the late Sen. Kennedy's seat](#) Washington Post
[UPI.com](#) - [Los Angeles Times](#) - [NewsOXY](#) - [CNN International](#)
[all 842 news articles »](#) ☑ [Email this story](#)



Benelux

From Wikipedia, the free encyclopedia

The **Benelux** is a union in Western Europe that comprises three neighbouring countries, **Belgium**, the Netherlands, and Luxembourg (Belgie/Belgique/Belgien - Les Pays Bas/Nederland/Die Nederlande - Lëtzebuerg/Luxemburg/Luxemburg), which lie in the



RMS Titanic

From Wikipedia, the free encyclopedia
(Redirected from *Titanic*)

"Titanic" redirects here. For other uses, see [Titanic \(disambiguation\)](#).

The **RMS Titanic** was an *Olympic-class* passenger liner owned by the White Star Line and was built at the Harland and Wolff shipyard in Belfast, Ireland. At the time of her construction, she was the largest passenger steamship in the world.



Shortly before midnight on 14 April 1912, four days into the ship's maiden voyage, *Titanic* struck an iceberg and sank two hours and forty minutes later, early on 15 April 1912. The sinking resulted in the deaths of 1,517 of the 2,223 people on board, making it one of the deadliest peacetime maritime disasters in history. The high casualty rate was due in part to the fact that, although complying with the regulations of the time, the ship did not have enough lifeboats for everyone aboard. The ship had a total lifeboat capacity of 1,178 people, although her maximum capacity was 3,547. A disproportionate number of men died due to the women and children first protocol that was followed.

The *Titanic* was designed by some of the most experienced some of the most advanced technology available, popularly believed to have been unsinkable.^[R] It was that, despite the extensive safety features, the *Titanic* part of the media about *Titanic*'s famous victims, the resulting changes to maritime law, and the disc contributed to the continuing interest in, and notoriety



WIKIPEDIA
The Free Encyclopedia

interesting  significant  relevant  feasible 



3D Brain MRI Image Segmentation

Advancing Computer Vision with Human in the Loop

Proposer: Ruogu Fang, rf294@cornell.edu

- Global Optimization for Tree Metrics
- Maximum Margin Learning to Learn Weight on Edges
 - Interesting: 3D volume segmentation on medical data
 - Significant: New energy function that can be globally optimized, $O(\log/L)$ running time
 - Relevant: Supervised learning for edge weight or unsupervised agglomerative clustering
 - Feasible: Graph cut on tree metrics has been implemented on RGB color image segmentation



SARAH IAMS

SM16@CORNELL.EDU

INSECT FLIGHT: WHERE (EXACTLY) ARE THEIR WINGS?



GOAL: SEGMENT THE IMAGE ANATOMICALLY

(> 40000 FRAMES OF HIGH SPEED VIDEO,
WITH MOSQUITOES IN VARIOUS CONTORTIONS)

Apprenticeship Learning for Chess

Proposer: Vasu Raman (vraman@cs.cornell.edu)

Single player's view of chess as an MDP

- Transitions known
- Board evaluation as unknown reward function

Apprenticeship learning (inverse reinforcement learning)!

- Using expert demonstrations (humans in the loop...)
- Find policy with performance comparable to or better than expert, on the expert's unknown reward function.

Pieter Abbeel and Andrew Y. Ng., Apprenticeship Learning via Inverse Reinforcement Learning, ICML, 2004.

Feasibility: grandmaster chess databases

Structured Learning for Object Detection

Proposer: Yimeng Zhang, yz457@cornell.edu



- Traditional Method

- Classifier: Input: a window,

- Output: 1/0

- Slide a window on the input image, and decide whether each window is the bounding box for the object

- classification

- Proposed Method

- Input: the whole image, Output: the likelihood to be a bounding box for each location

- Jointly give prediction to each position

- Take the context into consideration

- Related Works

- Using Joint kernel with SVM for object localization

- (Blaschko, M. B. and C. H. Lampert, BMVC 2009)

Learning in Robotic vision

- What we have:
 - Robot controlled by PC
 - Take images,
tracking/recognition



Proposed ideas:

- Learning to detect/avoid obstacle based on vision
 - Robot can get the feedback
 - Reinforcement learning
- Semantic Robot Vision Challenge
 - Given word (“orange”), surf the internet (Google image “orange”), and perform learning
 - Detect the objects in the environment