Where computer vision needs help from computer science

Bill Freeman
Electrical Engineering and Computer Science Dept.
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology

SODA Jan. 24, 2011

Outline

- How I got excited about computer vision
- Computer vision applications
- Computer vision techniques and problems:
  - High-level vision: combinatorial problems
  - Low-level vision: underdetermined problems
  - Miscellaneous problems

The Taiyuan University of Technology Computer Center staff, and me (1987)

The Taiyuan University of Technology Computer Center staff, and me (1987)

Me and my wife, riding from the Foreigners' Cafeteria

Me in my office at the Computer Center

Years ago, I read this book (re-issued by MIT Press in 2010), and got very excited about computer vision.

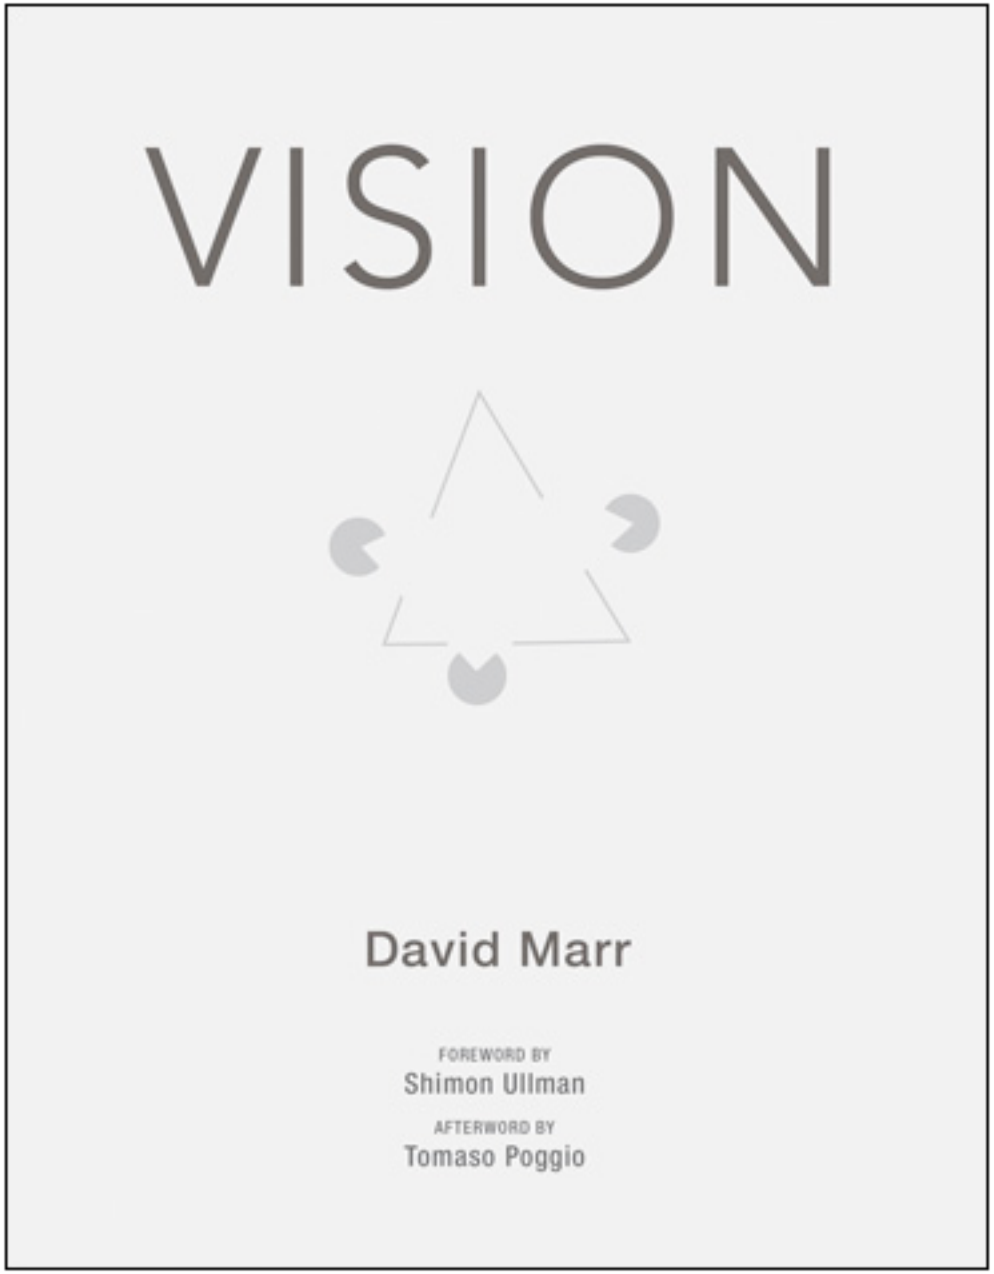

Goal of computer vision

Marr: "To tell what is where by looking".

Want to:

- Estimate the shapes and properties of things.
- Recognize objects
- Find and recognize people
- Find road lanes and other cars
- Help a robot walk, navigate, or fly.
- Inspect for manufacturing

Some particular goals of computer vision

- Wave a camera around, get a 3-d model out.
- Capture body pose of actor dancing.
- Detect and recognize faces.
- Recognize objects.
- Track people or objects
- Enhance images

Companies and applications

- Cognex
- Poseidon
- Mobileye
- Eyetoy
- Identix
- Google
- Microsoft
- Face recognition in cameras


ognex Machine Vision Systems and Machine Vision Sensors! - Microsoft Internet Explorer
Edit View Favorites Tools Help
Back Search Favorites Media
http://www.cognex.com/
Go
COGNEX
World Leader in Machine Vision
SITE MAP
PRESS
CONTACT US
PARTNERS
Search
INVESTOR'S BUSINESS
Corporate
Products
Support
Education
News | Events
Investors
Careers
-- Global Web Sites --
ognex ® is the world's leading supplier of machine vision systems, or computers that can "see". Our machine vision stems gauge, guide, inspect, and identify products on the fastest production lines. Our vision systems worldwide help anufacturers automate the production of a wide range of products .... from semiconductor chips to chocolate chip okies.
ur proven technology, application expertise, and worldwide support mean customers can rely on us to deliver machine sion solutions that work every time - even under the most difficult factory floor conditions. In the industrial machine sion market, Cognex has supplied more vision systems ... over 200,000 ... than any other company. Customers rely on ognex as a long-term partner, working alongside them to continually find new ways to improve the quality, productivity, d profitability of their manufacturing operations. Cognex ... the machine vision leader that industry relies on.
REGISTER
Machine Vision SEMINARS
LET US EVALUATE YOUR MACHINE VISION APPLICATION
FREE IN-SIGHT EXPLORER TRIAL SOFTWARE
NEW
Cognex Machine Vision News
3.03.04 Cognex Named Preferred Machine Vision Supplier by Leading German Automotive Engineering Firm
3.03.04 Machine Vision Bolsters Quality Control - PDF
3.01.04 Fortune Magazine Names Cognex Corporation's Founders "Heroes of Manufacturing"
Sign Up for our InFocus E-Newsletter
Machine Vision Industry Applications
Machine Vision Reports

Aquatic safety - Drowning problem - Poseidon - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Search Favorites Media

Address http://www.poseidon-tech.com/us/drowning.html Go Links

Poséidon

THE LIFEGUARD'S THIRD EYE

Aquatic Safety
The Drowning Problem

Drowning problem | Facts & Statistics | Lifeguarding challenges | Stages of a drowning

- Home
- The System
- Aquatic safety
- Installed sites
- About us
- Contact us
- News
- Events
- Site map

Select your country

According to the Centers for Disease Control, **9 people drown per day in the U.S**. For every person who drowns, **four times as many people nearly drown**. Many of these incidents happen **in pools staffed with certified professional lifeguards**.

If you've been to a pool recently, you've witnessed firsthand the **challenges that lifeguards face** in monitoring activity within a pool. Not only is it warm, but there are usually lots of swimmers, glare from the sun in some cases, and other distractions. The toughest part of a lifeguard's job is maintaining constant vigilance, and no human being can see everything all the time. But it only takes a second for someone to get into trouble and start to drown. Contrary to what most people think, drowning victims don't yell or wave their arms to alert someone that they are in trouble. They are in a state of shock, and are often silent.

It's vital that lifeguards reach a drowning victim before it's too late, and **every second counts**. To prevent death or lifelong injury, the resuscitation of drowning victims must be initiated as quickly as possible – ideally within 30 seconds.

The solution isn't just more lifeguards or better training. It's a better means of surveillance and detection. It's Poseidon. Poseidon helps lifeguards monitor what is happening in the pool, maintaining vigilance, and alerting them in seconds to a swimmer in trouble. Poseidon does not rescue drowning victims – lifeguards do – but it can help them more quickly initiate a rescue and save a life.

vimming pools with Poseidon - Microsoft Internet Explorer

Edit View Favorites Tools Help

Back Search Favorites Media

http://www.poseidon-tech.com/us/sites.html Go

quatic safety
nstalled sites
bout us
ontact us
ews
vents
ite map

ct your country

Here is a partial list of our installed sites:

USA

- McCoy Natatorium, Penn State University
  University Park
  Pennsylvania, USA

McCoy Natatorium
The Aquatics Facilities of Penn State University

- YMCA Southcoast, New Bedford Division
  New Bedford
  Massachusetts, USA
- Medina Community Recreation Center
  Medina
  Ohio , USA
- Metro Atlanta YMCA - Carl E. Sanders Family YMCA
  Buckhead
  Georgia , USA

Metro Atlanta YMCA
We build strong kids, strong families, strong communities.

- Heritage YMCA Group - 95th St. Family Center
  Naperville
  Illinois, USA

- The Sarasota Family YMCA Pools
  Florida, USA
- The Fort Wayne Community Schools' Pool
  Indiana, USA
- The St. Cloud School District Pools
  Minnesota, USA

EUROPE

Poseidon saved a life here!

From The Times

September 1, 2005

Saved by a computer lifeguard

Drowning girl is spotted on bottom of pool by new high-tech system that watches over swimmers Watch the rescue

By Russell Jenkins

RECOMMEND?

A YOUNG girl has been saved from drowning by an extraordinary computer system that keeps an eye on everybody in a swimming pool.

The girl was pulled unconscious from 12ft of water at the deep end of a public pool in Bangor, North Wales, when underwater cameras spotted that she was not moving and alerted a lifeguard. The lifeguard could not see the girl in the crowded pool but was able to respond to the alert within seconds.

It is the first time in Britain that the Poseidon surveillance system, manufactured by a French company, has helped lifeguards to save a swimmer from drowning. The campaign group Swimsafekids said last night that the rescue proved that the system could save many more lives if they were installed compulsorily.

The state-of-the-art system has been credited with saving three swimmers in France. Last year it helped to save a middle-aged German man who had a heart attack. So far, eight pools in Britain have installed the system.

The girl, from Rochdale, Greater Manchester, was on a camping holiday run by a charitable trust near Bangor. Along with her

Mobil Eye

Address http://www.mobileye.com/
Go
Home
Company
News
Jobs
Contact Us
MOBILEYE
Applications
Technologies
Products
Get Brochures
A Single-Camera
Driving Assistance
System on Chip
Copyright Mobileye N.V.

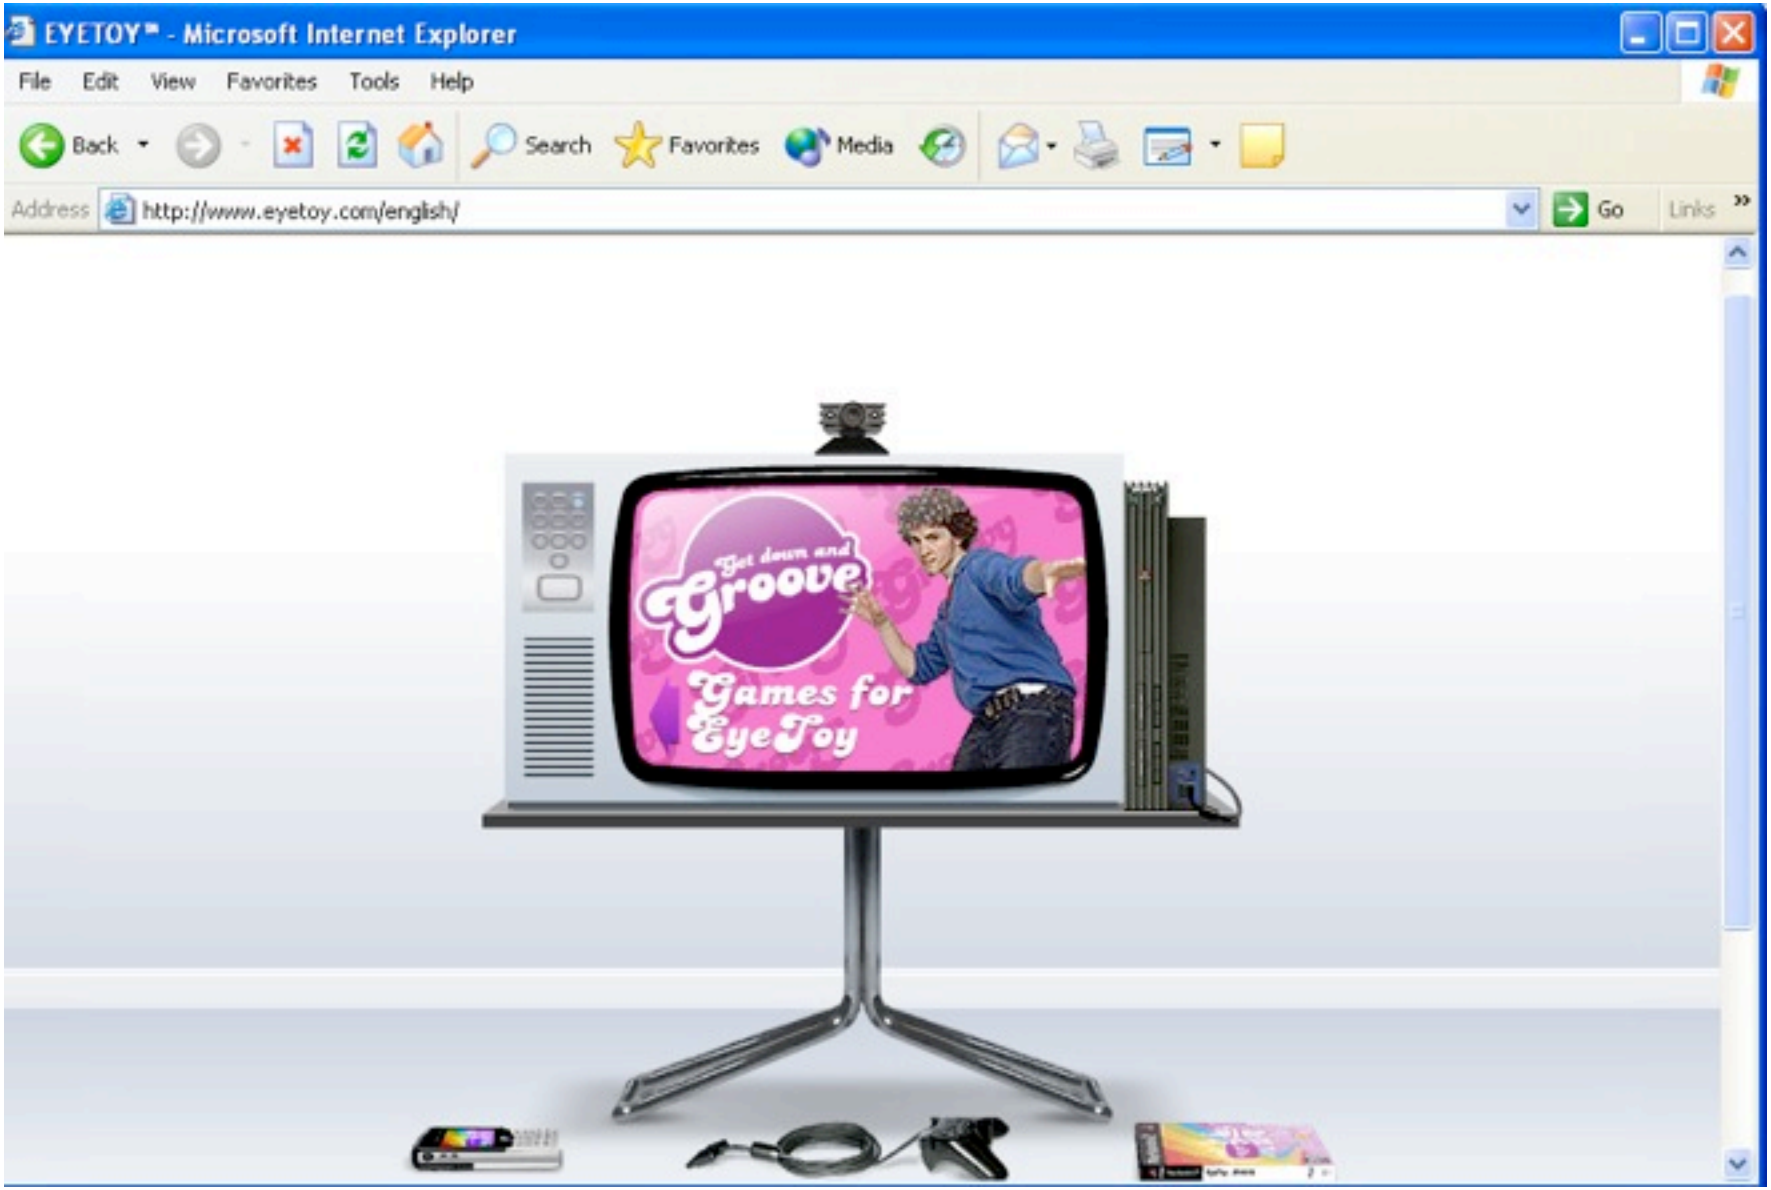
EYETOY™ - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://www.eyetoy.com/english/
Go Links
Get down and Groove
Games for EyeToy

Microsoft Kinect, 2010

File Edit View Favorites Tools Help

Back Search Favorites Media

Address http://www.identix.com/products/pro_security_bnp_argus.html Go Links

identix

SEARCH WEB SITE GO

EMPOWERING IDENTIFICATION™

Facial Surveillance

DOWNLOAD BROCHURES | REQUEST INFO | PRODUCT SUPPORT

Facial Surveillance

FaceIt® ARGUS

Configuration

FACIAL SURVEILLANCE

FaceIt® ARGUS

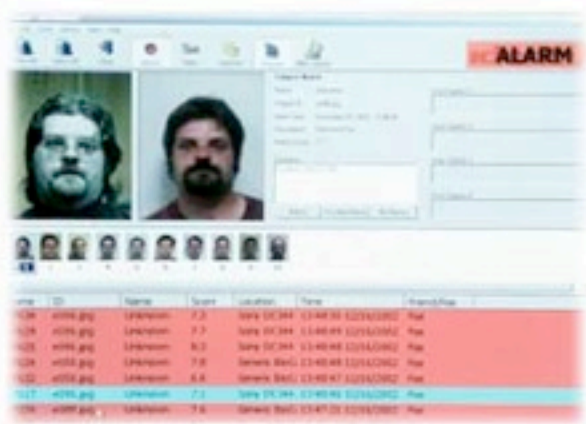

FaceIt® ARGUS is a scalable, off-the-shelf facial recognition system that detects and identifies humans as they pass through a camera's field of view. FaceIt ARGUS maximizes the value of CCTV by increasing deterrence, increasing active surveillance functionality, increasing investigative power and drastically diminishing or eliminating the challenges that exist when relying solely on operators to conduct surveillance.

FaceIt® ARGUS
Features and Benefits

Integrates into any size CCTV system and easily expands as cameras are added.

Central database administration provides ease of data sharing and management.

Overcomes human inability to recognize large numbers of unfamiliar faces and distraction in control room environments.

Visual and audible alert signals capture the operator's attention when a high confidence match is detected.

Adjustable threshold for confidence score maximizes correct alarms or minimizes the occurrence of false alarms depending on security requirements.

Google

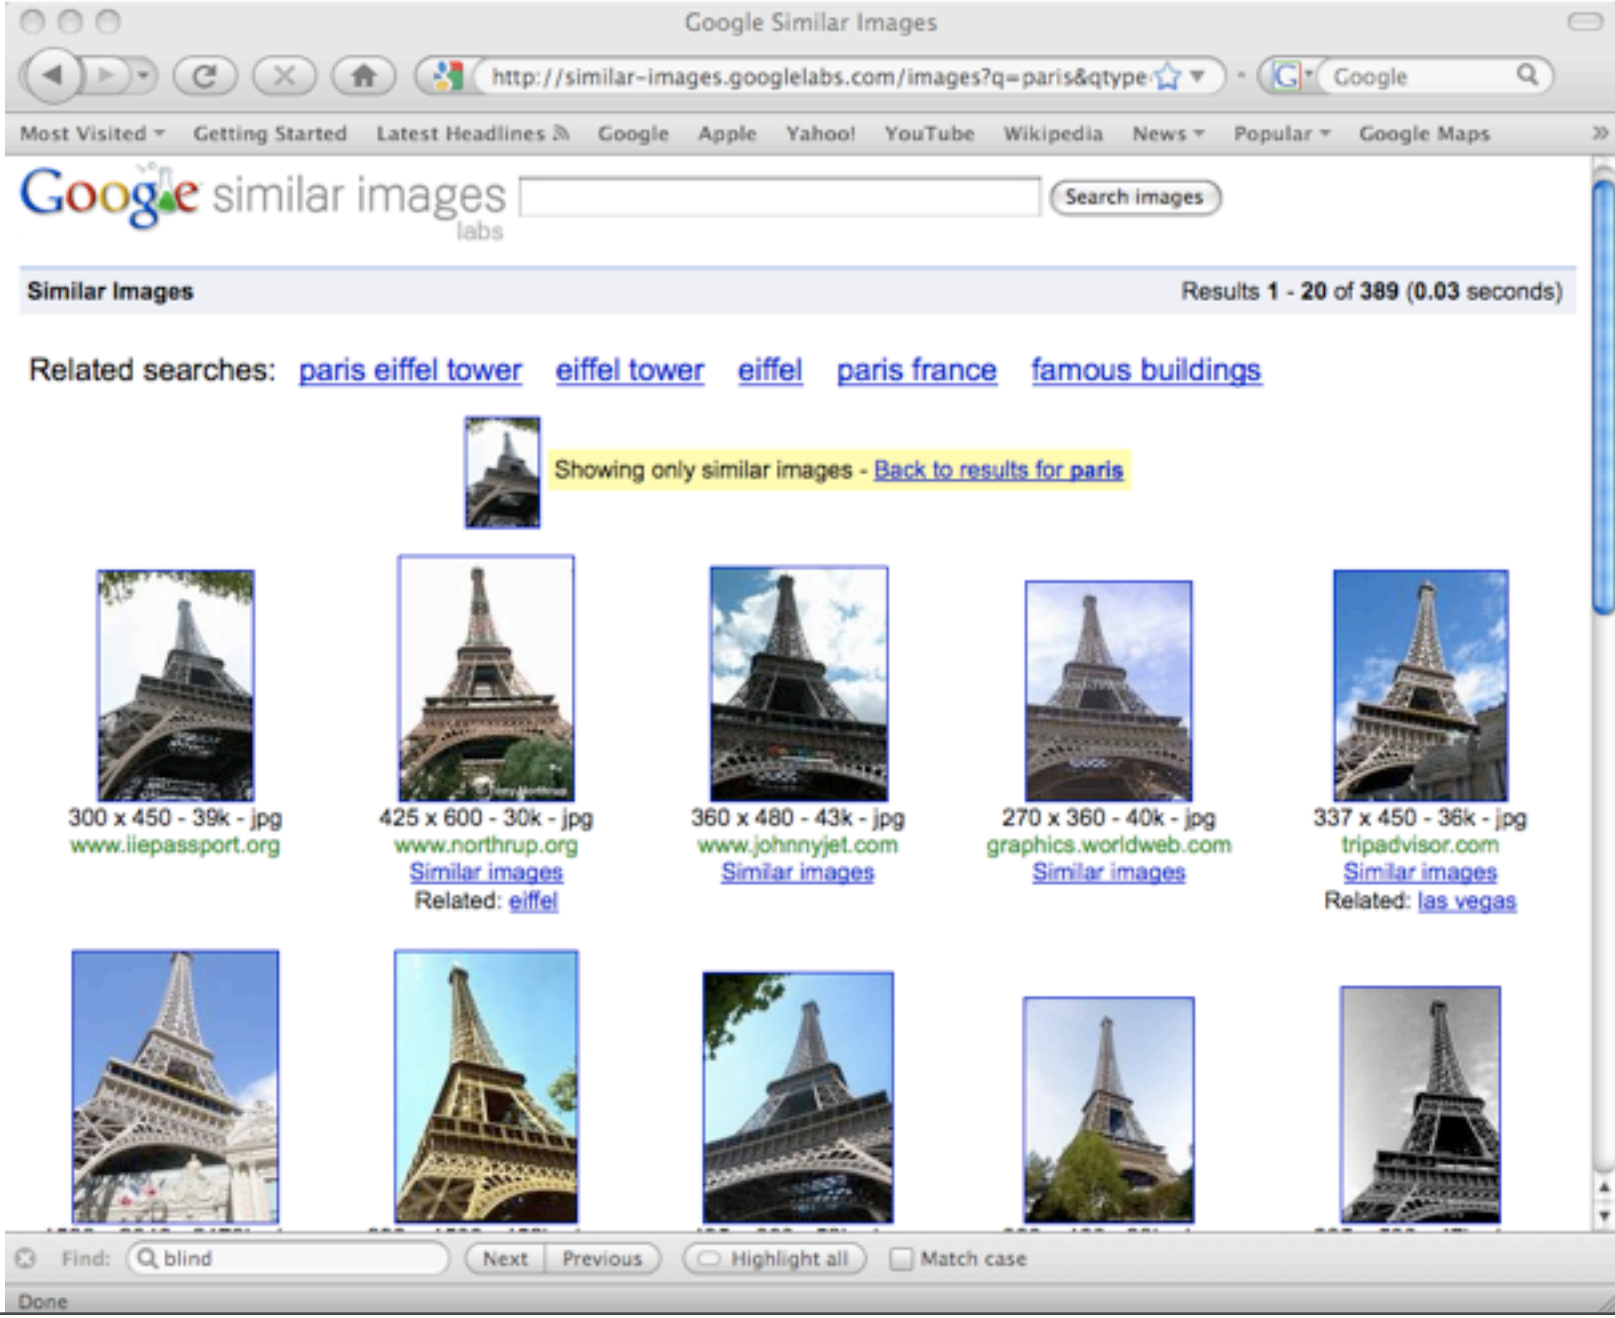

Google Similar Images
http://similar-images.googlelabs.com/images?q=paris&qtype
Google
Most Visited
Getting Started
Latest Headlines
Google
Apple
Yahoo!
YouTube
Wikipedia
News
Popular
Google Maps
Google similar images labs
Search images
Similar Images
Results 1 - 20 of 389 (0.03 seconds)
Related searches: paris eiffel tower eiffel tower eiffel paris france famous buildings
Showing only similar images - Back to results for paris
300 x 450 - 39k - jpg
www.iiepassport.org
425 x 600 - 30k - jpg
www.northrup.org
Similar images
Related: eiffel
360 x 480 - 43k - jpg
www.johnnyjet.com
Similar images
270 x 360 - 40k - jpg
graphics.worldweb.com
Similar images
337 x 450 - 36k - jpg
tripadvisor.com
Similar images
Related: las vegas
Find: blind
Next
Previous
Highlight all
Match case
Done

The Computer Vision Industry

David Lowe

This web page gives a listing of companies that develop computer vision products. Computer vision (also often referred to as "machine vision" or "automated imaging") is the automated extraction of information from images. This differs from image processing, in which an image is processed to produce another image. This page covers only products based on computer or machine vision, and it does not cover image processing or any of the many suppliers of sensors or other equipment to the industry.

Companies are categorized under their principal application area, and then listed alphabetically. Companies are listed only if they have web pages giving information about their products. Please let me know of any links that are missing.

Automobile driver assistance

Iteris (Santa Ana, California). Lane departure and collision warning systems for trucks and cars. Used in over 100,000 vehicles (2009). Also creates traffic monitoring systems.

MobilEye (Jerusalem, Israel). Vision systems that warn automobile drivers of danger, provide adaptive cruise control, and give driver assistance.

Digital Photography

Cloudburst Research (Vancouver, Canada). Develops fully automated image stitching for the iPhone platform (author of this list is a founder of the company).

Kolor (Challes les eaux, France). Develops the Autopano Pro software for automated panorama stitching of digital images. Also provides high-dynamic-range imaging by combining multiple exposures.

Eye and Head Tracking

Smart Eye (Göteborg, Sweden). Systems to track eye and gaze position. Applications include detection of drowsiness or inattention in drivers.

Games and Gesture Recognition

Canesta (Sunnyvale, California). Time-of-flight range sensors and software for gesture recognition. Acquired by Microsoft in 2010.

GestureTek (Toronto, Canada). Tracks human gestures for playing grames or interacting with computers.

PrimeSense (Tel-Aviv, Israel). Real-time projected infrared depth sensor and software for gesture recognition. Developed the sensing system in Microsoft's Xbox Kinect.

Reactrix (Redwood City, California). Interactive advertising for projected displays that tracks human gestures.

Sony EyeToy uses computer vision to track the hand and body motions of players to control the Sony Playstation. Sales were over 10 million units by 2008. (Wikipedia)

General purpose vision systems

Cognex (Natick, Massachusetts) is one of the largest machine vision companies (700 employees, 2009). Develops systems for inspection and localization tasks, people counting, and many other areas. (Hoover's)

Evolution Robotics (Pasadena, California). Vision systems for object recognition and navigation. Applications include mobile robotics and grocery retail. Develops integrated vision and navigation solutions for household robots.

Matrox Imaging (Dorval, Quebec, Canada). Software and hardware for machine vision applications.

National Instruments (Austin, Texas). Vision software and systems used for many applications, including inspection, biomedical, and security.

Neptec (Ottawa, Canada). Laser-based 3D vision systems for use on the space shuttles and other applications.

Newton Research Labs (Renton, Washington). Vision systems for high-speed tracking and mobile robots.

Point Grey Research (Vancouver, Canada). Real-time stereo vision systems, spherical vision systems, and imaging

Industrial automation and inspection: Electronics industry

KLA-Tencor (San Jose, California). Systems for inspection and process control in semiconductor manufacturing.

Orbotech (Yavne, Israel). Automated inspection systems for printed circuit boards and flat panel displays. (Hoover's)

Industrial automation and inspection: Food and agriculture

Montrose Technologies (Ottawa, Canada). Vision systems for the baked goods industry. Systems monitor bake color, shape, and size of bread, cookies, tortillas, etc.

Ellips (Eindhoven, The Netherlands). Vision systems for inspecting and grading fruits and vegetables.

Industrial automation and inspection: Printing and textiles

Advanced Vision Technology (Hod Hasharon, Israel). Systems to inspect output from high-speed printing presses.

Elbit Vision Systems Ltd. (Yoqneam, Israel). Vision systems for textile inspection and other applications.

Mnemonics (Mt. Laurel, New Jersey). Vision systems for print quality inspection and other applications.

Xiris Automation (Burlington, Ontario, Canada). Inspection for the printing and packaging industries.

Industrial automation and inspection: Other

Adept (Pleasanton, California). Industrial robots with vision for part placement and inspection.

Avalon Vision Solutions (Lithia Springs, Georgia). Vision systems for the plastics industry.

Basler (Ahrensburg, Germany). Inspection systems for optical media, sealants, displays, and other industries.

Object Recogntion for Mobile Devices

Kooaba (Zurich, Switzerland). Visual search for smartphones, photo management, and other applications.

SnapTell (Palo Alto, California). Image recognition and product search for camera phones. Owned by Amazon A9.

People tracking

Brickstream (Atlanta, GA). Tracking people within stores for sales, marketing, and security.

Reveal (Auckland, New Zealand). Systems for counting and tracking pedestrians using overhead cameras.

VideoMining (State College, PA). Tracking people in stores to improve marketing and service.

Safety monitoring

MG International (Boulogne, France). The Poseidon System monitors swimming pools to warn of accidents and drowning victims.

Security and Biometrics

Aimetis (Waterloo, Ontario, Canada). Systems for intelligent video surveillance.

Aurora (Northampton, UK). Systems for biometric face recognition.

AuthenTec (Melbourne, Florida). Fingerprint recognition systems with a novel sensor.

Cernium (Reston, Virginia). Systems for behavior recognition in real-time video surveillance.

Digital Persona (Redwood City, California). Fingerprint recognition systems.

Done

Three-dimensional modeling

Creative Dimension Software (Guildford, UK). Creates 3D models from a set of images. Objects are imaged on a calibration mat.

Eos Systems (Vancouver, Canada). PhotoModeler software allows creation of texture-mapped 3-D models from a small number of photographs. Uses some manual user input.

Eyetronics (Leuven, Belgium). Produces a 3-D scanner for the human body using structured light.

InSpeck (Quebec City, Canada). Uses projected light to create a full 3-D textured model of the human face or body in sub-second times.

Traffic and road management

Image Sensing Systems (St. Paul, Minnesota). Created the Autoscope system that uses roadside video cameras for real-time traffic management. Over 100,000 cameras are in use.

Yotta (Leamington Spa, UK). Imaging and scanning solutions for road network surveying.

Web Applications

Face.com (Tel Aviv, Israel). Image retrieval based on face recognition.

Incogna (Ottawa, Canada). Develops a system for image search on the web. Uses GPUs for increased performance.

LTU Technologies (Paris, France). Image retrieval based on content.

Photometria (San Diego, California). Virtual makeover website, TAAZ.com uses computer vision methods to allow users to try on makeup, hair styles, sunglasses, and jewelery.

Some particular goals of computer vision (status report)

- Wave a camera around, get a 3-d model out (almost)
- Capture body pose of actor dancing. Using multiple cameras (pretty well), using a single camera (not yet)
- Detect and recognize faces. (frontal, yes)
- Recognize objects. (working on it, lots of progress)
- Track people or objects (over short times)
- Enhance images (great image enlargements by 4x)

Outline

- About me
- Computer vision applications
- Computer vision techniques and problems:
  - High-level vision: combinatorial complexity
  - Low-level vision: underdetermined problems
  - Miscellaneous problems

Outline

- About me
- Computer vision applications
- Computer vision techniques and problems:
  - High-level vision: combinatorial complexity
  - Low-level vision: underdetermined problems
  - Miscellaneous problems

What makes computer vision hard?

What makes computer vision hard?

- variability.

What makes computer vision hard?

Mike Burton, http://www.psy.gla.ac.uk/~mike/averages.html

- variability.

intra-class variation

Slide from: Li Fei-Fei, Rob Fergus and Antonio Torralba, short course on object recognition, http://people.csail.mit.edu/torralba/shortCourseRLOC/

Object recognition issues

– Generative / discriminative / hybrid

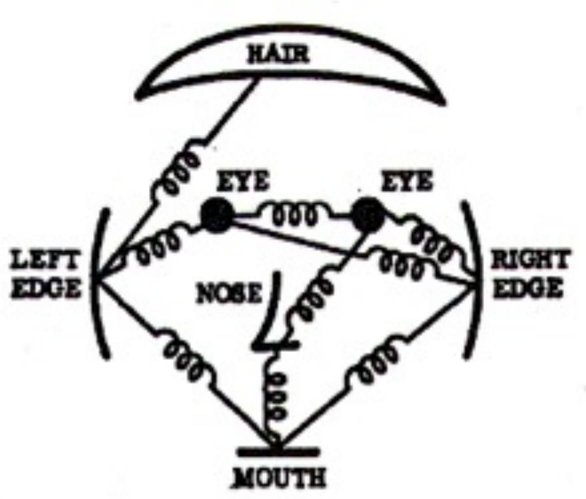

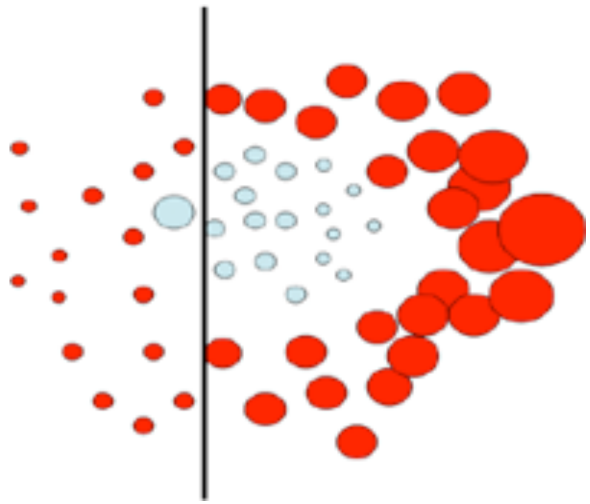

Slide from: Li Fei-Fei, Rob Fergus and Antonio Torralba, short course on object recognition, http://people.csail.mit.edu/torralba/shortCourseRLOC/

Object recognition issues

– Generative / discriminative / hybrid

– Appearance only or location and appearance

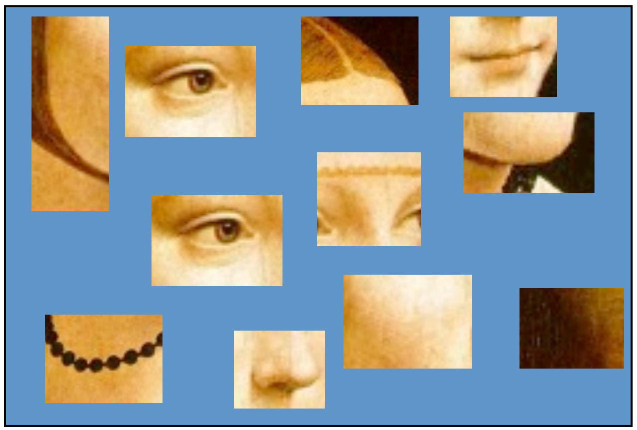

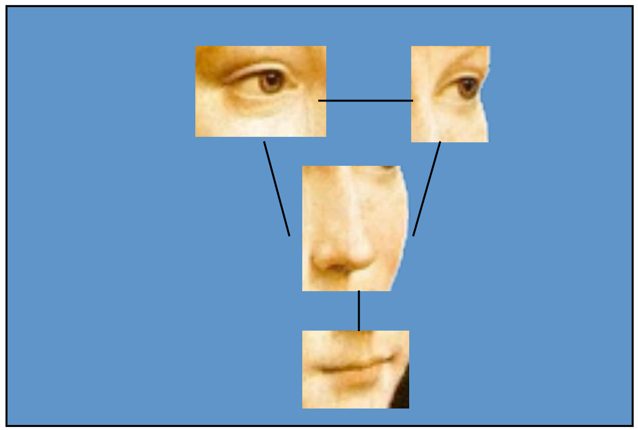

Slide from: Li Fei-Fei, Rob Fergus and Antonio Torralba, short course on object recognition, http://people.csail.mit.edu/torralba/shortCourseRLOC/

Object recognition issues

– Generative / discriminative / hybrid
– Appearance only or location and appearance
– Invariances
  • View point
  • Illumination
  • Occlusion
  • Scale
  • Deformation
  • Clutter
  • etc.

Slide from: Li Fei-Fei, Rob Fergus and Antonio Torralba, short course on object recognition, http://people.csail.mit.edu/torralba/shortCourseRLOC/

Let's go back in time, to the mid-1980's

What everyone looked like back then

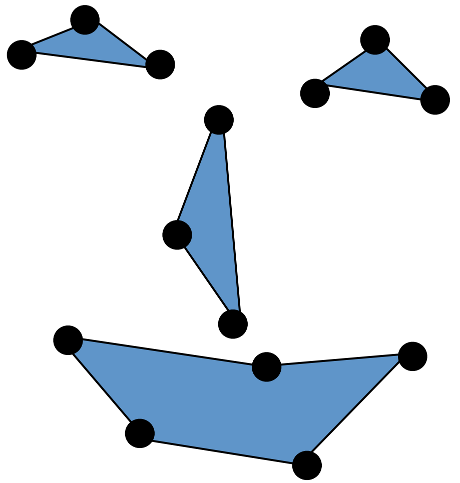

Features

- Points

but also,

- Lines
- Conics
- Other fitted curves

Objects

"blocks world"
A toy world in which to study image interpretation. All we have to do is to convert real world images to their blocks world equivalents and we're all set.

Yvan Leclerc and Martin Fischler, an optimization-based approach to the interpretation of single line drawings as 3-d wire frames.

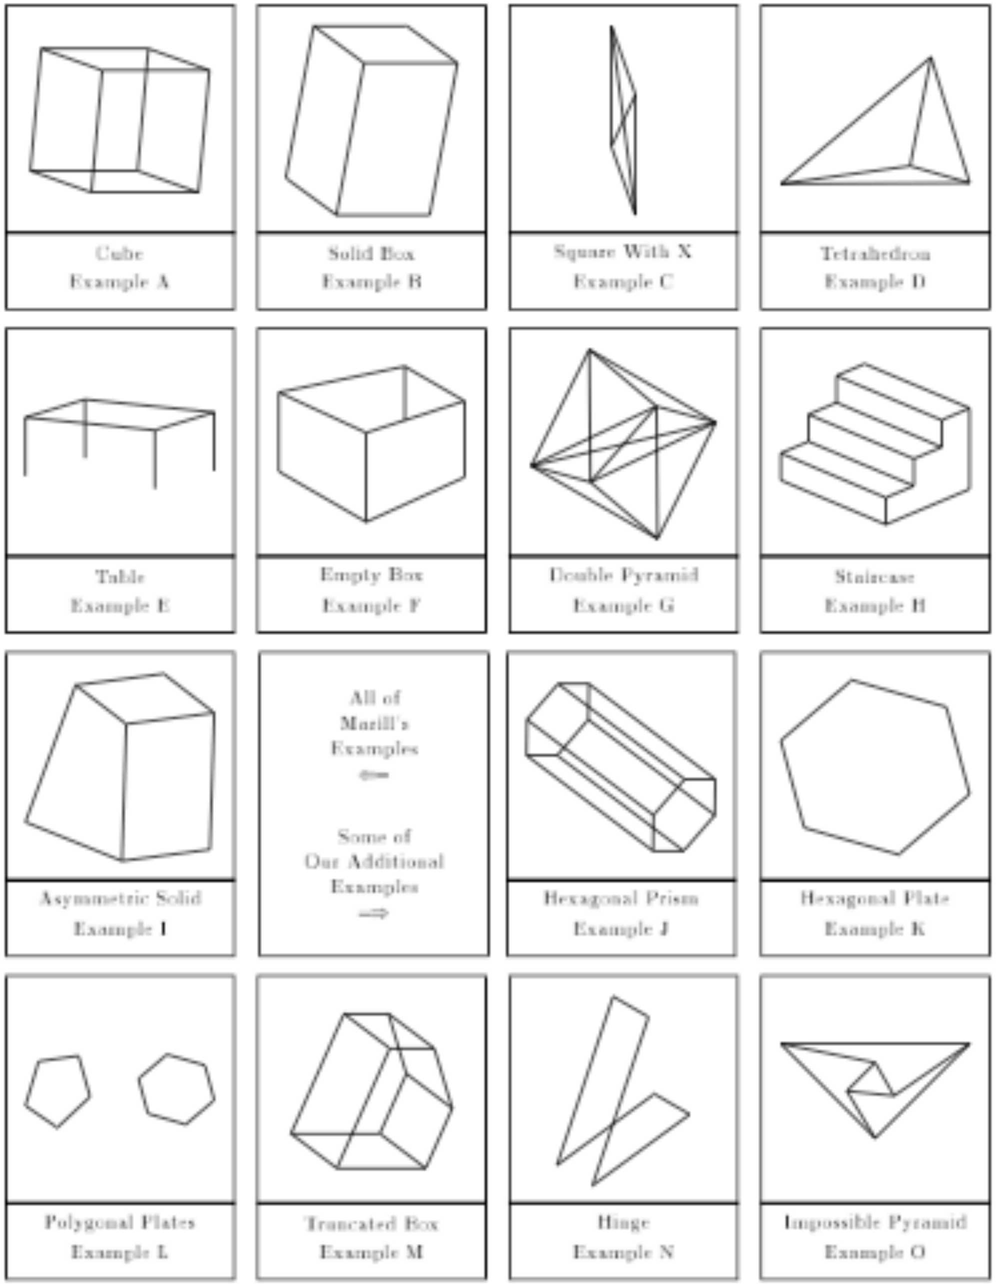

Figure 1: The line drawings examined in this paper. Examples A through I are taken from Marill's paper. Examples J through N are line drawings introduced in this paper for which Marill's algorithm failed to recover a psychologically plausible 3-D model. Example O is a line drawing for which a psychologically plausible 3-D model is not feasible.

Computer vision research results, 1986

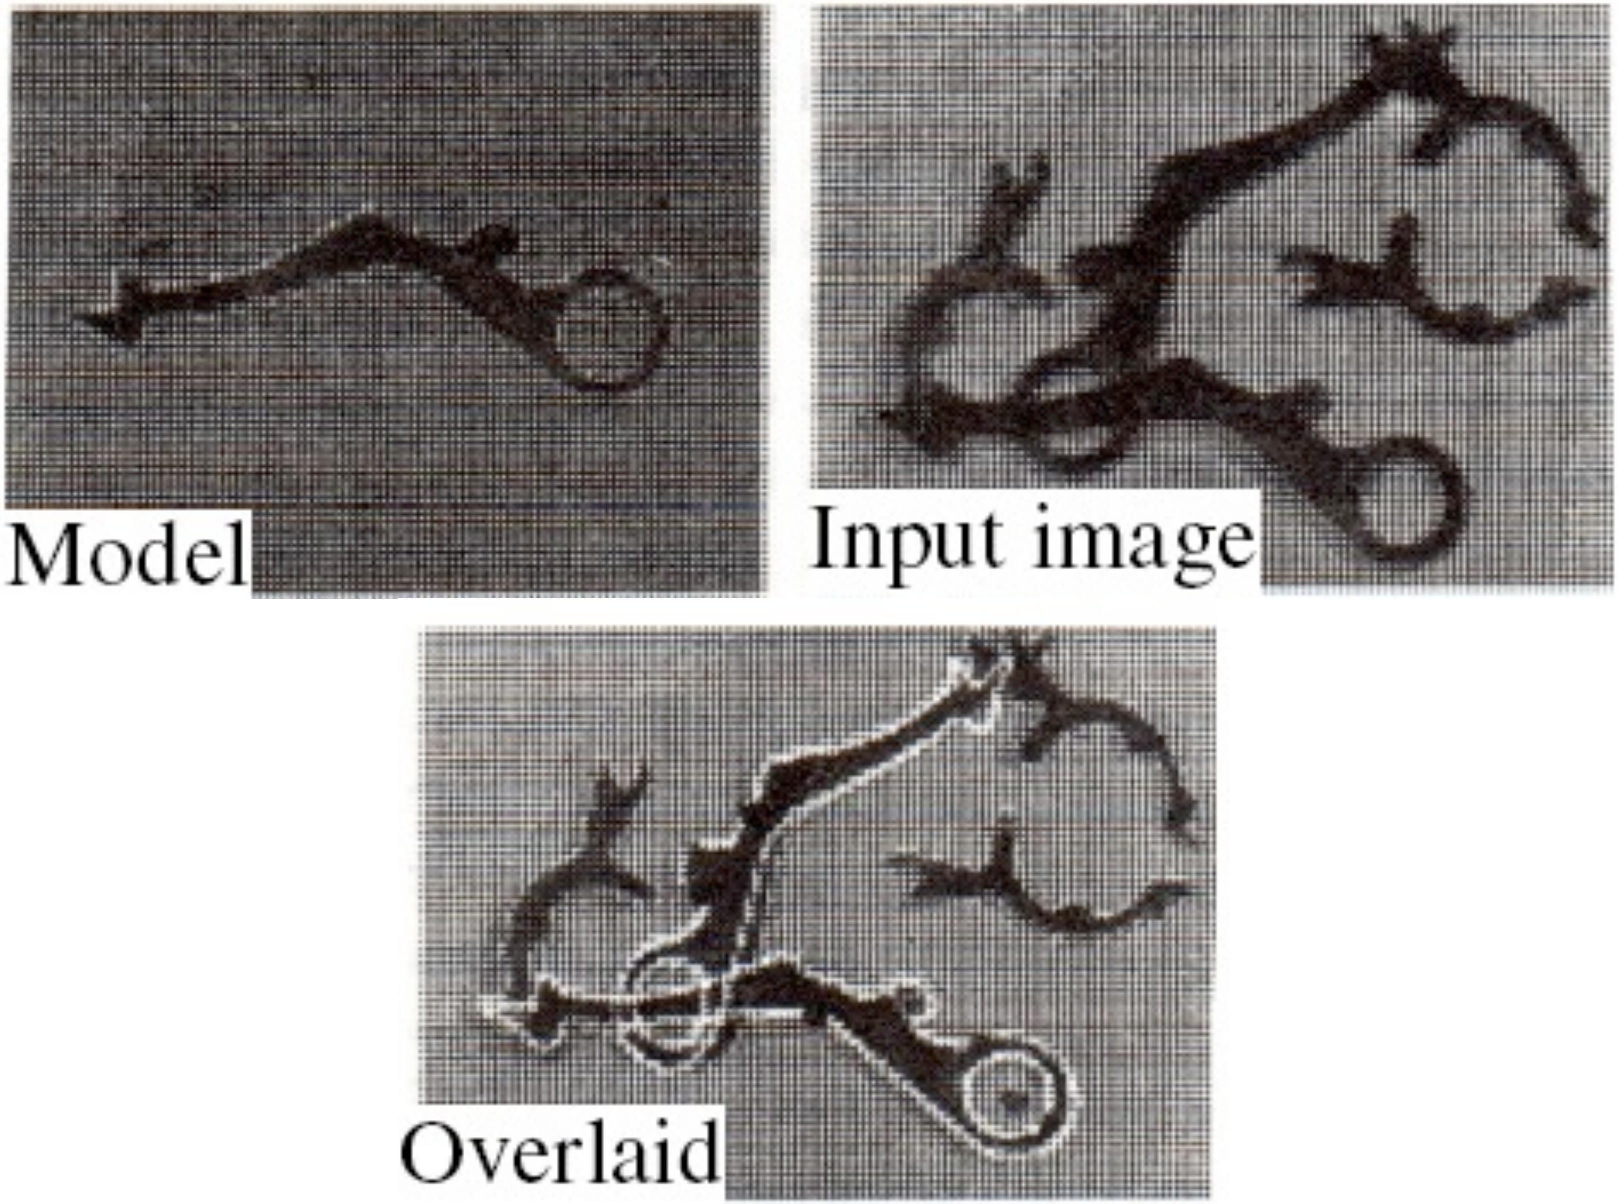

Huttenlocher and Ullman, Object recognition using alignment, ICCV, 1986

Computer vision research results, 1992

Input image

Edge points fitted with lines or conics

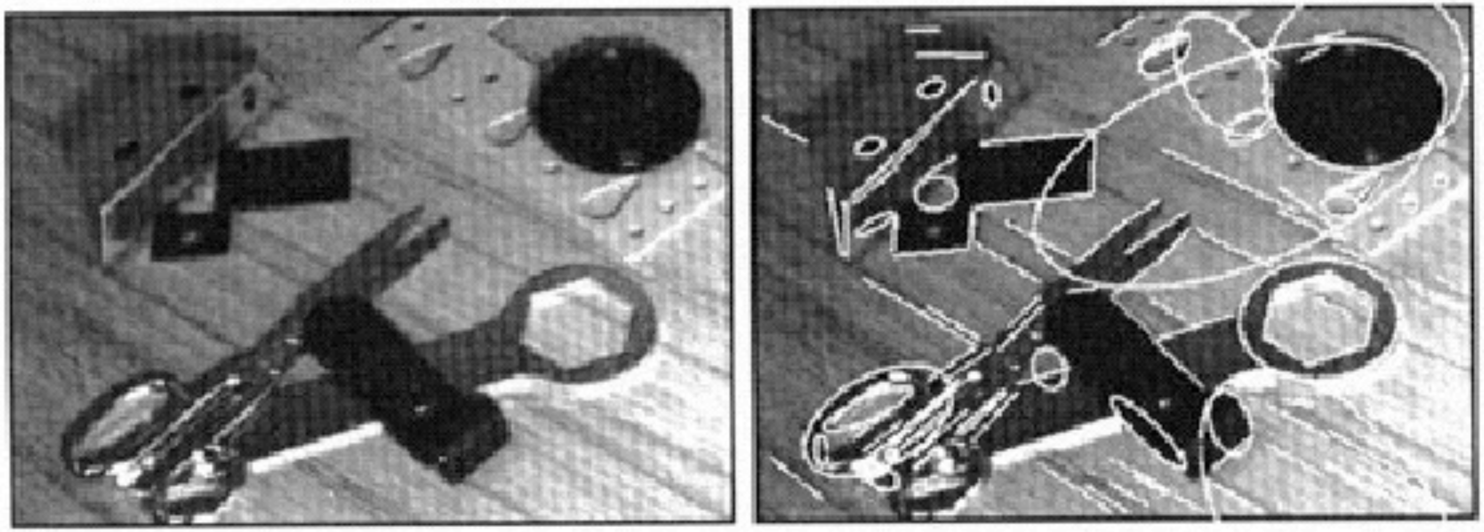

a b

6 years later:
Recognizing planar objects using invariants.

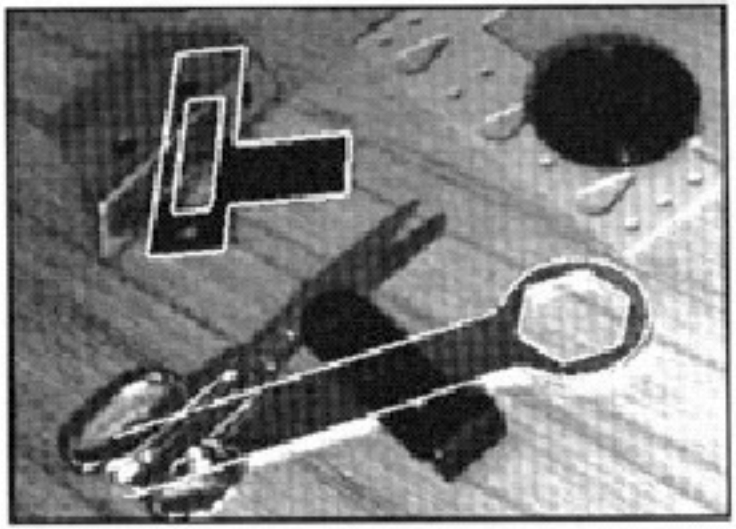

Objects that have been recognized and verified.

From Rothwell et al, Efficient model library access by projectively invariant indexing functions, CVPR 1992.

Back to the present…

What has allowed us to make progress?

- SIFT features
- Discriminative classifiers
- Bayesian methods
- Large databases

What has allowed us to make progress?

- SIFT features
- Discriminative classifiers
- Bayesian methods
- Large databases

CVPR 2003 Tutorial

Recognition and Matching Based on Local Invariant Features

David Lowe

Computer Science Department

University of British Columbia

http://www.cs.ubc.ca/~lowe/papers/ijcv04.pdf

CVPR 2003 Tutorial

Recognition and Matching Based on Local Invariant Features

David Lowe

Computer Science Department

University of British Columbia

http://www.cs.ubc.ca/~lowe/papers/ijcv04.pdf 9,000 citations!

Invariant Local Features

- Image content is transformed into local feature coordinates that are invariant to translation, rotation, scale, and other imaging parameters

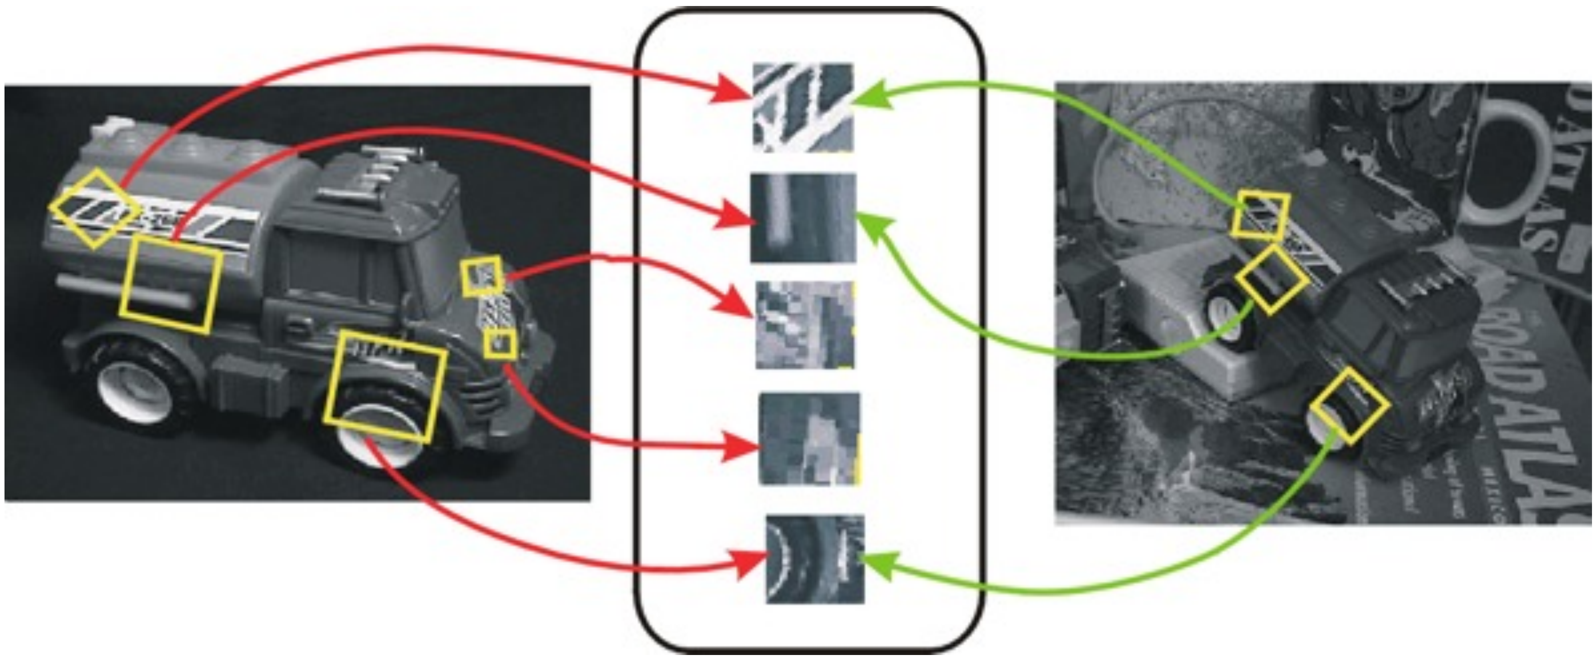

SIFT Features

B (1) and (2) A hand under two different lighting conditions (the pixel intensities vary greatly). (3) and (4) Orientation maps of those images are generally more robust to lighting changes than are the pixel intensities.

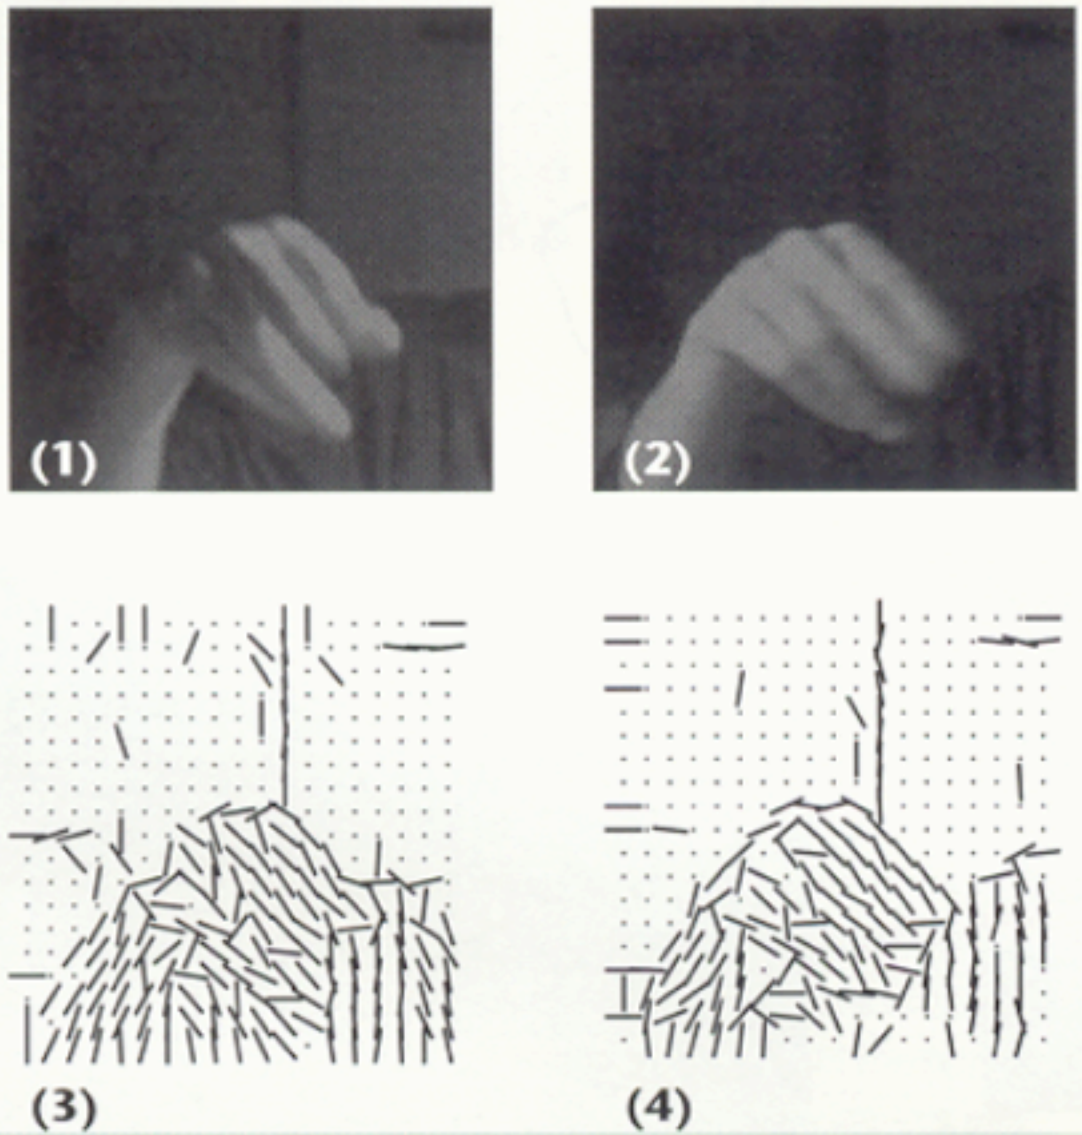

Freeman et al, 1998 http://people.csail.mit.edu/billf/papers/cga1.pdf

SIFT vector formation

- Computed on rotated and scaled version of window according to computed orientation & scale
  - resample a 16x16 version of the window
- Based on gradients weighted by a Gaussian of variance half the window (for smooth falloff)

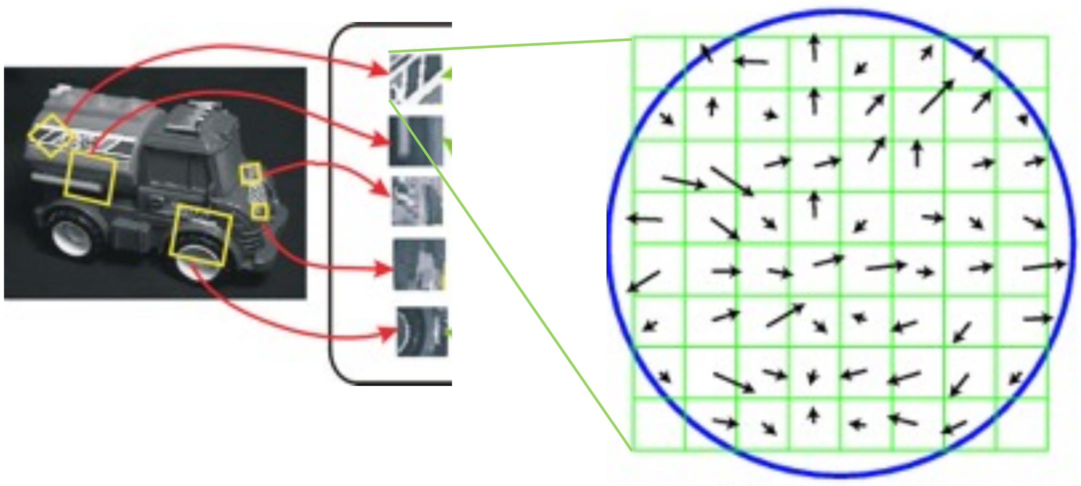

SIFT vector formation

- 4x4 array of gradient orientation histograms
  - not really histogram, weighted by magnitude
- 8 orientations x 4x4 array = 128 dimensions
- Motivation: some sensitivity to spatial layout, but not too much.

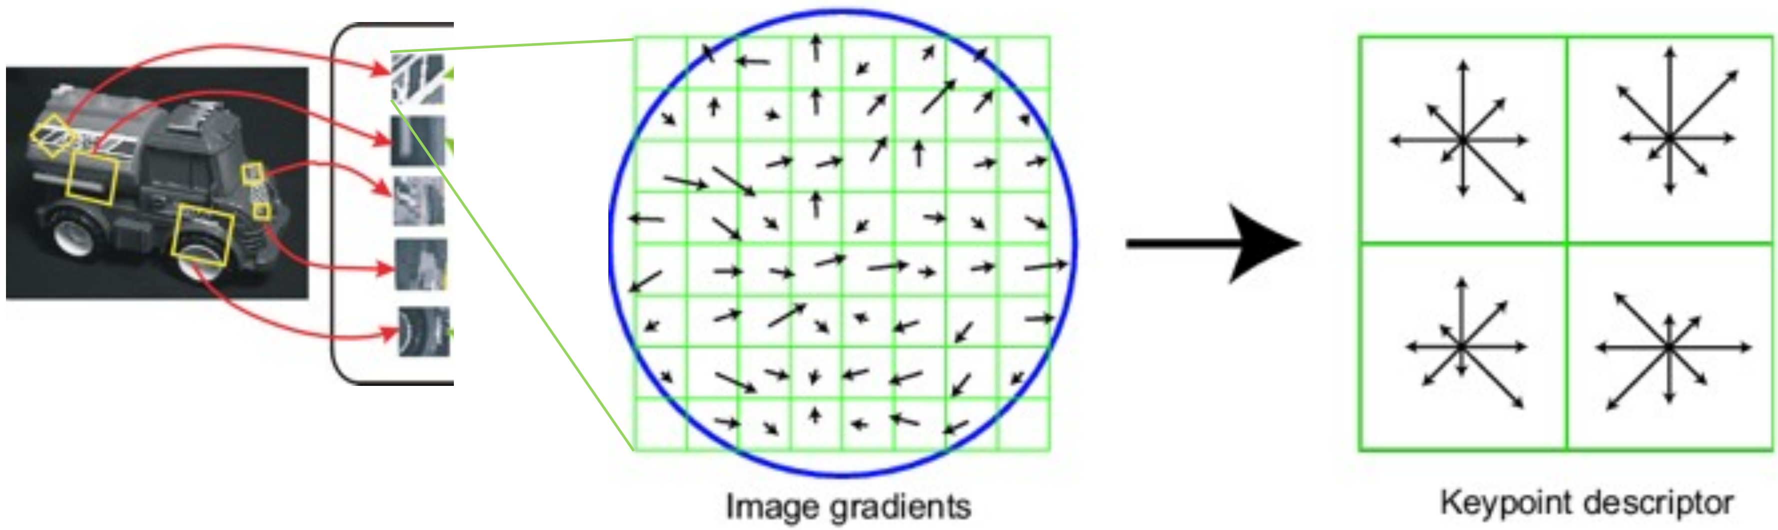

showing only 2x2 here but is 4x4

Feature stability to noise

- Match features after random change in image scale & orientation, with differing levels of image noise
- Find nearest neighbor in database of 30,000 features

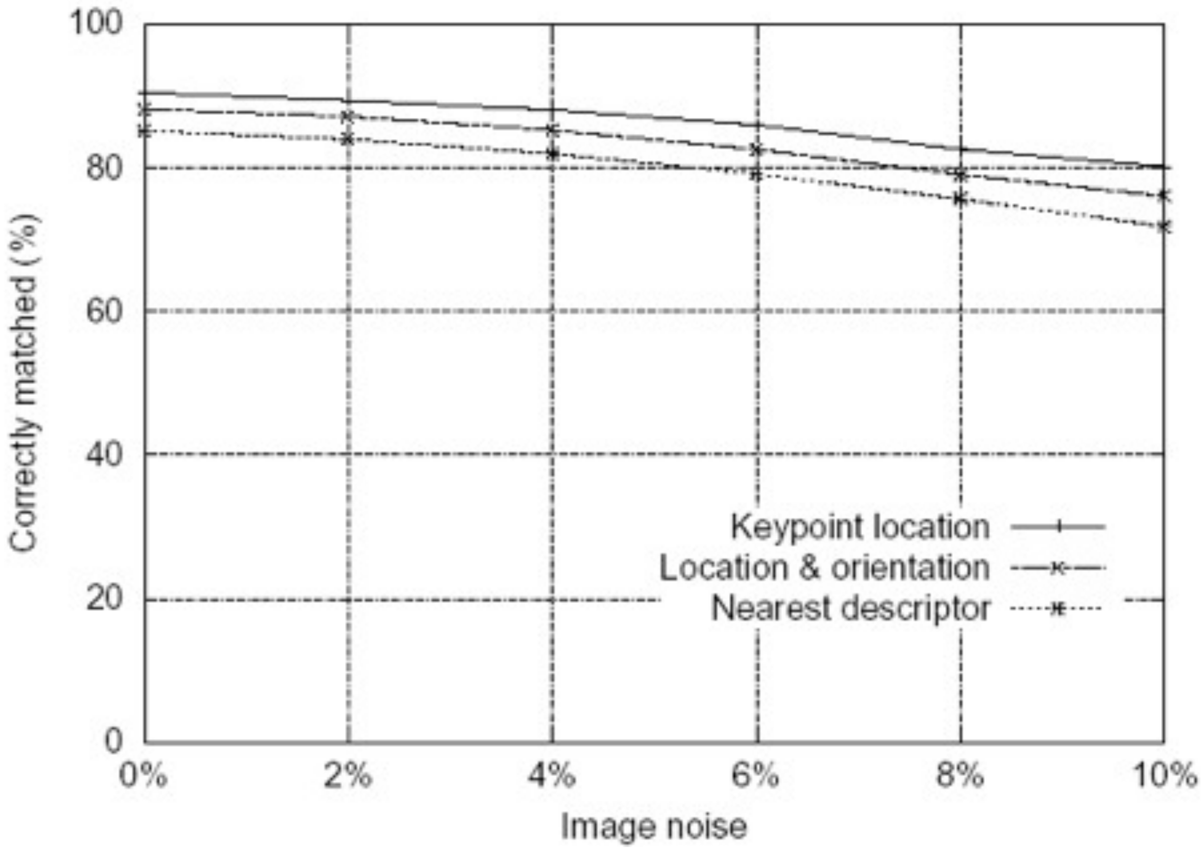

Feature stability to affine change

- Match features after random change in image scale & orientation, with 2% image noise, and affine distortion
- Find nearest neighbor in database of 30,000 features

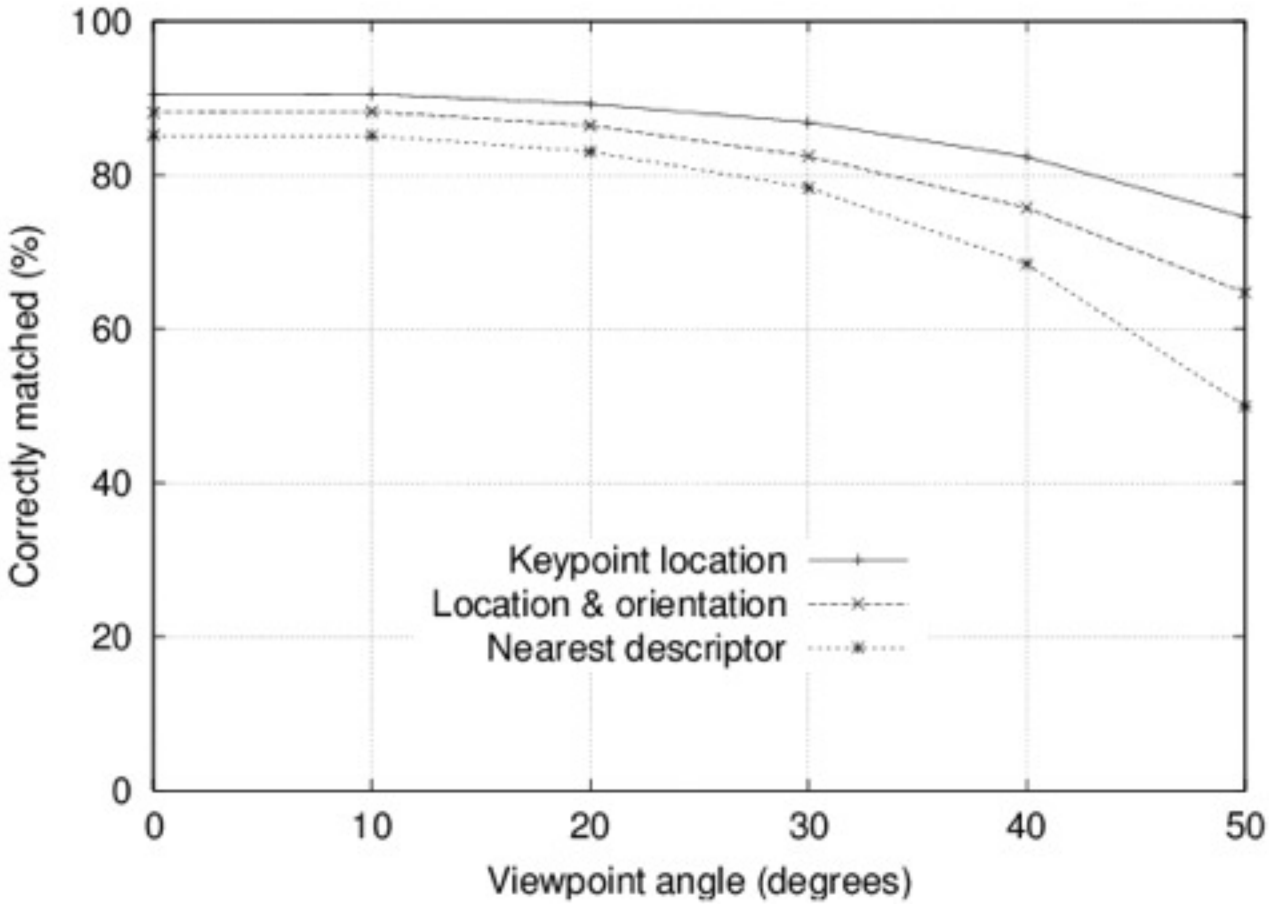

Distinctiveness of features

- Vary size of database of features, with 30 degree affine change, 2% image noise
- Measure % correct for single nearest neighbor match

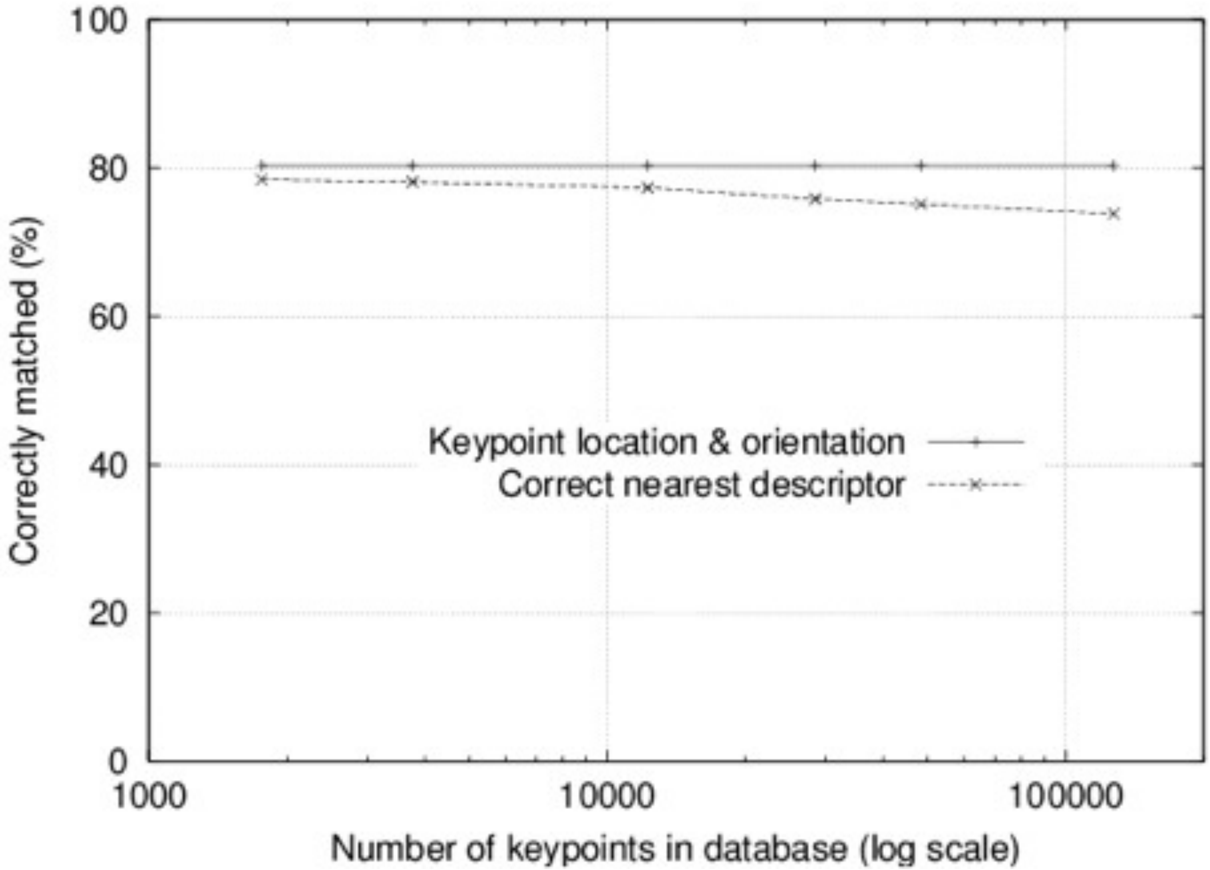

Figure 12: The training images for two objects are shown on the left. These can be recognized in a cluttered image with extensive occlusion, shown in the middle. The results of recognition are shown on the right. A parallelogram is drawn around each recognized object showing the boundaries of the original training image under the affi ne transformation solved for during recognition. Smaller squares indicate the keypoints that were used for recognition.

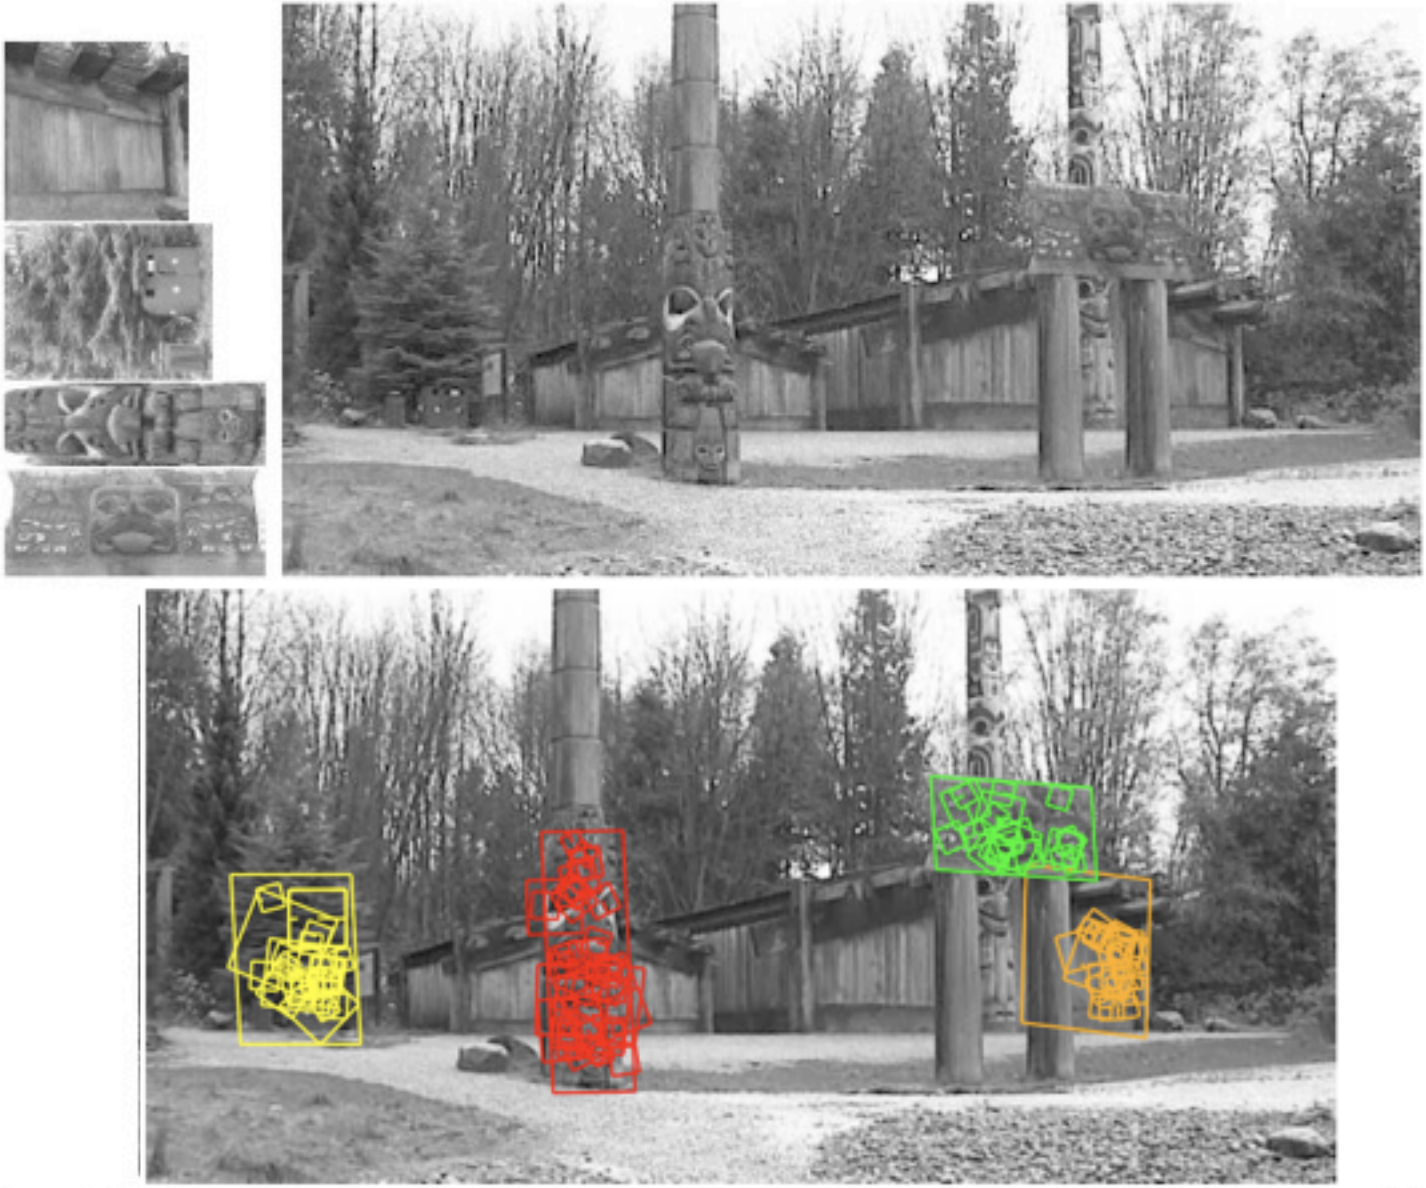

Figure 13: This example shows location recognition within a complex scene. The training images for locations are shown at the upper left and the 640x315 pixel test image taken from a different viewpoint is on the upper right. The recognized regions are shown on the lower image, with keypoints shown as squares and an outer parallelogram showing the boundaries of the training images under the affi ne transform used for recognition.

Building a Panorama

M. Brown and D. G. Lowe. Recognising Panoramas. ICCV 2003

Building a Panorama

M. Brown and D. G. Lowe. Recognising Panoramas. ICCV 2003

These feature point detectors and descriptors are the most important recent advance in computer vision and graphics.

- Feature points are used also for:
  - Image alignment (homography, fundamental matrix)
  - 3D reconstruction
  - Motion tracking
  - Object recognition
  - Indexing and database retrieval
  - Robot navigation
  - … other

More uses for SIFT features

SIFT features have also been applied to (categorical) object recognition

Extracting Words

- Find interest points using shape adapted (white) and maximally stable (yellow) regions

- Map ellipses to a circle

- Compute SIFT descriptor over circle

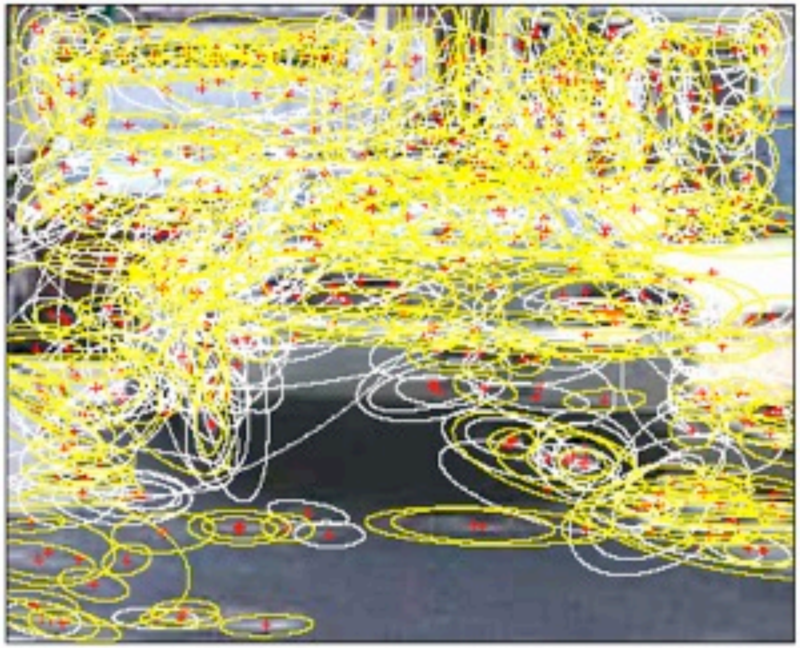

Visual words

- Vector quantize SIFT descriptors to a vocabulary of 2 or 3 thousand “visual words”.
- Heuristic design of descriptors makes these words somewhat invariant to:
  – Lighting
  – 2-d Orientation
  – 3-d Viewpoint

Object recognition using visual words

Find words

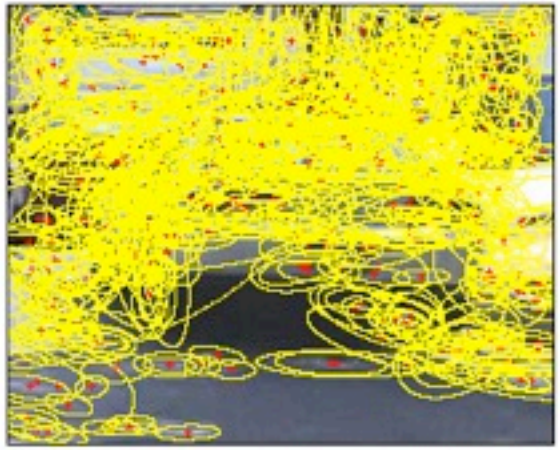

Form histograms

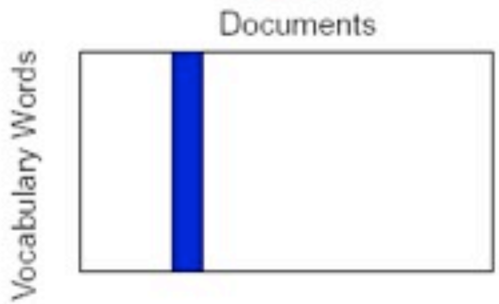

Compare with object class database

Now this starts to look like a discrete algorithms problem

m feature words in test image region, **n** possible matching features in a training database for each of **k** possible object classes. The feature word collections will have different sizes, and matching will be noisy.

Find the most probable object class.

My poll of the top researchers in computer vision (pictured here: participants in the BIRS Workshop on Computer Vision and the Internet)

“How do you think computer science can best help computer vision?”

“How do you think computer science can best help computer vision?”

Modal response:

“How do you think computer science can best help computer vision?”

Modal response:

“Fast, approximate nearest neighbor search in high dimensions.”

Nearest neighbor search in high dimensions

"Nearest neighbor search, but taking into account our particular data. or, tell us what questions we should be asking about our data in order to do nearest neighbor search well."

"Parallelism--where can we exploit it? kd tree high d search. Does LSH work as advertised? in practice not as well."

"Nearest neighbors in high-dimensions. category recognition. for instance recognition, nn for individual features works fine. but for category recognition, many times the local features are not, by themselves, a close match, due to within-class variations."

FAST APPROXIMATE NEAREST NEIGHBORS WITH AUTOMATIC ALGORITHM CONFIGURATION

Marius Muja, David G. Lowe
Computer Science Department, University of British Columbia, Vancouver, B.C., Canada
mariusm@cs.ubc.ca, lowe@cs.ubc.ca

International Conference on Computer Vision Theory and Applications (VISAPP), Lisbon, Portugal (Feb 2009)

Comparison of different algorithms

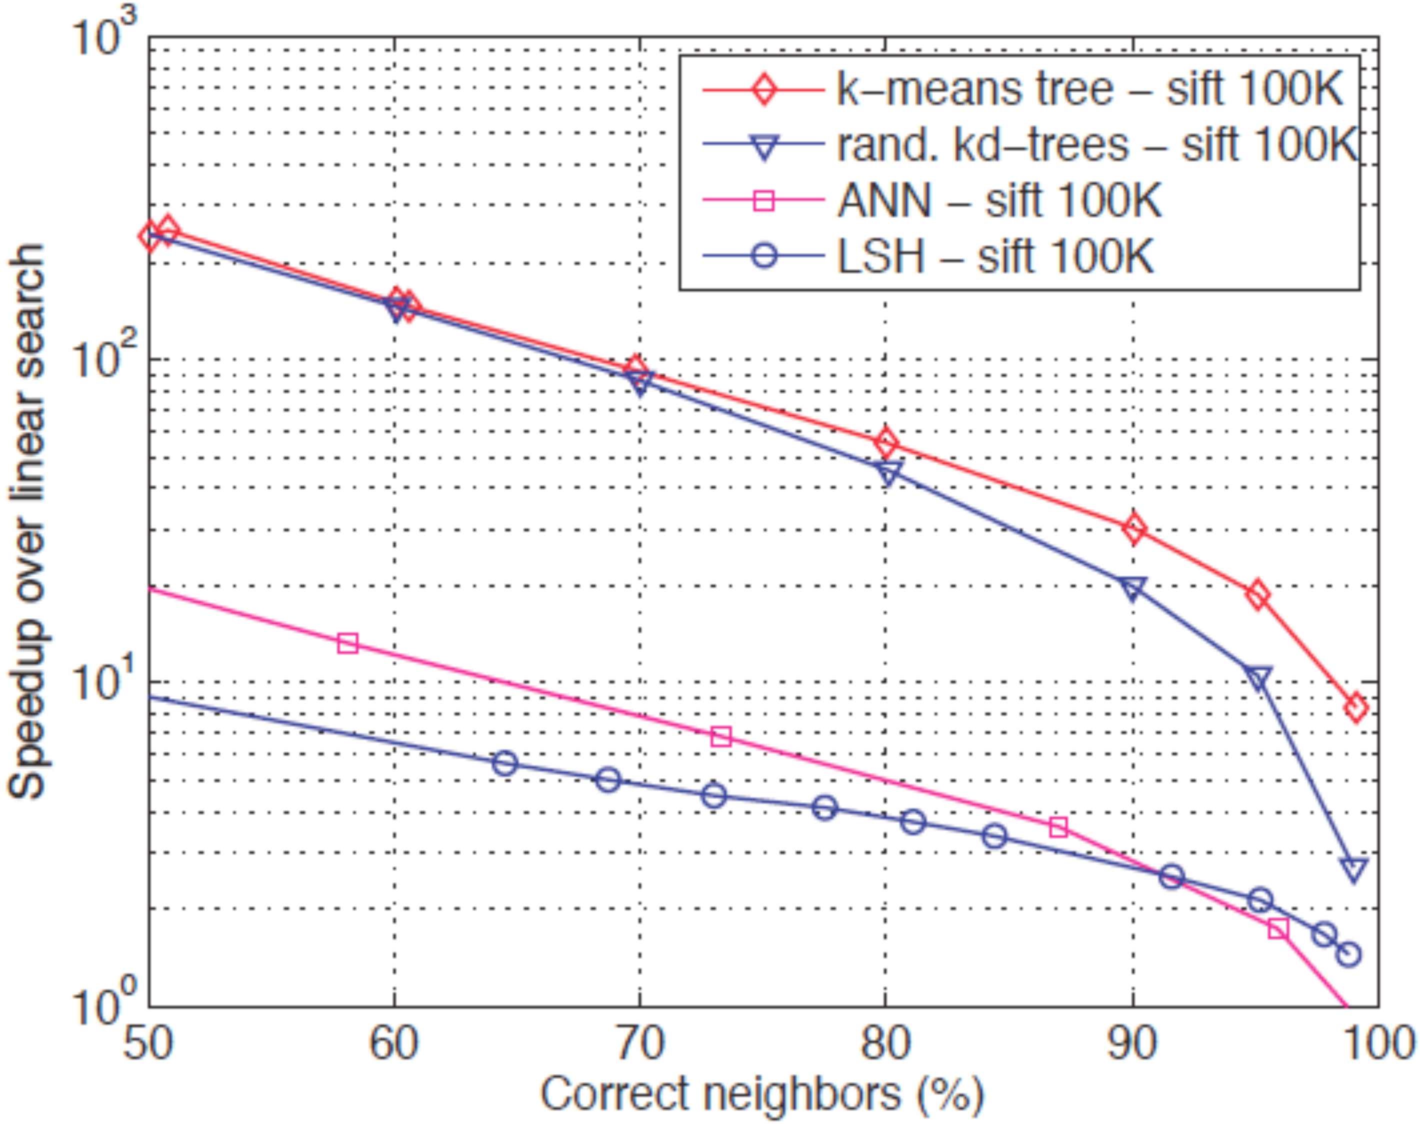

Additional structure present in NN problems for computer vision

m feature words in test image region, n possible matching features, in each of k possible object classes. The feature word collections will have different sizes, and matching will be noisy.

How can we quickly identify the most probable object categories? How handle feature variations in a principled way? How take feature positions into account efficiently?

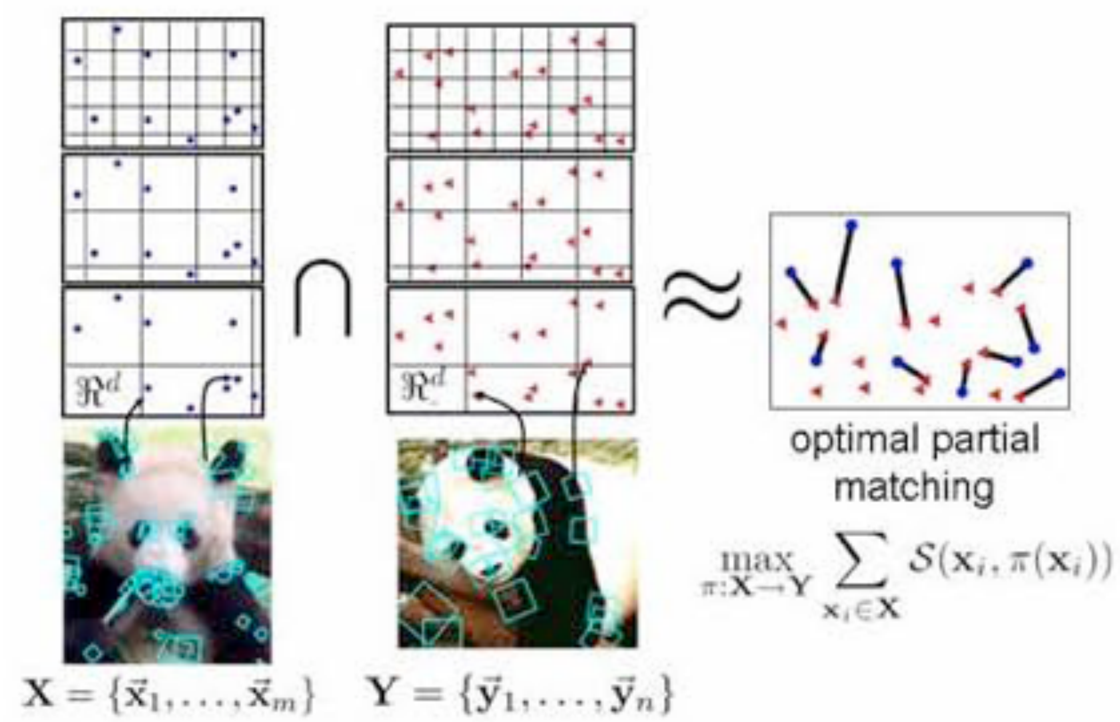

http://www.cs.utexas.edu/~grauman/research/projects/pmk/pmk_projectpage.htm

Another NN search problem, with structure: non-local means denoising

Image and movie denoising by nonlocal means

Antoni Buades Bartomeu Coll Jean-Michel Morel

Abstract

Neighborhood filters are image and movie filters which reduce the noise by averaging similar pixels. The object of the paper is to present a unified theory of these filters and reliable criteria to compare them to other classes of filters. First a CCD noise model will be presented justifying this class of algorithm. A classification of neighborhood filters is proposed, including classical image and movie denoising methods and a new one, the nonlocal-means (NL-means). In order to compare denoising methods three principles

International Journal of Computer Vision (IJCV) 76 (2008) 123–139

Non-local means denoising algorithm

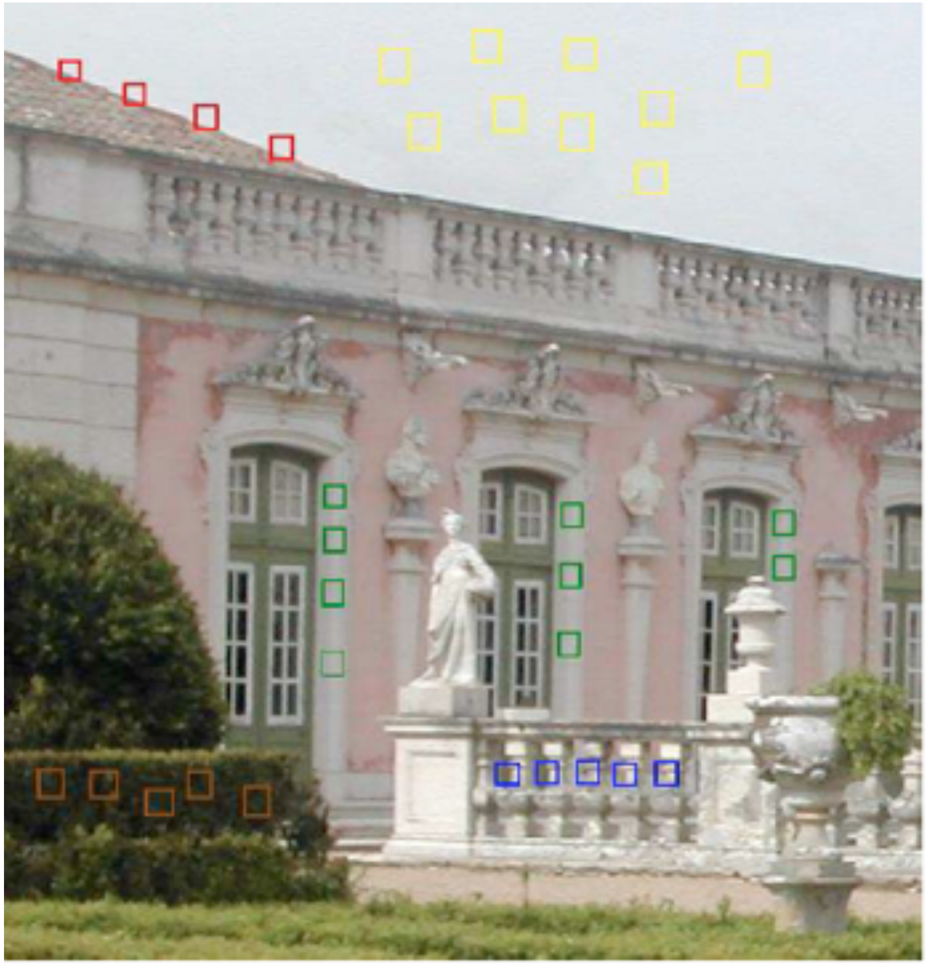

International Journal of Computer Vision (IJCV) 76 (2008) 123–139

An approx nearest-neighbor algo. that takes image spatial structure into account

- The complexity of exhaustive KNN search is O(NHM2log(K)) for each frome (N: number of pixels, H: temporal window size, M: spatial window size, K: the number of nearest neighbors).
- A solution: randomized search [Barnes et al. Siggraph '09]

- The complexity is O(NHKlog(K))<<O(NHM2log(K)) (typically M=40, K=7)

An approx nearest-neighbor algo. that takes image spatial structure into account

- The complexity of exhaustive KNN search is O(NHM2log(K)) for each frome (N: number of pixels, H: temporal window size, M: spatial window size, K: the number of nearest neighbors).
- A solution: randomized search [Barnes et al. Siggraph '09]

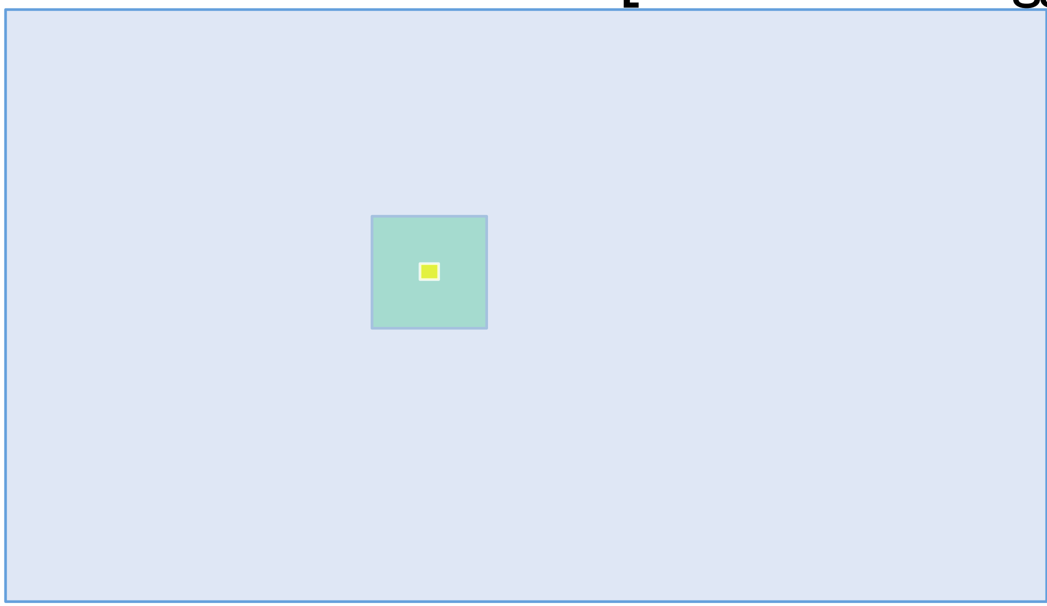

- The complexity is O(NHKlog(K))<<O(NHM2log(K)) (typically M=40, K=7)

An approx nearest-neighbor algo. that takes image spatial structure into account

- The complexity of exhaustive KNN search is O(NHM2log(K)) for each frome (N: number of pixels, H: temporal window size, M: spatial window size, K: the number of nearest neighbors).
- A solution: randomized search [Barnes et al. Siggraph '09]

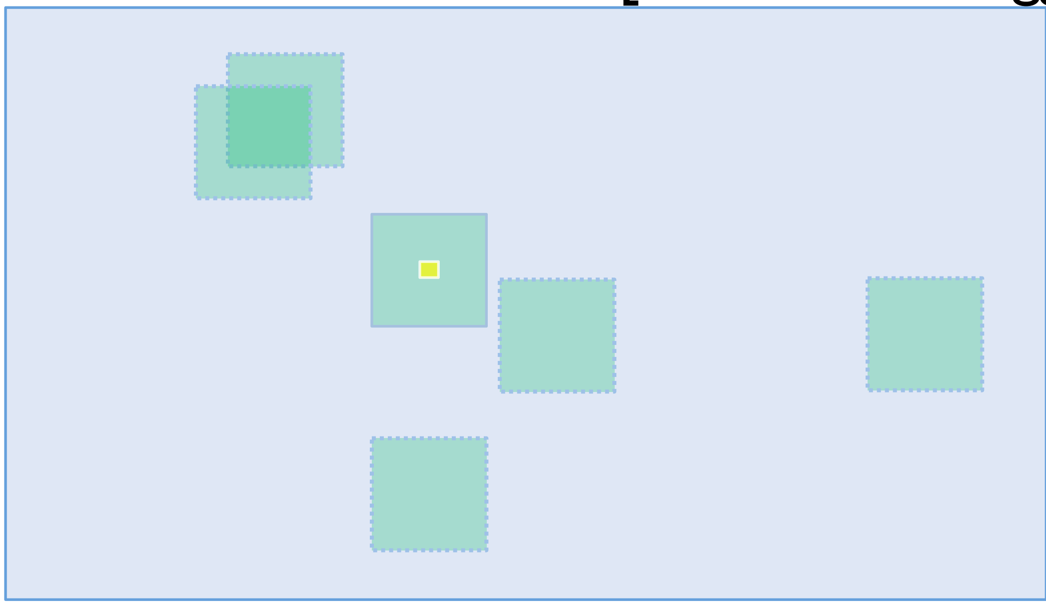

- The complexity is O(NHKlog(K))<<O(NHM2log(K)) (typically M=40, K=7)

An approx nearest-neighbor algo. that takes image spatial structure into account

- The complexity of exhaustive KNN search is O(NHM2log(K)) for each frome (N: number of pixels, H: temporal window size, M: spatial window size, K: the number of nearest neighbors).
- A solution: randomized search [Barnes et al. Siggraph '09]

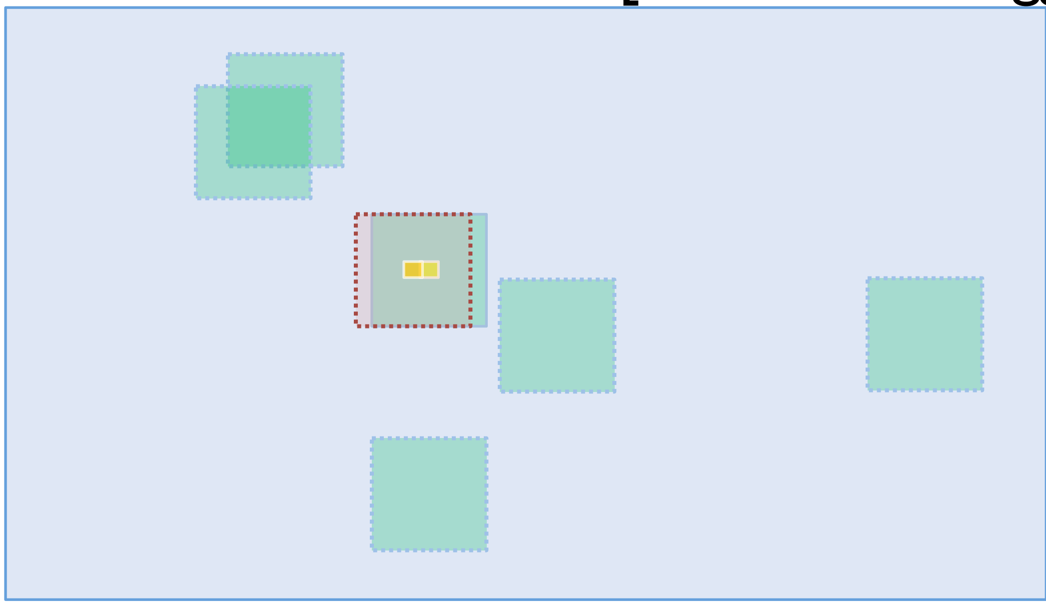

- The complexity is O(NHKlog(K))<<O(NHM2log(K)) (typically M=40, K=7)

An approx nearest-neighbor algo. that takes image spatial structure into account

- The complexity of exhaustive KNN search is O(NHM2log(K)) for each frome (N: number of pixels, H: temporal window size, M: spatial window size, K: the number of nearest neighbors).
- A solution: randomized search [Barnes et al. Siggraph '09]

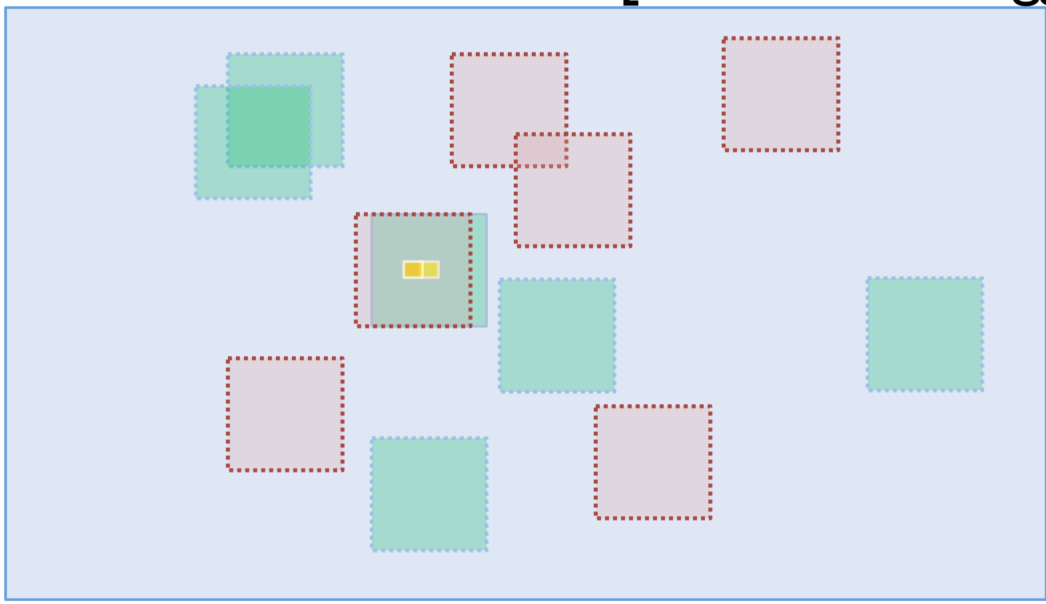

- The complexity is O(NHKlog(K))<<O(NHM2log(K)) (typically M=40, K=7)

An approx nearest-neighbor algo. that takes image spatial structure into account

- The complexity of exhaustive KNN search is O(NHM2log(K)) for each frome (N: number of pixels, H: temporal window size, M: spatial window size, K: the number of nearest neighbors).
- A solution: randomized search [Barnes et al. Siggraph '09]

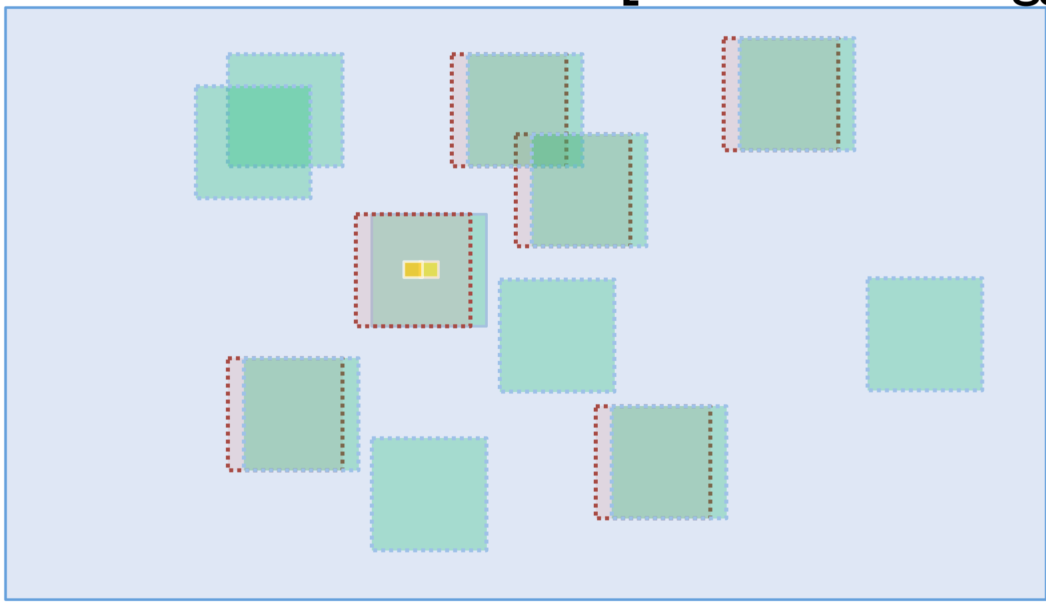

- The complexity is O(NHKlog(K))<<O(NHM2log(K)) (typically M=40, K=7)

An approx nearest-neighbor algo. that takes image spatial structure into account

• The complexity of exhaustive KNN search is O(NHM2log(K)) for each frome (N: number of pixels, H: temporal window size, M: spatial window size, K: the number of nearest neighbors).

• A solution: randomized search [Barnes et al. Siggraph '09]

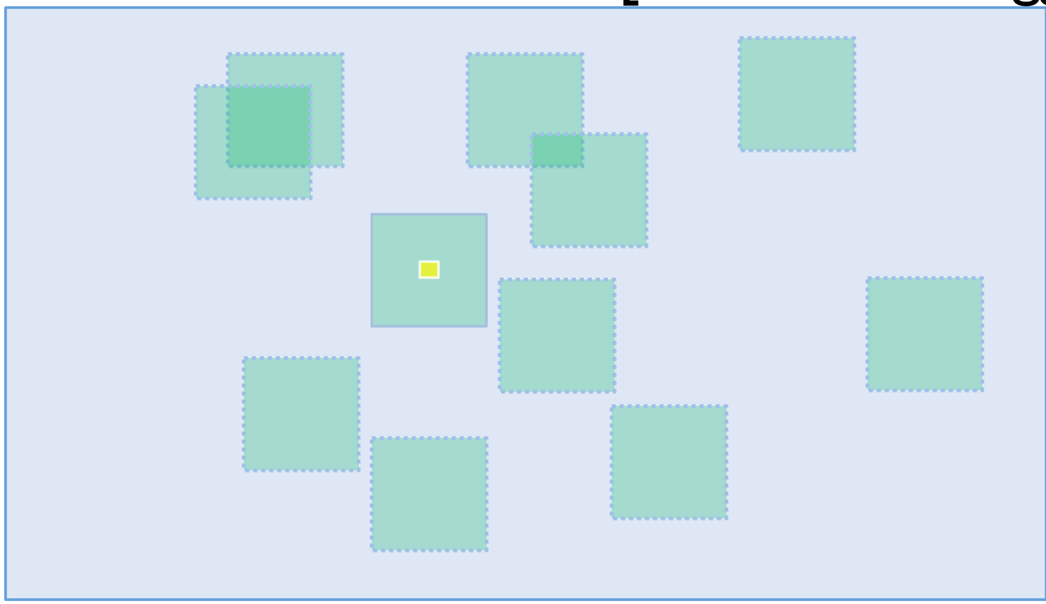

• The complexity is O(NHKlog(K))<<O(NHM2log(K)) (typically M=40, K=7)

An approx nearest-neighbor algo. that takes image spatial structure into account

• The complexity of exhaustive KNN search is O(NHM2log(K)) for each frome (N: number of pixels, H: temporal window size, M: spatial window size, K: the number of nearest neighbors).

• A solution: randomized search [Barnes et al. Siggraph '09]

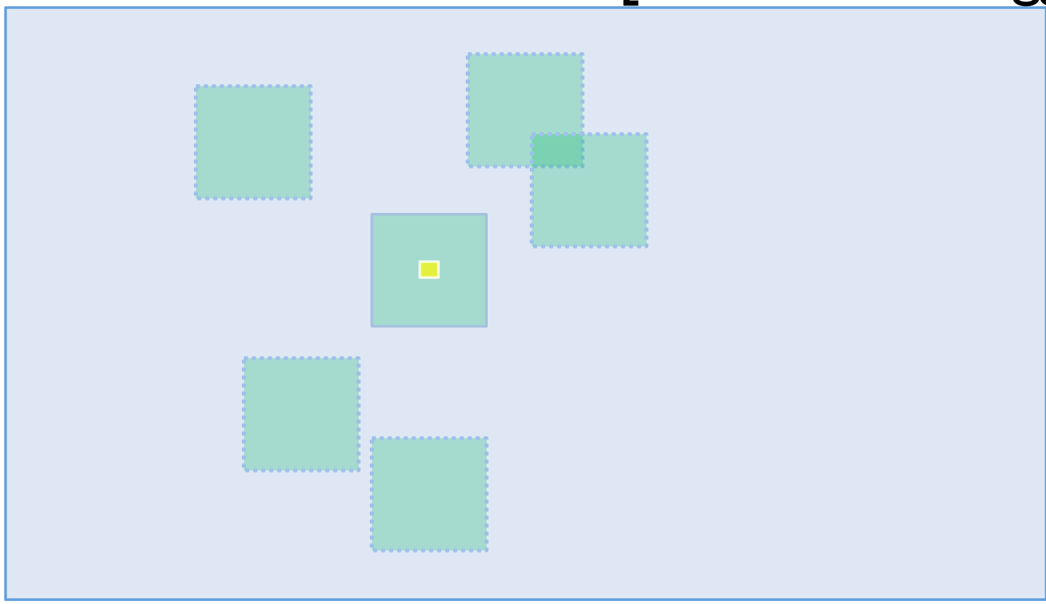

• The complexity is O(NHKlog(K))<<O(NHM2log(K)) (typically M=40, K=7)

An approx nearest-neighbor algo. that takes image spatial structure into account

• The complexity of exhaustive KNN search is O(NHM2log(K)) for each frome (N: number of pixels, H: temporal window size, M: spatial window size, K: the number of nearest neighbors).

• A solution: randomized search [Barnes et al. Siggraph '09]

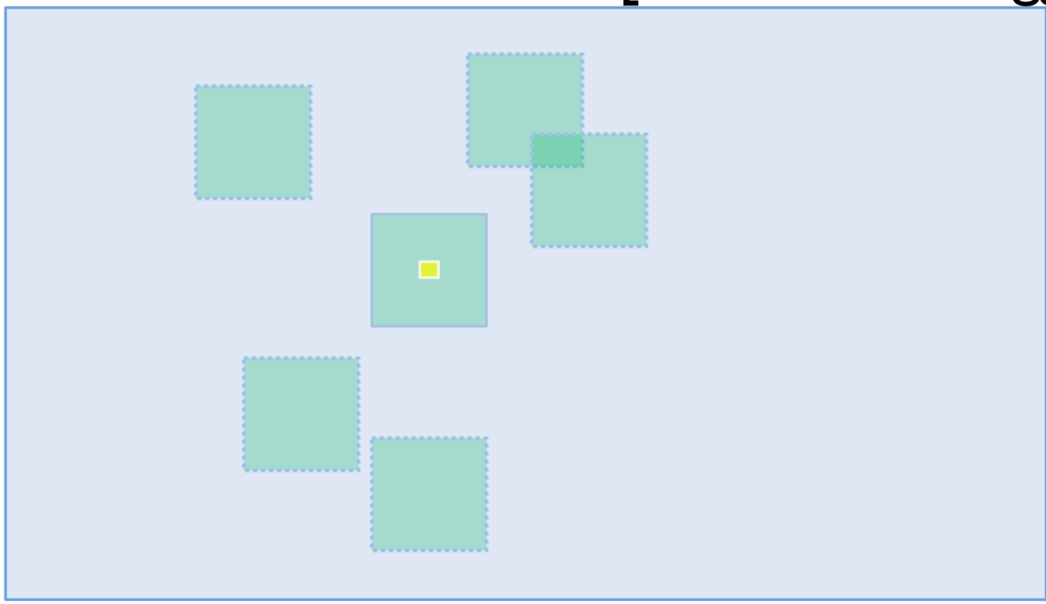

• The complexity is O(NHKlog(K))<<O(NHM2log(K)) (typically M=40, K=7)

This "patch match" algorithm is used in Photoshop's wildly successful "content aware delete" feature in release CS5 in 2010. Demo...

Problem:

Nearest neighbor search in high dimensions.

Applications:

Non-parametric texture synthesis and super-resolution. Image filling-in. Object recognition. Scene recognition.

References:

Fast approximate nearest neighbors with automatic algorithm configuration, Muja and Lowe, VISAPP 2009, http://www.cs.ubc.ca/~lowe/papers/09muja.pdf

PatchMatch: A Randomized Correspondence Algorithm for Structural Image Editing
ACM Transactions on Graphics (Proc. SIGGRAPH), August 2009
Connelly Barnes, Eli Shechtman, Adam Finkelstein,
Dan B Goldman, http://www.cs.princeton.edu/gfx/pubs/Barnes_2009_PAR/patchmatch.pdf

Another commonly expressed need: help in scaling up algorithms

Many vision problems lead to integer programs, linear programs, quadratic programs, and semi-definite programs for large amounts of high-dimensional data.

The standard solvers don't work and we need special purpose solvers that exploit the sparsity or structure of the problem, or develop online versions of the algorithms.

For multi-class categorization, want to recognize 1000's of object categories.
The categories live in taxonomies, and we want to exploit that structure.
The unlabeled training set can be huge.
Seek to generalize from the few labeled training examples.

Problem:

Scale up integer, linear, quadratic, semi-definite programs, exploiting sparsity or other structural characteristics.

Applications:

Vision problems with large-scale training sets.

References:

Pushmeet Kohli, Lubor Ladicky, Philip Torr
Robust Higher Order Potentials for Enforcing Label Consistency.
In: *IJCV 2009*.

Outline

- Computer vision applications
- Computer vision techniques and problems:
  - High-level vision: combinatorial problems
  - Low-level vision: underdetermined problems
  - Miscellaneous problems

Priors on images

original

with additive noise

Noise removed, using simple prior
http://www.cns.nyu.edu/pub/lcv/rajashekar08a.pdf

Noise removed, assuming
more complex prior model

Removing camera shake

Original

Removing camera shake

Original

Our algorithm

Close-up

Original

Close-up

Original

Naïve Sharpening

Close-up

Original

Naïve Sharpening

Our algorithm

Image formation process

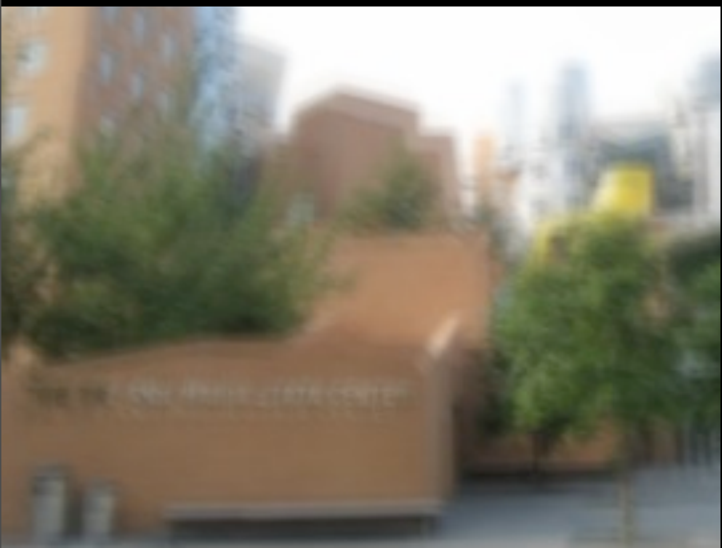

=

⊗

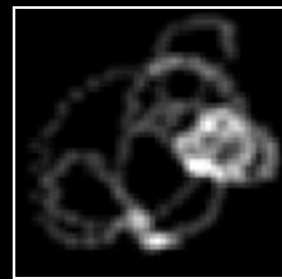

Blur kernel

Blurry image

Input to algorithm

Sharp image

Desired output

Convolution operator

Multiple possible solutions

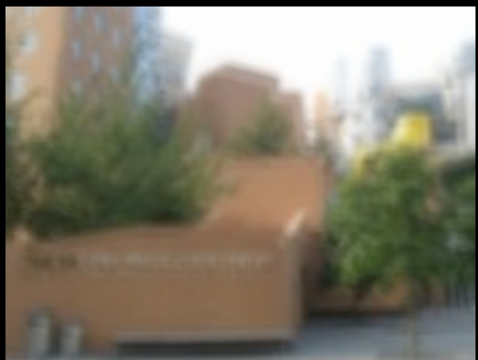

Blurry image

Multiple possible solutions

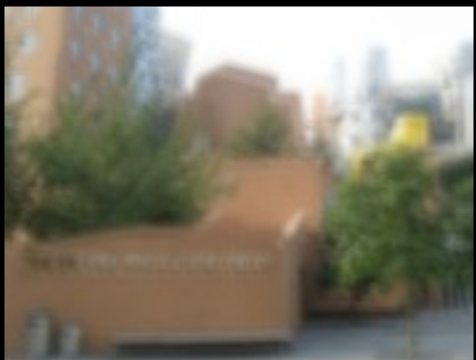

Blurry image

$=$

Sharp image

$\otimes$

Blur kernel

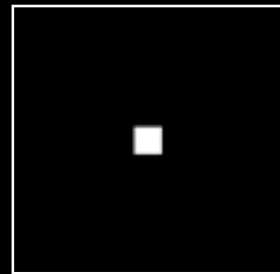

Multiple possible solutions

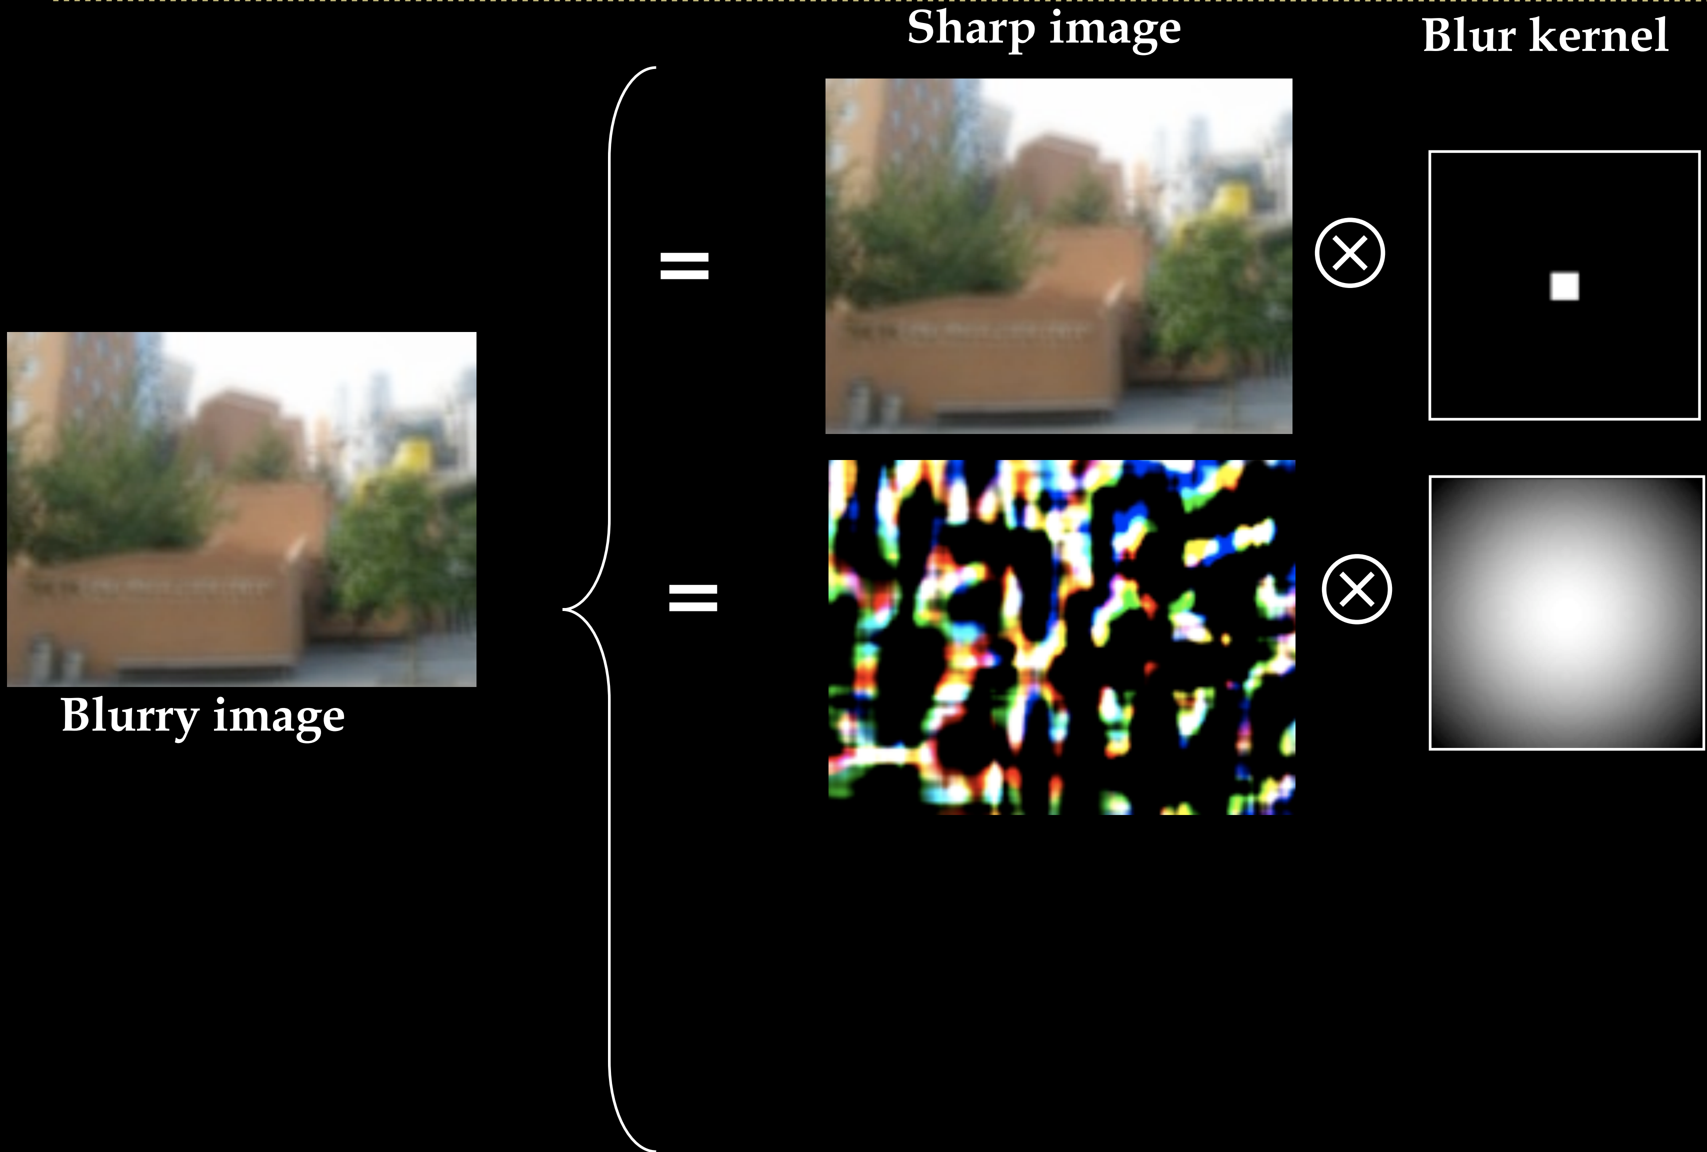

Multiple possible solutions

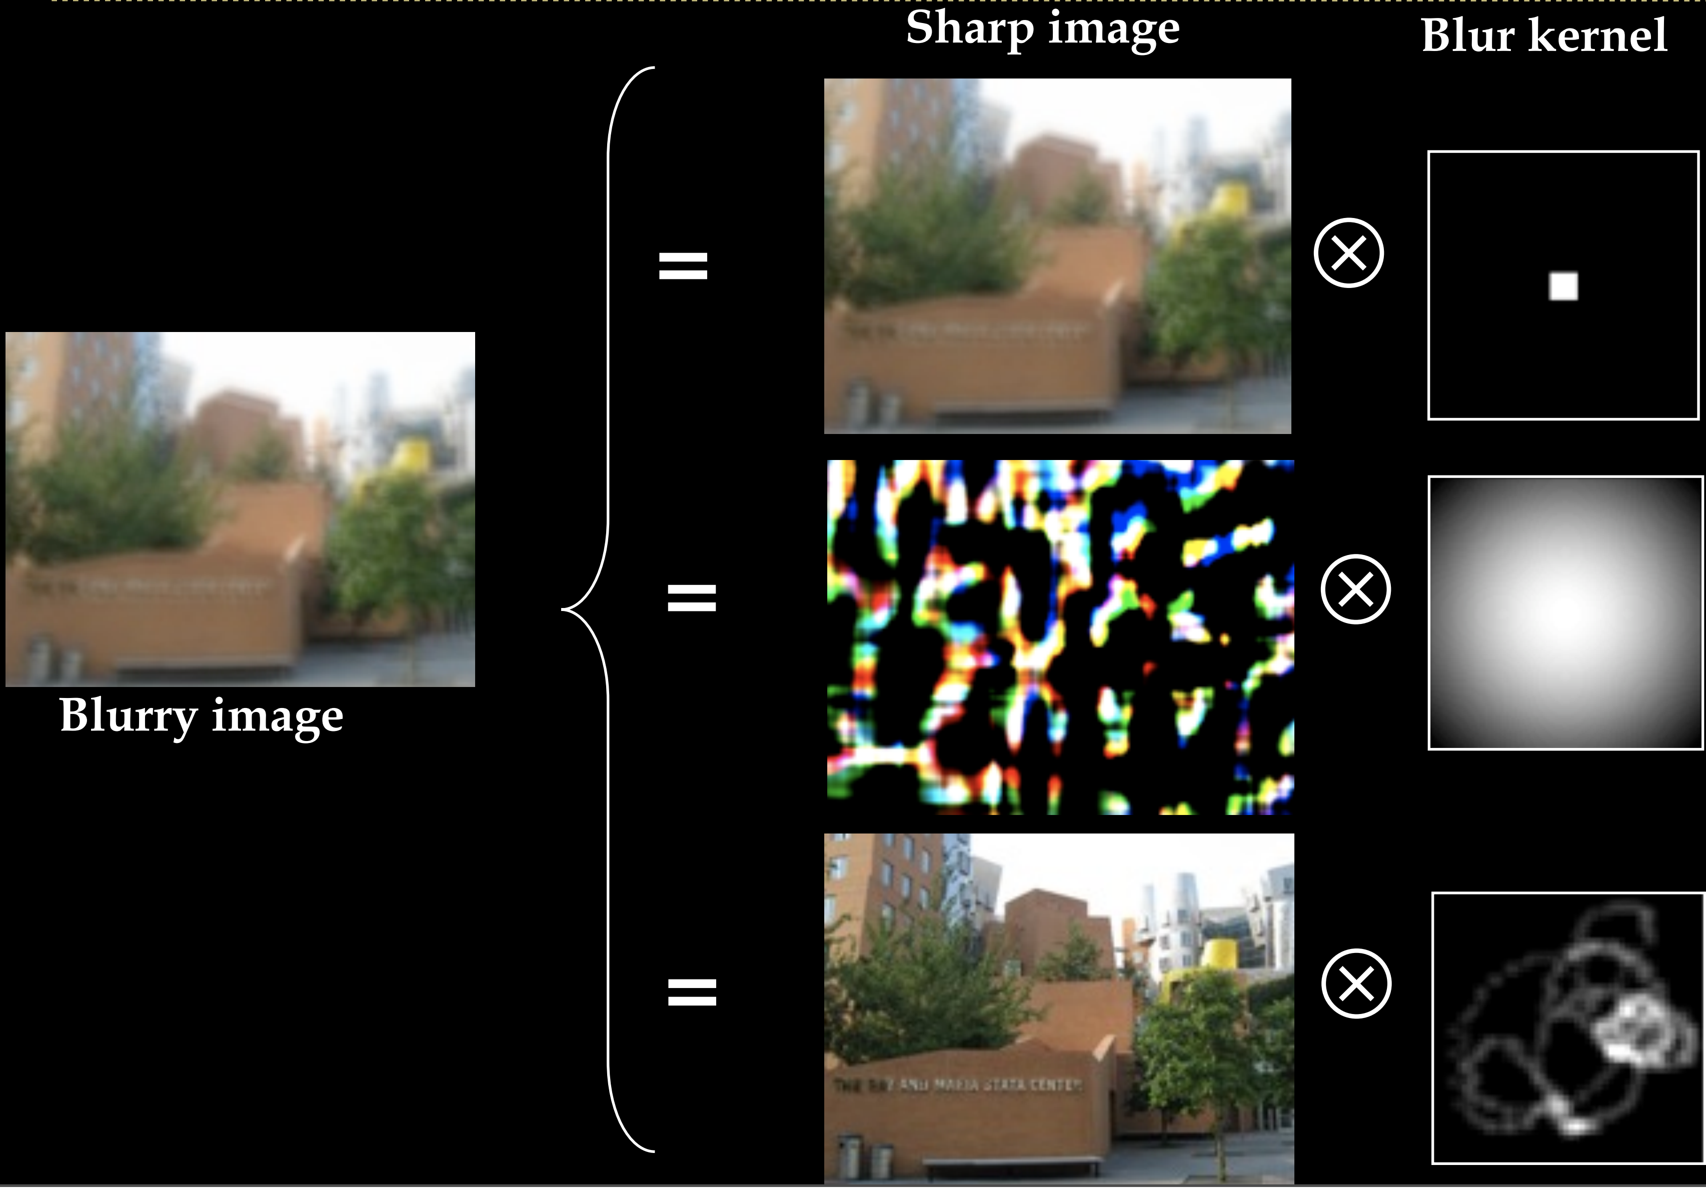

Multiple possible solutions

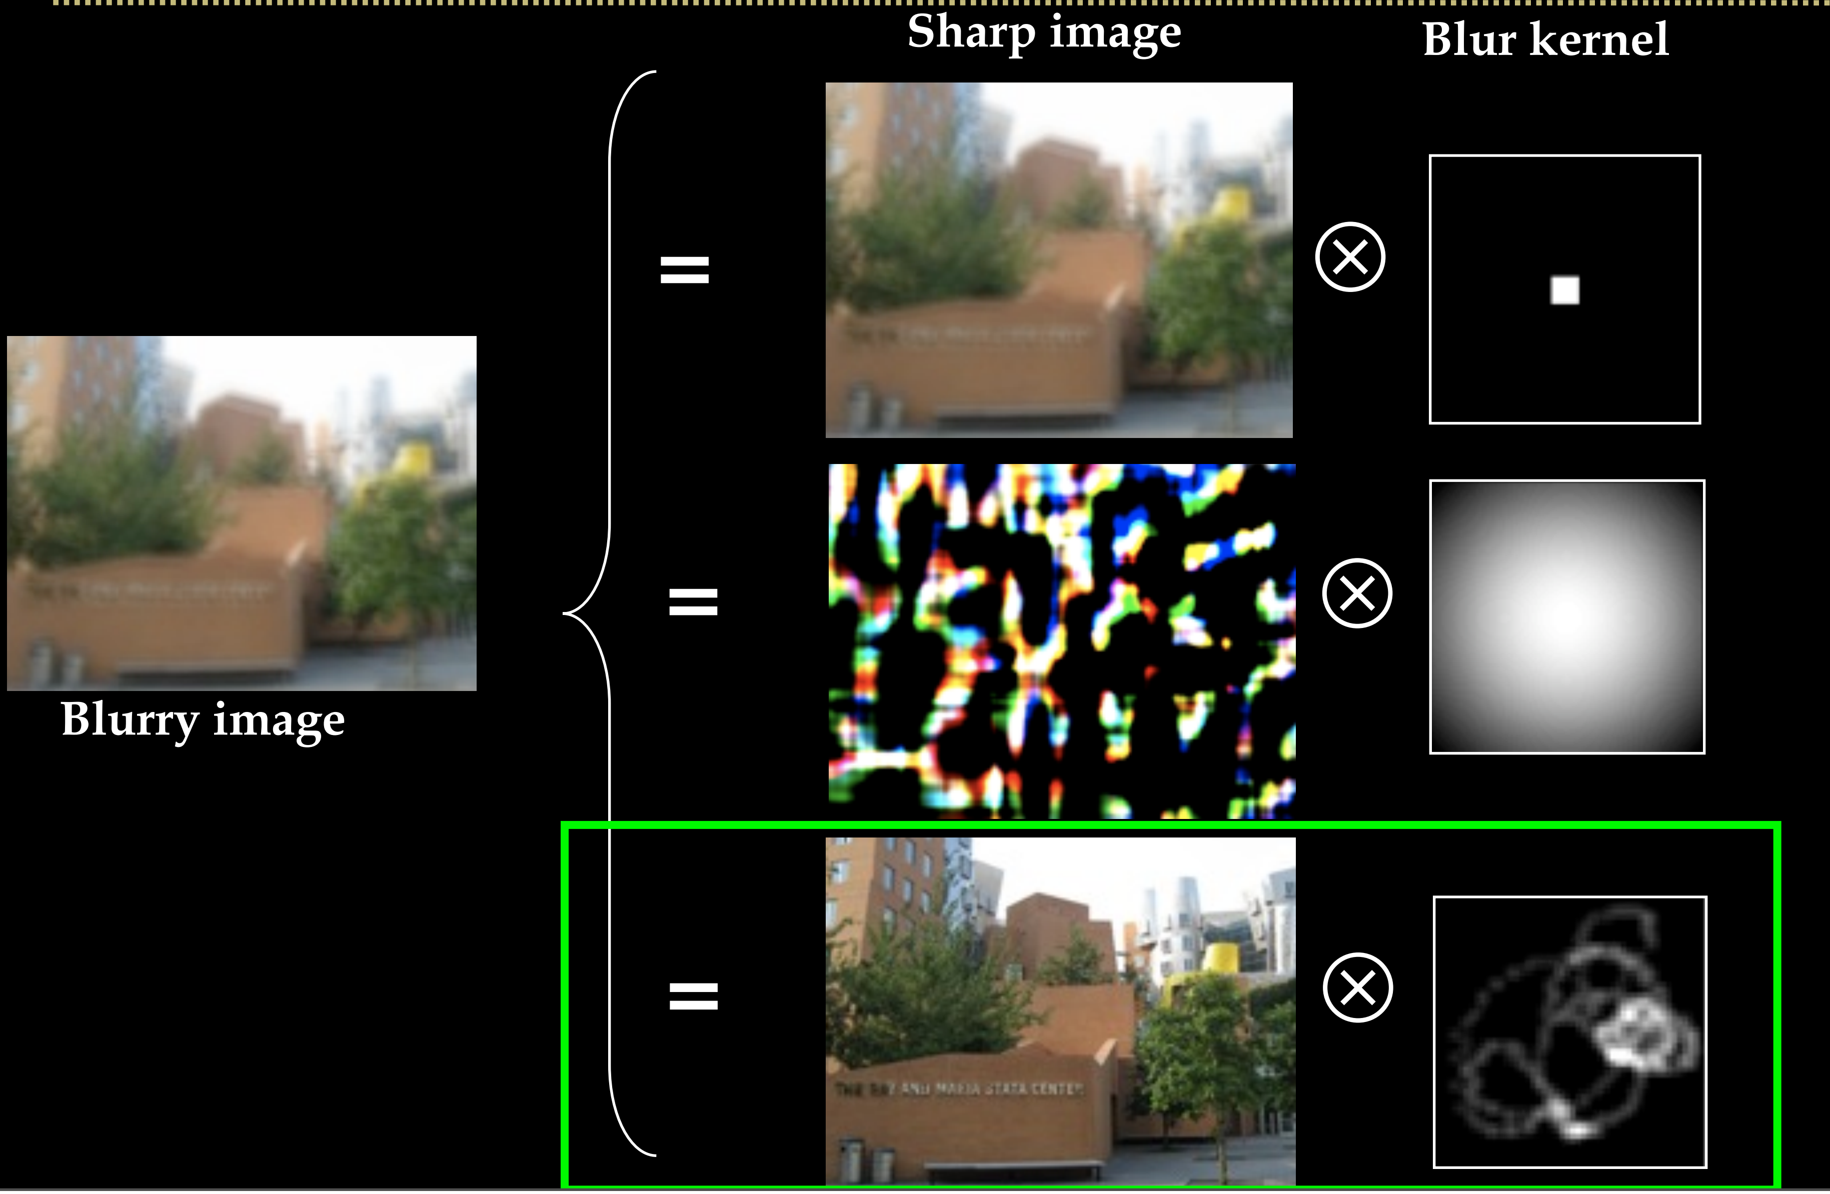

Is each of the images that follow sharp or blurred?

Natural image statistics

Characteristic distribution with heavy tails

Histogram of image gradients

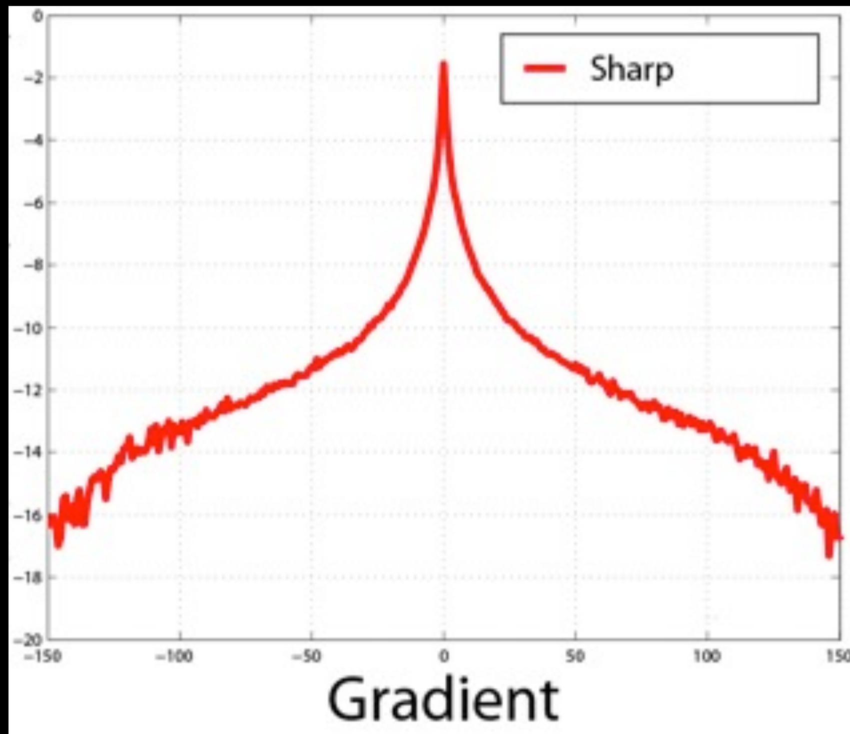

Blury images have different statistics

Histogram of image gradients

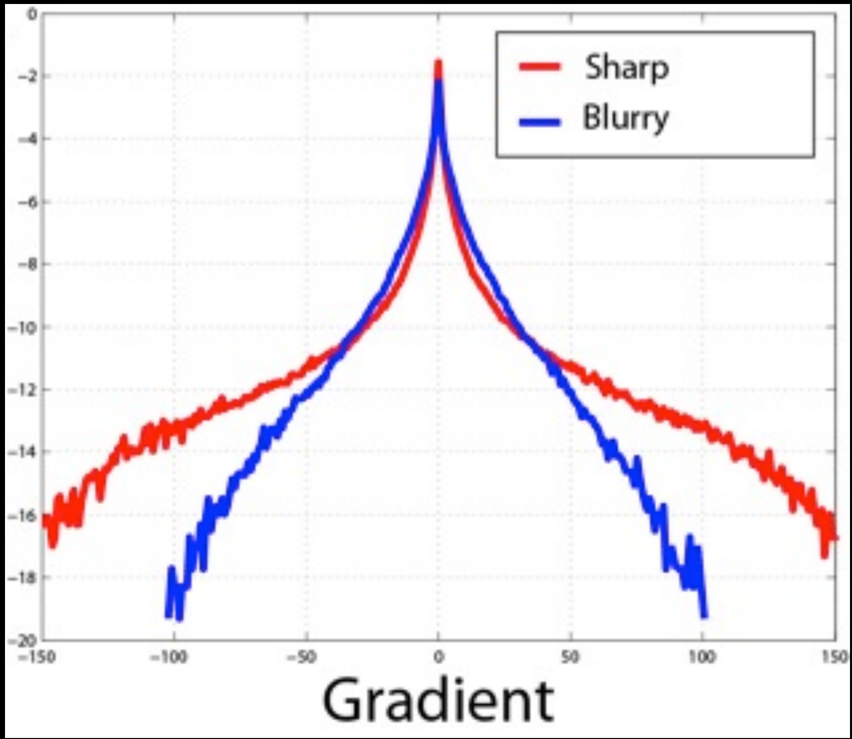

Parametric distribution

Histogram of image gradients

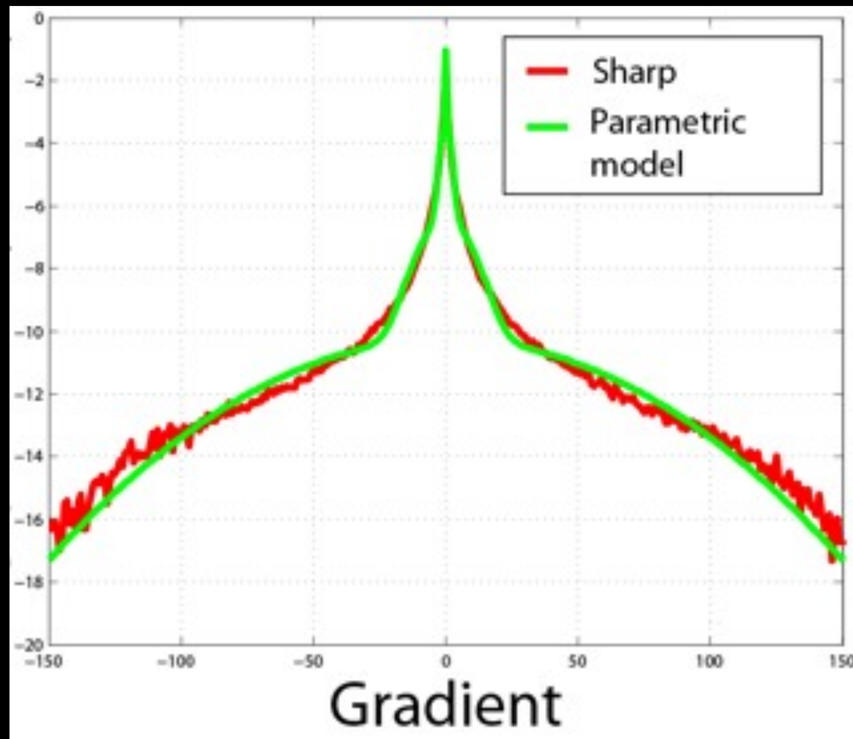

Use parametric model of sharp image statistics

Three sources of information

1. Reconstruction constraint:

Three sources of information

1. Reconstruction constraint:

Estimated sharp image

$\otimes$

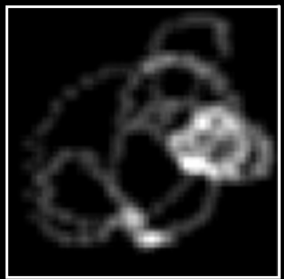

Estimated blur kernel

=

Input blurry image

Three sources of information

1. Reconstruction constraint:

Estimated sharp image

$\otimes$

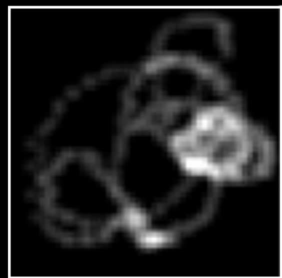

Estimated blur kernel

=

Input blurry image

2. Image prior:

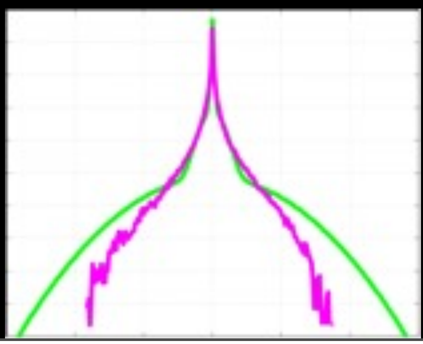

Distribution of gradients

Three sources of information

1. Reconstruction constraint:

Estimated sharp image

$\otimes$

Estimated blur kernel

=

Input blurry image

2. Image prior:

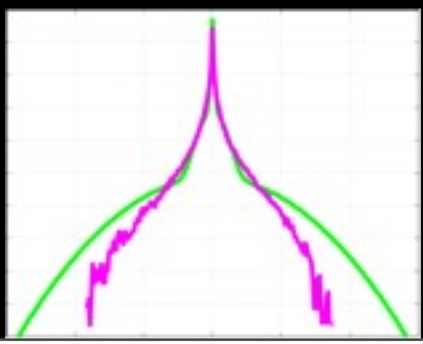

Distribution of gradients

3. Blur prior:

Positive & Sparse

Bayesian estimate of latent image, x, and blur kernel, b.

y – observed blurry image
x – unobserved sharp image
b – blur kernel

$$p(x, b|y) \propto p(y|x, b)p(x)p(b)$$

Likelihood

Latent image prior

Blur prior

Sparse and

$$e^{-\frac{\sum_i (x_i * b - y_i)^2}{2\sigma^2}} \qquad e^{-\lambda \sum_i |f(x_i)|} \qquad \geq 0$$

i – image patch index
f – derivative filter

Assumption: all pixels independent of one another

Original photograph

Blur kernel

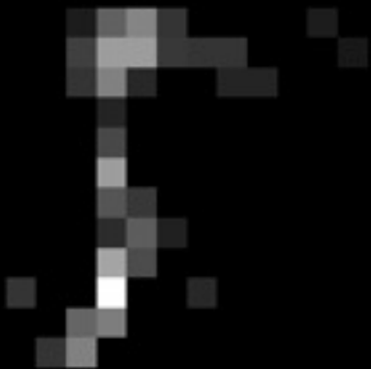

Our output


Matlab's deconvblind

Close-up of garland

Original

Matlab's `deconvblind`

Our output

Problem:

Statistical characterization of images.

Applications:

Low-level vision: noise removal, super-resolution, filling-in, texture synthesis.

References:

U Rajashekar and E P Simoncelli, Multiscale denoising of photographic images, in The Essential Guide to Image Processing,pages 241--261. Academic Press, Jul 2009. http://www.cns.nyu.edu/pub/lcv/rajashekar08a.pdf

R. Fergus, B. Singh, A. Hertzmann, S. Roweis, and W. T. Freeman, Removing camera shake from a single image, SIGGRAPH 2006. http://people.csail.mit.edu/billf/papers/deblur_fergus.pdf

Stefan Roth and Michael J. Black: Fields of Experts. International Journal of Computer Vision (IJCV), 82(2):205-229, April 2009. http://www.gris.informatik.tu-darmstadt.de/~sroth/pubs/foe-ijcv.pdf

Y. Weiss and W. T. Freeman, What makes a good model of natural images?, IEEE Computer Vision and Pattern Recognition (CVPR) 2007. http://people.csail.mit.edu/billf/papers/foe-final.pdf

IEEE International Conference on Computer Vision, Corfu, Greece, September 1999

Texture Synthesis by Non-parametric Sampling

Alexei A. Efros and Thomas K. Leung
Computer Science Division
University of California, Berkeley
Berkeley, CA 94720-1776, U.S.A.
{efros,leungt}@cs.berkeley.edu

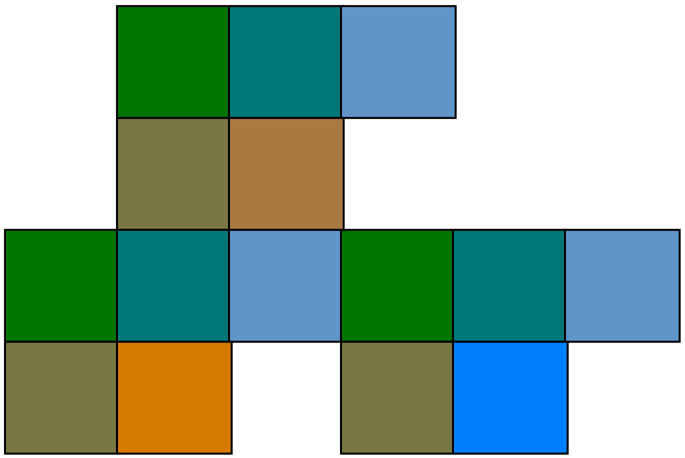

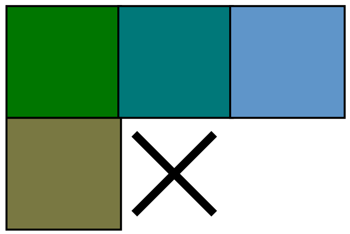

Algorithm

– Pick size of block and size of overlap
– Synthesize blocks in raster order

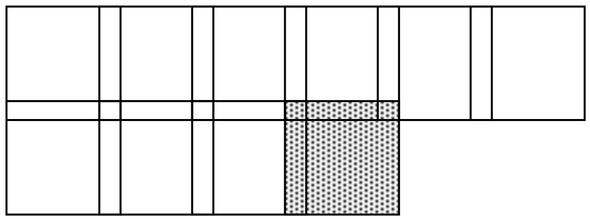

– Search input texture for block that satisfies overlap constraints (above and left)
  • Easy to optimize using NN search [Liang et.al., '01]
– Paste new block into resulting texture
  • use dynamic programming to compute minimal error boundary cut

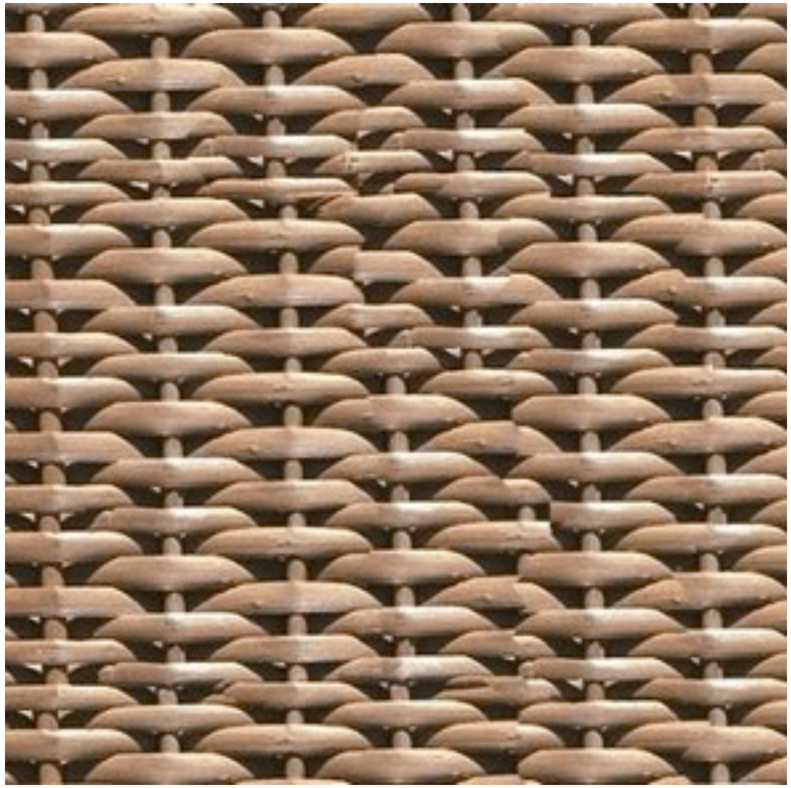

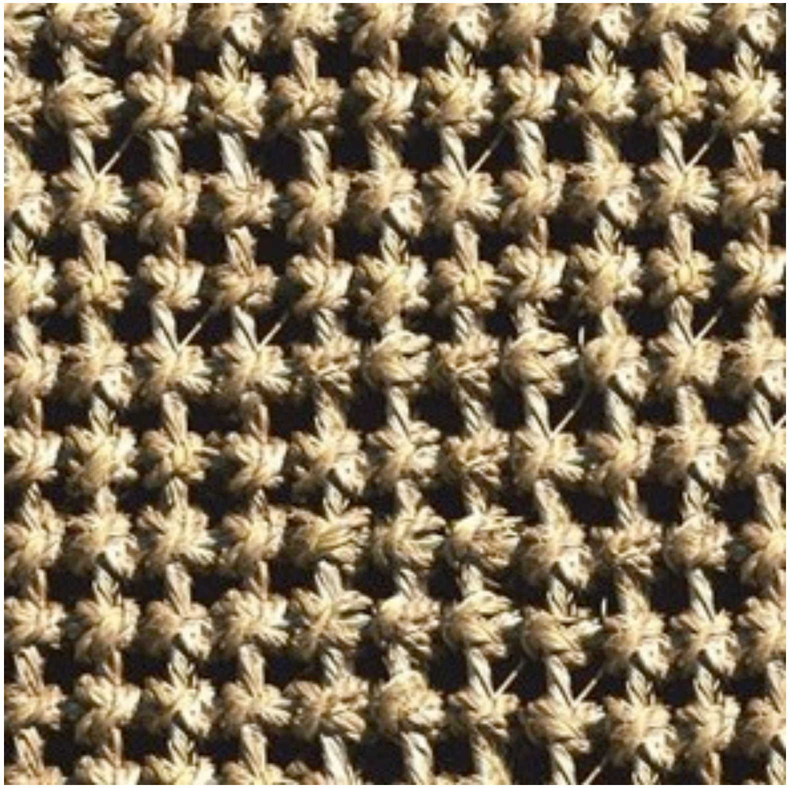

Problem:

How to construct and manage a non-parametric signal prior? How select the exemplars to use, how quickly find nearest neighbor matches?

Applications:

Low-level vision: noise removal, super-resolution, filling-in, texture synthesis.

References:

W. T. Freeman, E. C. Pasztor, O. T. Carmichael Learning Low-Level Vision International Journal of Computer Vision, 40(1), pp. 25-47, 2000. http://www.merl.com/reports/docs/TR2000-05.pdf

Alexei A. Efros and Thomas K. Leung, Texture Synthesis by Non-parametric Sampling, IEEE International Conference on Computer Vision (ICCV'99), Corfu, Greece, September 1999 , http://graphics.cs.cmu.edu/people/efros/research/NPS/efros-iccv99.pdf

Special case of an image prior: MRF

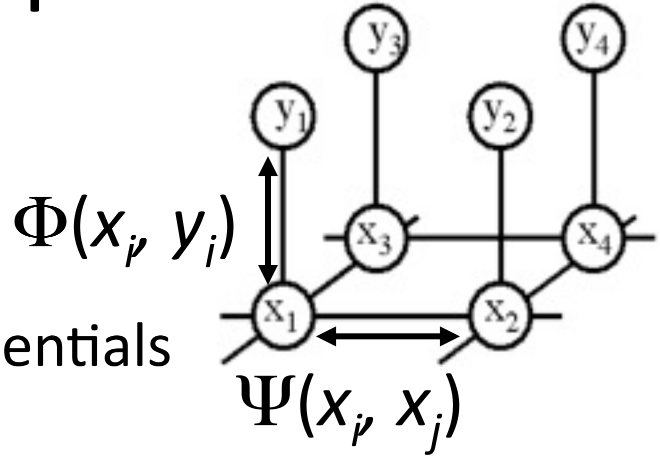

Network joint probability, pairwise potentials

$$P(x, y) = \frac{1}{Z} \prod_{i,j} \Psi\underbrace{(x_i, x_j)} \prod_i \Phi(x_i, y_i)$$

label

image

label-label compatibility function

neighboring label nodes

Image-label compatibility function

local observations

Some methods of approximate inference in MRF's

- loopy belief propagation
- graph-cuts (min-cut/max-flow)

MRF/CRF wishes

1. We'd like efficient algorithms for minimizing non-submodular functions, and which give bounds on the quality of the solution.

Input image

Segmentation using MRF with pairwise potentials

Segmentation using MRF with higher-order potentials

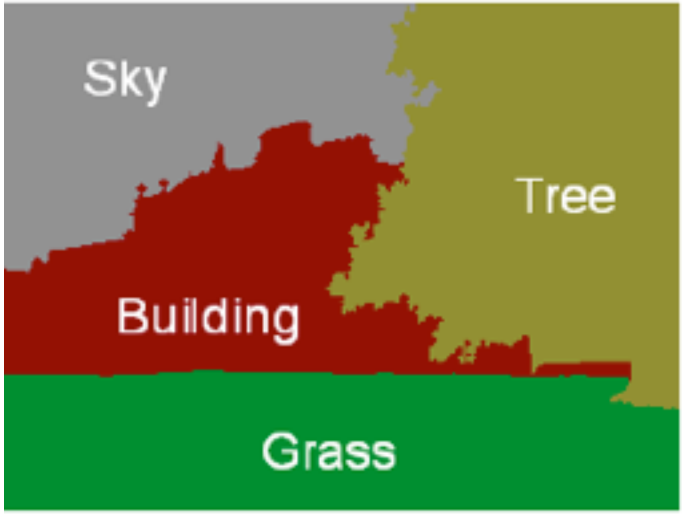

MRF/CRF wishes

2. There is real benefit to handling higher-order cliques. Need better ways to solve MRF's with such cliques, and provide performance bounds.
3. We often work with one of two kinds of constraints: (a) structure constraints, like planarity or treewidth and (b) language constraints, like submodularity or convexity. It would be useful to be able to combine the two.

Image Segmentation with A Bounding Box Prior

Victor Lempitsky, Pushmeet Kohli, Carsten Rother, Toby Sharp
Microsoft Research Cambridge

ICCV 2009

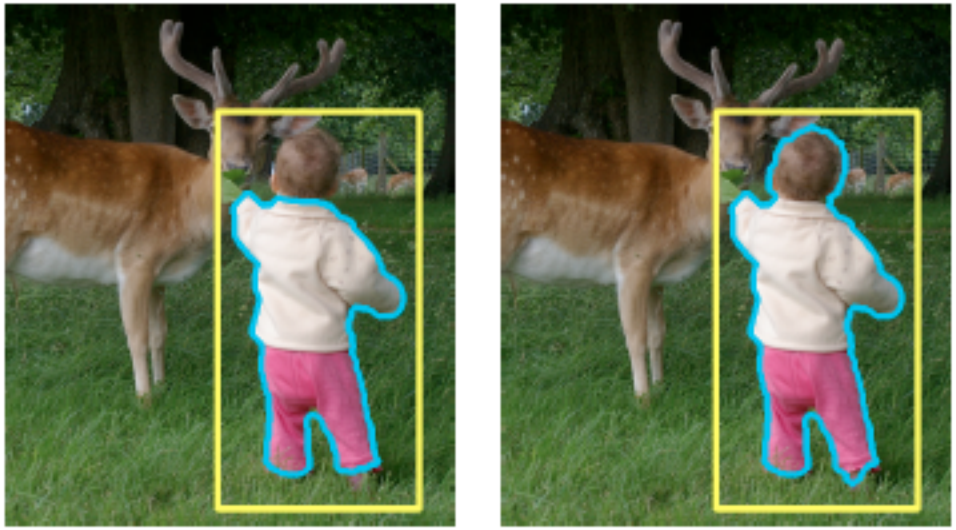

without the prior   with the prior

Figure 1. **Our tightness prior.** The segmentation on the left computed with graph cut is consistent with the low level image cues, yet inconsistent with the user input (in yellow) being too loose for this bounding box. By minimizing the same graph cut energy under a set of constraints, our method computes the segmentation that fits the bounding box in a sufficiently tight way, obtaining a better result (right).

Other constraints

- Topological constraints are often relevant to images, want to perform discrete optimization under such constraints.
- For example: specify that all states with some label within some neighborhood should be connected. Or that a user-specified founding box should somewhere touch a member of some label set. Lack efficient ways to solve that.

Problem:

Inference in Markov Random Fields. Want to handle higher order clique potentials, high-dimensional state variables, and real-valued state variables, language, structural, and topological constraints.

Applications:

Low-level vision: noise removal, super-resolution, filling-in, texture synthesis.

References:

Pushmeet Kohli, Lubor Ladicky, Philip Torr
Robust Higher Order Potentials for Enforcing Label Consistency.
In: International Journal of Computer Vision, 2009. http://research.microsoft.com/en-us/um/people/pkohli/papers/klt_IJCV09.pdf

Outline

- About me
- Computer vision applications
- Computer vision techniques and problems:
  – High-level vision: combinatorial complexity
  – Low-level vision: underdetermined problems
  – Miscellaneous problems

Compressed sensing

$$y = Wx + \eta$$

Compressed sensing

- Current sparsity assumptions are unrealistic for natural images.
- Is there a relaxed set of sparsity assumptions, which images meet, which would be useful for compressed sensing?
- Are there useful applications of compressed sensing in the domain of natural images?

Problem:

Is there a relaxed set of sparsity assumptions, met by natural images, useful for compressed sensing?

Applications:

Potential photographic applications.

References:

Y. Weiss, H. Sung-Chang and W. T. Freeman
Learning Compressed Sensing
Allerton Conference, 2008

Large, noisy datasets

- Relative importance of vision system components:
  - (1) datasets, (2) features, (3) algorithms
- Need progress handling large, noisy datasets.
  - we assume training and test distributions are the same; they rarely are. Under what circumstances can you break the assumption that the two distributions are the same?
  - What is the effect on algorithms when the IID assumption doesn't hold?
  - huge datasets make online learning important.

Problem:

Algorithms for large, inaccurately labeled datasets.

Applications:

Most modern algorithms use such datasets.

References:

Spectral Hashing

Y. Weiss, A. Torralba, R. Fergus.

Advances in Neural Information Processing Systems, 2008

Shai Avidan

Blind vision

Algorithm 4. Secure Classifier

Input: Alice has input test pattern $\mathbf{x} \in F^L$
Input: Bob has a strong classifier of the form $H(\mathbf{x}) = sign(\sum_{n=1}^{N} h_n(\mathbf{x}))$
Output: Alice has the result $H(\mathbf{x})$ and nothing else
Output: Bob learns nothing about the test pattern $\mathbf{x}$

1. Bob generates a set of N random numbers: $s_1, ..., s_N$, such that $s = \sum_{n=1}^{N} s_n$
2. For each $n = 1, ..., N$, Alice and Bob conduct the following sub-steps:
   (a) Alice and Bob obtain private shares a and b, respectively, of the dot product $\mathbf{x}^T \mathbf{y_n}$ using the secure-dot-product protocol.
   (b) Alice and Bob use the secure Millionaire protocol to determine which number is larger: a or $\Theta_n - b$. Instead of returning $\mathcal{A}$ or $\mathcal{B}$ the secure Millionaire protocol should return either $\alpha_n + s_n$ or $\beta_n + s_n$. Alice stores the result in $\mathbf{c_n}$.
3. Alice and Bob use the secure Millionaire protocol to determine which number is larger: $\sum_{n=1}^{N} \mathbf{c_n}$ or $\sum_{n=1}^{N} \mathbf{s_n}$. If Alice has a larger number then $\mathbf{x}$ is positively classified, otherwise $\mathbf{x}$ is negatively classified.

Problem:

Develop secure multi-party techniques for vision algorithms.

Applications:

Secure, distributed image analysis.

References:

S. Avidan and M. Butman
Blind Vision
European Conference on Computer Vision (ECCV), Graz, Austria, 2006.
http://www.merl.com/reports/docs/TR2006-006.pdf

Paper abstract:
Alice would like to detect faces in a collection of sensitive surveillance images she own. Bob has a face detection algorithm that he is willing to let Alice use, for a fee, as long as she learns nothing about his detector. Alice is willing to use Bob´ s detector provided that he will learn nothing about her images, not even the result of the face detection operation. Blind vision is about applying secure multi-party techniques to vision algorithms so that Bob will learn nothing about the images he operates on, not even the result of his own operation and Alice will learn nothing about the detector. The proliferation of surveillance cameras raises privacy concerns that can be addressed by secure multi-party techniques and their adaptation to vision algorithms.

me

Continuous to discrete representations

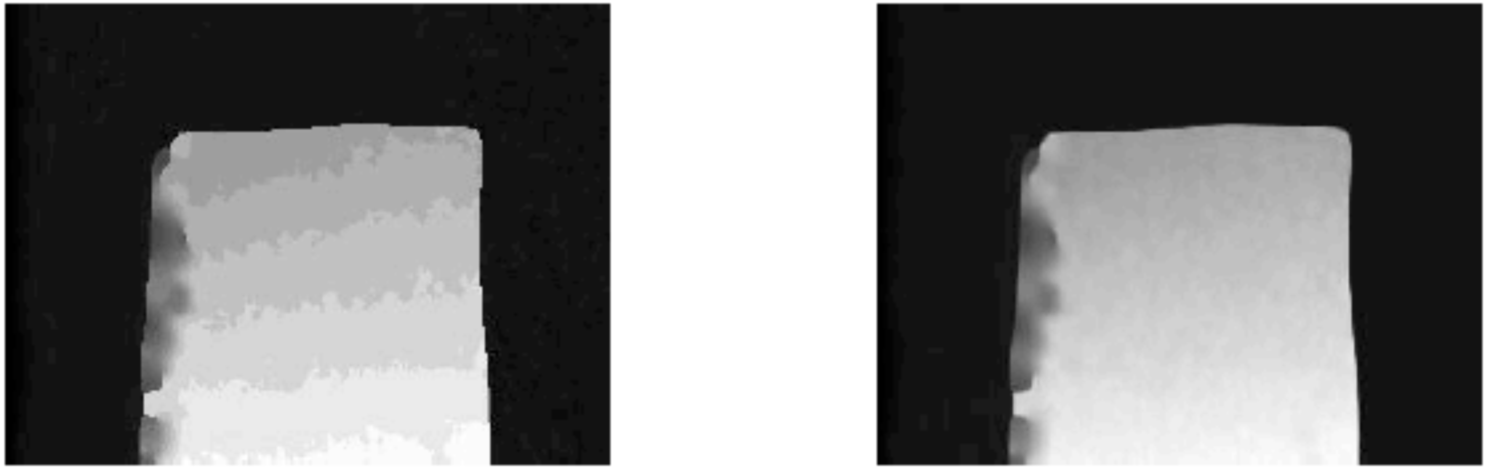

(a) MAP Estimate (b) MMSE Estimate

Figure 7. Comparison of MAP and MMSE estimates on a different MRF formulation. The MAP estimate chooses the most likely discrete disparity level for each point, resulting in a depth-map with stair-stepping effects. Using the MMSE estimate assigns sub-pixel disparities, resulting in a smooth depth map.

Problem:

A theory for how to optimally quantize and manipulate probabilities over a continuous domain.

Applications:

Probabilistic shape estimation.

Deva Ramanan

Evaluate easily over a powerset of all segmentations.

Deva Ramanan: wants a fast and efficient way to search over all possible segmentations of an image, scoring each one against some model.

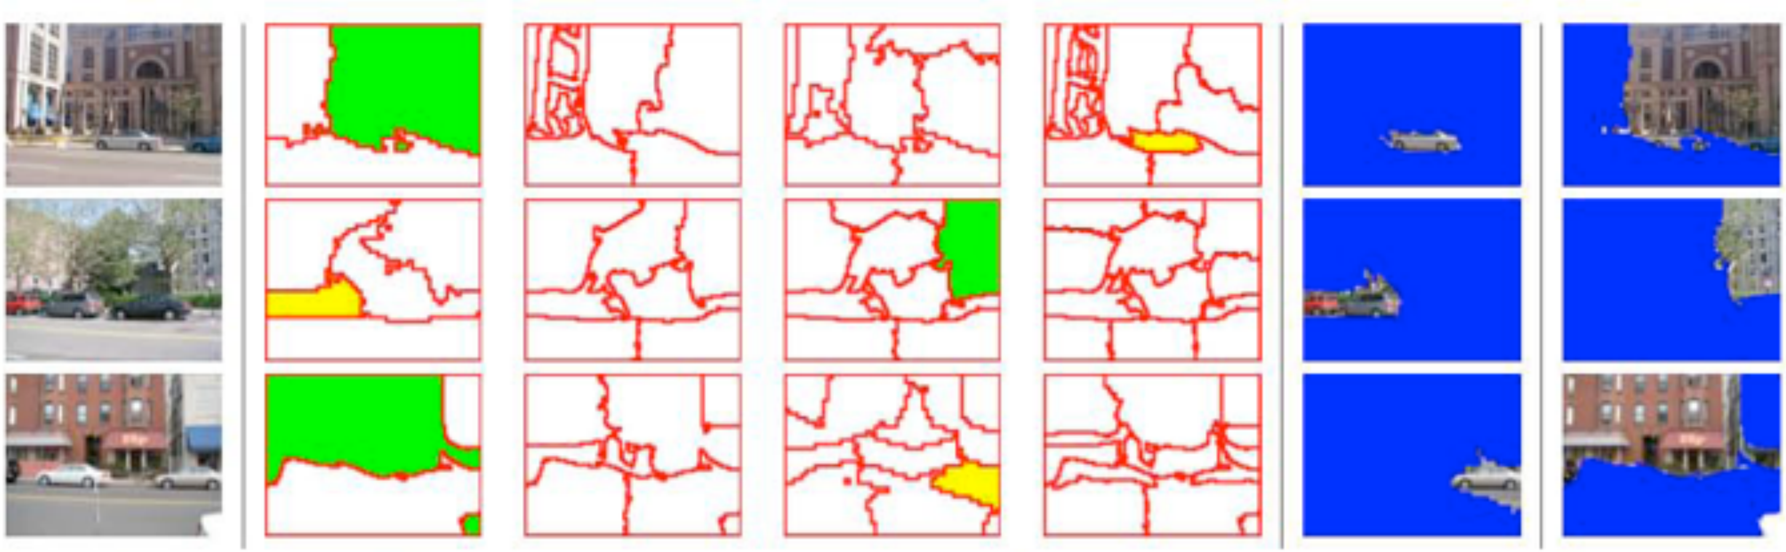

Figure 1. **Problem summary**. Given a set of input images (first column), we wish to discover object categories and infer their spatial extent (e.g. cars and buildings: final two columns). We compute multiple segmentations per image (a subset is depicted in the second through fifth columns; all of the segmentations for the first row are shown in Figure 4). The task is to sift the good segments from the bad ones for each discovered object category. Here, the segments chosen by our method are shown in green (buildings) and yellow (cars).

http://www.di.ens.fr/~russell/papers/Russell06.pdf

Problem:

Evaluate some segmentation-dependent function over (some approximation to) all possible segmentations. Note: different than bottom-up segmentation, which I would not recommend as a research project.

Applications:

Image understanding.

References:

Deva's home page: http://www.ics.uci.edu/~dramanan/

Using Multiple Segmentations to Discover Objects and their Extent in Image Collections, Bryan Russell, Alexei A. Efros, Josef Sivic, Bill Freeman, Andrew Zisserman in CVPR 2006, http://people.csail.mit.edu/brussell/research/proj/mult_seg_discovery/index.html


Alyosha Efros

Efros and Hoiem comments

Another representational issue relates to noisy evidence versus the lack of evidence. Presently, we usually treat those in the same way, but we might want to distinguish between there being a 10% similarity of something to a dog, versus a 10% probability that it's a dog, ie, to have a different description for weak relationship between things, and uncertainty about the relationship between things [10, 23]. It would be nice to allow for multiple different relationships between nodes–to have graphs with multi-colored edges [10].

Really, we'd like another breakthrough...

Algorithms for inference and classification,

– support vector machines

– boosting

– belief propagation

– graph cuts

have each led to much progress and creativity within the field. We're ready for the next algorithm...


David Lowe

David Lowe

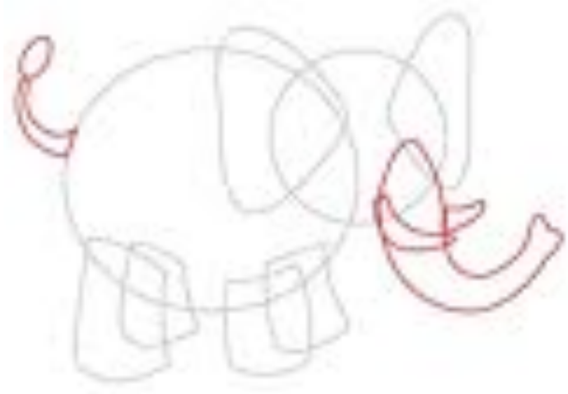

tlc.howstuffworks.com

need better features. an artist can draw then end of an elephant's trunk, and you know immediately what it is. but our features don't capture that similarity at all.

learning of features from images. what is a natural encoding of images? as a warning for what approach not to take: don't bother learning translation invariance, or rotation invariance. so a little bit of supervision is ok.

Most references are in citation list of SODA manuscript

Where computer vision needs help from computer science *

William T Freeman†

January, 2011

Abstract

This paper describes areas and problems where computer vision can use help from the discrete algorithms community.

1 Introduction

Computer vision is a good target for the discrete algorithms of computer science. While we are far from being able to interpret images reliably using a computer, it is clear that there will be many benefits when we do reach

eras use that capability to control exposure and focus settings. In the controlled conditions of a factory, computers routinely detect defects in manufactured parts and labels. For example, most manufactured diapers are visually inspected by computer [34]. Computers read license plates and digitized documents, monitor traffic, and track traffic lanes from cars in highways.

But if you look more closely, even those successes reveal where much more progress is needed. Face

Computer vision academic culture

No more "if only" papers

End-to-end empirical orientation.
There is a certain overhead in coming up to speed on the filters and representations.
Need dataset validation.
The competitive conferences have 20-25% acceptance rate. Other conferences have little impact. The competitive conferences: CVPR, ICCV, ECCV, NIPS.

Thus: best to collaborate with a computer vision researcher. We know that you can help us, and our doors are open.

A computer graphics application of nearest-neighbor finding in high dimensions

Kaneva, Sivic, Torralba, Avidan, and Freeman, Infinite Images, Proceedings of IEEE, 2010.

A computer graphics application of nearest-neighbor finding in high dimensions

Kaneva, Sivic, Torralba, Avidan, and Freeman, Infinite Images, Proceedings of IEEE, 2010.

The image database

•We have collected ~6 million images from Flickr based on keyword and group searches

– typical image size is 500x375 pixels

– 720GB of disk space (jpeg compressed)

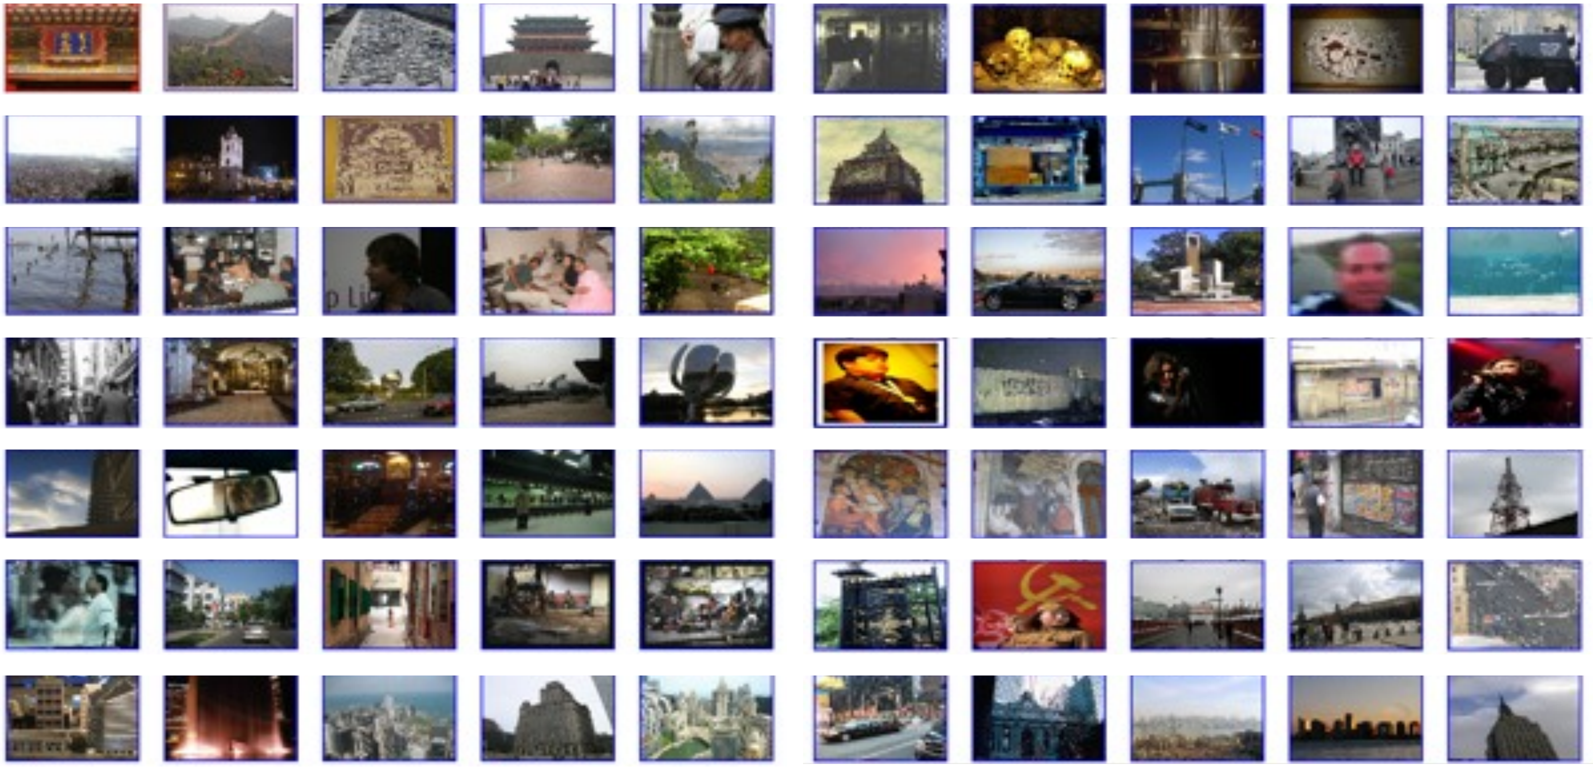

Image representation

Obtaining semantically coherent themes

We further break-up the collection into themes of semantically coherent scenes:

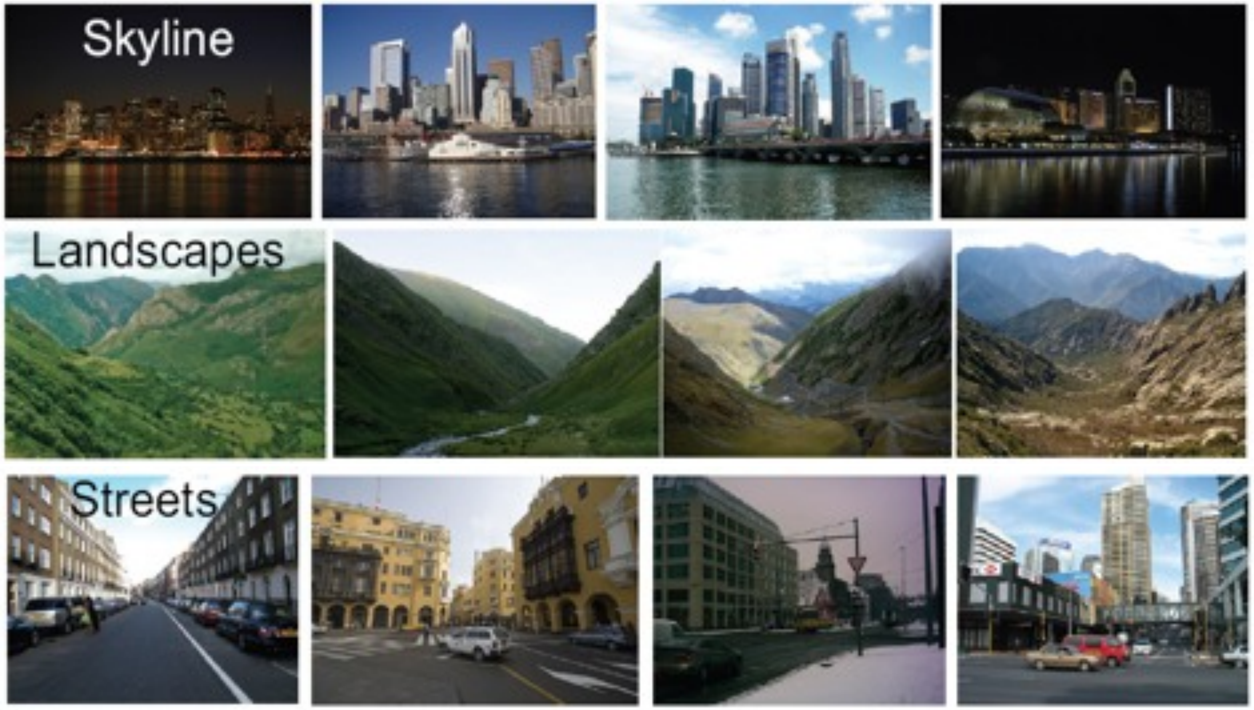

Train SVM-based classifiers from 1-2k training images [Oliva and Torralba, 2001]

Basic camera motions

Starting from a single image, find a sequence of images to simulate a camera motion:

Forward motion

Camera rotation

Camera pan

Scene matching with camera view transformations: Translation

1. Input image

Scene matching with camera view transformations: Translation

1. Input image

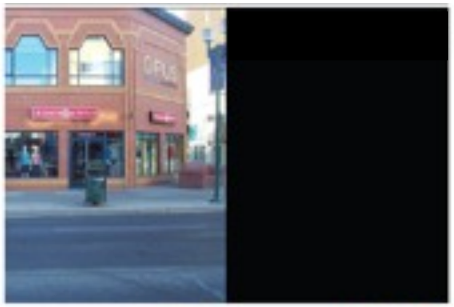

2. Move camera

Scene matching with camera view transformations: Translation

1. Input image

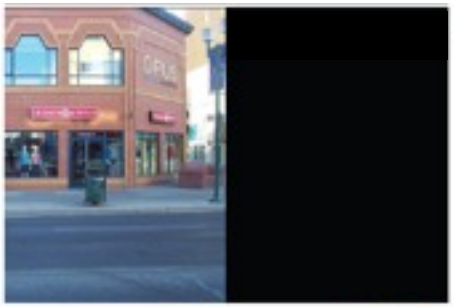

2. Move camera

3. Find a match to fill the missing pixels

Scene matching with camera view transformations: Translation

1. Input image

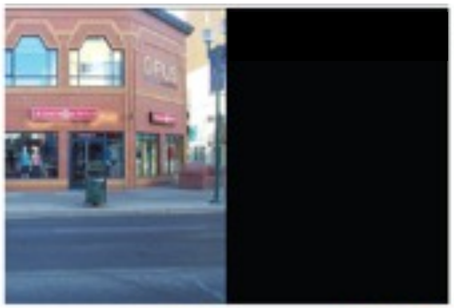

2. Move camera

3. Find a match to fill the missing pixels

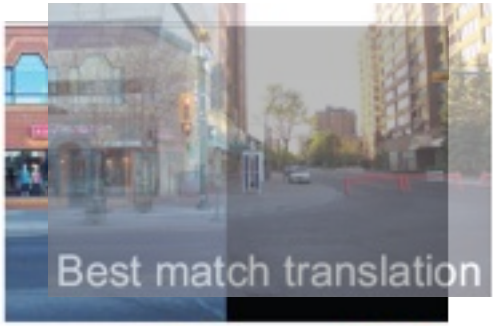

4. Locally align images

Scene matching with camera view transformations: Translation

1. Input image

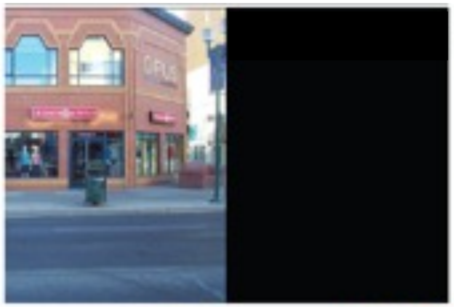

2. Move camera

3. Find a match to fill the missing pixels

4. Locally align images

5. Find a seam

Scene matching with camera view transformations: Translation

1. Input image

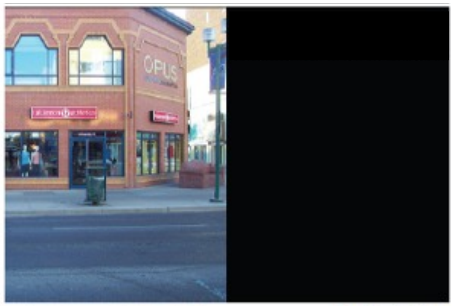

2. Move camera

3. Find a match to fill the missing pixels

4. Locally align images

5. Find a seam

6. Blend in the gradient domain

Scene matching with camera view transformations: Camera rotation

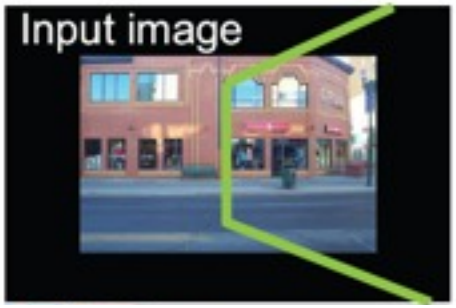

1. Rotate camera

Scene matching with camera view transformations: Camera rotation

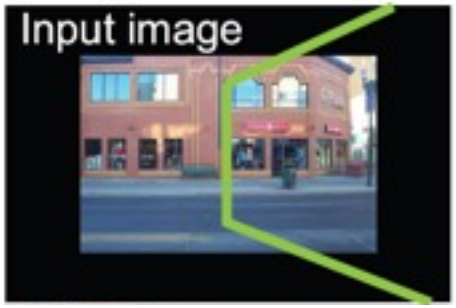

1. Rotate camera

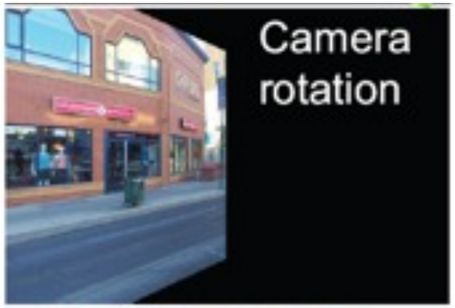

2. View from the
virtual camera

Scene matching with camera view transformations: Camera rotation

1. Rotate camera

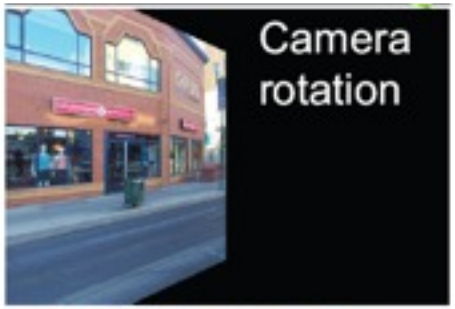

2. View from the virtual camera

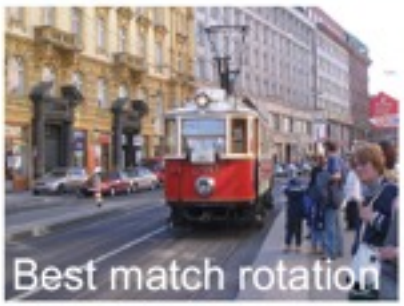

3. Find a match to fill-in the missing pixels

Scene matching with camera view transformations: Camera rotation

1. Rotate camera

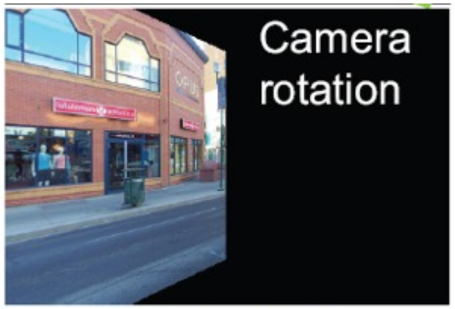

2. View from the virtual camera

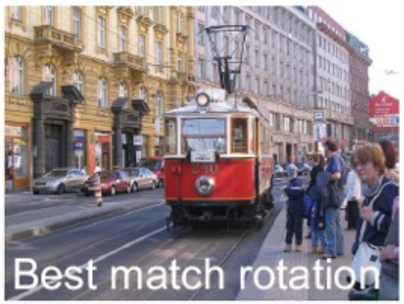

3. Find a match to fill-in the missing pixels

4. Stitched rotation

Scene matching with camera view transformations: Camera rotation

1. Rotate camera

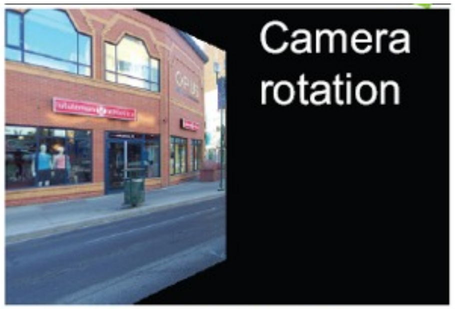

2. View from the virtual camera

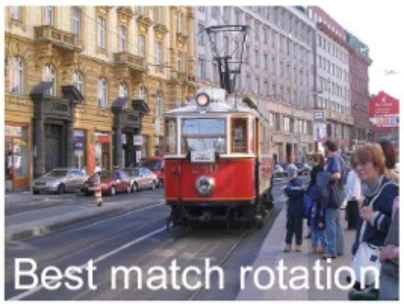

3. Find a match to fill-in the missing pixels

4. Stitched rotation

5. Display on a cylinder

More "infinite" images – camera translation

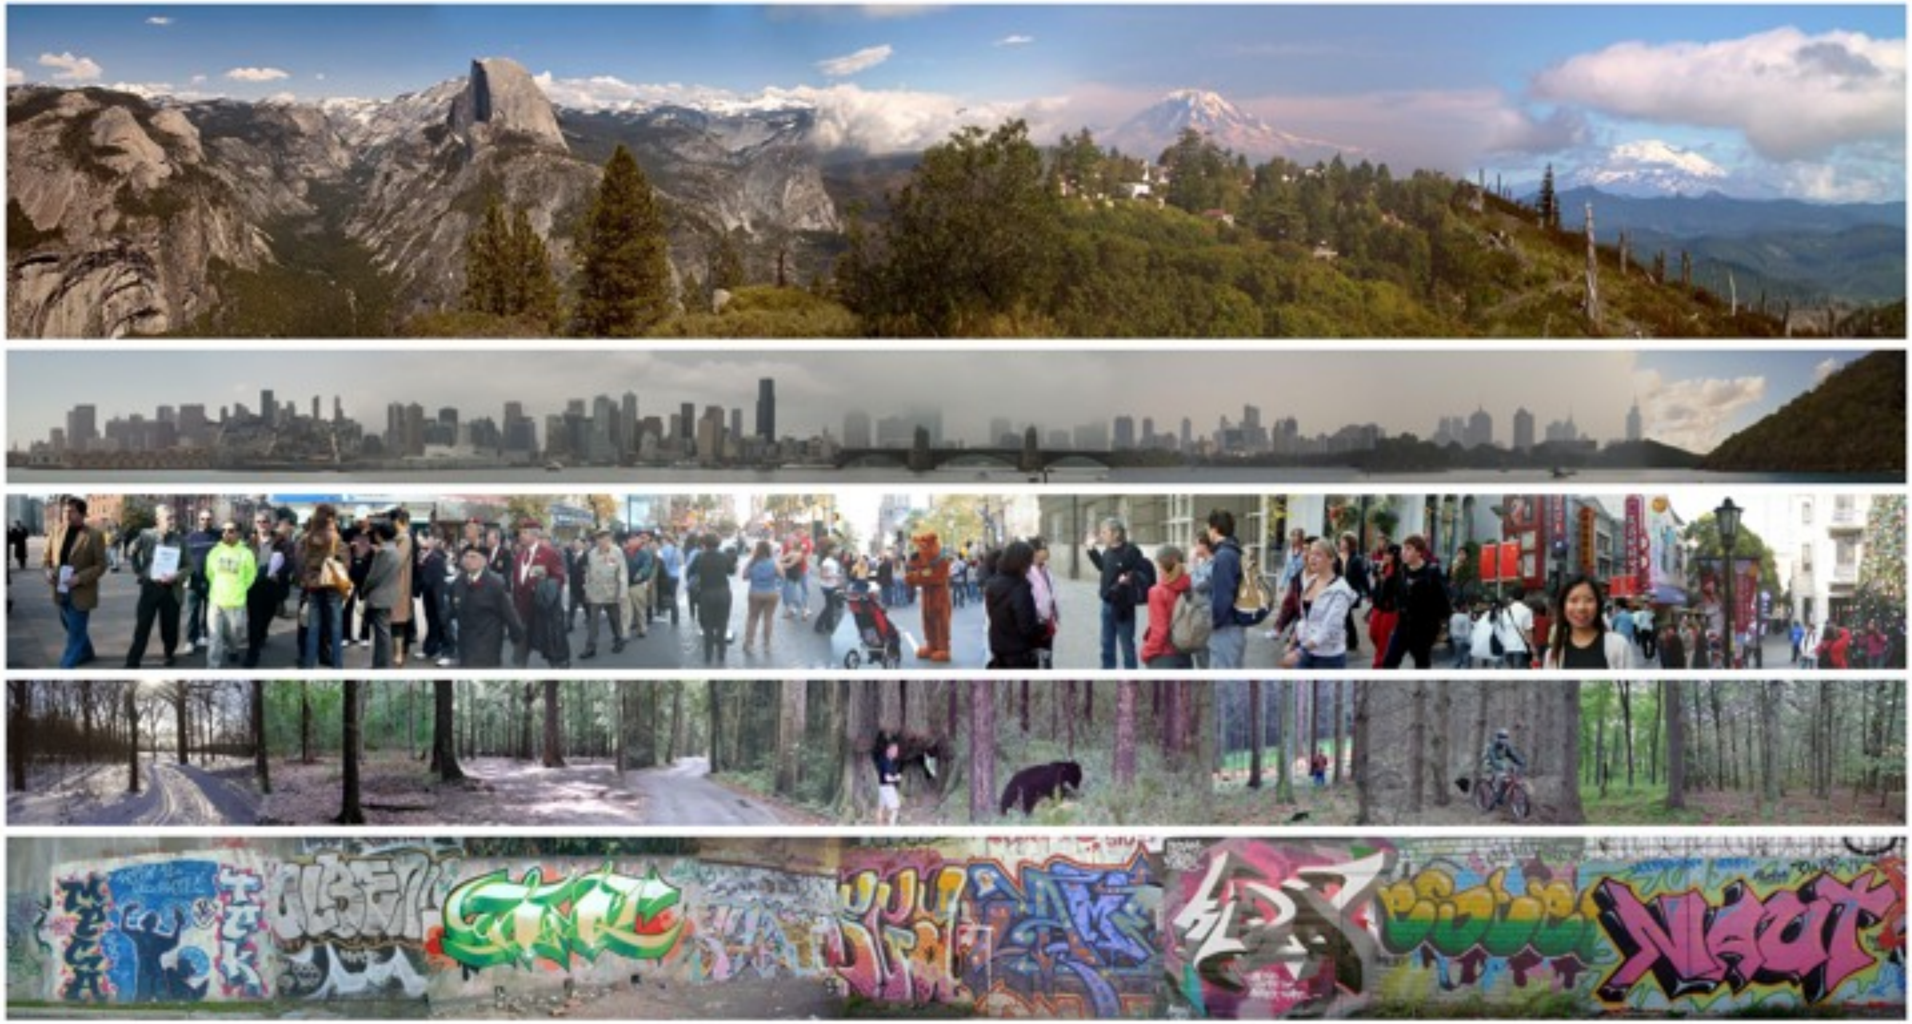


HOLLYWOOD

Image taxi

Image taxi

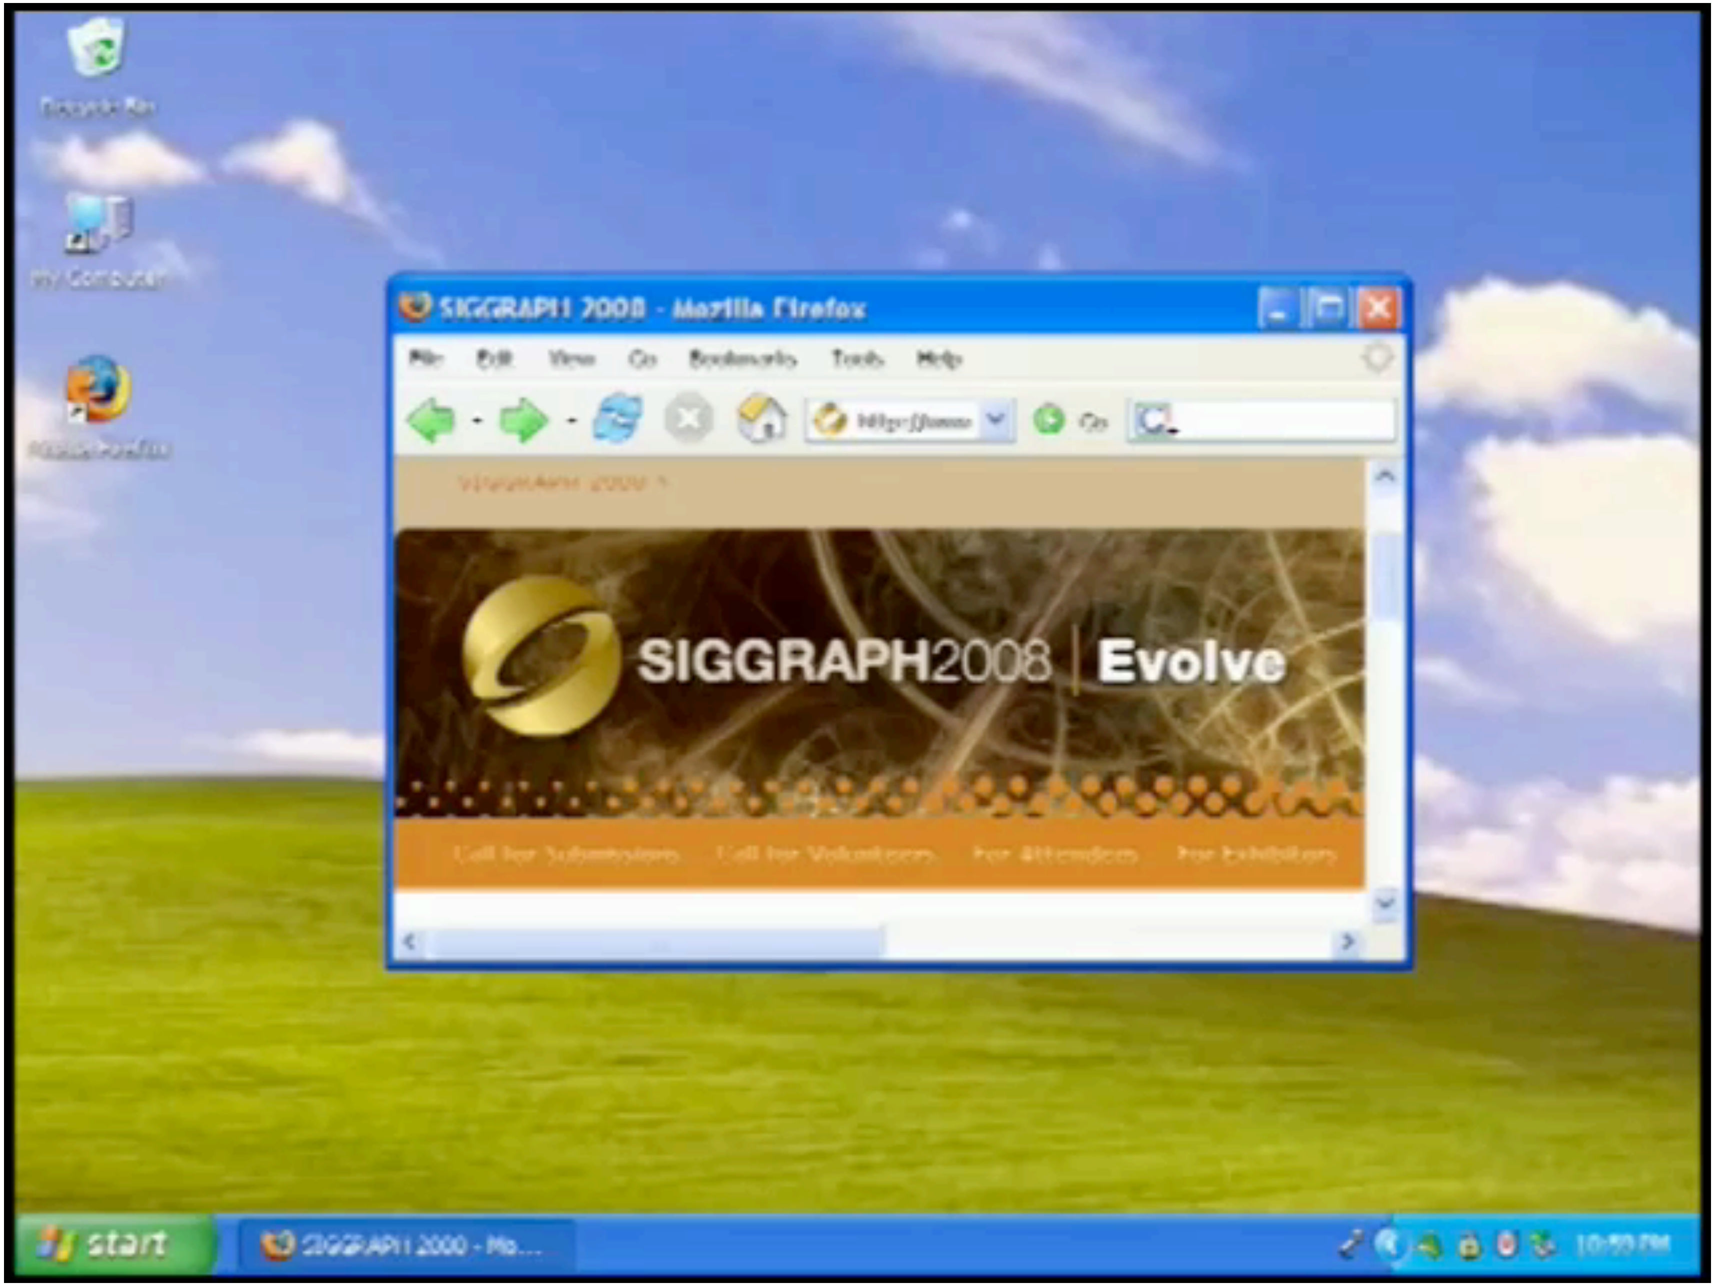
SIGGRAPH 2008 - Mozilla Firefox
SIGGRAPH2008
Evolve
start