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 ☟

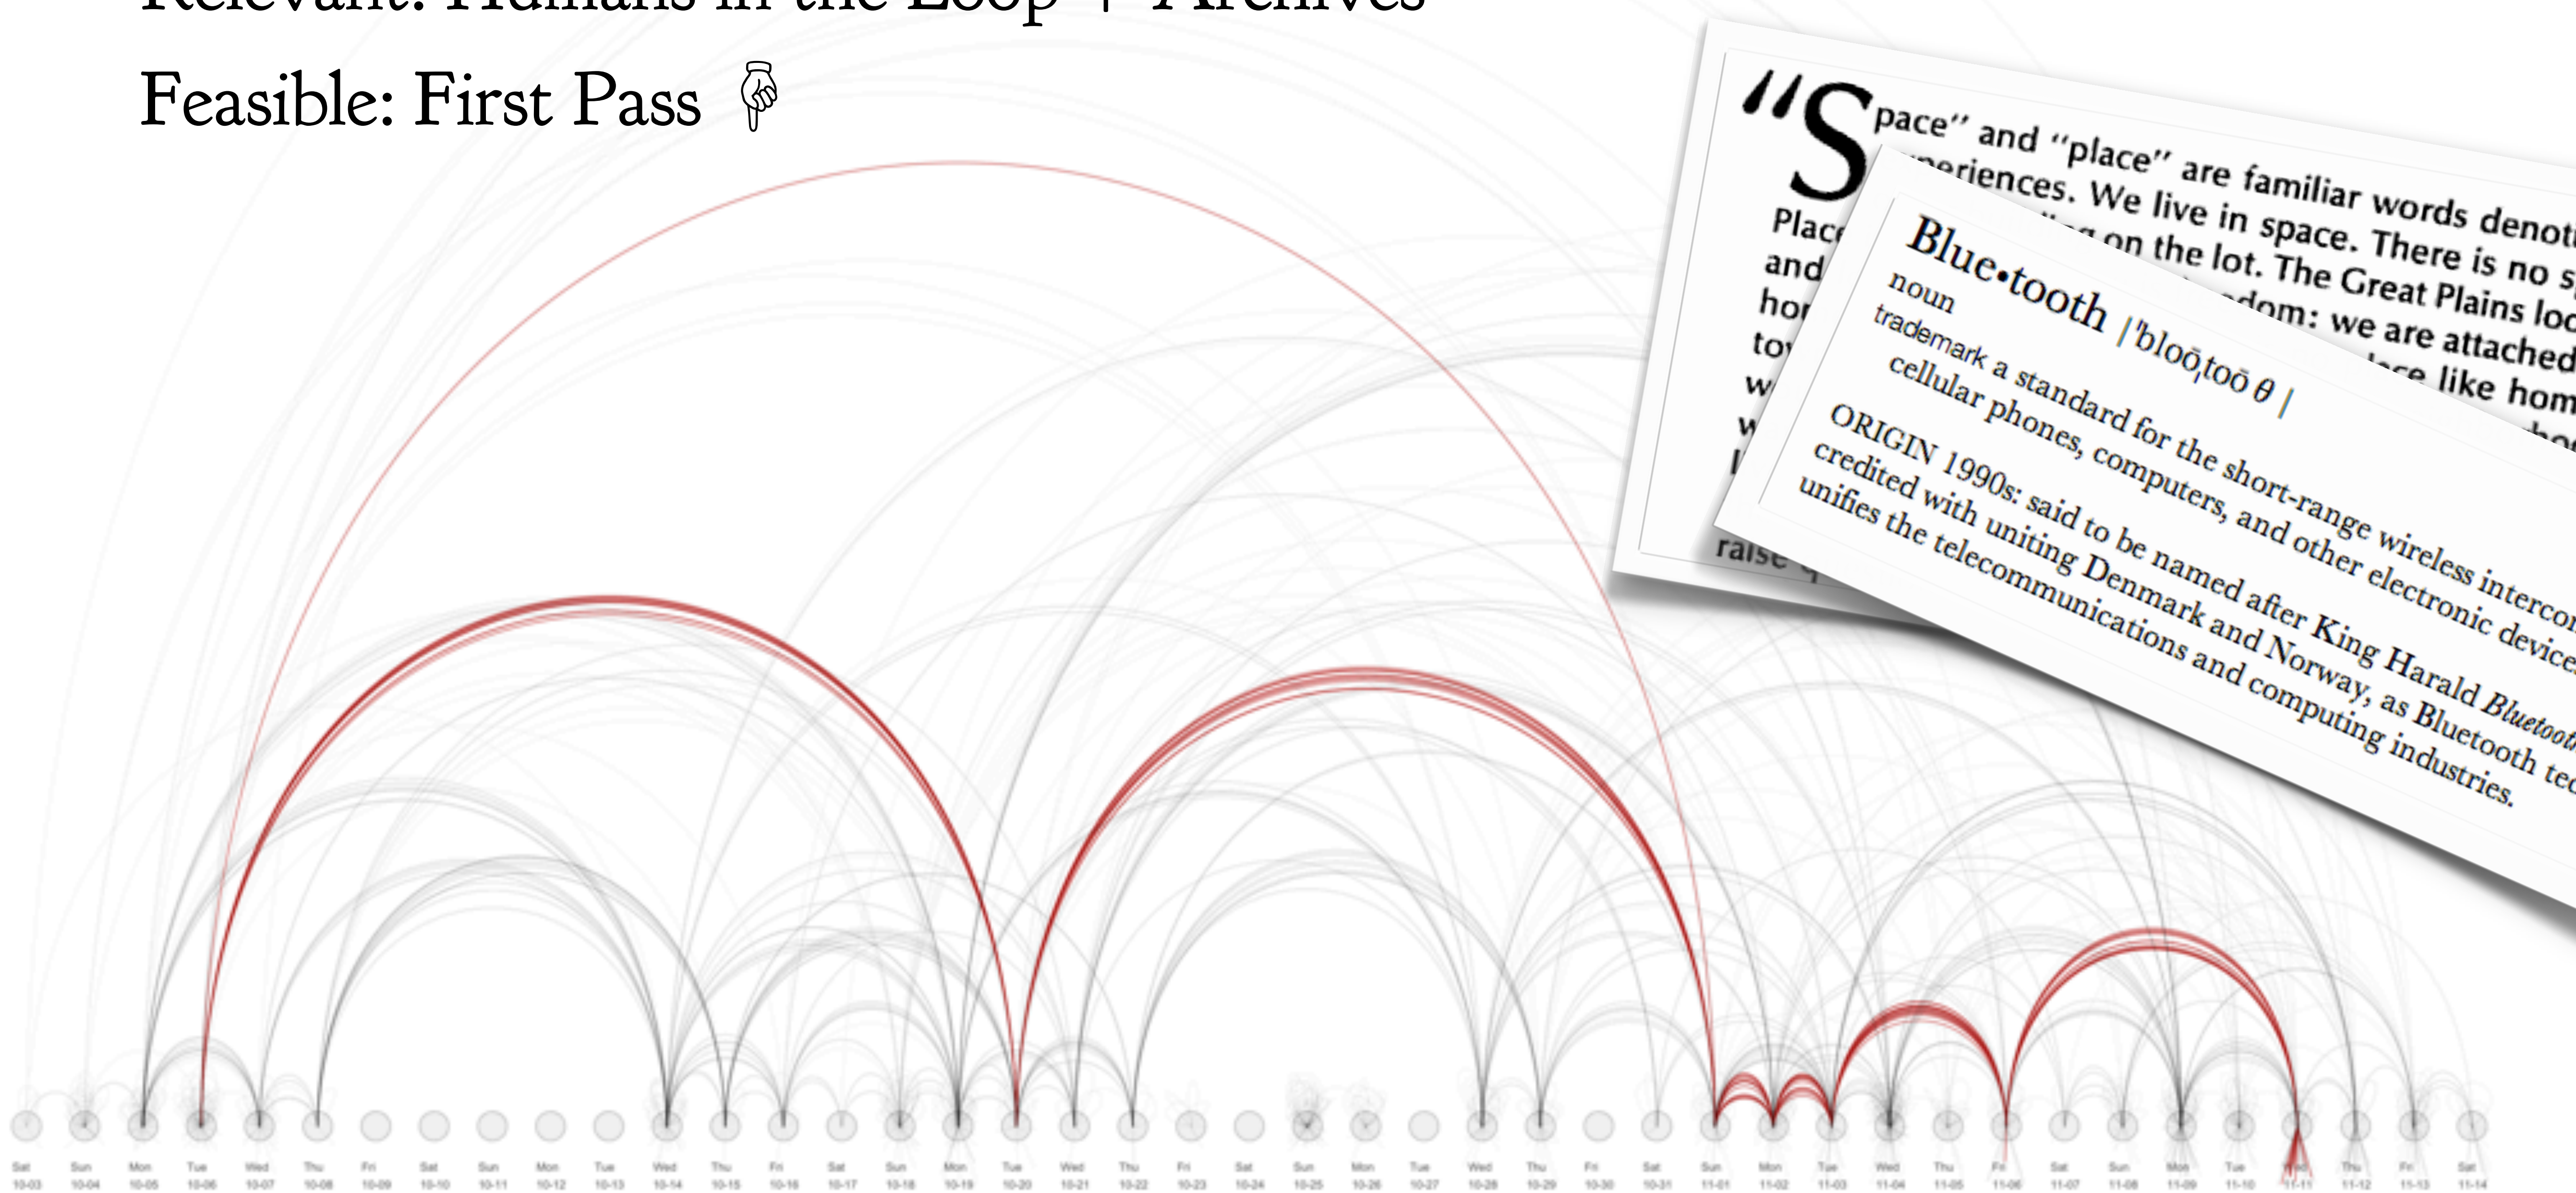

Andrejevic, M. (2007). ISpy: Surveillance and power in the interactive era. Lawrence, Kan: University Press of Kansas.

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

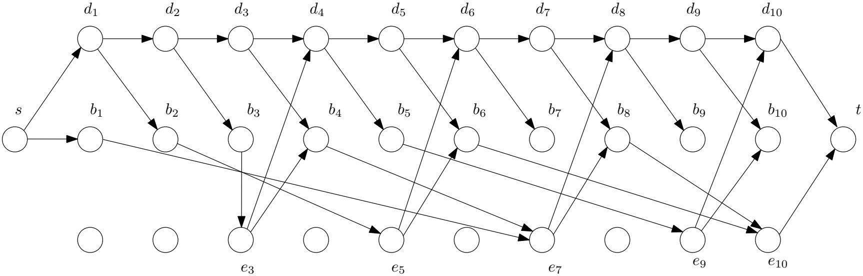

Scheduling Tasks using Natural Language

- 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

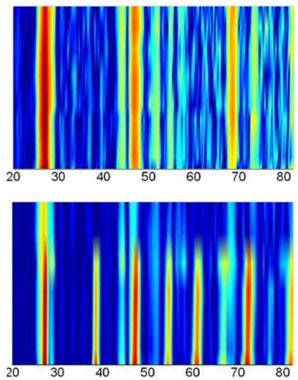

Discover new products in a new area using newly generated data

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

Significant

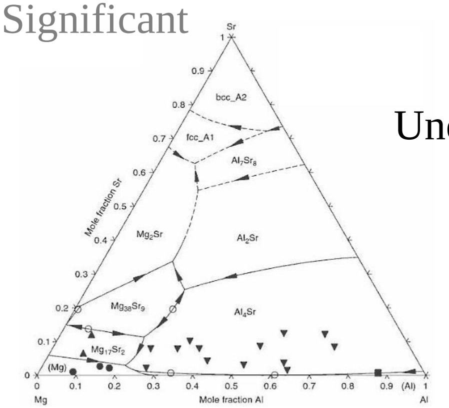

Material Property Understanding, Product Substitute, Resource Management…

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

Relevant

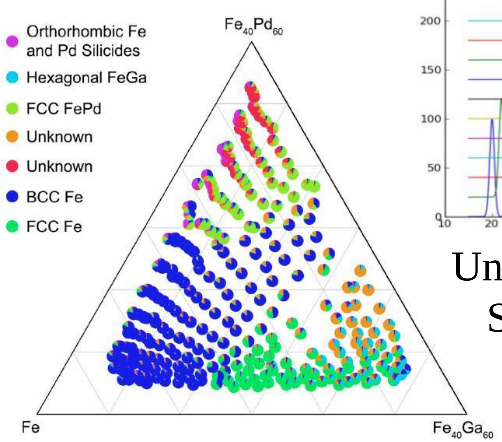

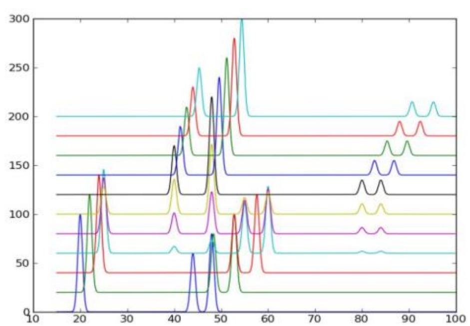

Unsupervised Learning, Structured Prediction, ICA, Clustering with Constraints…

[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

Cornell University
Cornell High Energy Synchrotron Source

Predicting Wikipedia updates

Ruben Sipos, rs@cs.cornell.edu

interesting | **significant** | **relevant** | **feasible**

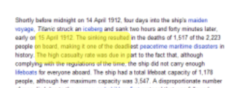

3D Brain MRI Image Segmentation

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

SMI6@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

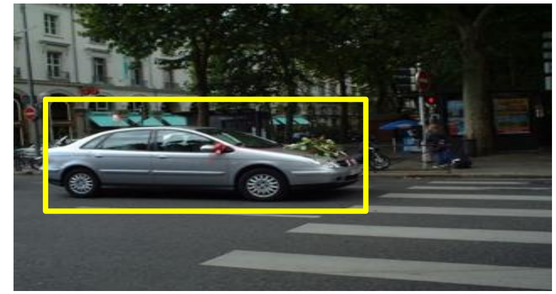

- Traditional Method
  - Classifier: Input: a window,
    Output: 1/0
  - Slide a window on the input image, and decide whether each window is the bounding box for the object →classification
- Proposed Method
  - Input: the whole image, Output: the likelihood to be a bounding box for each location
  - 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