

What do Image Generators Know?

D.A. Forsyth, UIUC

with Anand Bhattad, TTIC; Derek Hoiem, UIUC;

Lana Lazebnik, UIUC; Yuxiong Wang, UIUC; Shenlong Wang, UIUC


Sara Aghajanzadeh, UIUC; Daniel McKee, UIUC; Yuanyi Zhong, UIUC;

James Soole, UIUC

Structure

- Vision is flourishing
- What does vision do?
- Speculation tools are limited
- Big new player: Generative models
 - They know a lot of unexpected stuff
 - They're dumb about a lot of other stuff
- Agenda: Fix generative models and stop regressing things
 - Understand mistakes
 - Build sensitive evaluation procedures
 - Understand interactions between model architectures and weird behavior
 - “inductive bias” (?)

Vision group at Illinois



- **David Forsyth**
 - Fulton Watson Copp chair; Marr prize, 1993; 2 ex students with Marr prizes; IEEE Tech. Achievement, Fellow; ACM Fellow; EIC IEEE TPAMI; 4 patents



- **Jim Rehg**
 - HCESC director; multiple famous ex-students, best paper awards; 26 patents



- **Derek Hoiem**
 - best paper, CVPR 2006; ACM Doctoral Dissertation honorable mention; Sloan Fellow; PAMI-TC Young Researcher; major startup



- **Lana Lazebnik**
 - Microsoft Faculty Fellow; Sloan Fellow; Koenderink Prize (2016); EIC IJCV



- **Alex Schwing**
 - Visual learning, segmentation and GAN models



Saurabh Gupta

- Linking visual sensing to motion




Liangyan Gui

- Understanding human movement



Shenlong Wang



- Simulation and sensing for autonomous vehicles



Yuxiong Wang

- Learning to detect and classify with very little data

Vision group



Well-known ex-students:
Lana Lazebnik (UIUC)
Tamara Berg (UNC)
Pinar Duygulu (Hacettepe U.)
Ian Endres
Ali Farhadi (UW)
Varsha Hedau
Nazli Ikizler (Hacettepe U.)
Brett Jones
Kevin Karsch
Zicheng Liao
Deva Ramanan (CMU)
Raj Sodhi
Gang Wang (now Alibaba)
Amin Sadeghi
Zicheng Liao (Zhejiang U.)

The New Computer Vision
D.A. Forsyth

Likely about 2024 ←

Cover design opportunity!


WRONG!
(maybe next year)

Startups:

Lightform
Revery.ai
Reconstruct
Depix



- UIUC has autonomous vehicles class
- WITH ACTUAL VEHICLE!



Class project: brake for pedestrian



Top publications

Categories		English	
	Publication	<u>h5-index</u>	<u>h5-median</u>
1.	Nature	<u>488</u>	745
2.	IEEE/CVF Conference on Computer Vision and Pattern Recognition	<u>440</u>	689
3.	The New England Journal of Medicine	<u>434</u>	897
4.	Science	<u>409</u>	633
5.	Nature Communications	<u>375</u>	492
6.	The Lancet	<u>368</u>	678
7.	Neural Information Processing Systems	<u>337</u>	614
8.	Advanced Materials	<u>327</u>	420
9.	Cell	<u>320</u>	482
10.	International Conference on Learning Representations	<u>304</u>	584
11.	JAMA	<u>298</u>	498
12.	Science of The Total Environment	<u>297</u>	436
13.	IEEE/CVF International Conference on Computer Vision	<u>291</u>	484
14.	Angewandte Chemie International Edition	<u>281</u>	361
15.	Nature Medicine	<u>274</u>	474
16.	Journal of Cleaner Production	<u>272</u>	359
17.	International Conference on Machine Learning	<u>268</u>	424
18.	Chemical Reviews	<u>267</u>	461
19.	Proceedings of the National Academy of Sciences	<u>267</u>	405
20.	IEEE Access	<u>266</u>	364

.com

[Home](#) / Best Conferences - Computer Science

Best Computer Science Conferences

The ranking of leading conferences for Computer Science was prepared by Research.com, one of the primary websites for Computer Science research offering accurate data on scientific output since 2014.

The spot on the list is based on Impact Score data collected on 21-11-2023. It was based on a careful analysis of as much as 3,040 conference profiles and websites. [Show more](#)

All resear

All publi

All countr

Paper submission open


Rank Conference Details Impact Score

1 Rank	 Computer Vision and Pattern Recognition 18-06-2023 - 22-06-2023 - Vancouver	62.60
2 Rank	 International Conference on Computer Vision 11-10-2021 - 11-10-2021 - Montreal	43.80
3 Rank	 Neural Information Processing Systems 12-12-2023 - 14-12-2023 - New Orleans	40.60

4 Rank	 International Conference on Learning Representations 01-05-2023 - 05-05-2023 - Kigali	38.10
5 Rank	 AAAI Conference on Artificial Intelligence 07-02-2023 - 14-02-2023 - Washington DC	35.00
6 Rank	 International Conference on Machine Learning 17-07-2022 - 23-07-2022 - Baltimore	30.10
7 Rank	 European Conference on Computer Vision 29-09-2024 - 04-10-2024 - Milan	29.60
8 Rank	 Meeting of the Association for Computational Linguistics 22-05-2022 - 27-05-2022 - Dublin	28.90
9 Rank	 Empirical Methods in Natural Language Processing 06-12-2023 - 10-12-2023 - Singapore City	21.90
10 Rank	 International Joint Conference on Artificial Intelligence 23-07-2022 - 29-07-2022 - Vienna	21.10

Computer Vision Today

- **Vision in general is flourishing**
 - Greatest strengths
 - using data to construct effective, highly polished features
 - increasingly focused on discrimination and regression for huge datasets
 - a money game
 - vast flood of “important” papers



What should a young researcher do?

Find career making problems and solve them

Attributes of career making problems:

Change how others think and what they do

Broad scope

Technically hard, and require new insights

Unlike what everyone else is doing

Molesworth, you must learn the value of spiritual things!

G. Willans and R. Searle

“Molesworth”,

Penguin Classics reprint, 2000

Structure

- Vision is flourishing
- What does vision do?
- Speculation tools are limited
- Big new player: Generative models
 - They know a lot of unexpected stuff
 - They're dumb about a lot of other stuff
- Agenda: Fix generative models and stop regressing things
 - Understand mistakes
 - Build sensitive evaluation procedures
 - Understand interactions between model architectures and weird behavior
 - “inductive bias” (?)

What does vision do?

- Describe the world
 - to support:
 - movement, interaction, decision
- Predict the future
 - you can see what you might hit before you hit it
- Allow speculation
 - what would the world be like if:
 - I moved?
 - This moved?
 - how could I do this better?
 - etc.

Many aspects of classification and detection are largely closed.

Open questions:
what taxonomy?
change taxonomy with goal?
efficiency, etc.

Similar activities, but not the same thing.
Notice how extraordinary it
is that these problems seem do-able

What we should be able to do...

- Drag something around in its own scene
- Select something in one image/video, place in another
- Adjust scene by interaction with image/video
- Model:
 - perfect inverse graphics+interaction+perfect render
 - but this isn't likely the best mechanism

Speculation tools

- Tools

Remarkably well understood

Confused, but promising

Confused

Highly confusing

Highly confusing

- Object removal
- Relighting
- Object insertion
- Move camera, viewpoint
- Scene editing

Confusion: SOTA depends
A LOT on the
version of the task

Technologies

- Cut and paste
 - Works a bit, but sometimes not
- Inverse graphics
 - Spectacular at some cases, not great at others
- Image generators
 - New kid on the block
 - Erratic, but extremely promising

Cut and paste



Cut and paste



Cut and paste



Lalonde et al, 07

Illumination issues: bad match




Lalonde et al, 07


Lalonde et al, 07

“Inverse Graphics” Strategy


- Inverse graphics
 - Recover scene intrinsics from image
- Then forward graphics
 - Render scene using these intrinsics




(a) ORIGINAL SCENE




(b) DISTANCE



(c) REFLECTANCE



(d) ORIENTATION (VECTOR)



(e) ILLUMINATION

Barrow+Tenenbaum, 1978

Inverse graphics by regression

We're really good at this

- Make image-like things from images

- depth map from image
 - depth to closest surface at each pixel
- normal map from image
 - normal at viewed surface at each pixel
- albedo map from image
 - albedo at viewed surface at each pixel

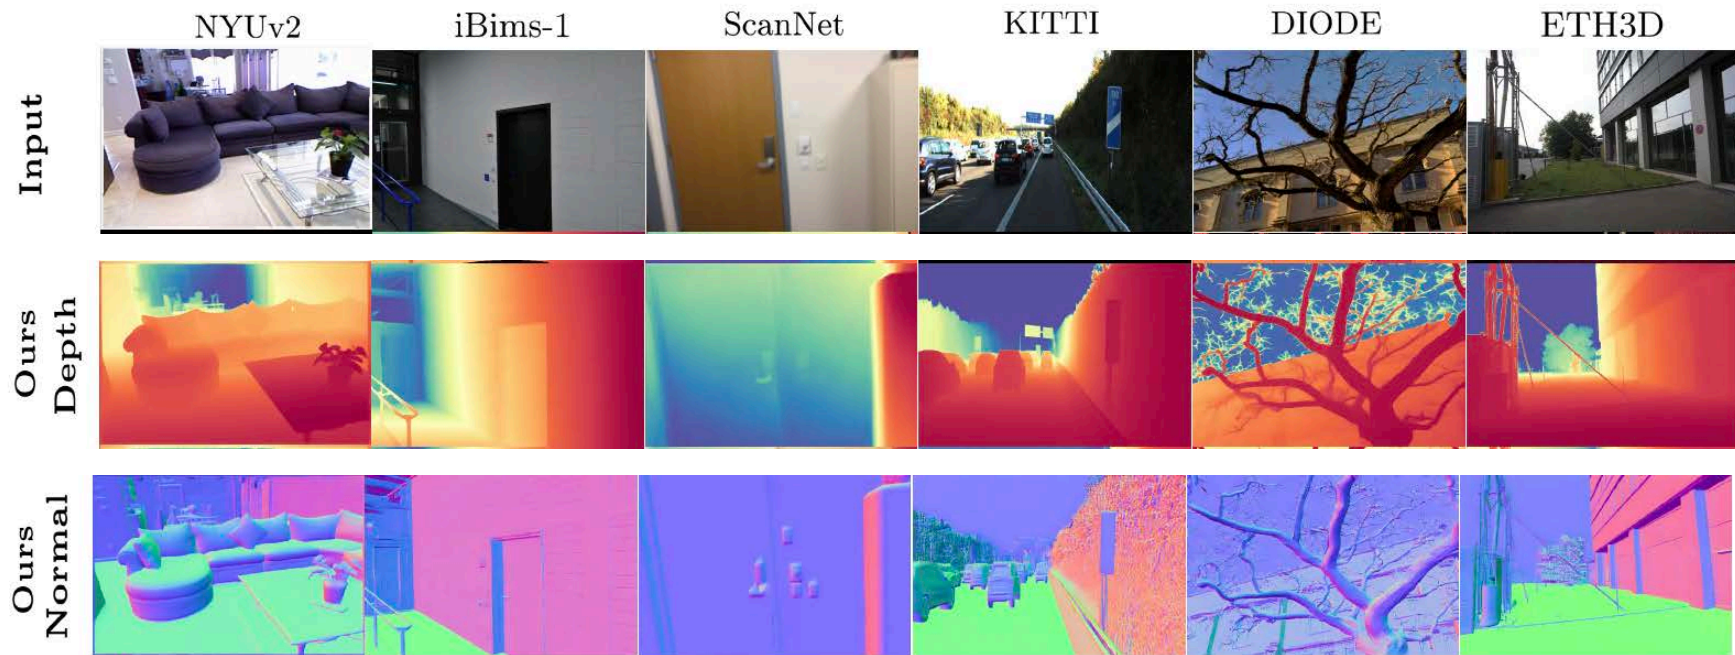
Highly developed

- Procedure

- encode image
 - with various methods!
- decode representation to (depth, normal, albedo)

Important recent trick: Use diffusion methods (below) as regression engine

(Likely) Best current single image intrinsics



(Likely) Best single image albedo



Input

Li et al. [27]

Zhu et al. [48]

Ours - Mean

GT

2020

2022

2024

Photograph



Relighted photograph



Yeh et al 2022: PhotoScene: Photorealistic Material and Lighting Transfer for Indoor Scenes

Inverse graphics problems

Photograph



Relighted photograph



Yeh et al 2022: PhotoScene: Photorealistic Material and Lighting Transfer for Indoor Scenes


Inverse graphics problems - II



Image



Image




Li et al 2022

Error control is really hard!

Notice

- One lighting suppresses details, the other picks them out
 - geometric errors inevitable



The difficulties

- Errors compound
 - and rendering procedures are elaborate, so error metrics are hard to adjust
 - differentiable rendering may help
 - differentiating global illumination is hard
- Important intrinsics may be missing from our model
 - Physical models aren't built to explain image variance

Structure

- Vision is flourishing
- What does vision do?
- Speculation tools are limited
- **Big new player: Generative models**
 - They know a lot of unexpected stuff
 - They're dumb about a lot of other stuff
- **Agenda: Fix generative models and stop regressing things**
 - Understand mistakes
 - Build sensitive evaluation procedures
 - Understand interactions between model architectures and weird behavior
 - “inductive bias” (?)

Images from random numbers

- Astonishing, recent fact:
 - We can generate images from random numbers
 - and conditioning information
- Two major procedures:
 - Generative adversarial network
 - Diffusion methods

In Limine: I'm ignoring text - image links to talk about scene representation and understanding

Applications: make commercial art cheaper



and political discourse nastier



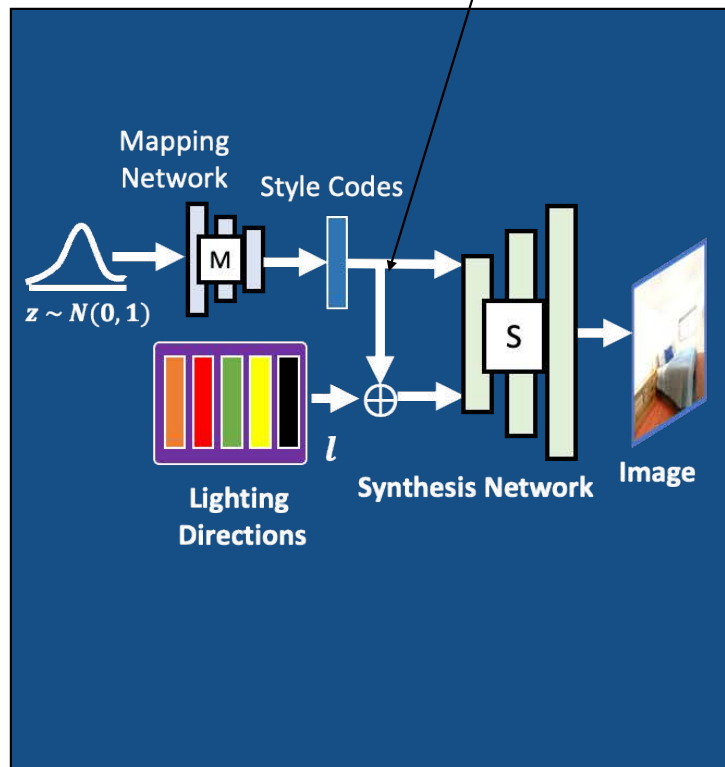
Structure

- Vision is flourishing
- What does vision do?
- Speculation tools are limited
- Big new player: Generative models
 - They know a lot of unexpected stuff
 - They're dumb about a lot of other stuff
- Agenda: Fix generative models and stop regressing things
 - Understand mistakes
 - Build sensitive evaluation procedures
 - Understand interactions between model architectures and weird behavior
 - “inductive bias” (?)

Why bother with image generators?

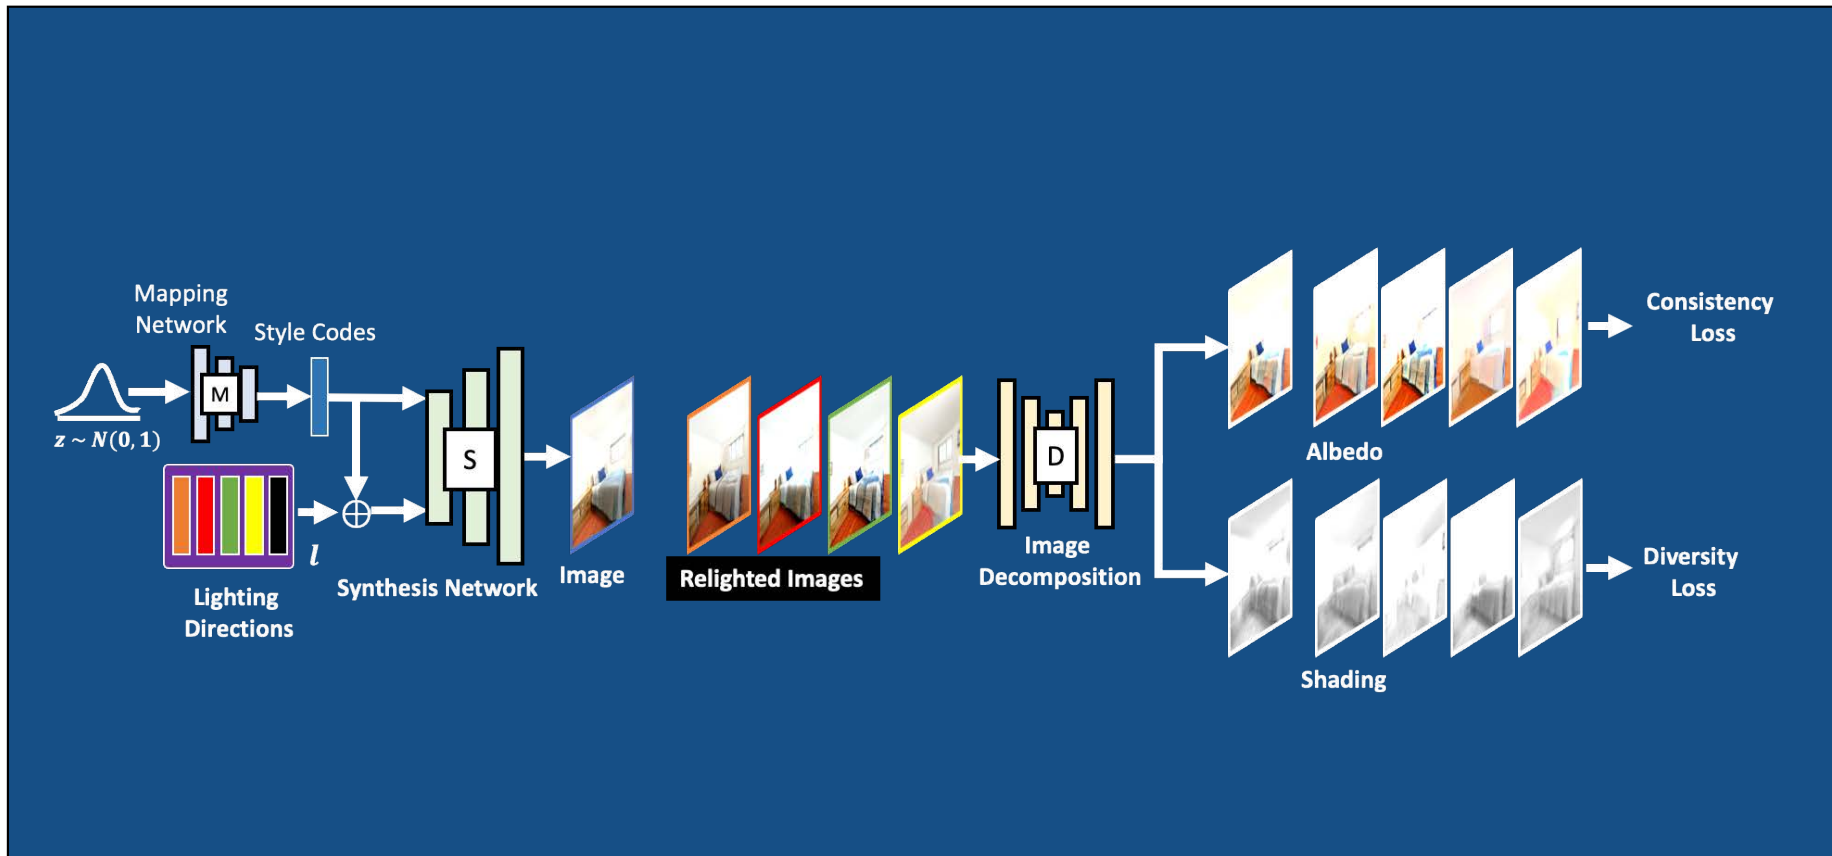
- Silly arguments
 - Money
 - Prurience (dogs playing poker)
- They know an awful lot about pictures, and data is easy
- Surprise: They know a lot about relighting, intrinsics, etc.
 - without being trained to do so
 - but have important areas of significant ignorance - why?
- Pose profound vision questions
 - what should we represent in what way?
 - Intrinsics - Latent; “Interaction hooks” - Explicit

- Add offset to StyleGAN latents
 - various effects by choice of offset
- Q: how to get desired result?
 - A (ish): search offsets



Yang et al. IJCV 2020; Epstein et al. ECCV 2022; Shi et al. NeurIPS 2022

Relighting: find directions that fix albedo



Bhattad et al, 23, 24: StyLitGAN: Prompting StyleGAN to Produce New Illumination Conditions

Relightings are realistic

relighting-1

relighting-2

relighting-3

relighting-4

relighting-5




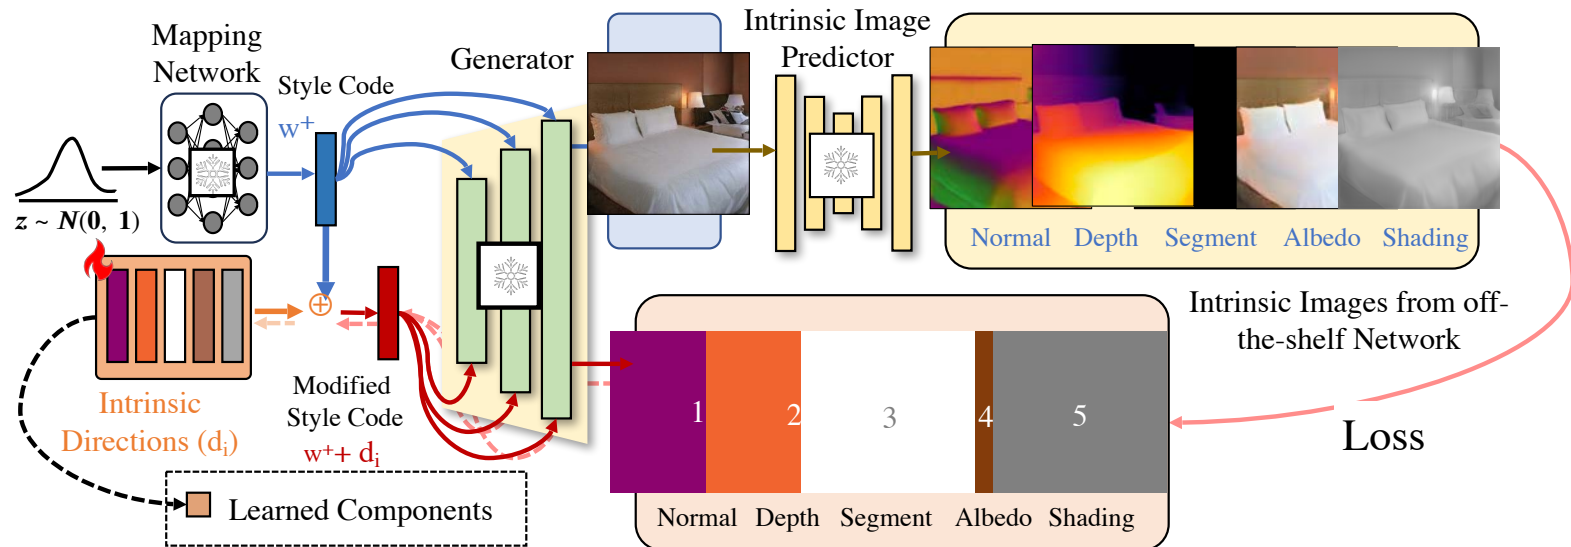
Brightness of the room is increased in relighting – 4

Brightness of the room is decreased in relighting – 5

If you can relight images

- you must know stuff about
 - depth
 - normal
 - surface properties
- Q: Does StyleGAN
- A1:
 - It should (kind of obvious)
- A2:
 - It can be made to produce this information (astonishing)

Recovering Intrinsic Images From StyleGAN




Recovered Intrinsic Images from StyleGAN


Recovered Intrinsic Images From StyleGAN



StyleGAN Image




StyleGAN Normals



StyleGAN Depth



StyleGAN Albedo




StyleGAN Shading



StyleGAN Segment

Recovered Intrinsic Images From SOTA Supervised Models



Normals




Depth



Albedo

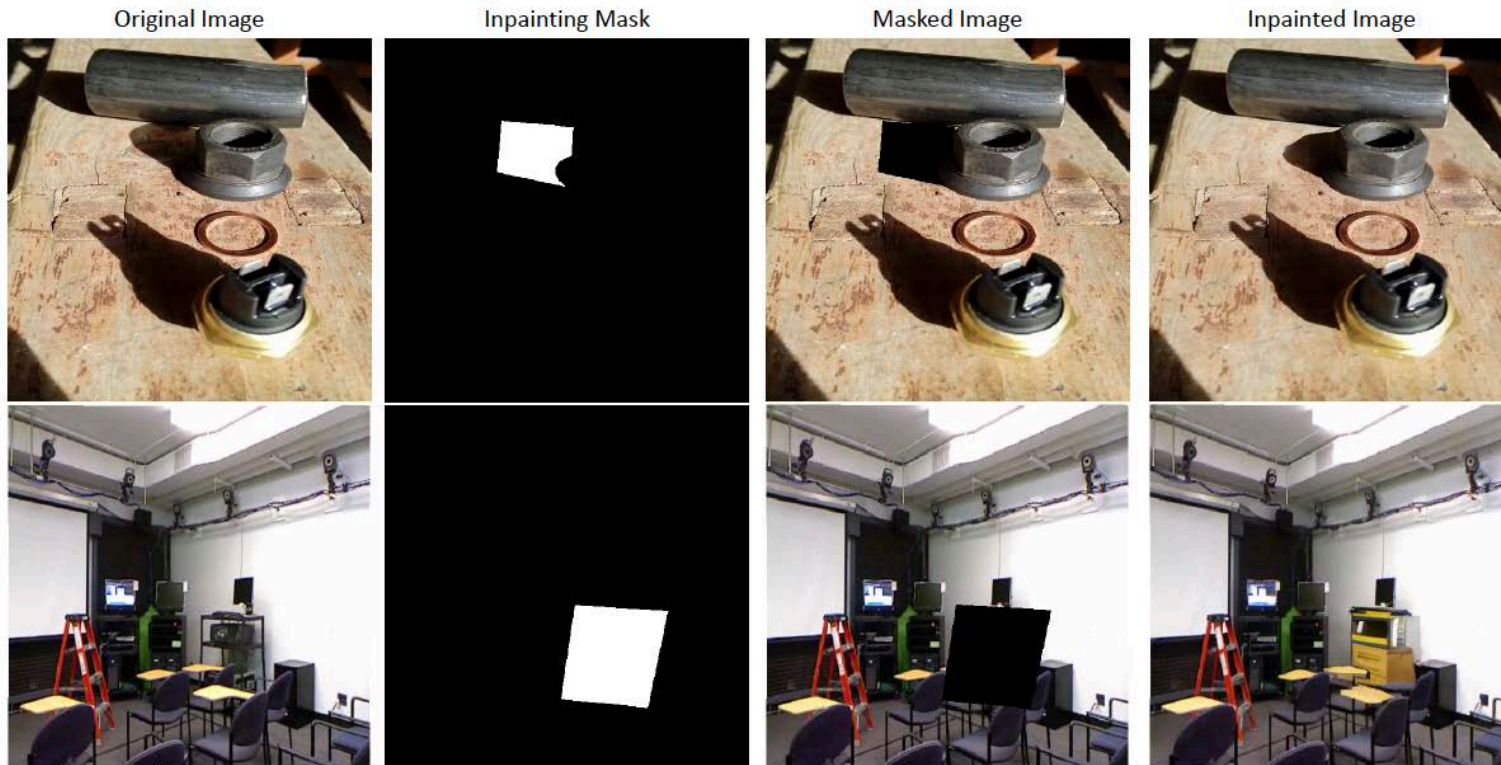


Shading



Segment

Stable diffusion “knows” some geometry




Zhan et al, 23: What does stable diffusion know about the 3D Scene?

Wallingford, Bhattad et al. NeurIPS 2024 360-1M: Learning to Imagine the world from 1M 360 videos

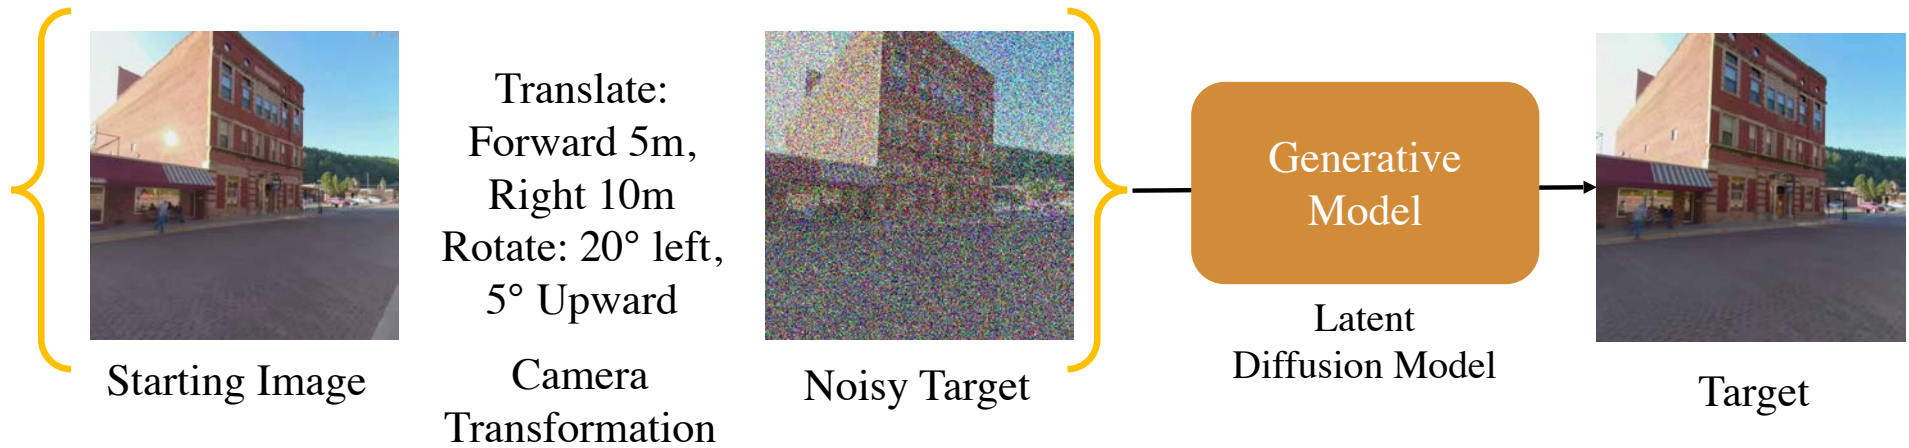


Aerial View




360-1M

Training



Comparison with SOTA


ZeroNVS (CVPR 2024)



Wallingford et al








This could change the game

- Forget building large labelled datasets
- Instead, build good generator, good inverter
 - and recover what you want from the generator
- Internal representation in generator is a geometric model
 - because it can answer geometric queries accurately (cf Malcolm Sabin)
 - and it's very good at making pictures


Structure

- Vision is flourishing
- What does vision do?
- Speculation tools are limited
- Big new player: Generative models
 - They know a lot of unexpected stuff
 - They're dumb about a lot of other stuff
- Agenda: Fix generative models and stop regressing things
 - Understand mistakes
 - Build sensitive evaluation procedures
 - Understand interactions between model architectures and weird behavior
 - “inductive bias” (?)


Original Image




Inpainting Mask



Masked Image



Inpainted Image



Zhan et al, 23: What does stable diffusion know about the 3D Scene?

Iffy projective geometry




<https://huggingface.co/spaces/stabilityai/stable-diffusion/discussions/1593>

Iffy projective geometry



<https://huggingface.co/spaces/stabilityai/stable-diffusion/discussions/1593>

Iffy projective geometry



<https://huggingface.co/spaces/stabilityai/stable-diffusion/discussions/1593>

Real or generated?



Are generated and real images different?

- Procedure:
 - build reduced geometric representation of image
 - lines; perspective field; object+shadow
 - using existing tools
 - train a classifier to distinguish between real and generated
 - prequalify images so that at test, classifier sees only “good” generated images.
- Result
 - Near perfect classification
- Conclusion:
 - generators are bad at lines; perspective; shadows

Evaluation and Results at Scale

Hard Test Set

Astonishing fact:

learn classifier with one generator's output,
apply to other generator output: very little change in performance
(they all seem to make the same kinds of mistake)

Curious fact: Deepfloyd is different

Curious fact, recent:
updated releases (bigger training sets, etc) aren't better.

Generators

Evaluated

Kandinsky

Stable Diffusion XL

PixArt-α

DeepFloyd IF

DALL-E 3

Firefly

1 Million+ Indoor

Images

500K+ Outdoor Images

Structure

- Vision is flourishing
- What does vision do?
- Speculation tools are limited
- Big new player: Generative models
 - They know a lot of unexpected stuff
 - They're dumb about a lot of other stuff
- Agenda: Fix generative models and stop regressing things
 - Understand mistakes
 - Build sensitive evaluation procedures
 - Understand interactions between model architectures and weird behavior
 - “inductive bias” (?)

What mistakes do generators make?

- (maybe one A): Problems with spatial or temporal scale
- **Communication:**
 - Hard to move fine scale information long distances over space or time
- **Volume:**
 - There's lots of fine scale information
- **Choice:**
 - What moves? and why?

Synthesizing clothing



Zhu et al 23: Tryondiffusion: A tale of two unets



Zhu et al 23: Tryondiffusion: A tale of two unets

Applications: make commercial art cheaper



Scale
relative to
bear



Scale
relative to
politician



He's got someone
else's nose!

What can you spot?





Did you see?

Soccer ball -> Tennis ball
Pentagon -> 4-gon
Frictional changes
(sliding, then spinning, etc)
Spin stopping, then starting
Weird path on tennis ball

....

What we should be able to do...

We can, but not well.


- Drag something around in its own scene Drag - the object geometry changes
- Select something in one image/video, place in another Insertion - lighting, aspect problems
- Adjust scene by interaction with image/video Adjust - tends to look muddy
- Model:
 - perfect inverse graphics+interaction+perfect render
 - but this isn't likely the best mechanism

Precision is everything - the something should not change identity

Editing is in a muddle

- Many different goals
 - some easy, some hard, many elided
- Easy (-ish!)
 - edit image with text
 - cause you don't need to preserve identity
- Super hard
 - when you want identity right
 - when you want geometry right

Editing - lively and broad



Text - looks like it works cause it's hard to be fussy

Editing Type: Object Addition (“Add a glass of milk next to the cookies”)



Editing Type: Object Removal (“Remove the wooden house and dock”)



Editing Type: Object Replacement (“Replace the llama toy with a floor lamp”)



Input

SDEdit [53]

P2P+NTI [35, 54]

Instruct-P2P [18]

MagicBrush [97]

MasaCTRL [19]


Stable Flow (ours)



“The cat is yelling and raising its paw”



“A rabbit toy sitting and wearing pink socks during the late afternoon”



“A rubber duck next to a purple ball during a sunny day”



“A rubber duck next to a purple ball during a sunny day”

Insertion - harder: you must preserve identity

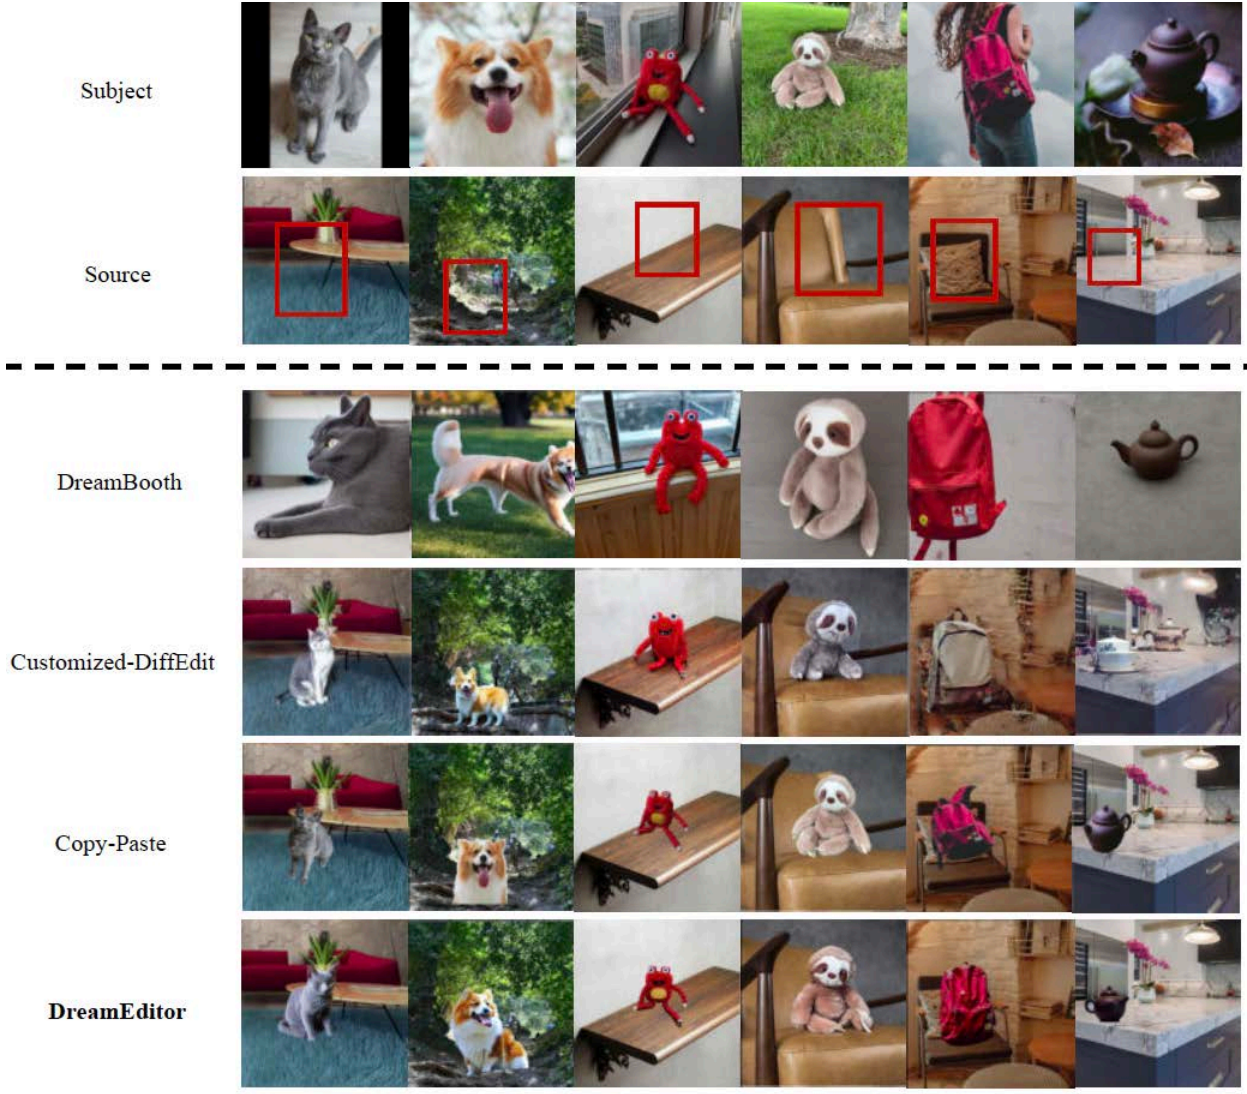



Figure 5: Results on Subject Addition Task Compared with Baselines



In some applications,
fussiness is crucial

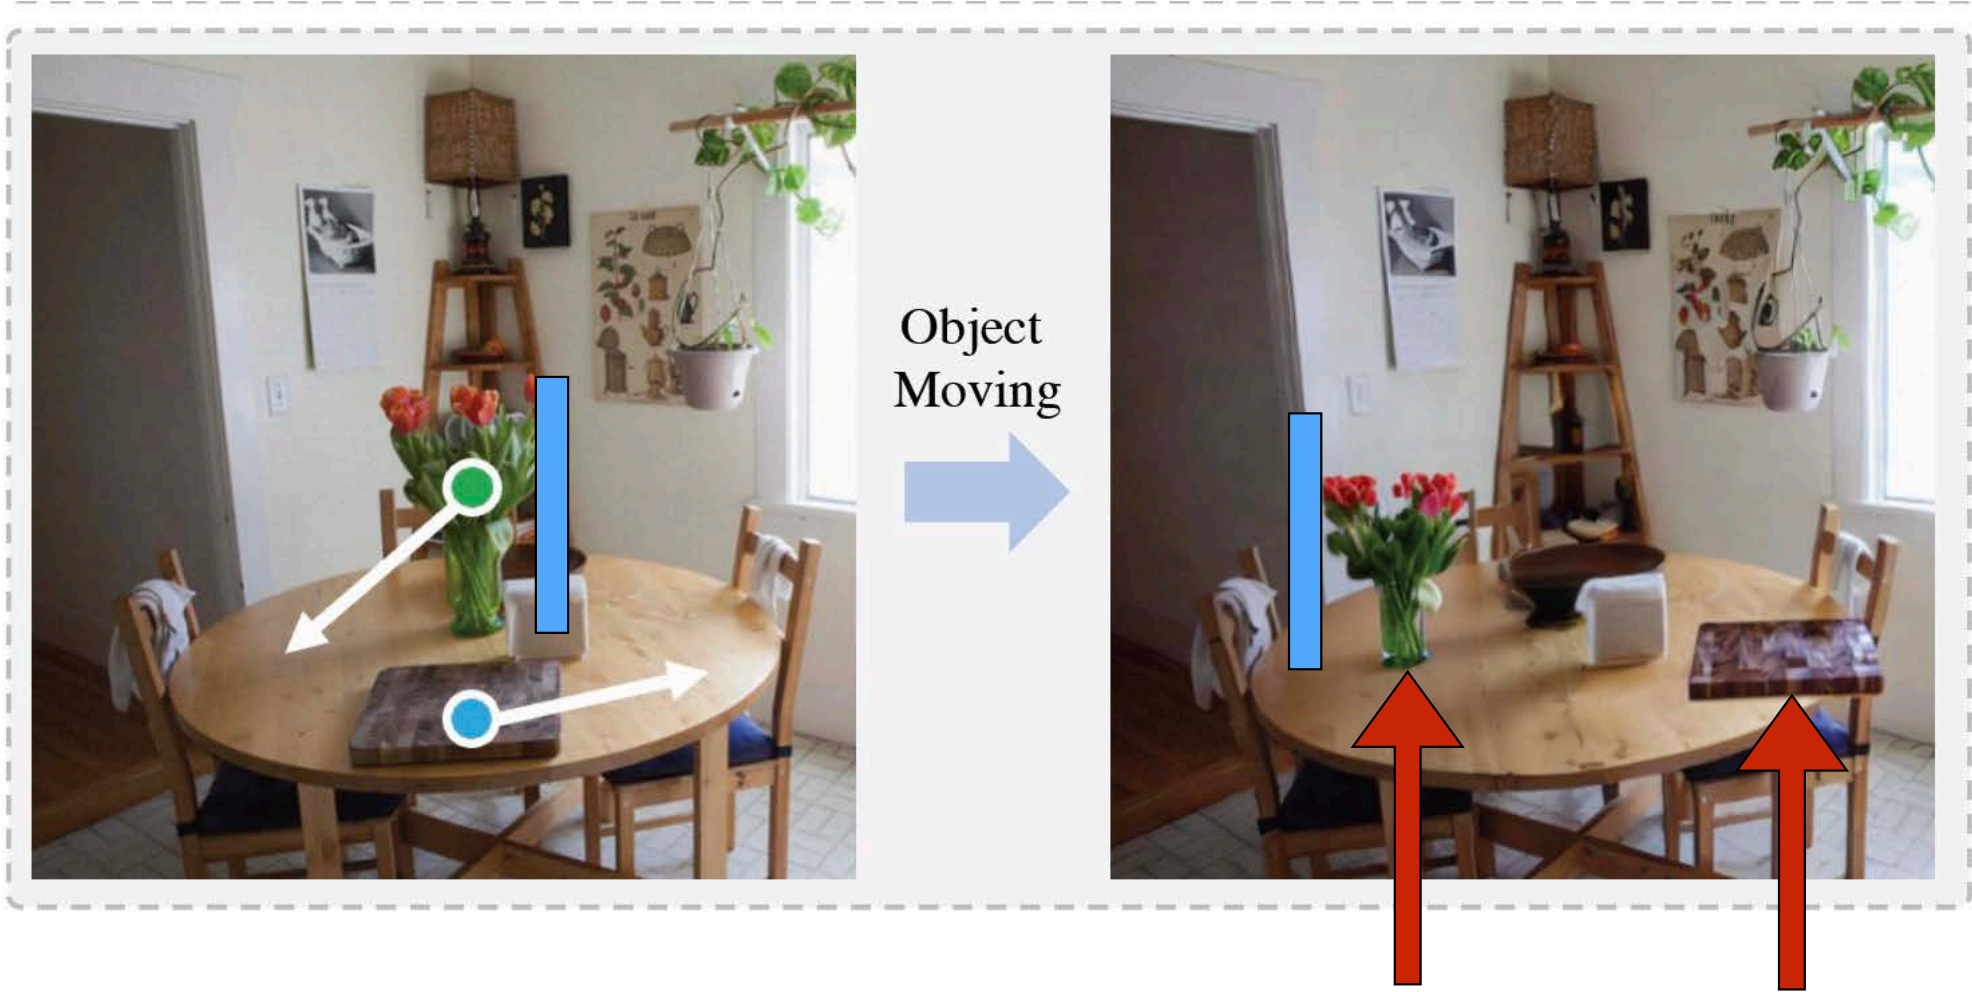
Chen et al 24: AnyDoor: Zero-shot Object-level Image Customization



Chen et al 24: AnyDoor: Zero-shot Object-level Image Customization



Chen et al 24: AnyDoor: Zero-shot Object-level Image Customization



Chen et al 24: AnyDoor: Zero-shot Object-level Image Customization

Structure

- Vision is flourishing
- What does vision do?
- Speculation tools are limited
- Big new player: Generative models
 - They know a lot of unexpected stuff
 - They're dumb about a lot of other stuff
- Agenda: Fix generative models and stop regressing things
 - Understand mistakes
 - Build sensitive evaluation procedures
 - Understand interactions between model architectures and weird behavior
 - “inductive bias” (?)

Evaluation must-dos

- How do you know if its working?
 - Is the image a “real image”? (currently, no)
 - Did the object preserve its identity? (currently no)
- What kinds of images are/are not generated?
 - and why?
- What kinds of error are typical?
 - tolerable? why?

Evaluating generators is very difficult

- Current procedures
 - mostly focused on short spatial scales (FID, LPIPS)
 - even these measures conceal the fact that generators are bad at short scales
 - or absurd:
 - “look at my nice pictures” (tough to measure progress like this)
 - anecdote (“look at their nasty pictures”...)
- Classifiers aren’t the answer
 - not sensitive to small improvements
- Two sample testing is not the answer
 - wildly unreliable in high dimension

Evaluation alternatives

- **CLIP-score:**
- **LMM-score:**
 - get a large multimodal model (LMM) to evaluate:
 - Editing accuracy
 - Context preservation
 - Visual quality
 - Logical realism
 - calibrate against user study
 - cf: GPT score, Chen+Wang
- **User studies**
 - but people aren't very good at this (as above!)

The field of image editing has traditionally relied on CLIPScore [304] as a major quantitative evaluation metric. Despite its effectiveness in assessing the alignment between images and corresponding textual prompts, CLIPScore may struggle in complicated scenes with many details and specific spatial relationships [301].

Chen+Wang 24: ProEdit: Simple Progression is All You Need for High-Quality 3D Scene Editing

Huang et al 25: Diffusion based image editing: a survey

Sort of agrees with people

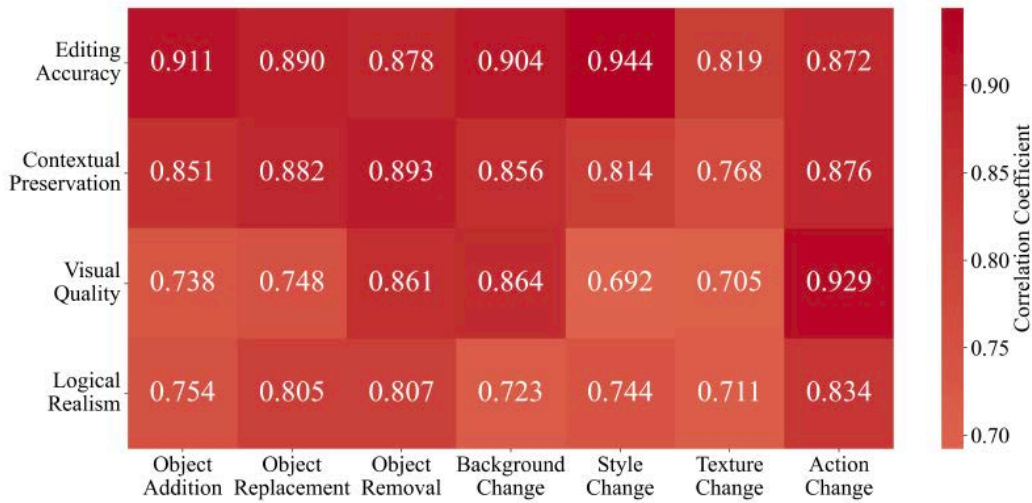


Fig. 10: Pearson correlation coefficients between LMM Score and the user study.




Fig. 11: Comparison of Pearson correlation coefficients between four metrics and the user study.

Structure

- Vision is flourishing
- What does vision do?
- Speculation tools are limited
- Big new player: Generative models
 - They know a lot of unexpected stuff
 - They're dumb about a lot of other stuff
- Agenda: Fix generative models and stop regressing things
 - Understand mistakes
 - Build sensitive evaluation procedures
 - Understand interactions between model architectures and weird behavior
 - “inductive bias” (?)

Why do generators make mistakes?


- Idea:
 - it is easy to build a path from image to noise
 - Denoising an image is relatively straightforward
 - we can build really good learned denoisers
 - so we can walk BACKWARDS along this path
 - Build a continuation path from noise to images
 - by
 - building a path
 - building a denoiser (some denoisers) that can cope with very small alpha

What happens if the denoiser makes errors that depend on some global properties of the input?

These errors would propagate along the path, and might very well accumulate.

Accumulated errors might lead to significant changes in output distribution

BIG Q: Does image regression make patterned errors?
BIG A: (well known) YES!




Zhong et al, 23: Improving Equivariance in State-of-the-Art Supervised Depth and Normal Predictors

[1] Forsyth et al., "Intrinsic Image Decomposition Using Paradigms." (TPAMI 2021)

What is to be done

- Understand what generators know and what they don't
 - and why...
- Improve evaluation practices
 - so we can tell when we're making progress
- Pay close attention to relations between scales
 - our architectures are missing something



What should a young researcher do? II

Smart, energetic, motivated young researchers can't stick out by doing what everyone else is doing

- Ignore your seniors
- Find something different and special

- Remember the wisdom of Molesworth:

“Akwire Kulture and Keep the Brane Kleen”