


Joint Alignment: What's It Good For?


Erik Learned-Miller, with
Vidit Jain, Andras Ferencz, Gary
Huang, Lilla Zollei, Sandy Wells,



Congealing (CVPR 2000, PAMI 2006)



MRI Bias Removal *(NIPS 2005, MICCAI 2005, PAMI 2006)*



Five Applications

- Image factorizations
 - For transfer learning, learning from one example
- Alignment for Data Pooling
 - 3D MR registration
 - EEG registration
- Artifact removal
 - Magnetic resonance bias removal
- Improvements to recognition algorithms
 - Alignment before recognition
- Defining anchor points for registration
 - Find highly repeatable regions for future registrations


Congealing

- Process of joint “alignment” of sets of arrays (samples of continuous fields).
- 3 ingredients
 - A **set of arrays** in some class
 - A parameterized family of **continuous transformations**
 - A criterion of **joint alignment**

Congealing Binary Digits



- 3 ingredients
 - A set of arrays in some class:
 - Binary images
 - A parameterized family of *continuous* transformations:
 - Affine transforms
 - A criterion of joint alignment:
 - Entropy minimization

Congealing



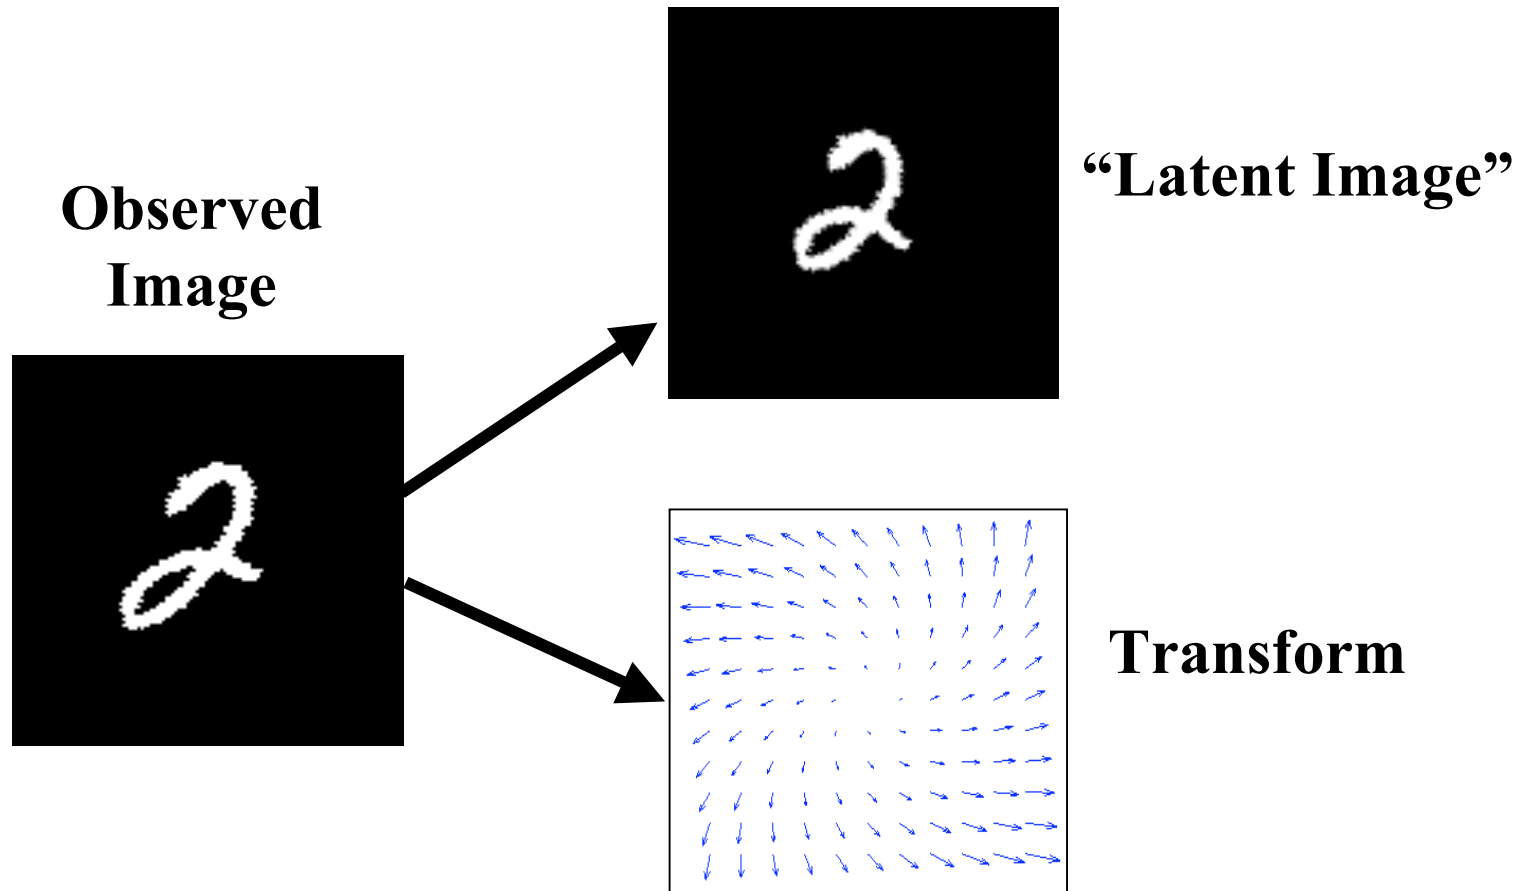
Criterion of Joint Alignment

- Minimize sum of pixel stack entropies by transforming each image.



Note: Mutual Information doesn't make sense here.

An Image Factorization



$$\arg \max_{\mathbf{T} \in \mathcal{T}} P(\mathbf{T} | \mathbf{I}) \stackrel{(a)}{=} \arg \max_{\mathbf{T} \in \mathcal{T}} P(\mathbf{I} | \mathbf{T}) P(\mathbf{T})$$

$$\stackrel{(b)}{=} \arg \max_{\mathbf{T} \in \mathcal{T}} P(\mathbf{I} | \mathbf{T})$$


$$\stackrel{(c)}{=} \arg \max_{\mathbf{T} \in \mathcal{T}} P(\mathbf{L}(\mathbf{I}, \mathbf{T}))$$

$$= \arg \max_{\mathbf{T} \in \mathcal{T}} \prod_{x,y} \prod_{i=1}^N p_{x,y}(L_i(x,y))$$

$$= \arg \max_{\mathbf{T} \in \mathcal{T}} \sum_{x,y} \sum_{i=1}^N \log p_{x,y}(L_i(x,y))$$

$$\stackrel{(d)}{\approx} \arg \min_{\mathbf{T} \in \mathcal{T}} \sum_{x,y} H(p_{x,y})$$

$$\stackrel{(e)}{\approx} \arg \min_{\mathbf{T} \in \mathcal{T}} \sum_{x,y} \hat{H}_{\text{Vasicek}}(L_1(x,y), \dots, L_N(x,y))$$



A pixel stack

Why Minimize Entropy?


- Negative entropy is just the average log likelihood of points under their own distribution.

Min entropy =
maximum non-parametric likelihood

The Independent Pixel Assumption

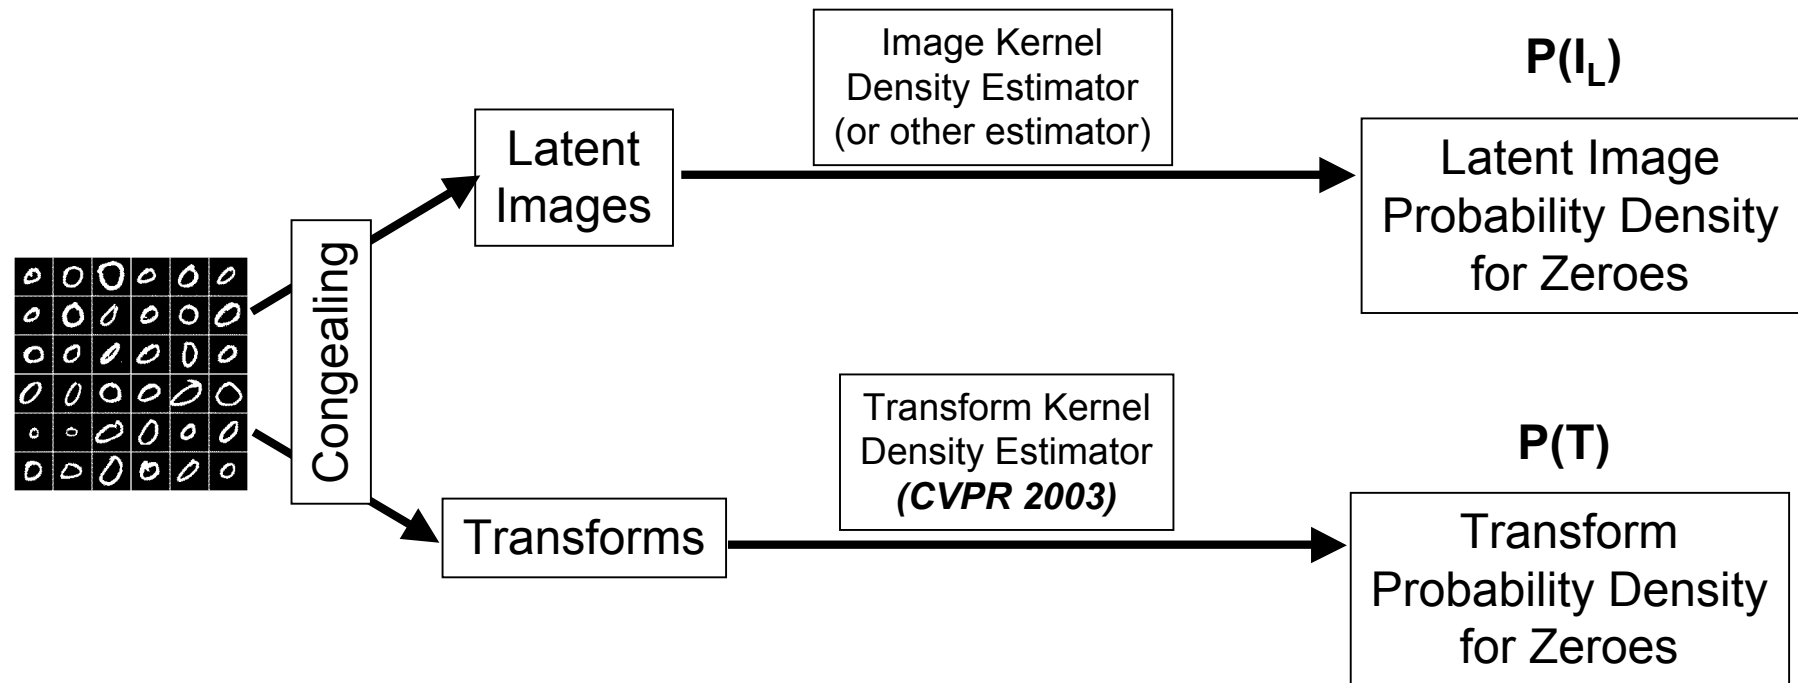
- Model assumes independent pixels
- A poor generative model:
 - True image probabilities don't match model probabilities.
 - Reason: heavy dependence of neighboring pixels.
- However! This model is great for alignment and separation of causes!
 - Why?
 - Relative probabilities of "better aligned" and "worse aligned" are usually correct.
- Once components are separated, a more accurate (and computationally expensive) model can be used to model each component.

Congealing




Each pair implicitly creates a sample of the transform T .

Character Models



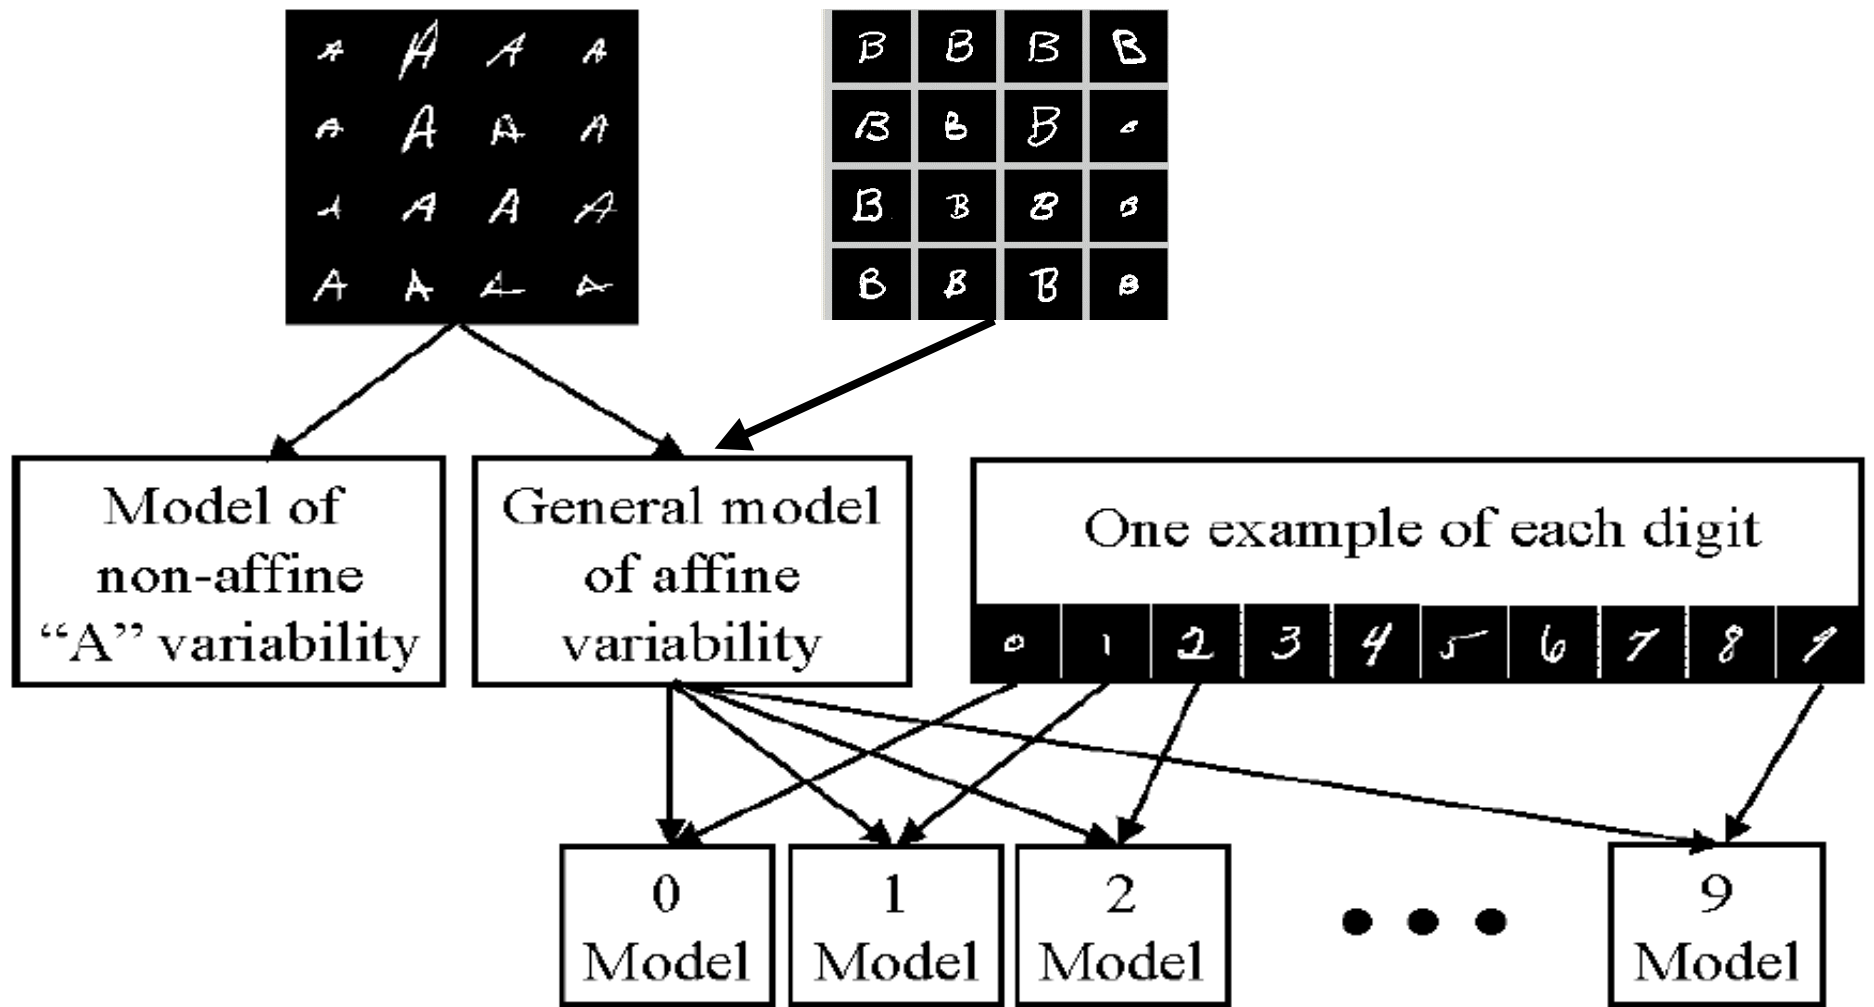
How do we line up a new image?

Sequence of successively “sharper” models



Take one gradient step with respect to each model.

Digit Models from One Example



Next Application: Alignment of 3D Magnetic Resonance Volumes

Lilla Zollei, Sandy Wells, Eric Grimson

Congealing MR Volumes: Joint Registration



- 3 ingredients
 - A set of arrays in some class:
 - Gray-scale MR volumes
 - A parameterized family of *continuous* transformations:
 - 3-D affine transforms
 - A criterion of joint alignment:
 - Grayscale entropy minimization

- Purposes:
 - Pooling data for fMRI studies
 - Building general purpose statistical anatomical atlases



Why Entropy?

- Drives volumes to having mass concentrated in a small number of tissues.
- Comparison to Transformed Mixture of Gaussians (Frey and Jojic).
- Convexity of entropy in distribution.


Congealing Gray Brain Volumes *(ICCV 2005 Workshop)*




Aligned Volumes



Validation: Synthetic Data




Unaligned input data sets



Aligned input data sets

Real Data



Unaligned input data sets

Aligned input data sets

Data set: 28 T1-weighted MRI; [256x256x124] with (.9375, .9375, 1.5) mm³ voxels

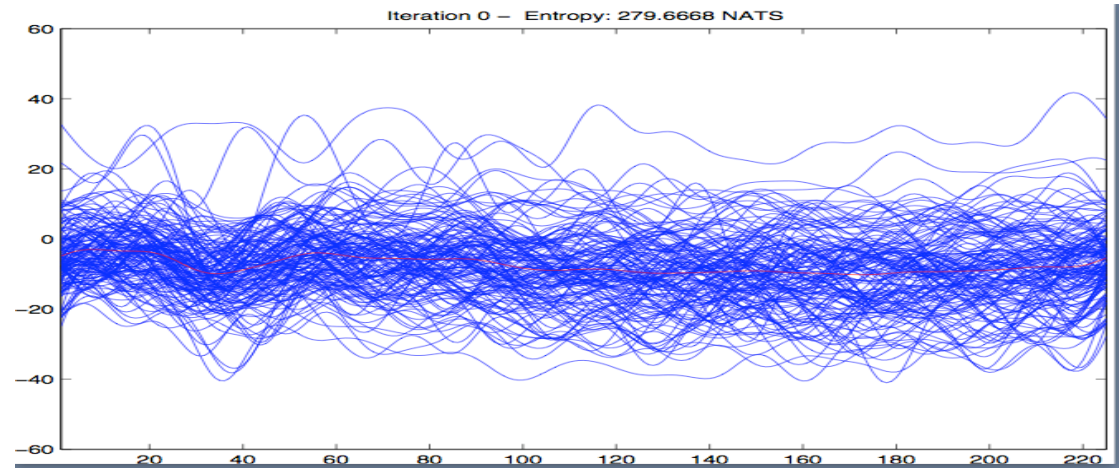
Experiment: 2 levels; 12-param. affine; N = 2500; iter = 150; time = **1209 sec!!**

MR Congealing Challenges


- Big data
 - 8 million voxels per patient
 - 100 patients
 - 12 transform *parameters*
 - 20 iterations
- Techniques:
 - Stochastic sampling
 - Multi-resolution techniques
 - Don't use visual basic

Signal to Noise in Event Related Potentials

Before
congealing




After
congealing



Next Application: Bias Removal in Magnetic Resonance Images

Parvez Ahammad, Vidit Jain

The Problem




Bias fields have low spatial frequency content

Bias Removal in MR as a Congealing Problem

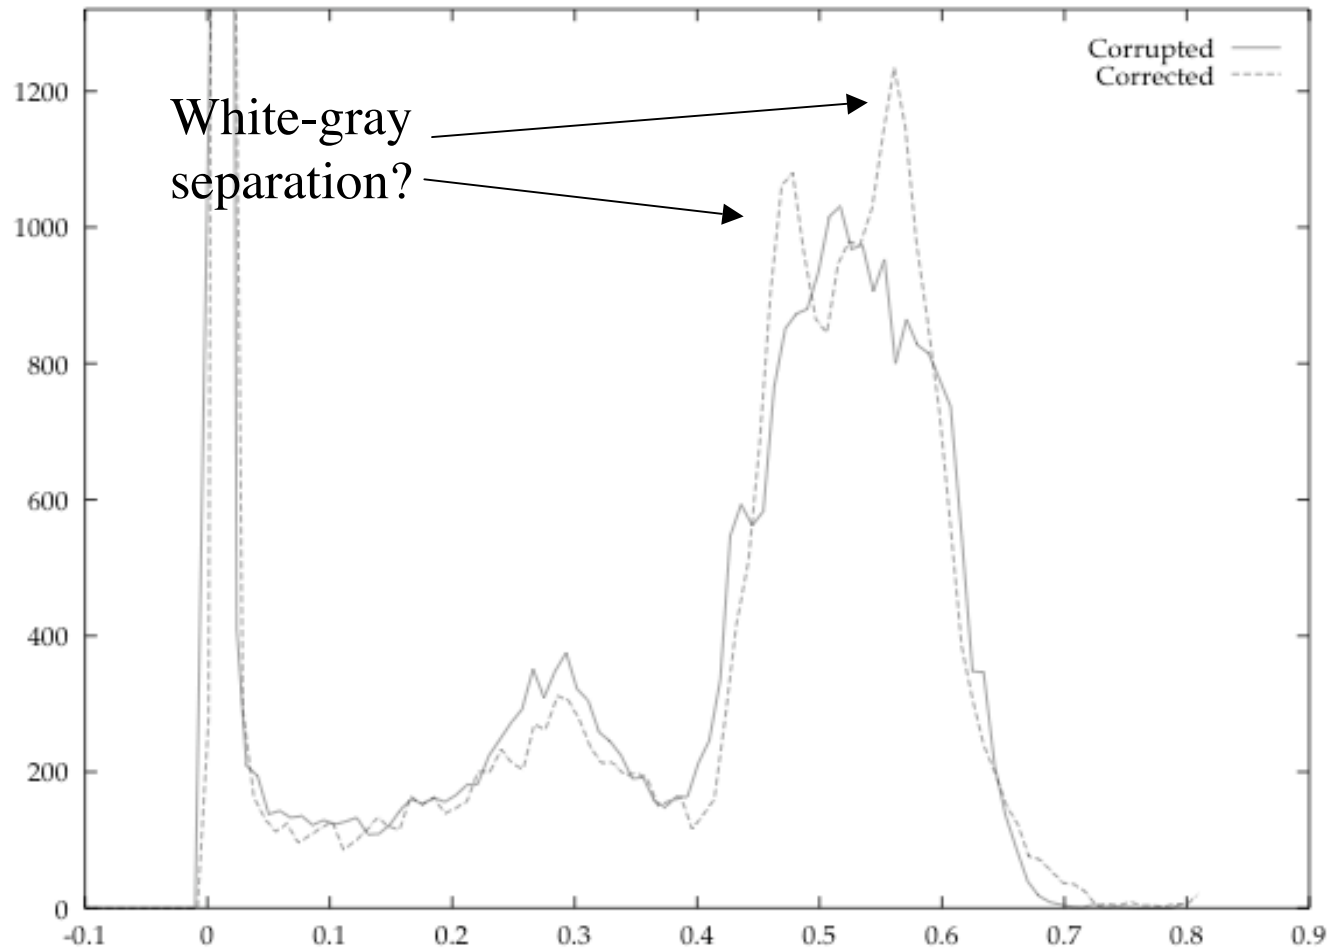
- 3 ingredients
 - A set of arrays in some class:
 - MR Scans of Similar Anatomy (2D or 3D)
 - A parameterized family of *continuous* transformations:
 - Smooth brightness transformations
 - A criterion of joint alignment:
 - Entropy minimization

Congealing with brightness transforms




Grayscale Entropy Minimization

Frequency of occurrence in image




Some Infant Brains

(thanks to Inder, Warfield, Weisenfeld)




- Pretty well registered (not perfect)
- Pretty bad bias fields

Fourier Basis for Smooth Bias Fields




Results

Original
Images



Bias
Corrected
Images



Assumptions


- Pixels in same location, across images, are independent.
 - When is this not true?
 - Systematic bias fields.
- Pixels in same image are independent, given their location.
 - Clearly not true, but again, doesn't seem to matter.
- Bias fields are truly bandlimited.


Some Other Recent Approaches

- Minimize entropy of intensity distribution in single image
 - Viola (95)
 - Warfield and Weisenfeld extensions (current)
- Wells (95)
 - Use tissue models and maximize likelihood
 - Use Expectation Maximization with unknown tissue type
- Fan (02)
 - Incorporate multiple images from different coils, but same patient.


Potential difficulties with single image method

- If there is a component of the brain that looks like basis set, it will get eliminated.
- Does this occur in practice?
 - Yes!





MRI Bias Removal



Next Application of Congealing: Improving recognition with alignment

Andras Ferencz, Vidit Jain, Jitendra Malik, Gary Huang

Better Alignment for Better Recognition

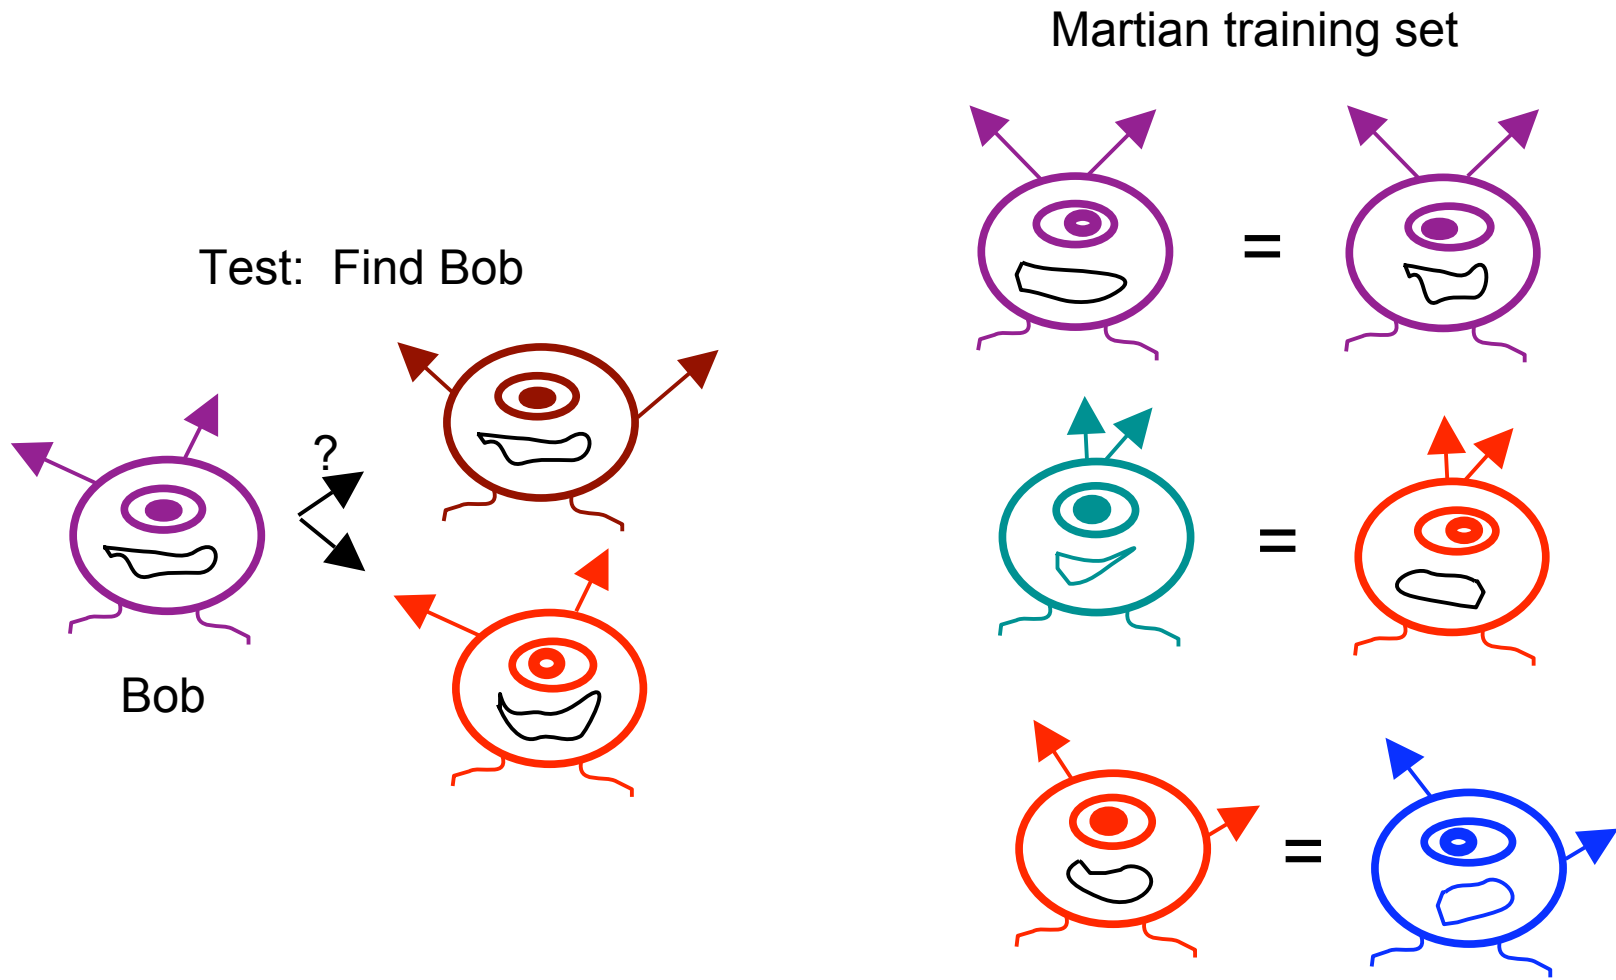
- Hyper-feature recognizer
 - Cars, faces, etc. (*ICCV 2005, NIPS 2005, BMVC 2006*)
- Depends heavily upon alignment:
 - Hyper-features are features conditioned on location and appearance of facial features
- Current alignment:
 - trained separately for each category, requires labeled parts
- Goals:
 - Develop automatic, parts free alignment methods
 - Improve alignments

Key Question for Identification

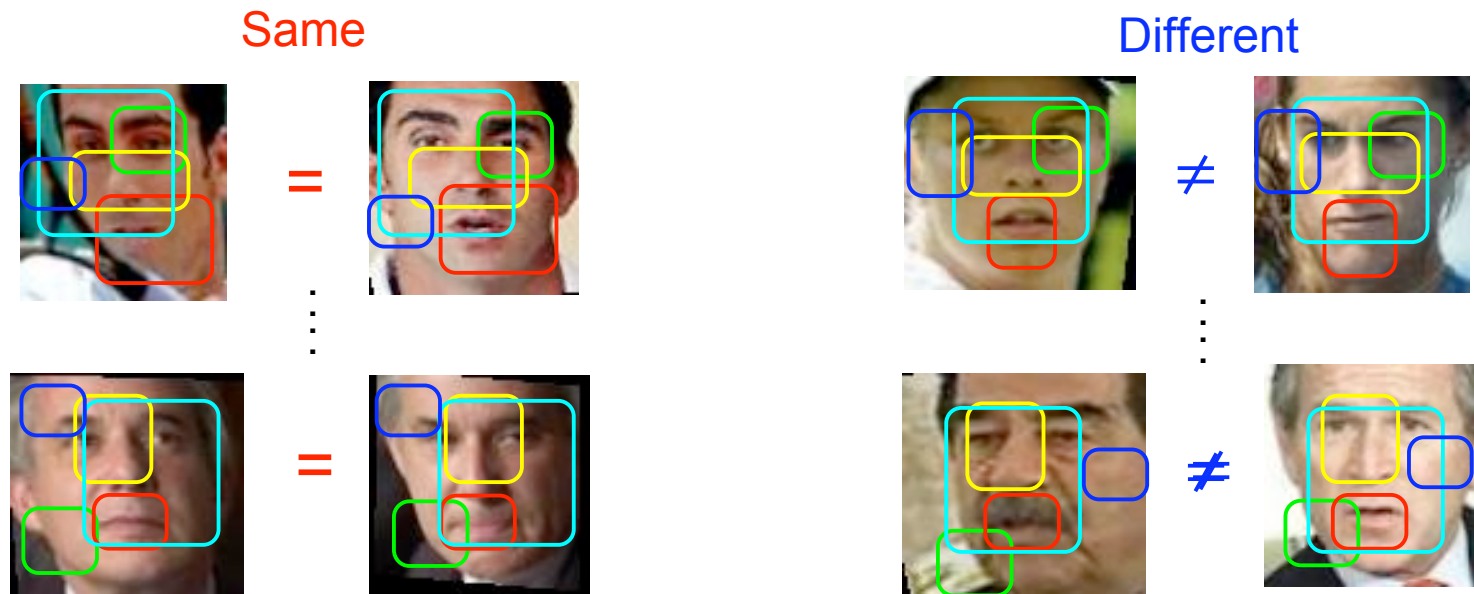
Which features are distinctive and persistent?



Crash Course on Martian Identification



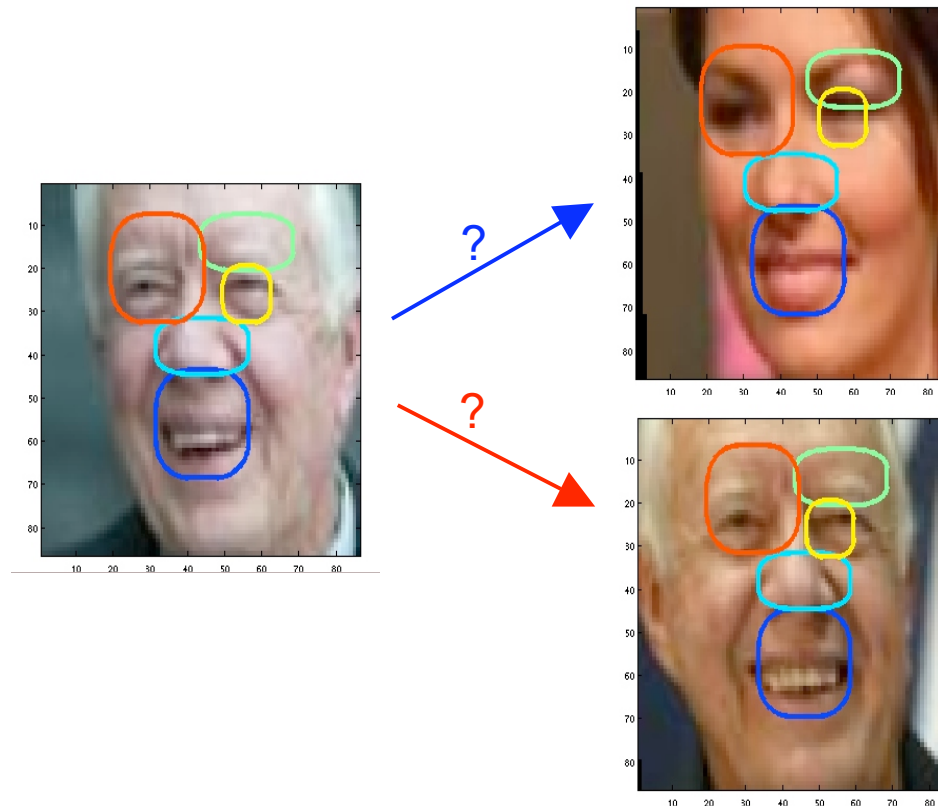
Learning About the Category



Training set: sample many patch pairs.

BETTER ALIGNMENT PRODUCES BETTER MODELS

An Example



Classification Results (correct)

same

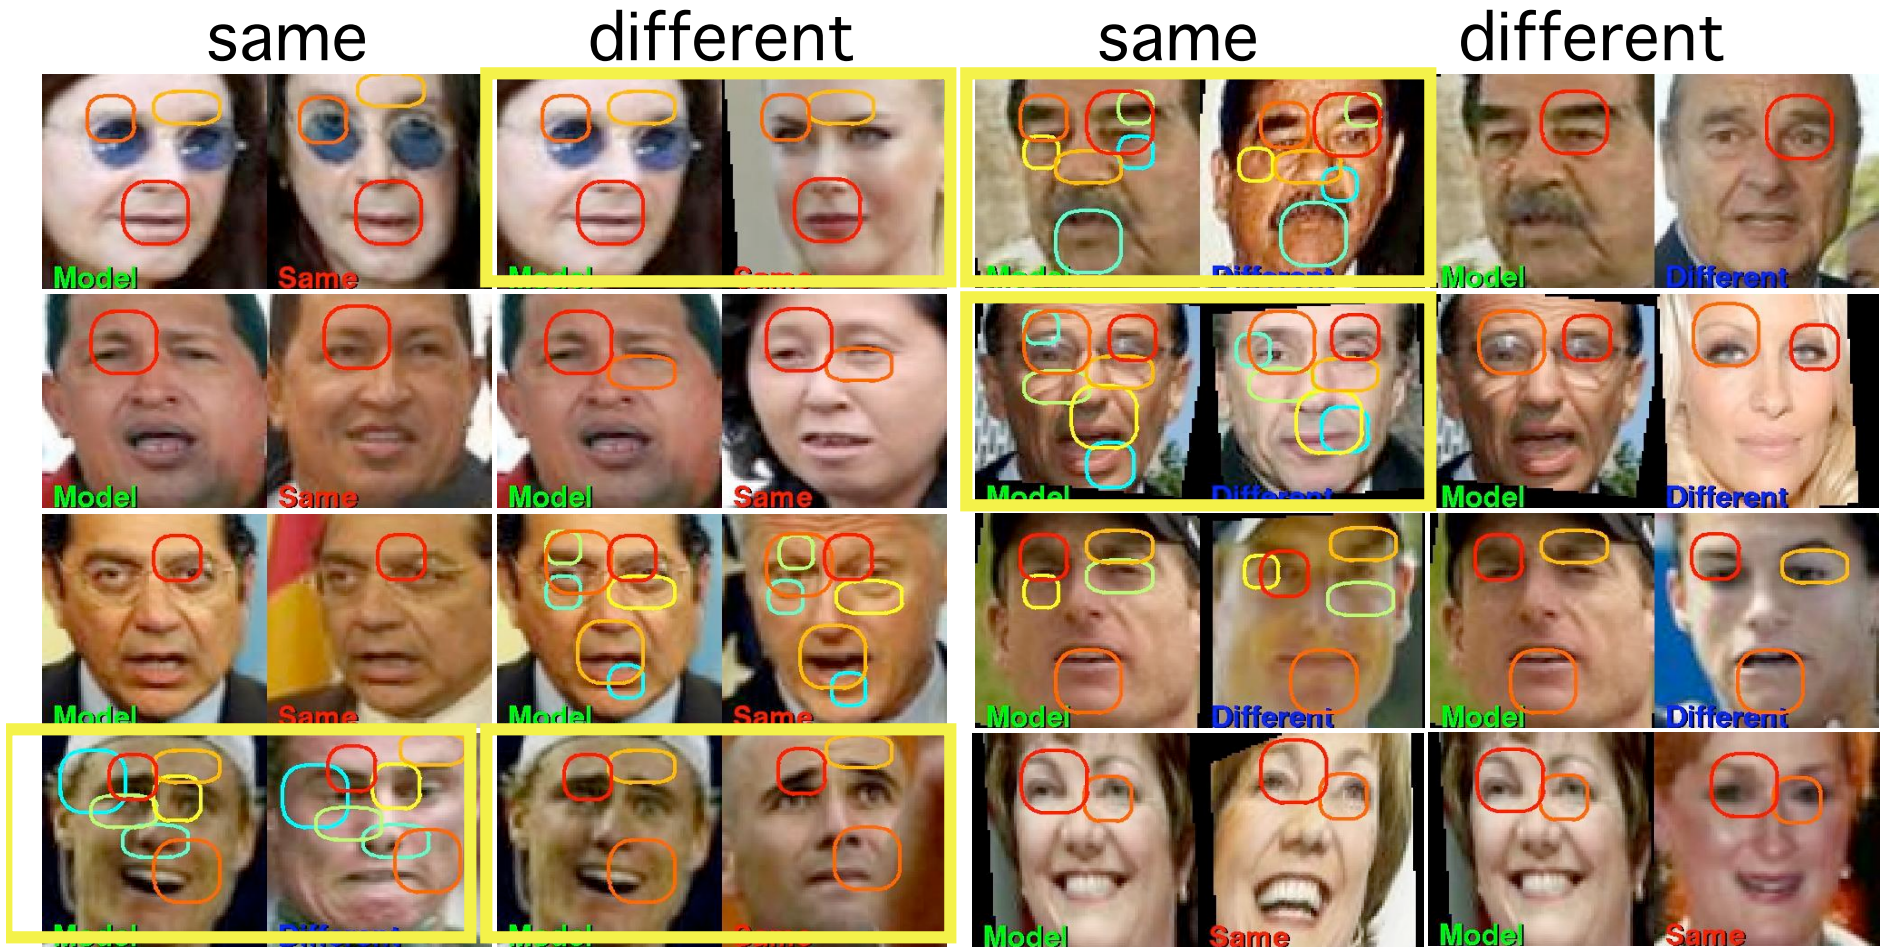
different

same

different



Classification Results (errors)



Congealing Faces

- Challenges:
 - High variability
 - Pixel values do not necessarily have low value when aligned
 - Lighting, hue may foil pixel-based method
- Use higher level-features
 - SIFT (what else?)

Convert face images to arrays of multinomials

- Start with data set of faces
- Compute SIFT at each pixel
- Cluster SIFT vectors (16 clusters)
- At each pixel, form posterior (multinomial) over clusters
- Distribution of pixel stack is mean of multinomial vectors
- Now, do congealing over these multinomial vectors



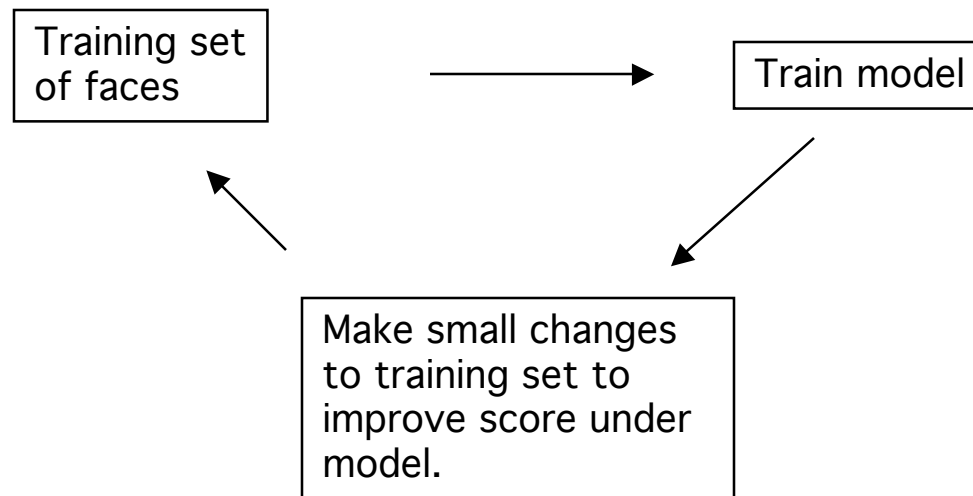






Converting any Model into a Congealer


- Congealing as cascade of independent pixel models.
 - Why not use other models?
 - For example, CMU face detector?



Next Application: Finding Anchor Points for Registration

Just started...


How Do We Register This?




Answer: Anchor Points

- Radiologists use highly robust, highly informative reference points to determine local coordinate systems in the brain.
- Idea: Use congealing to determine what these anchor points should be.

How Do We Register This?



Rerun algorithm on 3D windows



Unaligned input data sets

Aligned input data sets

Data set: 28 T1-weighted MRI; [256x256x124] with (.9375, .9375, 1.5) mm³ voxels

Experiment: 2 levels; 12-param. affine; N = 2500; iter = 150; time = **1209 sec!!**

Properties of Good Anchor Points

- High reliability
 - Low across patient entropy
- High distinctiveness
 - High within patient entropy

Summary

- Remove source of variability
 - MR bias removal
 - MR anatomical alignment
 - ERP signal alignment
 - Better alignment for recognition (hyper-features)
- Model a source of variability
 - Form factorized models (learning from one example)
- Define points of high saliency and repeatability (anchor points) for difficult registration problems