

Apologia pro vita sua

D.A. Forsyth, UIUC

MP: Aye. In them days, we'd a' been glad to have the price of a cup o' tea.

GC: A cup ' COLD tea.

EI: Without milk or sugar.

TJ: OR tea!

MP: In a filthy, cracked cup.

EI: We never used to have a cup. We used to have to drink out of a rolled up newspaper.

EI: Right. I had to get up in the morning at ten o'clock at night, half an hour before I went to bed, (pause for laughter), drink a cup of sulphuric acid, work twenty-nine hours a day down mill, and pay mill owner for permission to come to work, and when we got home, our Dad and our mother would kill us, and dance about on our graves singing 'Hallelujah.'

MP: But you try and tell the young people today that... and they won't believe ya'.

History of vision, IMHO

- Particle size distribution analysis
- Grouping
- Color constancy
- Mutual Illumination
- Invariance
- Algebraic surfaces
- Naked people
- Shading
- People
- Words and pictures
- Attributes
- Sentences
- Intrinsic images


PSDA

- No pix, likely didn't work, but I got an MS, so...


Grouping

- Nothing to see here

Color constancy



Color Figure 2. Images of the first Mondriaan, imaged under (clockwise, from top left) white, blue-green, blue, and yellow lights.



Color Figure 7. Outputs of **Crule** for the first Mondriaan derived from images under (clockwise, from top left) white, blue-green, blue, and yellow lights. **Crule** performs well, because these images look similar.




Fig. 10. Cumulative distribution of the statistic described in the text for sixty descriptors computed by **Crule**, operating on images of the first Mondriaan without borders (plotted as a solid line), and on images of the first Mondriaan both with and without borders (plotted as a dashed line). The statistic measures the scatter of the descriptors computed for the chips in images under different lights—the larger the statistic, the wider the scatter, and the poorer the algorithm. Note that adding colored borders to the Mondriaan imaged does not significantly affect the descriptors that **Crule** computes, as expected.




Fig. 1. The convex hull of the gamut obtained by observing 180 color chips under white light. This formed the observed canonical gamut for the experimental work. If t represents the green light for the gamut shown in figure 2, $\Psi(\cdot; t)$ would take this set to a superset of that gamut.






Fig. 3. The convex hull of the gamut observed for a Mondriaan image under blue light. Notice that there are significant differences between this gamut, and those of figures 1 and 2. Crule uses this skewing effect to infer possible illuminants.


Mutual Illumination




17. Image of a black room, containing black objects. The black line indicates the section represented by figure 19.



18. Image of a white room, containing white objects. The black line indicates the section represented by figure 20.



19 Section of image intensity for figure 17.



20. Section of image intensity for figure 18. Notice just how pronounced the effects of the reflexes are.

Invariants

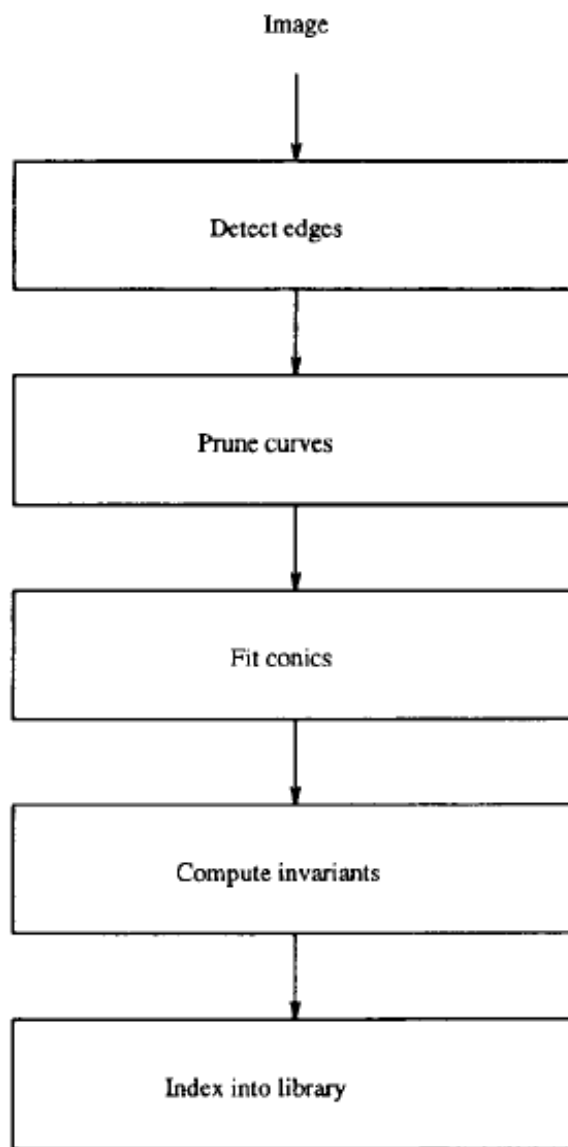


Fig. 4. The Block diagram of a model-based vision system for objects consisting of pairs of general coplanar curves.





Fig. 3. Two views of labels consisting of pairs of ellipses on card in cluttered scenes. These labels are conceptually similar to Nielsen's [34] labels but use different invariants. The labels come from a model base of 15 labels. These labels are positioned at different angles and distances to the camera and are recognized by their invariant descriptors. Model instances are found by detecting edges, pruning edges that are too short or are obviously not conics, fitting conics to all edges, and computing invariant descriptors for every pair of fitted conics. These descriptors are then used to index into the model base. The labels are correctly identified in each case, and their corresponding number is superimposed on the image.

Algebraic surfaces








Fig. 1. The outline and contour generator of a curved object, viewed from a perspective camera.



Note that a similar fact appears as an exercise in (Hartshorne, 1977) (p. 188, ex. 8.4).

Consider the following exact sequence of sheaves associated with the curve:

$$0 \rightarrow \mathcal{I} \rightarrow \mathcal{O}_{P^3} \rightarrow \mathcal{O}_C \rightarrow 0$$

where the symbols have their usual meaning as in, for example, (Hartshorne, 1977). Taking the associated cohomology sequence, and twisting by 1, we obtain the following long exact sequence:

$$0 \rightarrow H^0(P^3, \mathcal{I}(1)) \rightarrow H^0(P^3, \mathcal{O}_{P^3}(1)) \rightarrow H^0(C, \mathcal{O}_C(1)) \rightarrow H^1(P^3, \mathcal{I}(1)) \rightarrow \dots$$

for $0 < i < n$ and for all $j \in \mathbb{Z}$ gives that $H^1(P^3, \mathcal{I}(1))$ is empty, and so we have:

$$0 \rightarrow H^0(P^3, \mathcal{O}_{P^3}(1)) \rightarrow H^0(C, \mathcal{O}_C(1)) \rightarrow 0$$

that is, the two are isomorphic. ■

Finding Naked People

Table 1. Overall classification performance of the system, in various configurations, to 4289 control images and 565 test images. Configuration F is the primary configuration of the grouper, fixed before the experiment was run, which reports a nude present if either a girdle, a limb-segment girdle or a spine group is present, but not if a limb group is present. Other configurations represent various permutations of these reporting conditions; for example, configuration A reports a person present only if girdles are present. There are fewer than 15 cases, because some cases give exactly the same response.

System configuration	Response ratio	Test response	Control response	Test images marked	Control images marked	Recall	Precision
Skin filter	7.0	79.3%	11.3%	448	485	79%	48%
A	10.7	6.7%	0.6%	38	27	7%	58%
B	12.0	26.2%	2.2%	148	94	26%	61%
C	11.8	26.4%	2.2%	149	96	26%	61%
D	9.7	38.6%	4.0%	218	170	39%	56%
E	9.7	38.6%	4.0%	218	171	39%	56%
F (primary)	10.1	42.7%	4.2%	241	182	43%	57%
G	8.5	54.9%	6.5%	310	278	55%	53%
H	8.4	55.9%	6.7%	316	286	56%	52%





Figure 6. Typical control images wrongly classified as containing nudes. These images contain people or skin-colored material (animal skin, wood, bread, off-white walls) and structures which the geometric grouper mistakes for spines or girdles. The grouper is frequently confused by groups of parallel edges, as in the industrial images. Note that regions marked as skin can contain texture at a larger scale than that measured by the texture filter. An ideal system would require that limbs not have texture at the scale of the limb, and would be able to automatically determine an appropriate scale at which to search for limbs.

Words and pictures



Linking words and pictures



Brown, Della Pietra, Della Pietra & Mercer 93; Melamed 01

- In its simplest form, missing variable problem
- Caveats
 - might take a lot of data; symmetries, biases in data create issues

Object recognition as machine translation: Learning a lexicon for a fixed image vocabulary -
P Duygulu, K Barnard, JFG de Freitas, DA Forsyth ECCV 2002



Finding People







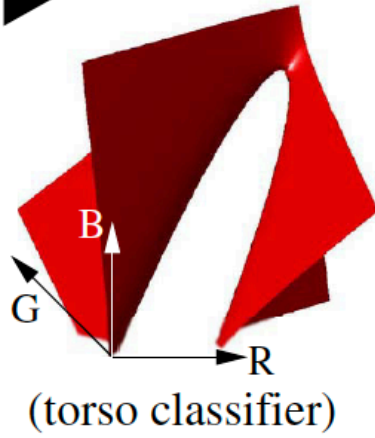
Figure 16. Examples showing representatives for images with two people; these representatives give quite a good guide to the person's configuration (top row); the bottom row shows some cases where the configurations are not represented correctly. Images have been airbrushed so they can be shown salve pudore.




Learn
Limb
Classifiers



Label
Pixels



General
Pose
Pictorial
Structure






Ramanan, Forsyth and Zisserman CVPR05

More Words and Pictures

It was there and we didn't



sky, sun, clouds, sea, waves, birds, water



tree, people, sand, road, stone, statue, temple, sculpture, pillar



tree, birds, snow, fly




sky, water, tree, plane, elephant, herd



mountain, sky, water, clouds, tree



sky, sun, jet, plane




mountain, sky, water, tree, grass, plane, ground, giraffe



water, people, pool, swimmers



tree, people, shadows, road, stone, statue, sculpture, pillar




people, buildings, stone, temple, sculpture, pillar, mosque

It was there and we predicted it

It wasn't and we did

Scene Discovery by Matrix Factorization, N Loeff, A Farhadi, ECCV 2008



Loeff Farhadi Forsyth??

Rooms and (inevitably) Lighting



V. Hedau, D. Hoiem, D.A. Forsyth, "Recovering the layout of cluttered rooms", ICCV 2009

Infer then render

- Kind of circular
- Record
 - good for CGI insertion
 - shakey for scenes




Results



Results







Shown at 2x speed





Describing Objects



Coping with the unfamiliar



What is an object like?

Professor Piehead

and his assistant, TIM



52

Viz comic, issue 101

Describing

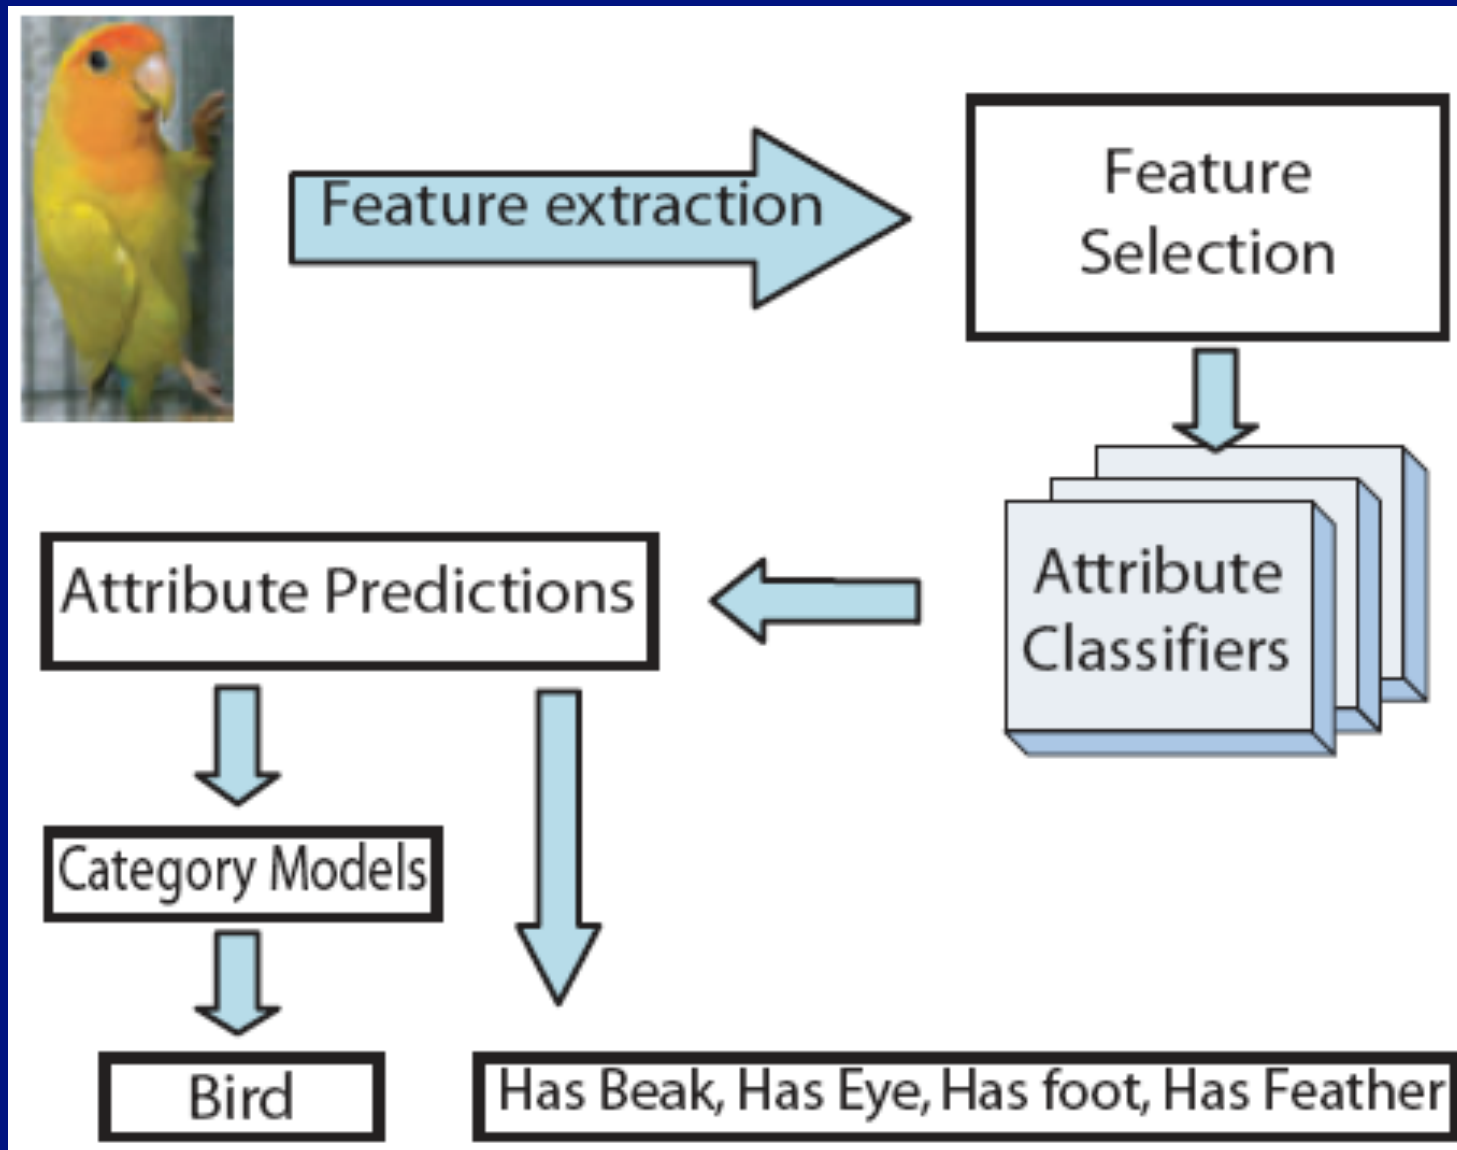


“Man+dog”

“Man in camouflage clothes restraining a scary attack dog with a leash.”

- **Attributes**
 - describe things by properties
 - again, a small “vocabulary” describes many different objects
- **Primitives**
 - a small “vocabulary” makes up many different objects
 - typically, shapes, but that isn’t compulsory
 - eg shared parts; texture encodings; deep learning

General architecture






'is 3D Boxy'
 'is Vert Cylinder'
 'has Window' ~~'has Screen'~~
 'has Row Wind'
~~'has Headlight'~~



'has Hand'
 'has Arm'
~~'has Plastic'~~
 'is Shiny'



'has Head'
 'has Hair'
 'has Face'
~~'has Saddle'~~
 'has Skin'



'has Head'
 'has Torso'
 'has Arm'
 'has Leg'
~~'has Wood'~~




'has Head'
 'has Ear'
 'has Snout'
 'has Nose'
 'has Mouth'




'has Head'
 'has Ear'
 'has Snout'
 'has Mouth'
 'has Leg'



~~'has Furniture Back'~~
~~'has Horn'~~
~~'s Screen'~~
 'has Plastic'
 'is Shiny'



'is 3D Boxy'
 'has Wheel'
 'has Window'
 'is Round'
 'has Torso'



'has Tail'
 'has Snout'
 'has Leg'
~~'has Text'~~
~~'has Plastic'~~



'has Head'
 'has Ear'
 'has Snout'
 'has Leg'
 'has Cloth'



'is Horizontal Cylinder'
~~'has Beak'~~
~~'has Wing'~~
~~'has Side mirror'~~
 'has Metal'



'has Head'
 'has Snout'
 'has Horn'
 'has Torso'
~~'has Arm'~~

Sentences from Images

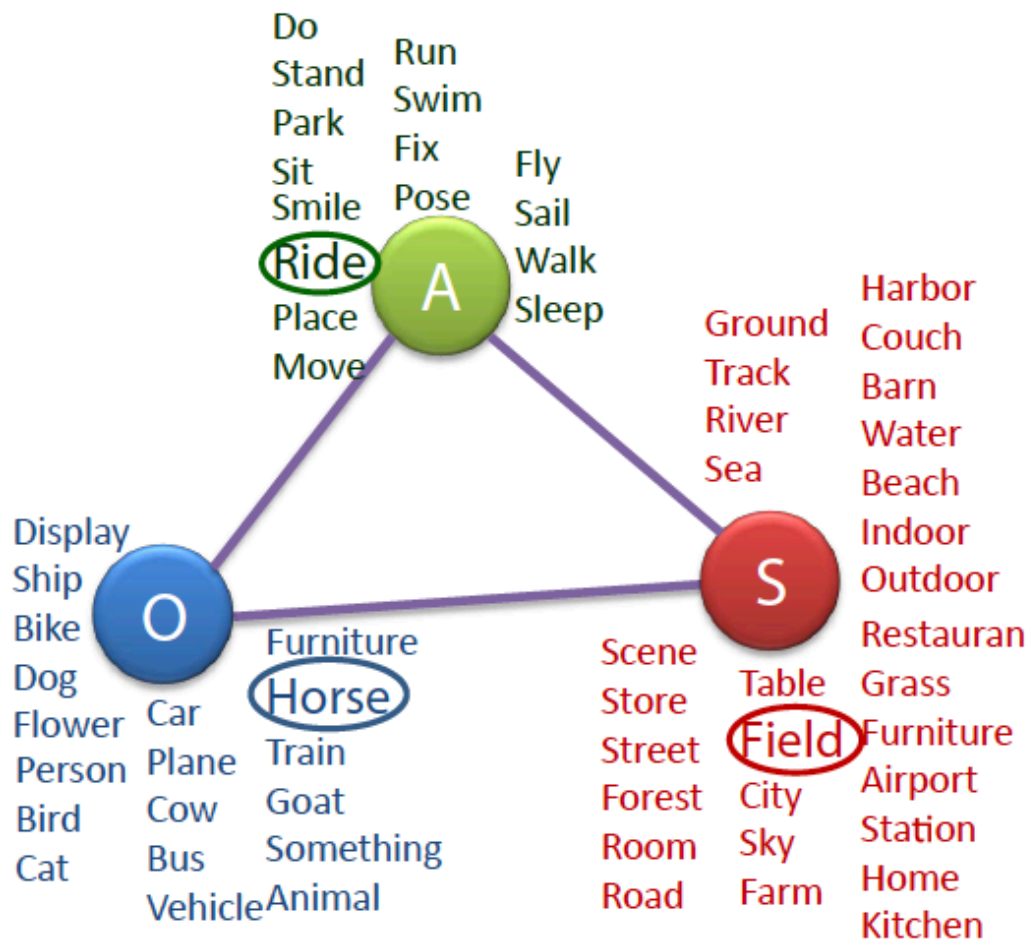


Fig. 2. We represent the space of the meanings by triplets of $\langle \text{object}, \text{action}, \text{scene} \rangle$. This is an MRF. Node potentials are computed by linear combination of scores from several detectors and classifiers. Edge potentials are estimated by frequencies. We have a reasonably sized state space for each of the nodes. The possible values for each nodes are written on the image. “O” stands for the node for the object, “A” for the action, and “S” for scene. Learning involves setting the weights on the node and edge potentials and inference is finding the best triplets given the potentials.

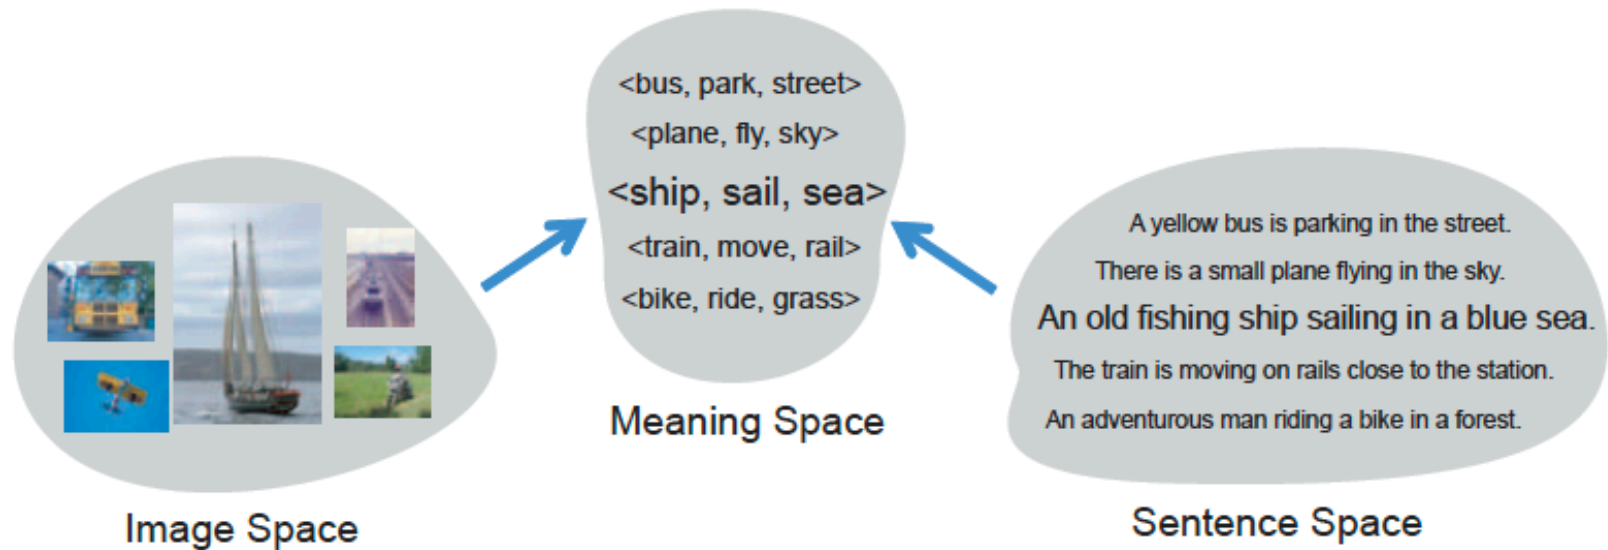


Fig. 1. There is an intermediate space of meaning which has different projections to the space of images and sentences. Once we learn the projections we can generate sentences for images and find images best described by a given sentence.

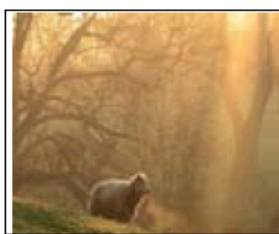
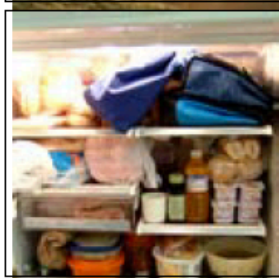
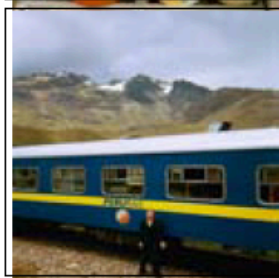
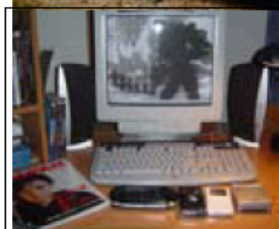

	<p>(pet, sleep, ground) (dog, sleep, ground) (animal, sleep, ground) (animal, stand, ground) (goat, stand, ground)</p>	<p>see something unexpected. Cow in the grassfield. Beautiful scenery surrounds a fluffly sheep. Dog hearing sheep in open terrain. Cattle feeding at a trough.</p>
	<p>(furniture, place, furniture) (furniture, place, room) (furniture, place, home) (bottle, place, table) (display, place, table)</p>	<p>Refrigerator almost empty. Foods and utensils. Eatables in the refrigerator. The inside of a refrigerator apples, cottage cheese, tupperwares and lunch bags. Squash apenny white store with a hand statue, picnic tables in front of the building.</p>
	<p>(transportation, move, track) (bike, ride, track) (transportation, move, road) (pet, sleep, ground) (bike, ride, road)</p>	<p>A man stands next to a train on a cloudy day A backpacker stands beside a green train This is a picture of a man standing next to a green train There are two men standing on a rocky beach, smiling at the camera. This is a person laying down in the grass next to their bike in front of a strange white building.</p>
	<p>(display, place, table) (furniture, place, furniture) (furniture, place, furniture) (bottle, place, table) (furniture, place, home)</p>	<p>This is a lot of technology. Somebody's screensaver of a pumpkin A black laptop is connected to a black Dell monitor This is a dual monitor setup Old school Computer monitor with way to many stickers on it</p>

Fig. 3. Generating sentences for images: We show top five predicted triplets in the middle column and top five predicted sentences in the right column.


Intrinsic Images

Big technical points from the distant past


- Intrinsic images = maps of scene properties




(a) ORIGINAL SCENE




(b) DISTANCE



(c) REFLECTANCE



(d) ORIENTATION (VECTOR)



(e) ILLUMINATION

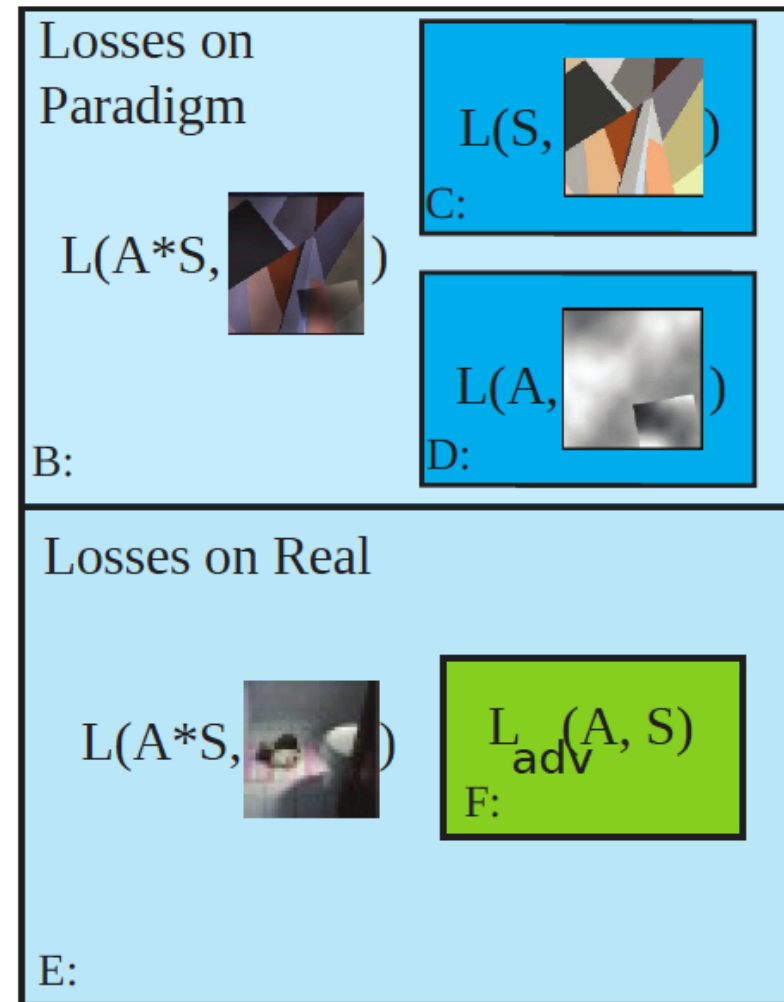
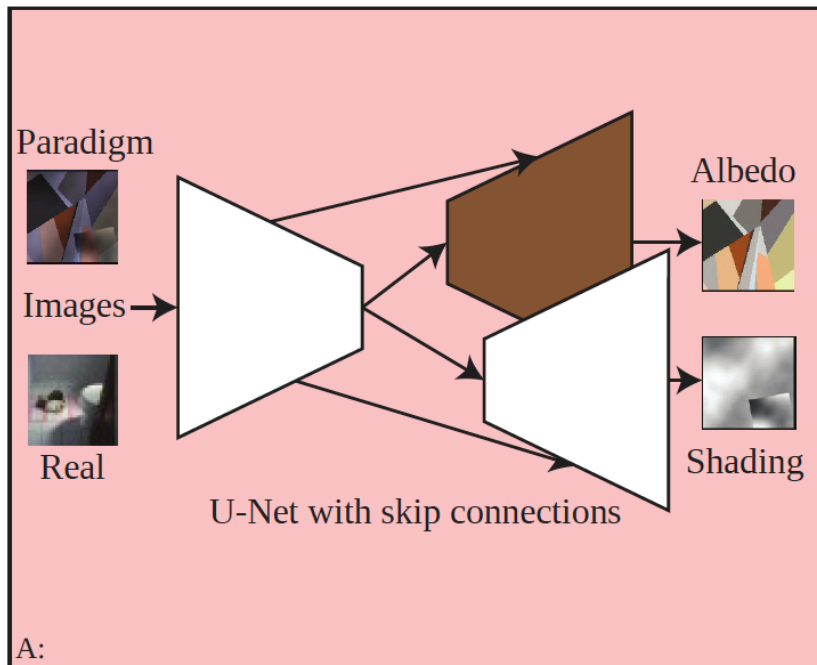
Barrow+Tenenbaum, 1978

Intrinsic images

- Intrinsic
 - shape, and affordances that follow
 - surface properties, and affordances that follow
 - volume properties, and affordances that follow
- What doesn't change when
 - object moves from image to image?
 - light changes?
- Often dumbed down to albedo estimation

Easy losses

- Paradigms should be correctly decomposed
 - with small residual
- Composing decomposed images
 - should have small residual



Nasty problem

Image



- Translate, rotate, scale image
 - albedo for translated (etc) image should be translated albedo
 - shading for translated (etc) image should be translated shading
- But the network doesn't know that...

BR



Rescale



Flip



TL



Model 1



Model 0



Averaging very strongly suppresses error

Image


BBAF



BR




Rescale



Flip



TL



Model 1



Model 0





Categories ▾






English ▾

	Publication	<u>h5-index</u>	<u>h5-median</u>
1.	Nature	<u>376</u>	552
2.	The New England Journal of Medicine	<u>365</u>	639
3.	Science	<u>356</u>	526
4.	The Lancet	<u>301</u>	493
5.	IEEE/CVF Conference on Computer Vision and Pattern Recognition	<u>299</u>	509
6.	Advanced Materials	<u>273</u>	369
7.	Nature Communications	<u>273</u>	366
8.	Cell	<u>269</u>	417
9.	Chemical Reviews	<u>267</u>	438
10.	Chemical Society reviews	<u>240</u>	368

Top Computer Science Conferences

Ranking is based on *Conference H5-index* ≥ 12 provided by Google Scholar Metrics

Show Due only All Categories
All Countries Search by keyword

Rank	Publisher	Conference Details	H5-index	Impact Score
1	 IEEE	CVPR : IEEE/CVF Conference on Computer Vision and Pattern Recognition Jun 21, 2021 - Jun 24, 2021 - Nashville , United States http://cvpr2021.thecvf.com/	299	51.98
2	 NeurIPS	NeurIPS : Neural Information Processing Systems (NIPS) Dec 6, 2021 - Dec 14, 2021 - Online , Online https://nips.cc/	198	33.49
3	 IEEE	ICCV : IEEE/CVF International Conference on Computer Vision Oct 11, 2021 - Oct 17, 2021 - Montreal , Canada http://iccv2021.thecvf.com/home	176	32.51
4	 Springer	ECCV : European Conference on Computer Vision Oct 11, 2021 - Oct 17, 2021 - Montreal , Canada http://iccv2021.thecvf.com/	144	25.91
5	 AAAI	AAAI : AAAI Conference on Artificial Intelligence Feb 2, 2021 - Feb 9, 2021 - Vancouver , Canada https://aaai.org/Conferences/AAAI-21/	126	25.57

Vision

Vision

Vision

State of our discipline - I

- Fantastic successes
 - “selling shit” (© G. Gkioxari)
 - ex-students earn much more than they used to
 - huge conferences offer rich intellectual experiences
 - startups raising absurd sums of money
 - regular complete revisions of what is known
- Impacting people way outside our community
- Many very deep problems remain open
 - but may be open to attack

State of our discipline - II

- Utter intellectual disorder
 - We can do things, but mostly don't know why
 - “It works better this way”
- Fanatical adherence to experiment
 - mostly hilariously poorly conducted
 - no error bars, repeated trials, no significance on ablations, etc, etc
- Minimal collective values or vision

Mood

- Sadness
 - whatever worked wasn't mine
- Willful ignorance
 - "I don't read papers before 2015" !
- Anxiety
 - will I get a job/be promoted/etc
- Fear
 - Risk-taking may be very costly
- Exclusion
 - "I'm not <X> so I'm not welcome"
- Laughable deference to authority
 - "**** said **** on twitter!" (so it must be true)
- Historical perspective
 - 99% of the old vision literature was crap



Mood

- Sadness
 - whatever worked wasn't mine
- Willful ignorance
 - "I don't read papers before 2015" !
- Anxiety
 - will I get a job/be promoted/etc
- Fear
 - Risk-taking may be very costly
- Exclusion
 - "I'm not <X> so I'm not welcome"
- Laughable deference to authority
 - "**** said **** on twitter!" (so it must be true)
- Historical perspective
 - 99% of the old vision literature was crap



Big reveal: 99% of our literature is too, and there's more of it!

CVPRWN: Non coping strategies



<bad things, go far far away>



<to the point of nausea>



<virtue by work>

CVPR workshop

- Scholars and big models: how can academics adapt?
 - Organizers: Anand Bhattad, Unnat Jain, Angjoo Kanazawa and Sara Beery
- Numerous discussants: I kept summary notes
 - CVPRWN

CVPRWN: Our field is very hard

- Hardness measure:
 - Smaller half-life for knowledge -> harder field
- Ours was hard, now outrageously hard
 - but there are new effects
 - ignore some ideas, and they'll go away quickly
- Challenges:
 - by the time you've picked up new stuff, it's irrelevant
 - someone has already done what you're doing
- Responses:
 - finely detailed differentiation, often absurd
 - excessive framing and citation practices
 - massive exaggerations of significance - everything is disruptive!


CVPRWN: What will happen

- A lot needs to be done
- It won't be like the stuff we're used to
 - Robotics
 - What does vision do?
 - Questions keep changing
 - Properly conducted experiments
 - Speed and power
 - Ethics
 - etc...
- A great decoupling between Eng AI and Science AI?


A Scary Movie



Depth (omnimap, current best depth est)



Normal (omnimap, current best normal est)




Points:

- Vision is flourishing
- Intrinsic images are scene maps
 - predictors learned with regression from labelled data
 - depth, normal, albedo - but albedo is not like normal, depth
- Lighting messes up intrinsic image estimates
 - MIT dataset gives suggestive evidence that relighting could help
- StyleGAN yields a procedure to relight in general
- Generative models are very well informed
 - and can be made to produce depth, normal, albedo!
- but have some fascinating gaps in knowledge

but not all...


Original Image




Inpainting Mask



Masked Image




Inpainted Image



Iffy projective geometry



Iffy projective geometry



Iffy projective geometry




Synthesizing clothing








Tuning-based Methods



Reference-based Methods





Object
Moving

