

Mentoring
Undergraduates in
Applied Mathematics
and Computer Science
Research

Fred Park

Whittier College

AMS Sectional Meeting, UC Riverside

November 9th, 2019

My Background

Undergrad in Research at UCLA with 3 students

Team lead at UCI Interdisciplinary Computational Applied Math (iCAMP) REU for 2 years

Undergraduate Research at various levels at Whittier College
e.g. fellowships, senior seminar, or otherwise

Problem Formulation

Coming up with an interesting problem is the first step

Main thing I found

Problem should be :

- interesting to both the mentee and mentor
- problem should challenge both
- problem should be engaging to both

Problem Formulation

Problem should be :

- interesting to both the mentee and mentor
- problem should challenge both
- problem should be engaging to both

example 1: student interested in illusory contours → match

example 2: student interested in modeling drums in music.

me = interested in ML to mine patterns in music → no match.

Don't be afraid to suggest a different direction.

Goal is to find a project that interests the both of you above all!

Problem Formulation

keep a list of potential undergrad research problems either of your own curiosity or from your research.

flexibility is usually key as research tends to different directions quickly.

I know I'm just stating the obvious! :)

Case study 1: Image Segmentation

What is Segmentation:

Computer Vision Task: find the boundaries of salient regions in an image





Image Segmentation: Motivation

Recall the Los Angeles Riots of 1992



Image Segmentation: Some Motivation

High Point of Riots: Reginald Denny Beaten Mercilessly on Nat'l. TV



Public Outrage!

Perpetrators at large!

Calculus Based Image Processing Used to Enhance Footage

Cognitech and UCLA Image Processing Group Help LAPD

THE WALL STREET JOURNAL.

California Company Uses Calculus to Pin The Crime on the Criminal Who Did It

By ROBERT LANGRETH


Staff Reporter of THE WALL STREET JOURNAL

The two men were on trial for murder. Convictions might have been easy: A gas station security camera had filmed the whole tussle, culminating in fatal gunshots. But the videotape was so blurry that no one could really tell who attacked whom. The two argued self-defense, and the Los Angeles County jury hung.

So local detectives turned to Cognitech Inc., a tiny company armed with a powerful new technique for enhancing fuzzy images. Cognitech's improved video clearly showed the suspects pinning the victim face down against the ground and firing into his skull. Both defendants eventually pleaded guilty.

In the past two years, analyzing crime and accident videotapes has blossomed into a full-time business for Cognitech, based in Santa Monica, Calif. It is among a handful of companies applying sophisticated mathematics to clearing up crime and accident videotapes.

Before these companies existed, police trying to enhance poor videos had to buy commercial "Photoshop" software, which generally processes one frame at a time and is limited to simple operations such as improving contrast. Or they could send their



This computer-generated image is the first step in a process that allowed Cognitech to identify a tattoo (circled) on the arm of Reginald Denny's attacker

frame and the nature of distortions caused by poor focus, atmospheric refraction, electronic noise, and

work is done on computer workstations at employees' desktops.

On a typical day, Mr. Rudin prowls the office looking over employees' shoulders as they analyze videotapes, asking questions and making suggestions. For particularly stubborn videotapes, the indefatigable Mr. Rudin often stays late into the night adapting computer programs to do that type of image better.

Cognitech isn't alone in the expanding video-enhancement field. Another small company, Trec Inc. in Huntsville, Ala. sells software for enhancing videotapes to the FBI and other law enforcement agencies. And Aerospace Corp., a nonprofit military-research agency, recently started a small unit to analyze crime videotapes. It now handles a couple of dozen cases per year.

Neither the FBI nor Trec will comment on their enhancement techniques. Aerospace, for its part, says a variety of standard mathematical methods for improving images serve it just fine. "Standard image enhancement is a whopping field," agrees Massachusetts Institute of Technology electrical engineer Alan Willsky. He adds that "the jury's still out on the overall impact" of Cognitech's method.

In any case, all sides expect their caseloads to increase as more businesses buy sophisticated video-

Los Angeles Times

LOCAL U.S. WORLD BUSINESS SPORTS ENTERTAINMENT HEALTH LIVING TRAVEL OPINION

Search

GO

MONEY & CO.

TECHNOLOGY

PERSONAL FINANCE

SMALL BUSINESS

COMPANY TOWN

JOBS

REAL ESTATE

CARS

YOU ARE HERE: LAT Home → Collections → Photographs

The Region : SOUTHERN CALIFORNIA ENTERPRISE : Cognitech Thinks It's Got a Better Forensic Tool : The firm uses complex math in video image-enhancing technology that helps in finding suspects.

f Recommend



0



0

Tweet

Submit


+1

September 05, 1994 | KAREN KAPLAN | TIMES STAFF WRITER



It started out as just a speck on a photograph of a man who threw a brick at truck driver Reginald Denny at Florence and Normandie avenues in the opening hours of the 1992 Los Angeles riots.

But when Leonid Rudin subjected it to a complicated computer algorithm and a slew of complex mathematical equations, that speck--originally less than 1/6,000th the size of the total photograph--was revealed to be a rose-shaped tattoo on the arm of the man, later identified in court as Damian Monroe Williams.

Reginald Denny Beating Investigation



Original Image



Perpetrator's tatoo superresolution



Comparison with the suspect's real tatoo

Image Processing Crime Footage

← Original Crime Scene Video

← Tattoo Superresolution

← Comparison w/ Real Tatoo

**Outcome:
All 3 Criminals Convicted!**


Case study 1: Image Segmentation

my prior work --> image segmentation and shape priors


Prior UCI REU → lots of shape modeling

Shape Priors → shape info to aid in tough segmentation problems

Shape Prior Segmentation




With Shape




Without Shape

Case study 1: Image Segmentation

Common interest with student → illusory contours



upside down triangle




bright disk

Case study 1: Image Segmentation

Common interest with student → illusory contours



cash money sign



pacman and square

Case study 1: Image Segmentation

common interest → illusory contours

Goal:

create model for Shape Prior Segmentation & Illusory Contour Capture!

My previous work on shape prior segmentation

→ does not work on Illusory Contour Capture

→ Model would be Geodesic active contours + shape

Application: Active Contours

- given an image $f : \Omega \rightarrow \mathfrak{R}$
- evolve a curve C to detect objects in f
- the curve has to stop on the boundaries of the objects

Initial Curve → Evolutions → Detected



What intrinsic quality of curve can minimize to make it move inward??

Case study 1: Image Segmentation

Map out the project
step by step with hurdles

hurdle 1: Software considerations

Use high level scripting language in initial phase of research

e.g. Matlab, Python, R, etc.

Tutorials online

If new student → work on a mini project

e.g. loading an image (selfie) and doing basic manipulation

Then build up. Smaller steps add up faster to progress than a big one bc less likely to get stuck

Case study 1: Image Segmentation

hurdle 2: code parametric curve and get it to move inward

wait. Must minimize:

$$\inf_C F(C) = \int_0^1 |C'(s)| ds$$

Involves the Calculus of Variations!!

Case study 1: Image Segmentation

hurdle 2: code parametric curve and get it to move inward

wait. Must minimize:

$$\inf_C F(C) = \int_0^1 |C'(s)| ds$$

Involves the Calculus of Variations!!

Solution: 1-2 weeks of reading Peter Olver's notes on this & calculating the Gateaux Derivative of $F(C)$

- Get functional gradient of $F(C)$
- Gradient Descent to minimize $F(C)$
- Curve moves inward to min arc length

Case study 1: Image Segmentation

hurdle 3: how to stop the curve?

Case study 1: Image Segmentation

hurdle 3: how to stop the curve?

edge detector!

Boundary detection: stopping edge-function (external forces)

$$g \geq 0, \quad g \downarrow, \quad \lim_{t \rightarrow \infty} g(t) = 0$$

Example:
$$g(|\nabla u_0|) = \frac{1}{1 + |\nabla G_\sigma * u_0|^p}$$

$g \sim 0$ on edges

$g \sim 1$ in flatter regions



Case study 1: Image Segmentation

hurdle 4: Put it together → GAC Model

Geodesic model (Caselles, Kimmel, Sapiro '95)

$$\inf_C F(C) = 2 \int_0^1 |C'(s)| g(|\nabla I(C(s))|) ds$$

g: edge detector

NRG is 0 at edges → curve doesn't move

point is, we progressively build up

Case study 1: Image Segmentation

hurdle 5: Put it together → GAC Model + Shape

Householder and Park 2019

$$\inf_C F(C) = \int_0^1 |C'(s)| g(|\nabla I(C(s))|) ds + \textit{Shape}$$

Move a curve that stops at edges in image
while enforcing shape

Case study 1: Image Segmentation

Householder and Park 2019

$$\inf_C F(C) = \int_0^1 |C'(s)| g(|\nabla I(C(s))|) ds + Shape$$



(a) Initialization



(b) Evolving




(c) Evolving




(d) Result


Case study 1: Image Segmentation




(a) Initialization



(b) Evolving




(c) Evolving




(d) Result

Case study 1: Image Segmentation




without shape




with shape


Case study 1: Image Segmentation




(a) Initialization



(b) Evolving



(c) Evolving



(d) Result

Case study 1: Image Segmentation

Important Discussion with the Student:

- where does model work. where it doesn't → model limitations
- advantages over other methods
- no method is a be all end all

all have advantages and disadvantages regardless of what author says
(for the most part)

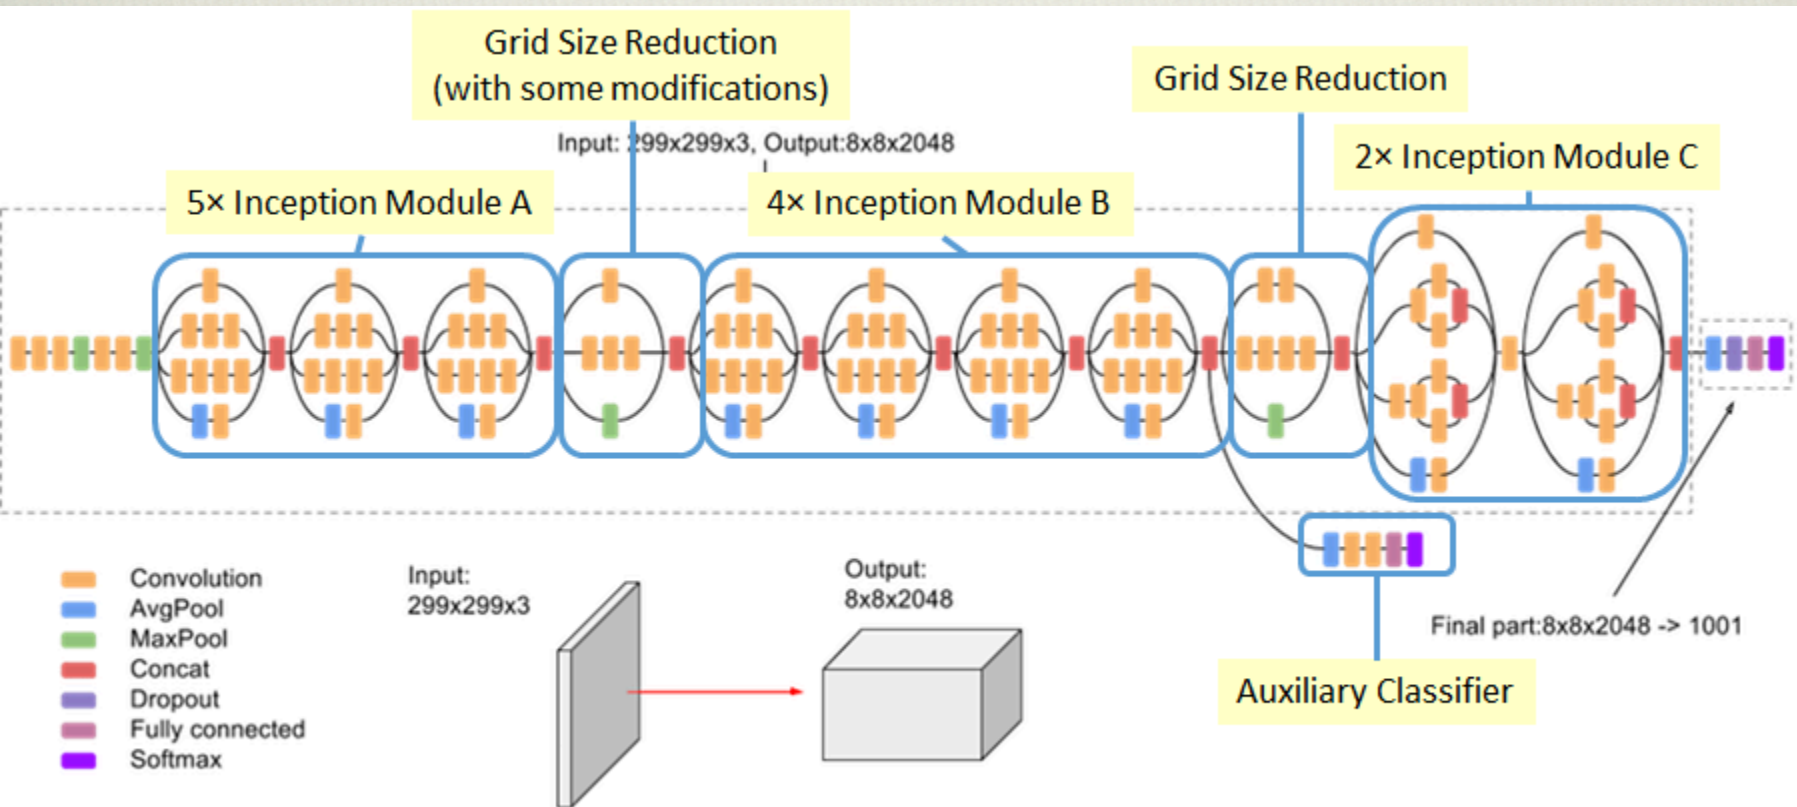
Undergrads really need to see this point!

Case study II: Deep Learning

Motivate by interest

Student was interested in Deep Learning Architectures

e.g. resnets, inception, wide-resnets



Case study II: Deep Learning

Motivate by interest

Student was interested in Deep Learning Architectures

I was doing work on sparsity promoting norms in computer vision

→ try sparsity promoting regularizers in Deep Networks

Case study II: Deep Learning

Motivate by interest

Student was interested in Deep Learning Architectures

I was doing work on sparsity promoting norms in computer vision


→ try sparsity promoting regularizers in Deep Networks

→ Instead of creating new network, systematically prune existing ones

→ Compressed networks have huge advantages!

Case study II: Deep Learning

Example of a compressed network via pruning



Case study II: Deep Learning

Project: sparse neural networks

hurdle 1:

ML research bottleneck = computing power

Cannot do deep learning without GPU acceleration!
(GPU = graphics processing unit)

GPUs are expensive!

Cloud based solutions are even more so!

Solution??

Case study II: Deep Learning

option 1: hardware

- nvidia: GPU research grants
- internal faculty research and development grants
- fellowship funding
- external research grant funding

option 2: cloud computing

- Google: cloud computing research grants. \$5k in credits is easy to obtain and fast
- AWS: cloud comp. grants. More involved but bigger awards. \$15k

We got both → \$20k cloud computing credits

Case study II: Deep Learning

One tip for Cloud Computing Grants:

- Apply for grants for cloud computing credits with google or AWS
- Tie undergrad research into your own research. Makes proposal stronger than if only undergrad research.

For projects:

- extension of previous work
- tweak of previous work
- new direction of previous work
- something new altogether that interests the both of you

Case study II: Deep Learning

Hurdle 2: Coding.

- Use Pytorch (facebook AI)
- Tensorflow/Keras (Google)
- or create a neural net from scratch in matlab or python (no GPU)
- Can also use Scikitlearn for python (no GPU)

Hurdle 2: testing models

Use standard net e.g LeNet5 on standard dataset
e.g. MNIST handwritten digits

LeNet5 trainable on most laptops.

Other sets like CIFAR 10/100, ImageNet require GPU acceleration.

Case study II: Deep Learning

Steep Learning curve: software, gpu, cloud computing etc.

Use github to share code.

block off multi-hour time slots to sit down and code and run training of models with your student.

the importance of coding and figuring things out together in a team environment with ML cannot be stressed enough.

Case study II: Deep Learning

Get standard results on network

add regularization e.g. L1, L2, L1-L2 etc. and see what happens.

Only regularize certain layers. See what happens.

vary the regularization parameter

check sparsity and accuracy

try it on groups/neurons for conv nets.

show effectiveness

again, no one method does it all

advantages and disadvantages in detail

Case study II: Deep Learning

Mathematical:

can prove anything?

convergence?

any links to other papers?

can prove results on a smaller network even if results may or may not hold for larger one

disseminate:

Conference:

undergrad SCURR, College Ugrad Conferences

ML/AI/IP conferences: ICIP, CVPR, etc.

Journal:

SIAM SIURO

Institutional undergrad research journals


Thank You for
Your Attention!



Application: “active contour”

- giving an image $f : \Omega \rightarrow \mathfrak{R}$
- evolve a curve C to detect objects in f
- the curve has to stop on the boundaries of the objects

Initial Curve \longrightarrow Evolutions \longrightarrow Detected




Basic idea in classical active contours

Curve evolution and deformation (internal forces):

$$\text{Min } Length(C) + Area(inside(C))$$

Boundary detection: what is it? What is stopping criteria for curve?

Initial Curve \longrightarrow Evolutions \longrightarrow Detected



Data Fidelity Term

$$\int_{\text{inside}(C)} |f - c_1|^2 dx dy + \int_{\text{outside}(C)} |f - c_2|^2 dx dy$$

where

$c_1 = \text{average}(f)$ inside C

$c_2 = \text{average}(f)$ outside C

Fit > 0

Fit > 0


Fit > 0

Fit ~ 0



Minimize: (Fitting + Regularization)

Fitting not depending on gradient

detects  "contours without gradient"

Chan-Vese (CV) Model

Fitting + Regularization terms (length, area)

$$\inf_{c_1, c_2, C} F(c_1, c_2, C) = \mu \cdot |C| + \nu \cdot \text{Area}(\text{inside}(C)) \\ + \lambda \int_{\text{inside}(C)} |u_0 - c_1|^2 dx dy + \lambda \int_{\text{outside}(C)} |u_0 - c_2|^2 dx dy$$





C = boundary of an open and bounded domain

$|C|$ = the length of the boundary-curve C

- ❖ P.W. Constant Version of Mumford Shah Model
- ❖ Fit constant homogeneous regions while enforcing regularity on boundary of C
- ❖ Active Contours without Edges

Experimental Results

Evolution of C



Advantages

Automatically detects interior contours!

Works very well for concave objects

Robust w.r.t. noise


Detects blurred contours

The initial curve can be placed anywhere!

Allows for automatic change of topologies

Convex Relaxed CV Model

- Partition boundaries in MS & Potts model rep'd by L_0 Norm: $J(u) = \|\nabla u\|_0$
- Gradient distribn's mostly vertical and horizontal in natural images.
- $L_1 - \alpha L_2$ TV norm is better approx'n to L_0 via level lines than L_2 TV.
- α chosen based on gradient distrib's

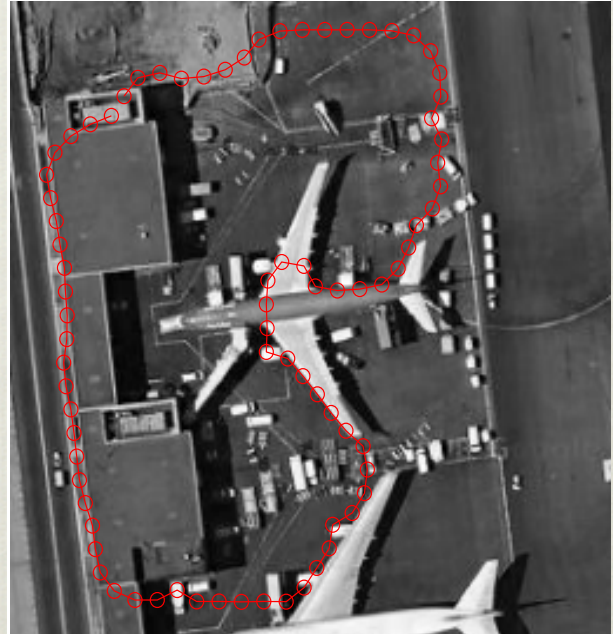




$L_1 - 0.5L_2$ closer to L_0 than:
 L_1 , L_2 , and $L_1 - L_2$

Level lines plot

Shape Prior Segmentation

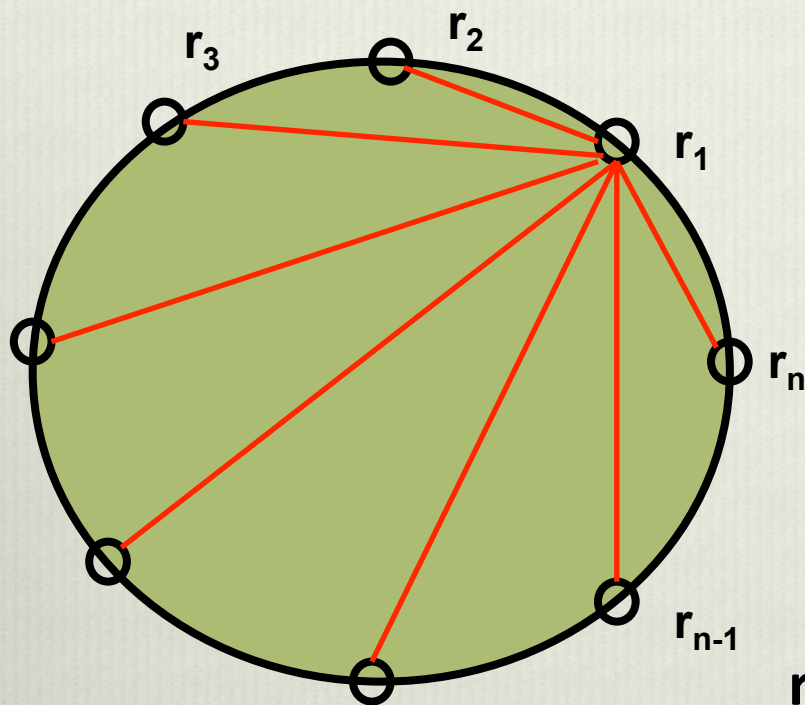
Why are Shape Priors Needed?



Example of MS Segmentation Without Shape!

- Difficult Cases: Clutter, Regions w/ non-uniform intensities, Occluded Objects
- Prior Must be compatible with Segmentation Models i.e. both can be minimized

Cliques Invariant Signature



Motivatation:

Bending Invariant Signatures

Elad and Kimmel 03'

Intervortex Distances:

$$\sum_{i,j} (\|p_i - p_j\|^2 - \|r_i - r_j\|)^2$$

r_i : pts. lying on reference shape

p_i : pts. lying on some evolving contour

Incorporation into:

- Geodesic Active Contours (Snakes)
- Polygonal Implementation of the P.W. Constant MS Model

Cliques Shape Matching Energy

S Polygonal Rep. of a Reference Shape: $S = \{\vec{r}_i\}_{i=1}^N$

$[d_{ij}]$: Symmetric Matrix of Intervertex



Distances $d_{ij} = |\vec{r}_i - \vec{r}_j|$

Σ an Evolving Polygonal Contour: $\Sigma = \{\vec{p}_j\}_{j=1}^N$

Shape Matching Energy:

$$\inf_{\Sigma, s} \left\{ E_c(\vec{p}_1, \vec{p}_2, \dots, \vec{p}_N, s) = \sum_{i,j=1}^N (|\vec{p}_i - \vec{p}_j|^2 - s d_{i,j}^2)^2 \right\}$$

- 's': scale parameter to be min'd over as well
- Invariance to Rigid Motion
- Scale Invariance



Proposed Model

$$E(\Sigma, c_1, c_2, s) = E_{MS} + E_C$$

$$= \min \left\{ \text{Per}(\Sigma) + \lambda \int_{\text{In}(\Sigma)} (c_1 - f)^2 dx dy + \lambda \int_{\text{Out}(\Sigma)} (c_2 - f)^2 dx dy + \alpha \sum_{i,j=1}^N \left(|\vec{p}_i - \vec{p}_j|^2 - sd_{ij}^2 \right)^2 \right\}$$

α : shape strength


Σ : Evolving Polygonal Curve

Polygonal CV Model + Shape!


Best approx of 'f' in L^2 sense taking 2 values c_1 and c_2

While Enforcing Σ matches reference shape

Shape Prior Segmentation




With Shape



Without Shape

Shape Prior Segmentation Example



Learned Ref Shape








Image to be Segmented



Prior Juxt'd on Image




Increasing
shape strength

$\alpha = 0.1$

$\alpha = 0.5$

$\alpha = 1.0$

Final Seg'd. Image!



Disocclusion



$$E(\Sigma, c_1, c_2, s) = E_{MS} + E_C$$

$$= \min \left\{ \text{Per}(\Sigma) + \lambda_R \int_{\text{In}(\Sigma)} (c_1 - f)^2 dx dy + \lambda_R \int_{\text{Out}(\Sigma)} (c_2 - f)^2 dx dy + \alpha \sum_{i,j=1}^N \left(|\vec{p}_i - \vec{p}_j|^2 - s d_{ij}^2 \right)^2 \right\}$$


$$\lambda_R = \begin{cases} 0 & \text{if } \vec{x} \in R \\ \lambda & \text{otherwise.} \end{cases}$$

R: Occluded Region

Don't fit data in R




Disocclusion Example




144x144

Learned reference shape




500x500


Occluded Image



Scale difference



Initial contour




Final evolved
contour



Disoccluded contour
w/ true image


Shape Prior Segmentation: Very Difficult Case



Initial curve




With Prior



No Prior!

- Fuselage matches tarmac.
- 2 completely different intensities in plane

Very difficult segmentation!



prior

Ongoing and Future Work

- ❖ Convergence proof of the DCA algorithm for proposed model
- ❖ Other ways of achieving/exploiting directional sparsity
- ❖ Shape Prior Segmentation: Modeling both occluders and shapes
- ❖ Neural Network \rightarrow semantic segmentation. Interplay between a trained model and a mathematical one. Not a 2 step approach but a synergistic one.
- ❖ Using CNN's for applications to spatially varying blind deconvolution.
- ❖ Stochastic Primal Dual methods for Neural Network Optimization
- ❖ Convex relaxation techniques for NN's. Some work done on quantizing weights. Xin et al. '18

Thank You for
Your Attention!

