

## Topics in Multimedia Research

Dr. Michael S. Lew

Director of the Media Research Group  
co-Head of the Imagery and Media Research Cluster  
Leiden University  
E-mail: [mlew@liacs.nl](mailto:mlew@liacs.nl)

## Multimedia Research Community

- *Major Scientific Conferences*
  - *ACM International Conference on Multimedia*
  - *ACM SIGGRAPH - computer graphics*
- *Large Influential Research Labs*
  - *Google*
  - *Yahoo Research*
  - *Microsoft Research*

## LIACS Media Research Group

*We are interested in the frontiers of multimedia research, technology which will often be commonplace 2-5 years in the future.*

## Some Collaborators

(about 10 fulltime members in total)

- **Dr. Erwin M. Bakker**, [erwin@liacs.nl](mailto:erwin@liacs.nl)
  - <http://www.liacs.nl/~erwin/>
  - Audio/Speech search and modeling
  - BioInformatics
- **Dr. Nies Huijsmans**, [huijsman@liacs.nl](mailto:huijsman@liacs.nl)
  - <http://www.liacs.nl/~huijsman/>
  - Computer vision
  - Cultural heritage search and browsing
- **Dr. Michael Lew**, [mlew@liacs.nl](mailto:mlew@liacs.nl)
  - <http://www.liacs.nl/~mlew/>
  - Machine learning
  - Graphics & Multimedia Information Retrieval



## Some Project Topics

- *Topics*
  - *Social Computing and Recommendation Systems*
  - *Social Media Analysis*
  - *Computer Vision*
    - *Visual Concept Detection*
    - *Face Detection and Tracking*
  - *Computer Graphics and Infinite Worlds*
  - *Searching large image/video collections*

## Social Computing & Recommendation Engines

- Recommendation engines are at the intersection of machine learning and societal usage.
- Just about every Internet shop wants to recommend new items to buy (Bol, Amazon)
- or People to connect to
- or Groups to join
- or What movie should you see next?

## Social Computing & Recommendation Engines

- List of  $m$  users and a list of  $n$  Items .
- Each user has a list of items he/she expressed their opinion about (can be a null set).
- Explicit opinion - a rating score (numerical scale).

## Collaborative Filtering (CF)

- The task of a CF algorithm is to find item likeliness of two forms :

**Prediction** – a numerical value, expressing the predicted likeliness of an item the user hasn't expressed his/her opinion about.

**Recommendation** – a list of  $N$  items the active user will like the most (Top- $N$  recommendations).

## Collab Filter - No Problems

	Star Wars	Sense and Sensibility	Inception	Pride and Prejudice	The Empire Strikes Back
Alex	9	6	9	3	10
Jamie	7	8	7	10	7
Terry	9	5	9	?	?

## Problem: User Conflict

What do you do when user ratings conflict?

	Star Wars	Sense and Sensibility	Inception	Pride and Prejudice	The Empire Strikes Back
Alex	10	5	9	3	10
Jamie	9	5	8	10	7
Terry	9	5	9	?	?

## Project - Video Recommendation

MovieLens Project has complete ratings databases of varying size: 100K, 1M, etc.

<http://www.grouplens.org/node/73>

### Plan:

- (1) Download a MovieLens DB
- (2) Write a simple prediction algorithm
- (3) Design and write a more sophisticated algorithm
  - indirect links
  - other machine learning algorithms
  - etc.

## Reading

*A large-scale evaluation and analysis of personalized search strategies*

Zhicheng Dou, Ruihua Song, Ji-Rong Wen

In Proceedings WWW '07 Proceedings of the 16th International Conference on World Wide Web

*Evaluating collaborative filtering recommender systems*

Jonathan L. Herlocker, J.A. Konstan, L.G. Terveen, J.T. Riedl

ACM Transactions on Information Systems (TOIS), Volume 22 Issue 1, January 2004

## Social Media Analysis

- For a social post or a new product, there are often reviews and comments.
- How can one automatically summarize the positive/negative view toward the post or item?
- Starting point: analyze occurrence of certain categories of words which have positive or negative meanings.

## Computer Vision

- Visual Concept Detection



## Visual Concept Detection




## Computer Vision

- Visual Concept Detection
- **Start:** Use OpenCV to extract salient points and perform initial classification
- Instead of creating test sets with thousands of positive and negative examples, how can one interactively give feedback to the system to reduce training time?
- Begin by allowing user to select incorrectly labeled patches and re-train with correct labels



## Computer Vision

- Could we link a face tracker to a game character?
- **Start:** Use OpenCV (public software) to do tracking



## CG: Infinite 3D Worlds

- What are effective methods for creating 3D worlds of unlimited size?
- How would one render an entire planet?
- **Start:** Look at combining functions like *sin/cos* or *random number generators* to get infinite depth maps

## World Building



## 3D Worlds Influenced By Natural Phenomena

Computer graphics worlds are often designed manually  
What about evolving or creating them by some rough  
knowledge of physics or environmental phenomena?

- How to create a 3D world based on the continental fault lines?
- How to create a 3D world based on the movement of storms?
- **Start:** Model initial world using height map and alter using basic physics:  $\text{Force} = \text{mass} * \text{acceleration}$

## Graphical Crowd Models

- Recent work has modeled crowds typically ignores different "agent" size, effects (children, trampling,...).

**Start:** Model more realistic crowd flow using rules



## Searching Media Collections

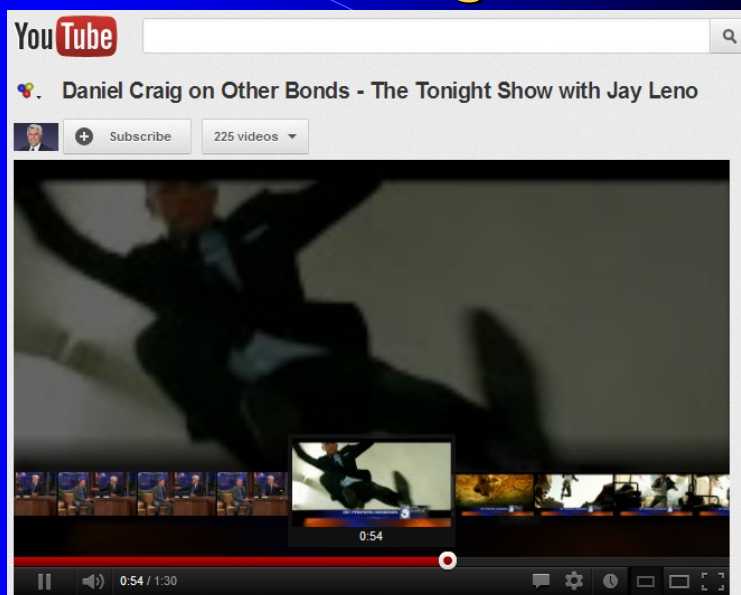
- Suppose you have a million images and video, how do you find anything?
- Image Near-Copy Detection
  - How do you find copies of images which are cropped, have additional text written on them, scaled, rotated, etc?

**Start:** Human eye works quite well. Model a new feature descriptor after the retina

## High Performance Search

- In distributed web downloading (i.e. Google), you do not want the overhead of cross checking signatures (*cryptographic hash function*) between download nodes.
  - webpage -> signature (hopefully unique)
  - minimize number of bits in signature
- Which signatures algorithm (i.e. MD5, CRC) performs the best for web downloads at a particular number of bits?
- How do they perform against a modern pseudo-random number generator?
- uniqueness (or collisions) vs. cpu generation time

## Video Browsing/Search



## Video Browsing/Search

- YouTube represents each video shot as a thumbnail.
- What other representations of video shots are there?
- How effective would it be to have immediately responsive seeking in long videos or other interfaces for extremely fast video browsing/search?

## SQL and Image Search

- How can one convert image search to an SQL search problem?
- `SELECT ImgID FROM ImgTB WHERE color="green" AND texture="trees" AND position="upper"`
- Stage 1: Start by extracting straightforward color and texture features from images and storing in SQL database (open source public code)
- Stage 2: Look deeper into what features would be most useful to store in SQLDB and what queries would be necessary?



## Image Search in 3D

What about 3D interfaces? What would be an effective interface?



## Concluding Remarks

The project descriptions are only *starting points*.

The research question/goal and approaches will evolve with discussion.

If these subjects are interesting to you, please email me for a chat about what you would like to do.