Learning High Dynamic Range from Outdoor Panoramas

Jinsong Zhang and Jean-François Lalonde

Université Laval
Québec City, Canada

Outdoor lighting

Image credit: Barrett&MacKay | Getty Images

Dynamic range

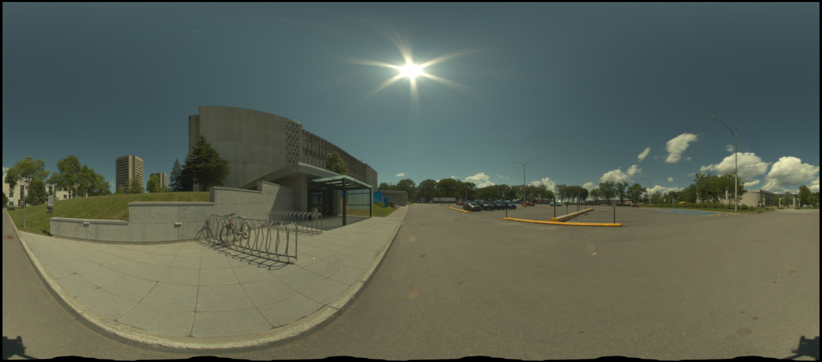

Dynamic range

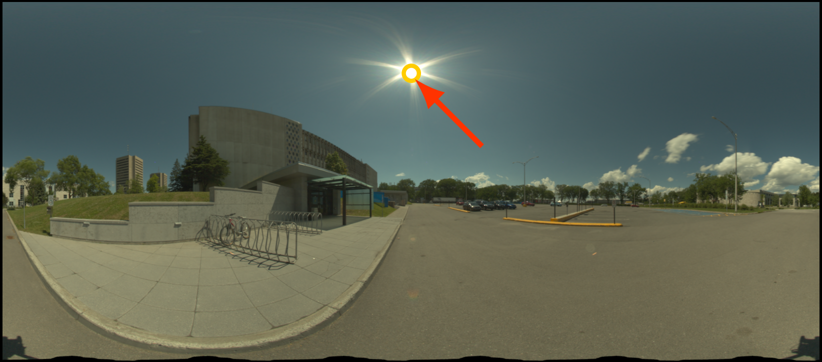

Dynamic range

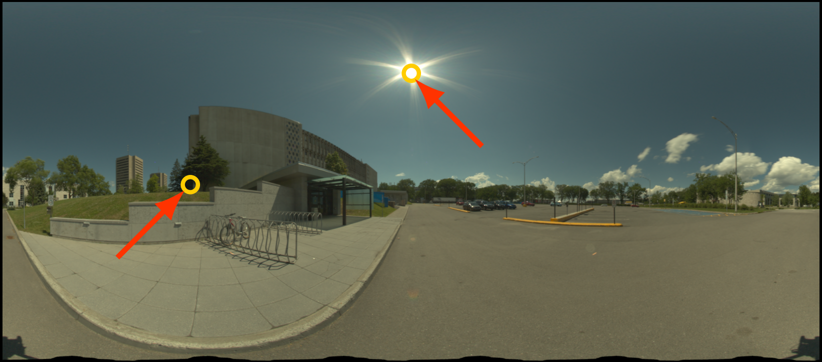

Dynamic range

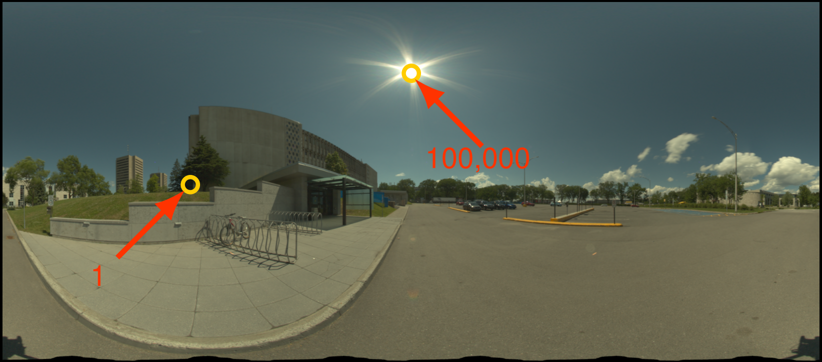

High Dynamic Range (HDR) image

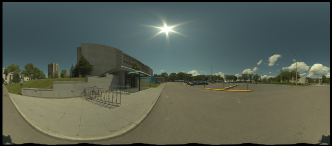

High Dynamic Range (HDR) image

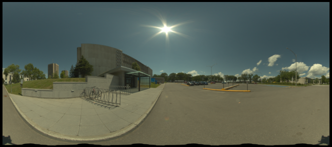

UNIVERSITÉ LAVAL

High Dynamic Range (HDR) image

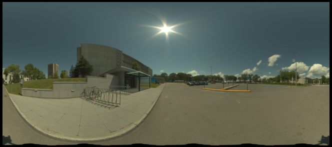

High Dynamic Range (HDR) image

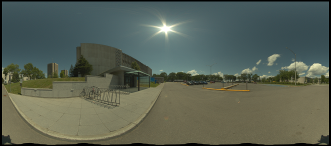

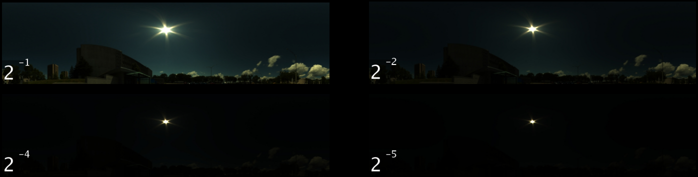

High Dynamic Range (HDR) image

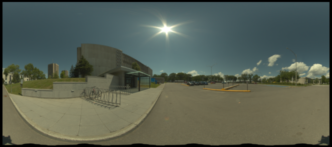

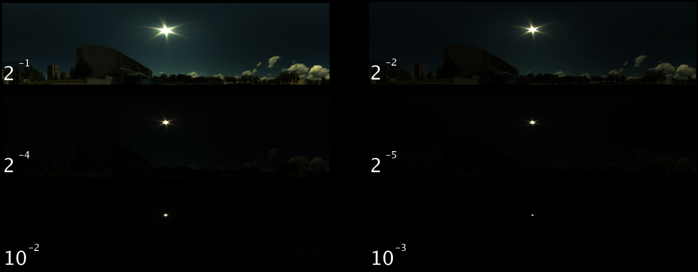

High Dynamic Range (HDR) image

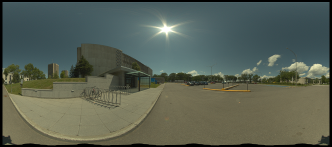

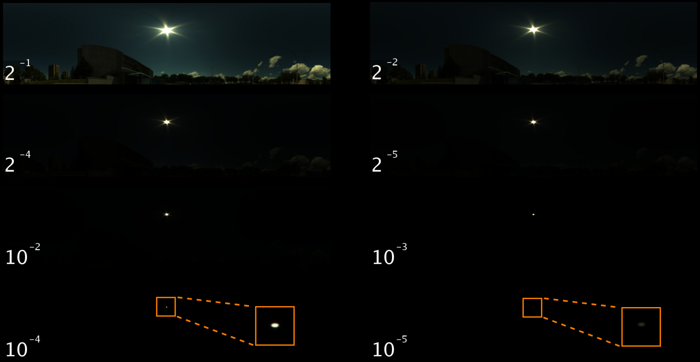

Outdoor lighting capture

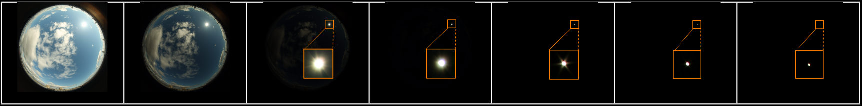

Stumpfel, et al. Direct HDR capture of the sun and sky. AFRIGRAPH, 2004.

UNIVERSITÉ LAVAL

Outdoor lighting capture

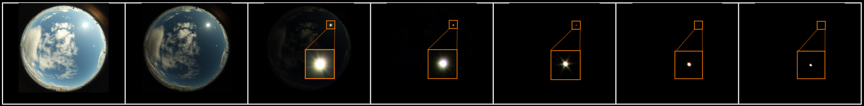

Combine 7 LDR images at different exposures into HDR

Stumpfel, et al. Direct HDR capture of the sun and sky. AFRIGRAPH, 2004.

Outdoor lighting capture

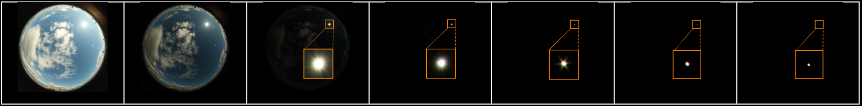

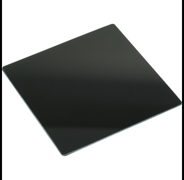

Combine 7 LDR images at different exposures into HDR
Neutral density filter

Stumpfel, et al. Direct HDR capture of the sun and sky. AFRIGRAPH, 2004.

Outdoor lighting capture

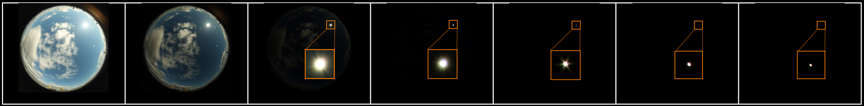

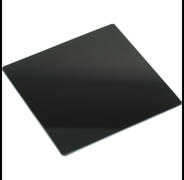

Combine 7 LDR images at different exposures into HDR
Neutral density filter
Complicated capture procedure

Stumpfel, et al. Direct HDR capture of the sun and sky. AFRIGRAPH, 2004.

Learn to predict the extremely HDR outdoor lighting from a single, LDR panorama.

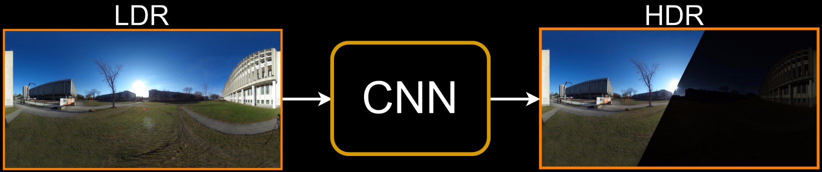

Learn to predict the extremely HDR outdoor lighting from a single, LDR panorama.

LDR

HDR

Key idea: Large dataset of synthetic panoramas lit by **real skies**

Learn to predict the extremely HDR outdoor lighting from a single, LDR panorama.

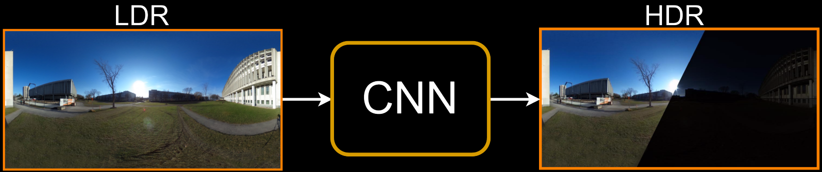

Key idea: Large dataset of synthetic panoramas lit by **real skies**
Assumption: Sun is centered in the middle

Hold-Geoffroy, et al. Deep Outdoor Illumination Estimation. CVPR, 2017.

Learn to predict the extremely HDR outdoor lighting from a single, LDR panorama.

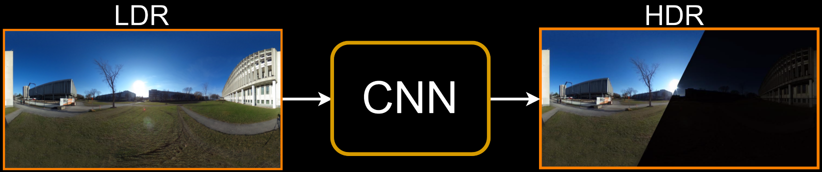

Key idea: Large dataset of synthetic panoramas lit by **real skies**
Assumption: Sun is centered in the middle
Contributions:

1. Frame the LDR to HDR problem as a deep learning problem
2. Provide a novel real LDR/HDR panorama dataset
3. Introduce three novel applications

Hold-Geoffroy, et al. Deep Outdoor Illumination Estimation. CVPR, 2017.

Synthetic panoramas lit by real skies

Synthetic panoramas lit by real skies



Real sky

Laval HDR Sky Database

- 38,000 images
- 103 days over 3 years

UNIVERSITÉ LAVAL

Synthetic panoramas lit by real skies

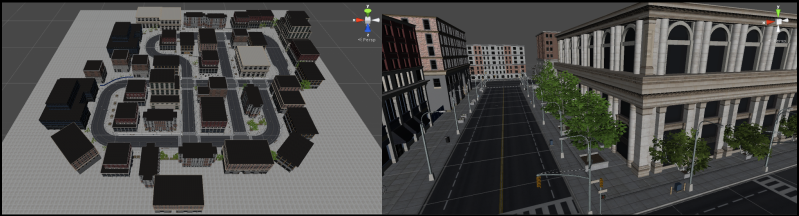

Real sky

Laval HDR Sky Database

- 38,000 images
- 103 days over 3 years

City model

- Realistic
- Detailed
- Random position
- Different exposures

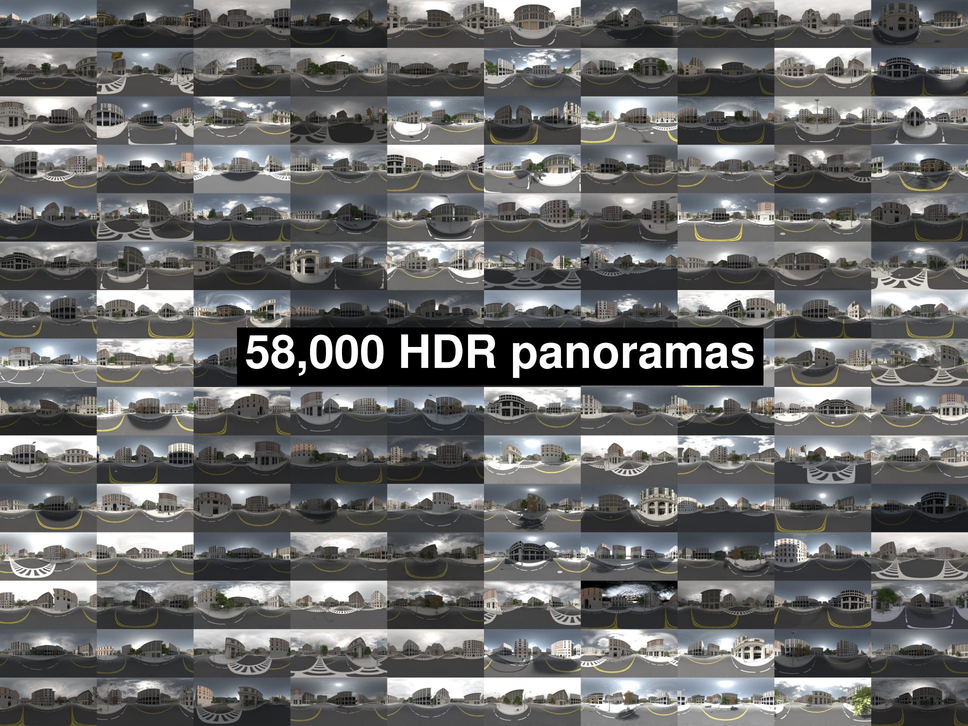
58,000 HDR panoramas

CNN structure

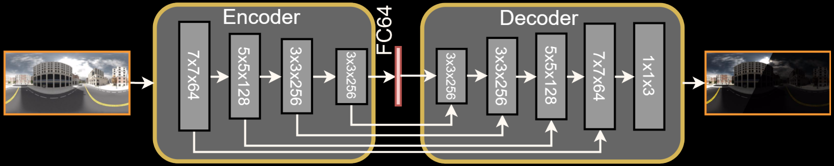

CNN structure

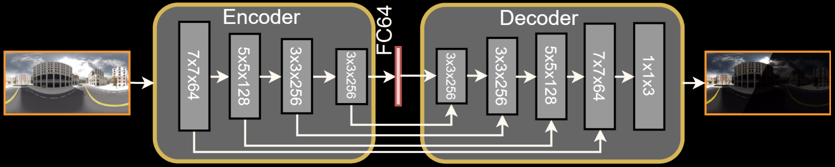

$$\mathcal{L}_{\mathrm{HDR}}(\mathbf{y}, \mathbf{t}) = ||\mathbf{y}_{\mathrm{HDR}} - \mathbf{t}_{\mathrm{HDR}}||_1 , \mathbf{t}_{\mathrm{HDR}} = \alpha(\mathbf{t}^*_{\mathrm{HDR}})^{1/\gamma}$$

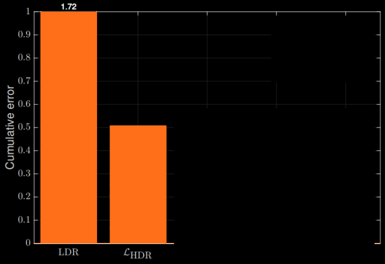

CNN structure

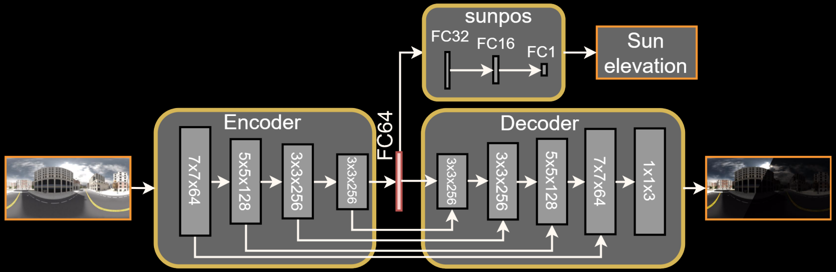

$$\mathcal{L}_{\mathrm{HDR}}(\mathbf{y},\mathbf{t}) = ||\mathbf{y}_{\mathrm{HDR}} - \mathbf{t}_{\mathrm{HDR}}||_1, \mathbf{t}_{\mathrm{HDR}} = \alpha(\mathbf{t}^*_{\mathrm{HDR}})^{1/\gamma}$$

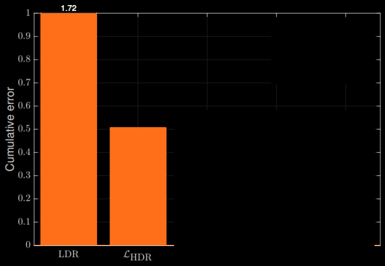

UNIVERSITÉ LAVAL

CNN structure

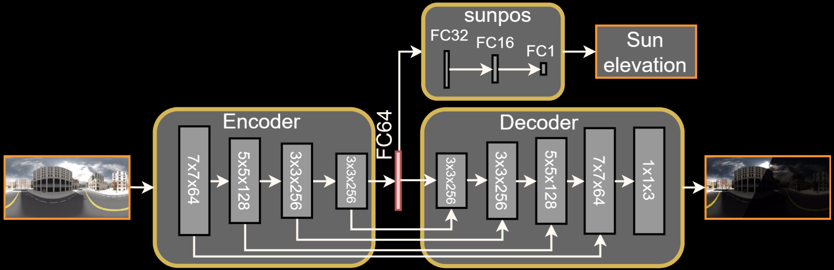

$$\mathcal{L}_{\mathrm{HDR}}(\mathbf{y}, \mathbf{t}) = ||\mathbf{y}_{\mathrm{HDR}} - \mathbf{t}_{\mathrm{HDR}}||_1, \mathbf{t}_{\mathrm{HDR}} = \alpha(\mathbf{t}^*_{\mathrm{HDR}})^{1/\gamma}$$

$$\mathcal{L}_{\theta}(\mathbf{y}, \mathbf{t}) = ||\mathbf{y}_{\theta} - \mathbf{t}_{\theta}||_2$$

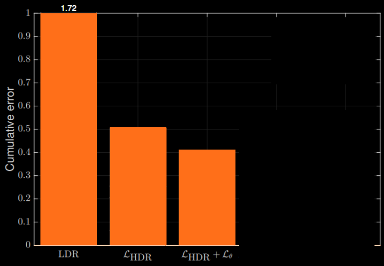

CNN structure

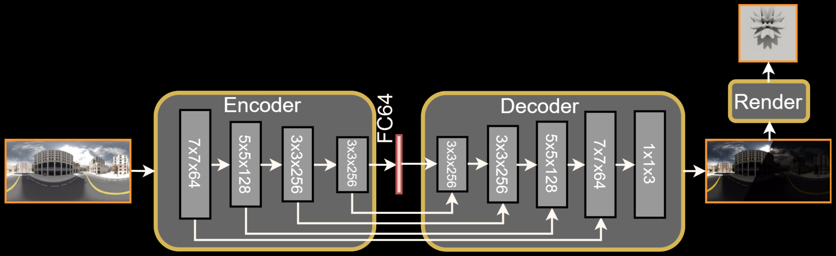

$$\mathcal{L}_{\mathrm{HDR}}(\mathbf{y}, \mathbf{t}) = ||\mathbf{y}_{\mathrm{HDR}} - \mathbf{t}_{\mathrm{HDR}}||_1, \mathbf{t}_{\mathrm{HDR}} = \alpha(\mathbf{t}^*_{\mathrm{HDR}})^{1/\gamma}$$

$$\mathcal{L}_{\theta}(\mathbf{y}, \mathbf{t}) = ||\mathbf{y}_{\theta} - \mathbf{t}_{\theta}||_2$$

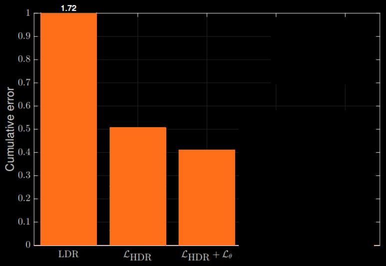

UNIVERSITÉ LAVAL

CNN structure

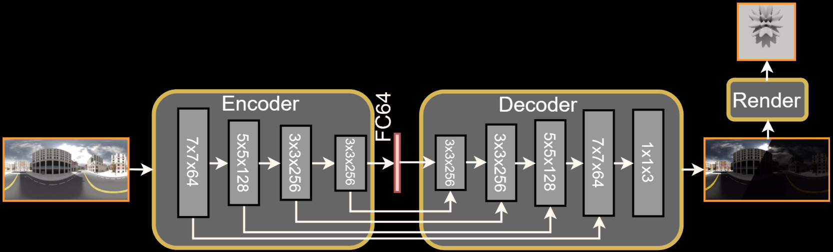

$$\mathcal{L}_{\text{HDR}}(\mathbf{y},\mathbf{t}) = ||\mathbf{y}_{\text{HDR}} - \mathbf{t}_{\text{HDR}}||_1 , \mathbf{t}_{\text{HDR}} = \alpha(\mathbf{t}^*_{\text{HDR}})^{1/\gamma}$$

$$\mathcal{L}_{\theta}(\mathbf{y},\mathbf{t}) = ||\mathbf{y}_{\theta} - \mathbf{t}_{\theta}||_2$$

$$\mathcal{L}_{\text{render}}(\mathbf{y},\mathbf{t}) = ||\mathbf{T}\mathbf{y}_{\text{HDR}} - \mathbf{T}\mathbf{t}_{\text{HDR}}||_2$$

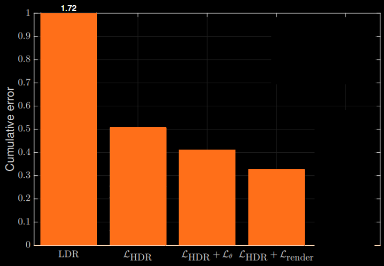

UNIVERSITÉ LAVAL

CNN structure

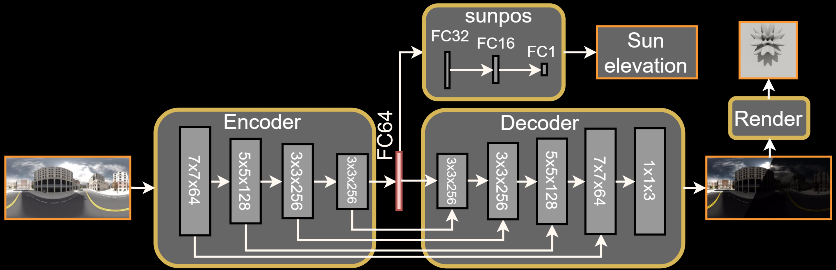

$$\mathcal{L}_{\mathrm{HDR}}(\mathbf{y},\mathbf{t}) = ||\mathbf{y}_{\mathrm{HDR}} - \mathbf{t}_{\mathrm{HDR}}||_1,\ \mathbf{t}_{\mathrm{HDR}} = \alpha(\mathbf{t}^*_{\mathrm{HDR}})^{1/\gamma}$$

$$\mathcal{L}_{\theta}(\mathbf{y},\mathbf{t}) = ||\mathbf{y}_{\theta} - \mathbf{t}_{\theta}||_2$$

$$\mathcal{L}_{\mathrm{render}}(\mathbf{y},\mathbf{t}) = ||\mathbf{T}\mathbf{y}_{\mathrm{HDR}} - \mathbf{T}\mathbf{t}_{\mathrm{HDR}}||_2$$

$$\mathcal{L}_{\mathrm{all}}(\mathbf{y},\mathbf{t}) = \mathcal{L}_{\mathrm{HDR}}(\mathbf{y},\mathbf{t}) + \lambda_1\mathcal{L}_{\theta}(\mathbf{y},\mathbf{t}) + \lambda_2\mathcal{L}_{\mathrm{render}}(\mathbf{y},\mathbf{t})$$

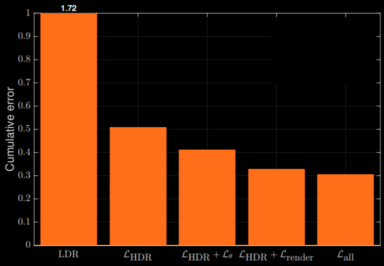

UNIVERSITÉ LAVAL

Results on synthetic data

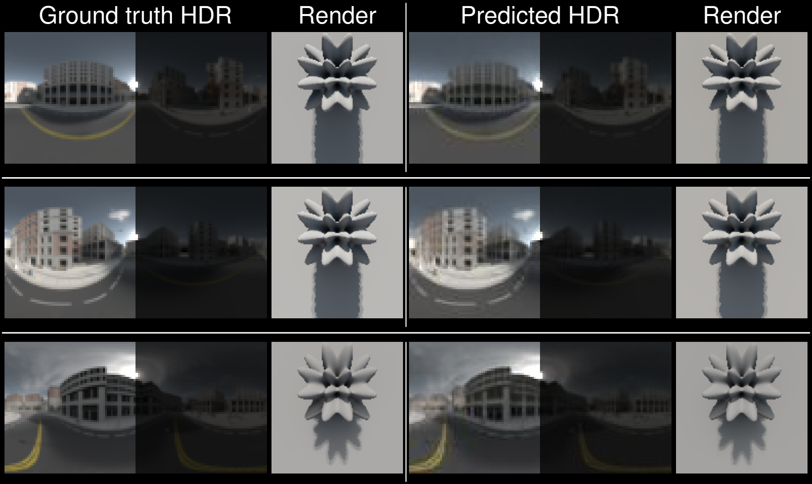

Results on synthetic data

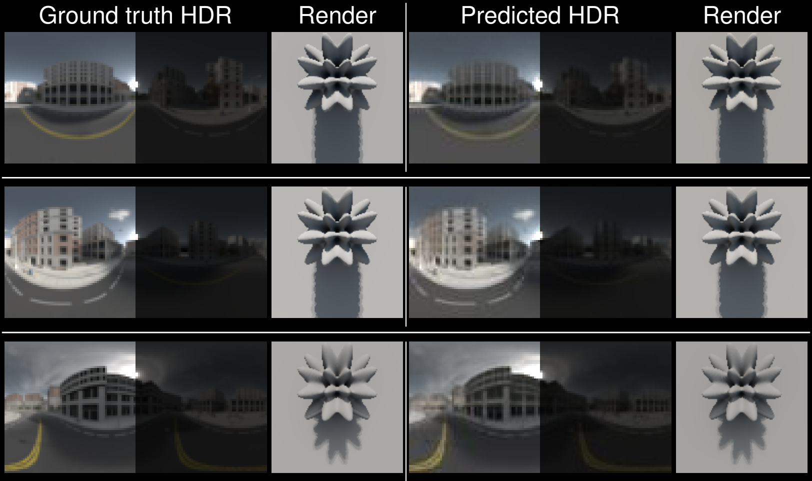

How well does it work with real images?

Novel real data

LDR panorama

HDR sky

Ricoh Theta S

Canon 5D Mark III

Novel real data

Ricoh Theta S

Canon 5D Mark III

Towards a real scenario

Method	Error metrics			
	HDR	Render	Sun elevation	Sun intensity
LDR	5.30	1.34	0.21	0.54

UNIVERSITÉ LAVAL

Towards a real scenario

Method	Error metrics			
	HDR	Render	Sun elevation	Sun intensity
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	**1.19**	**0.10**	**0.43**

Towards a real scenario

Method	Error metrics			
	HDR	Render	Sun elevation	Sun intensity
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	**1.19**	**0.10**	**0.43**

No access to camera

Towards a real scenario

Method	Error metrics			
	HDR	Render	Sun elevation	Sun intensity
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	**1.19**	**0.10**	**0.43**

No access to camera

With access to camera

Experiments on real data

No access to camera

Method	Error metrics			
	HDR	Render	Sun elev.	Sun inten.
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	1.19	0.10	0.43

Experiments on real data

No access to camera

Method	Error metrics			
	HDR	Render	Sun elev.	Sun inten.
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	1.19	0.10	0.43

Augmenting synthetic dataset

Experiments on real data

No access to camera

Method	Error metrics			
	HDR	Render	Sun elev.	Sun inten.
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	1.19	0.10	0.43

Augmenting synthetic dataset

- Camera response function (CRF)

Scene 1 Scene 2

CRF

CRF

CRF

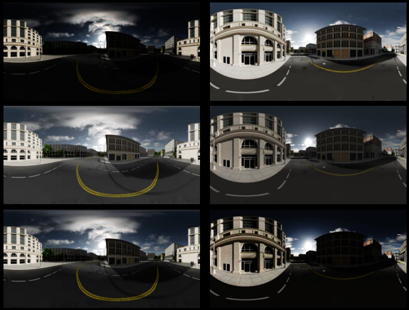

Grossberg and Nayar. What is the Space of Camera Response Functions? CVPR, 2003

Experiments on real data

No access to camera

Method	Error metrics			
	HDR	Render	Sun elev.	Sun inten.
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	1.19	0.10	0.43

Augmenting synthetic dataset
- Camera response function (CRF)
- White balance (WB)

Grossberg and Nayar. What is the Space of Camera Response Functions? CVPR, 2003

Experiments on real data

No access to camera

Method	Error metrics			
	HDR	Render	Sun elev.	Sun inten.
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	1.19	0.10	0.43
WB+CRF	**3.59**	**1.06**	**0.08**	**0.31**

Augmenting synthetic dataset

- Camera response function (CRF)
- White balance (WB)

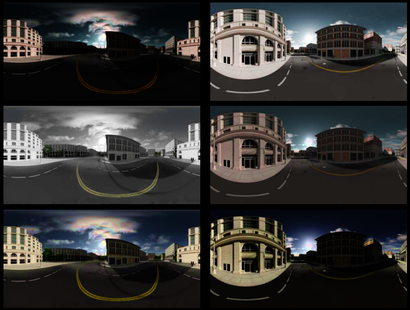

Grossberg and Nayar. What is the Space of Camera Response Functions? CVPR, 2003

Experiments on real data

With access to camera

Method	Error metrics			
	HDR	Render	Sun elev.	Sun inten.
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	1.19	0.10	0.43
WB+CRF	3.59	1.06	0.08	0.31

Experiments on real data

With access to camera

Method	Error metrics			
	HDR	Render	Sun elev.	Sun inten.
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	1.19	0.10	0.43
WB+CRF	3.59	1.06	0.08	0.31

1. Camera calibration
   - Inverse response function
   - White balance transformation

Experiments on real data

With access to camera

Method	Error metrics			
	HDR	Render	Sun elev.	Sun inten.
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	1.19	0.10	0.43
WB+CRF	3.59	1.06	0.08	0.31

1 Camera calibration

- Inverse response function
- White balance transformation

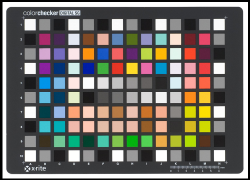

Experiments on real data

With access to camera

Method	Error metrics			
	HDR	Render	Sun elev.	Sun inten.
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	1.19	0.10	0.43
WB+CRF	3.59	1.06	0.08	0.31
Calibration	**2.99**	**1.03**	**0.08**	**0.30**

1. Camera calibration
   - Inverse response function
   - White balance transformation

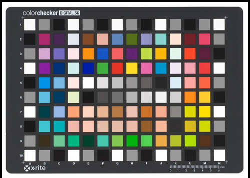

Experiments on real data

With access to camera

Method	Error metrics			
	HDR	Render	Sun elev.	Sun inten.
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	1.19	0.10	0.43
WB+CRF	3.59	1.06	0.08	0.31
Calibration	**2.99**	**1.03**	**0.08**	**0.30**

1. Camera calibration
   - Inverse response function
   - White balance transformation
2. Finetuning with real HDR data

Experiments on real data

With access to camera

Method	Error metrics			
	HDR	Render	Sun elev.	Sun inten.
LDR	5.30	1.34	0.21	0.54
Baseline (trained on synthetic)	5.34	1.19	0.10	0.43
WB+CRF	3.59	1.06	0.08	0.31
Calibration	2.99	1.03	0.08	0.30
Fine-tuning	**2.55**	**0.64**	**0.07**	**0.22**

1. Camera calibration
   - Inverse response function
   - White balance transformation
2. Finetuning with real HDR data

Experiments on real data

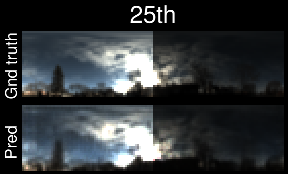

Experiments on real data

Experiments on real data

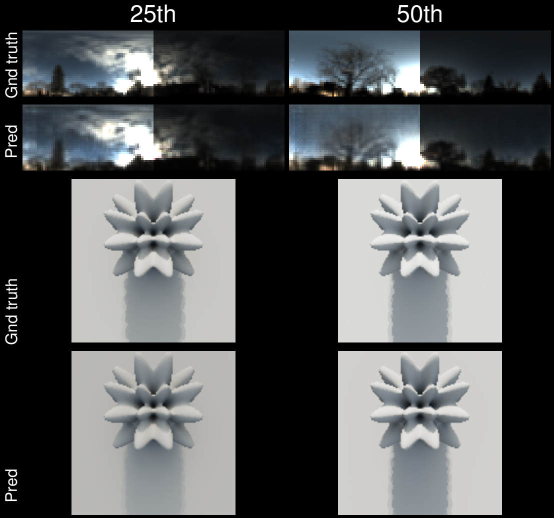

Experiments on real data

UNIVERSITÉ LAVAL

Single shot outdoor light probe

Single shot outdoor light probe

Single shot outdoor light probe

Render with LDR

UNIVERSITÉ LAVAL

Single shot outdoor light probe

Render with LDR

Render with Ours

Results on the finetune model

Single shot outdoor light probe

UNIVERSITÉ LAVAL

Single shot outdoor light probe

UNIVERSITÉ LAVAL

Single shot outdoor light probe

Single shot outdoor light probe

UNIVERSITÉ LAVAL

Single shot outdoor light probe

Single shot outdoor light probe

Visualization in Google Street View imagery

Google Street View imagery

UNIVERSITÉ LAVAL

Visualization in Google Street View imagery

Google Street View imagery

Render with LDR

UNIVERSITÉ LAVAL

Visualization in Google Street View imagery

Google Street View imagery

Render with LDR

Render with Ours

Visualization in Google Street View imagery

Google Street View imagery | Render with LDR | Render with Ours

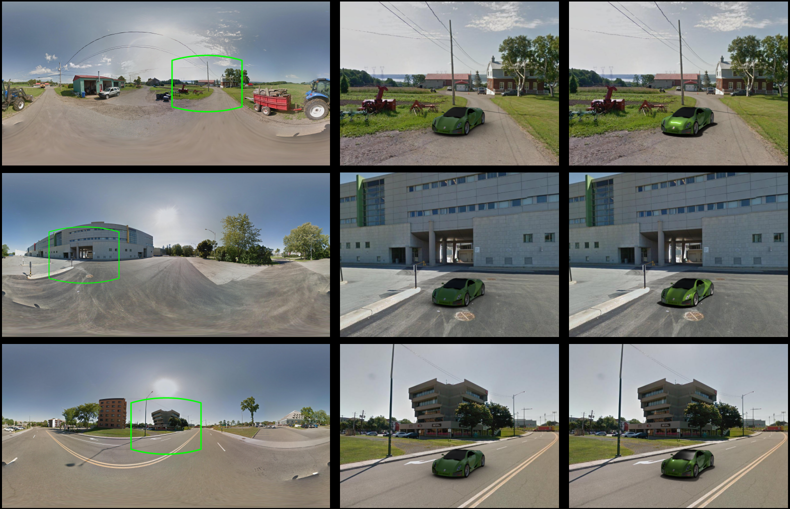

UNIVERSITÉ LAVAL

Conclusion

Learn to predict the extremely HDR outdoor lighting from a single, LDR panorama.

Conclusion

Learn to predict the extremely HDR outdoor lighting from a single, LDR panorama.

Contributions:

1. Frame the LDR to HDR problem as a deep learning problem

Conclusion

Learn to predict the extremely HDR outdoor lighting from a single, LDR panorama.

Contributions:

1. Frame the LDR to HDR problem as a deep learning problem
2. Provide a novel real dataset

Conclusion

Learn to predict the extremely HDR outdoor lighting from a single, LDR panorama.

Contributions:

1. Frame the LDR to HDR problem as a deep learning problem
2. Provide a novel real dataset
3. Introduce several novel applications

Conclusion

Learn to predict the extremely HDR outdoor lighting from a single, LDR panorama.

Contributions:

1. Frame the LDR to HDR problem as a deep learning problem
2. Provide a novel real dataset
3. Introduce several novel applications

Code and data are available on our website:
jflalonde.ca/projects/learningHDR