
COMP 875
Machine Learning Methods
in Image Analysis

What the class is about


- “Applied” machine learning and statistical methods
- Applications are primarily, though not exclusively, to computer vision and medical imaging
- Students from other research areas are welcome
- Exact list of topics to be determined by *you!*

Who should take this class?


- This is meant as an “advanced” graduate course
- Ideally, you should have taken COMP 665, 775, 776, or Data Mining (or similar courses elsewhere)
- You should be comfortable reading and understanding papers in recent conferences such as CVPR, ICCV, MICCAI, NIPS, ICML, etc.
- You should have some experience doing research presentations
- If you have questions or doubts about your background, please talk to me after this class

Why Machine Learning?


- Image analysis early on: simple tasks, few images




(a) Original picture.



(b) Differentiated picture.



(c) Line drawing.

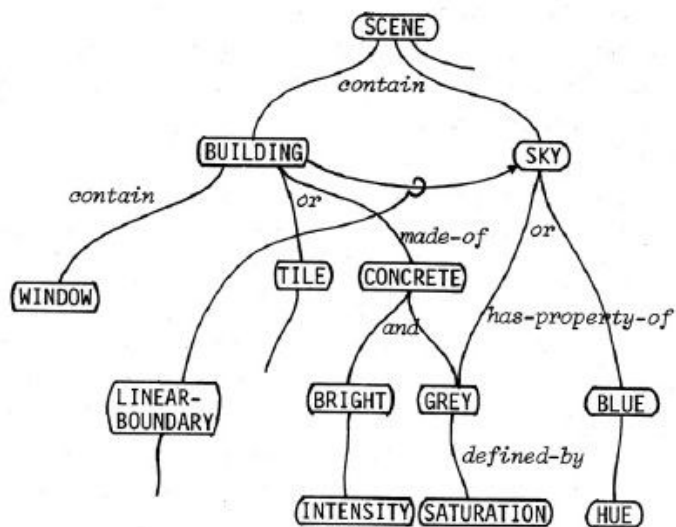


(d) Rotated view.

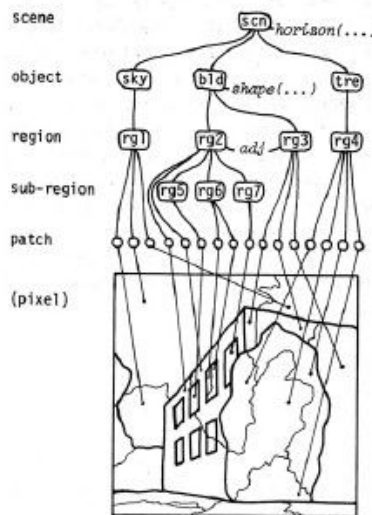
L. G. Roberts, *Machine Perception of Three Dimensional Solids*, Ph.D. thesis, MIT Department of Electrical Engineering, 1963.

Why Machine Learning?

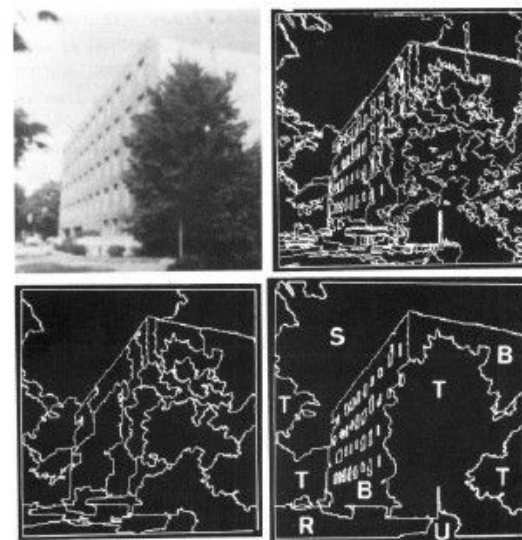
- Image analysis early on: try to program a computer directly using rules and symbolic representations



(a) Bottom-up process



(b) Top-down process



(c) Result

Figure 3. A system developed in 1978 by Ohta, Kanade and Sakai [33, 32] for knowledge-based interpretation of outdoor natural scenes. The system is able to label an image (c) into semantic classes: S-sky, T-tree, R-road, B-building, U-unknown.

Y. Ohta, T. Kanade, and T. Sakai, "An Analysis System for Scenes Containing objects with Substructures," *Proceedings of the Fourth International Joint Conference on Pattern Recognition*, 1978, pp. 752-754.

Why Machine Learning?

- Today: Lots of data, complex tasks




Internet images,
personal photo albums



Movies, news, sports



Surveillance and security



Medical and scientific images

Why Machine Learning?

- Today: Lots of data, complex tasks
- Instead of trying to encode rules directly, learn them from examples of inputs and desired outputs

Not Just Image Analysis

- Speech recognition
- Document analysis
- Spam filtering
- Computer security
- Statistical debugging
- Bioinformatics
-

Topics (tentative)

- Classifiers: linear models, boosting, support vector machines
- Kernel methods
- Bayesian methods, Expectation Maximization
- Random field models
- Sampling techniques such as Markov Chain Monte Carlo
- Unsupervised learning: density estimation, clustering
- Manifold learning and dimensionality reduction
- Distance metric learning
- Semi-supervised learning
- Online and active learning
- Sequential inference (i.e., tracking)
- Large-scale learning

Class requirements

- Class format: lectures and student presentations
- Grading:
 - Presentation: 35%
 - Project: 35%
 - Participation: 30%

Presentation

- You are “professor for a day”: you need to give a one-hour lecture that would be interesting and accessible to all the students in the class
- You are responsible for selecting your own topic and paper(s)
 - Look at the [list of reading materials](#) on the class webpage
 - Look through recent conference proceedings
 - Pick a topic of interest based on your own research

Presentation Guidelines

- **Evaluation criteria**
 - **Integration:** utilize multiple sources
 - **Critical thinking:** separate the essential from the non-essential; critique the papers you present; think of alternative applications and future research directions
 - **Interactivity:** try to involve the rest of the class
- **Structuring the presentation**
 - Will depend on your focus
 - Broadly speaking, you may want to focus either on a particular learning topic, or a particular application

Sample Presentation Outline

- Introduction
 - Problem definition
 - Problem formulation
 - Significance
- Survey of methods for solving this problem
- Detailed presentation of one or more specific methods
- Discussion
 - Pluses and minuses of different methods
 - Compare and contrast different approaches
 - Ideas for improvement and future research
 - Alternative applications
 - Alternative methods for solving the same problem
 - Connect your topic to other topics discussed earlier in class

Presentation Timeline

- **Reading list:** **due next Thursday, September 3rd**
- **Preliminary slides:** due Monday the week before your scheduled presentation
- **Practice meeting:** scheduled for the week before your presentation
- **Final slides:** due by the end of the day after your presentation
- All of the above are part of your presentation grade (35% of total class grade)
- A note on slides: you must explicitly credit all sources

Project

- Your project topic may be the same as your presentation topic
 - Not required, but may make your life easier
- Two options: implementation or survey paper

Implementation

- Implement one or more methods from literature
- Conduct a comparative evaluation
- Implement your own ideas or extensions of existing methods

- Deliverable: an “informal” final report and (possibly) a short presentation
- Students may collaborate, but each must submit his/her deliverables
- You can use existing code and/or software, provided you document all your sources and it doesn't make your project trivial

Survey Paper

- Comprehensive tutorial, literature review
- A “formal” academic paper
- Typeset in LaTeX, 10-15 pages (single-spaced, 11pt font)
- Must be individual

Project timeline (tentative)

- **Project proposal:** due end of September (details to follow)
- **Progress report** (for implementation) or **draft paper** (for survey, ~5 pages): due end of October
- **Final report or paper:** due last day of class (December 8th)
- All of the above are part of your project grade (35% of total class grade)

Participation (30% of the grade)

- Class attendance, being on time
- Answer questions in review sessions at the beginning of class
- Be prepared
 - Read all the material before the class and come up with ~3 questions for discussion
 - I may call on anyone at any time
- Participate in discussions
- Send email to me and/or the class mailing group links to material that may be of interest
- *If you never speak up in class, the best grade you can get is P+!*

What's next?

- First few weeks: lectures on the basics of machine learning
- **Reading lists due next Thursday, September 3rd**
 - Also send any date constraints/preferences
 - Topics may require some conflict resolution
 - “Bonus points” for the first two students to present
- Class schedule finalized by the end of third week of class
- If you have any questions about whether you should take this class, talk to me now!