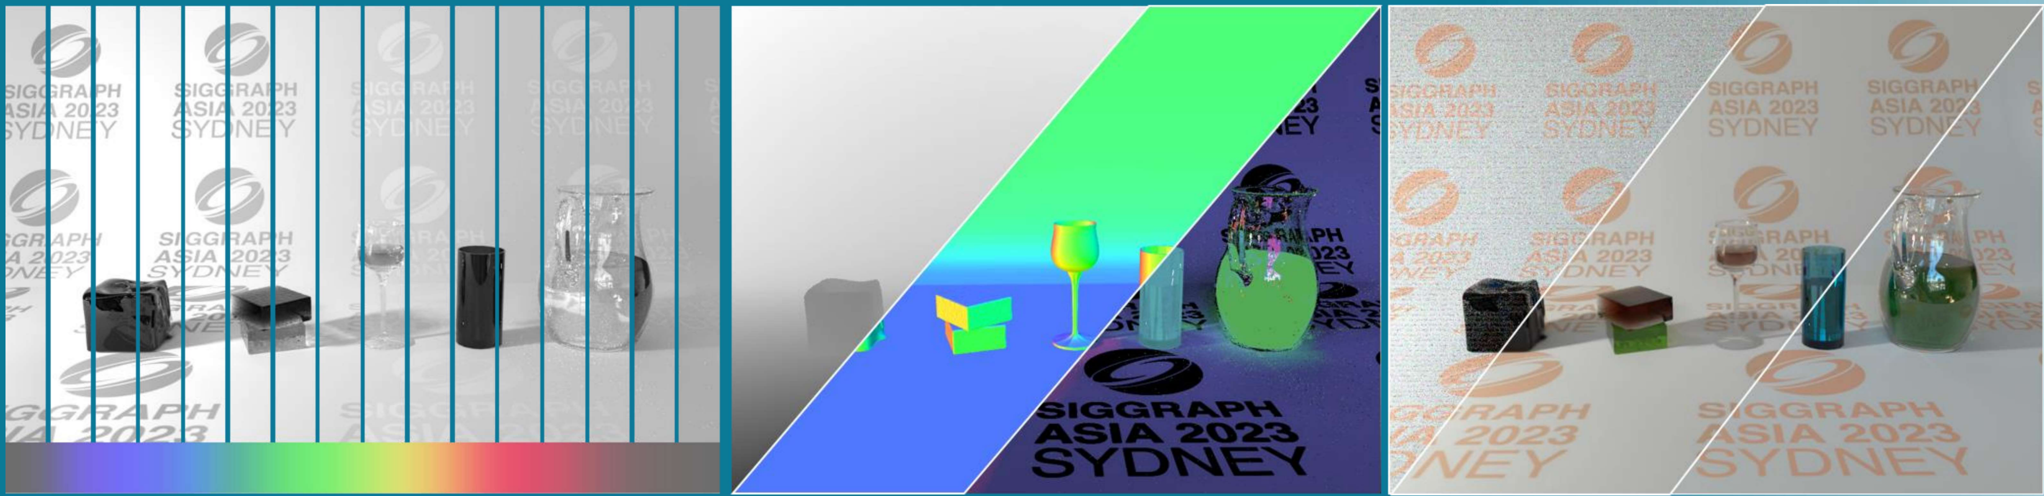

Spectral Monte Carlo Image Denoising

M. Noizet[1], R. Rouphael[1,2], S. Prévost[1], H. Deleau[1], L. Steffenel[1] and L. Lucas[1]

[1]Université de Reims Champagne-Ardenne, LICIIS, France
[2]United Visual Researchers, France

JCAD 2023

SUMMARY

- General context
- Previous works
- Contributions and results
- Conclusion

General context

- Predictive visualisation – in interactive time – of complex materials for industry (CA^{2}O)
  - ANR LUCE PRCE 2021-2024
  - Optical simulation with Spectral information
    - Generate predictive image for virtual moke-up
    - Time consuming
    - All light phenomena (metarism, polarization, etc.)

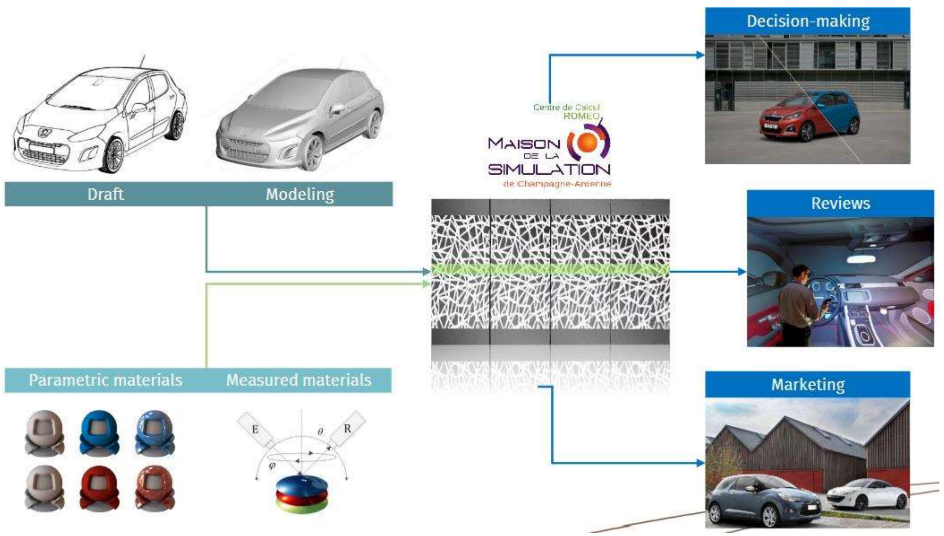

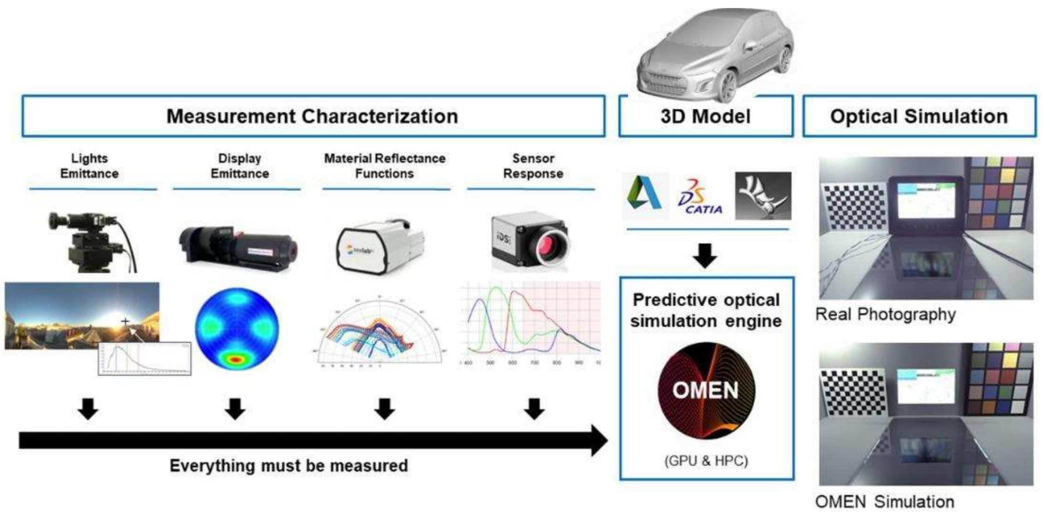

METARISM

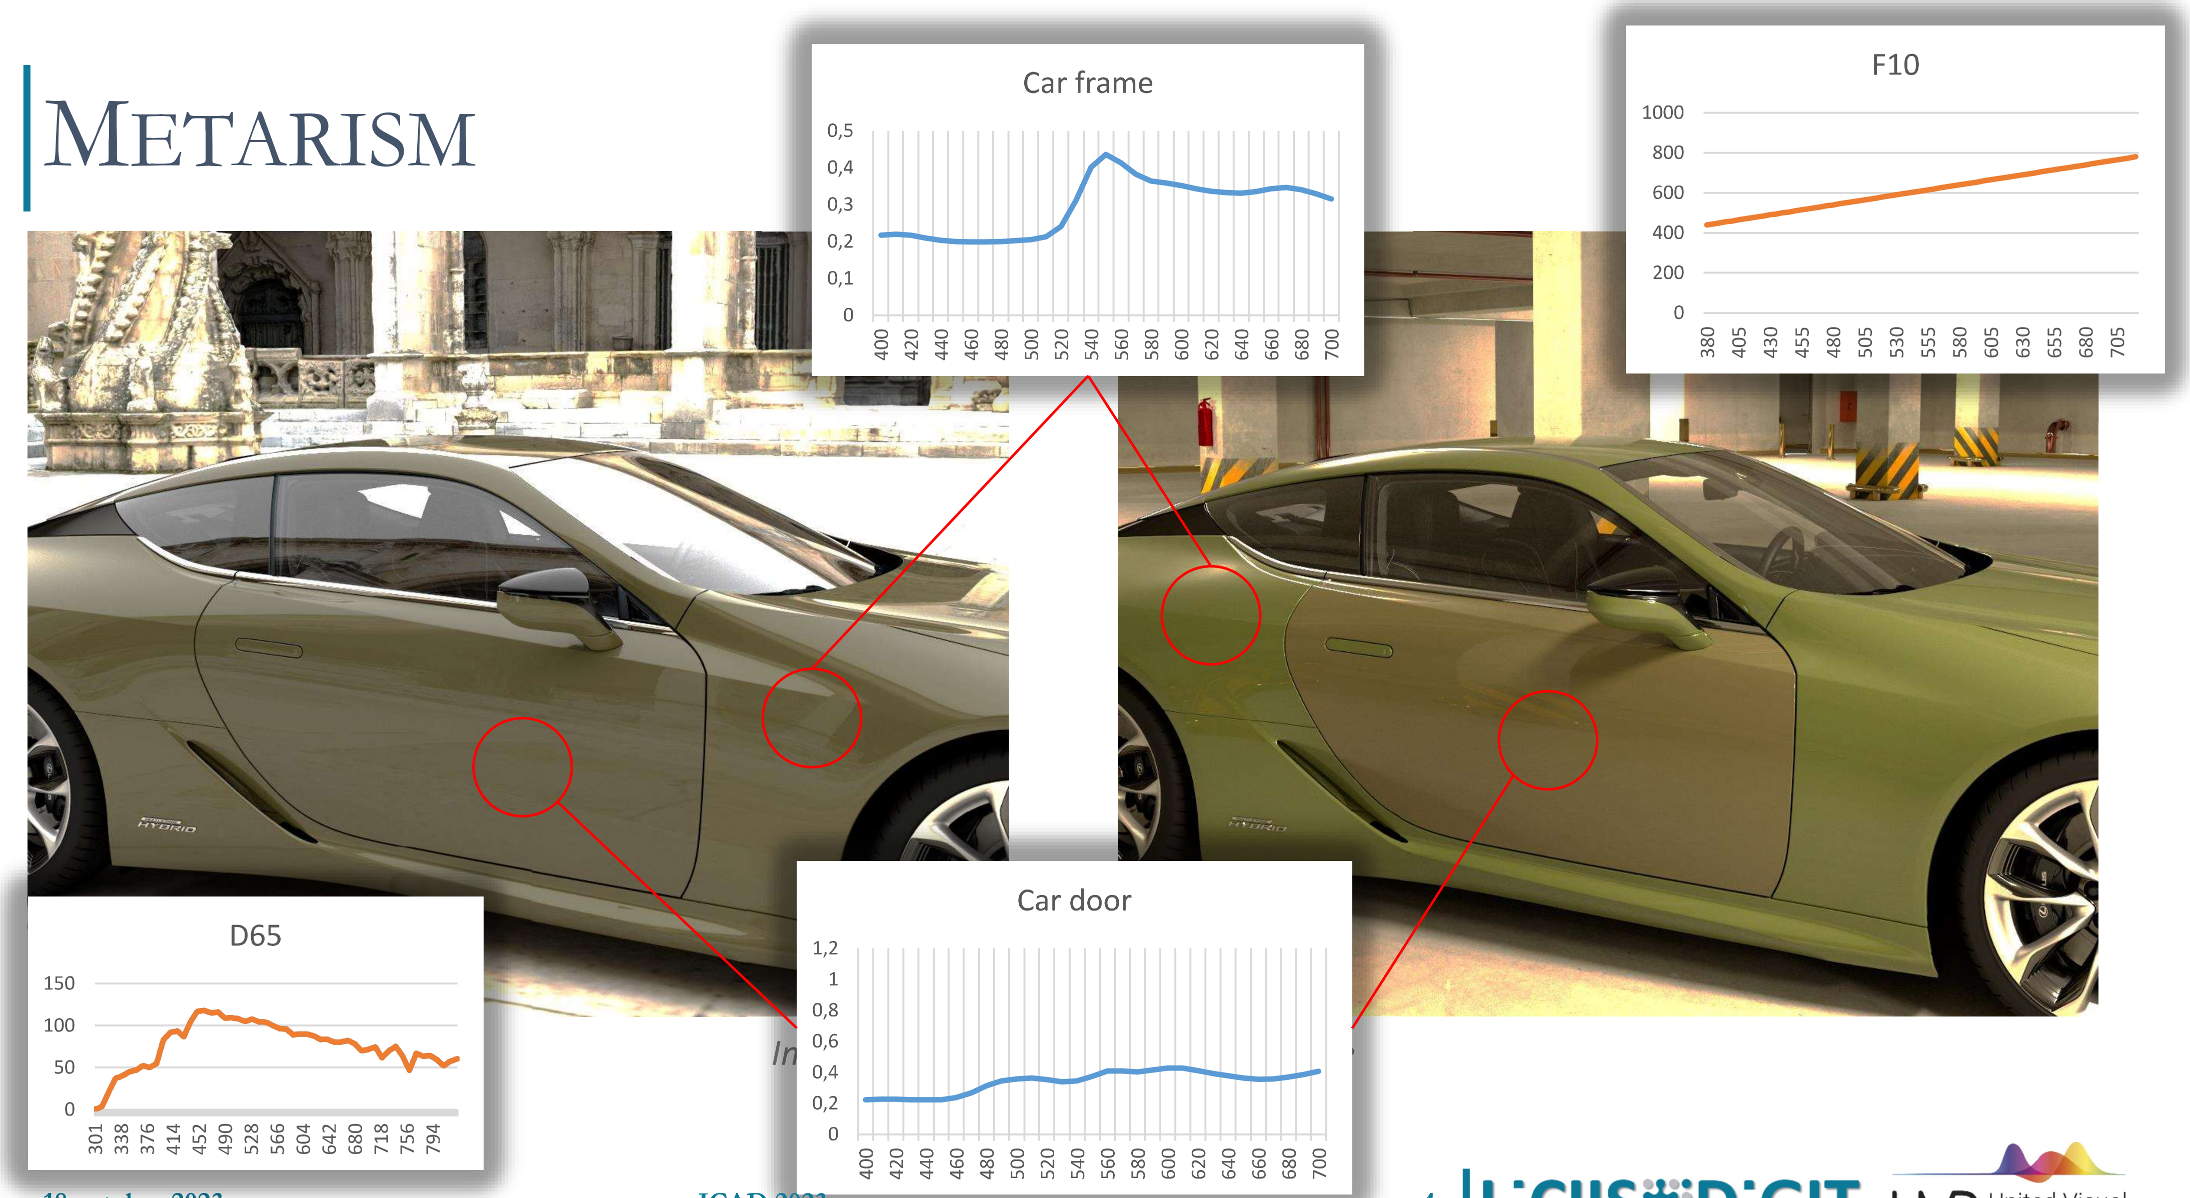
Car frame
0,5 0,4 0,3 0,2 0,1 0
400 420 440 460 480 500 520 540 560 580 600 620 640 660 680 700
F10
1000 800 600 400 200 0
380 405 430 455 480 505 530 555 580 605 630 655 680 705
D65
150 100 50 0
301 338 376 414 452 490 528 566 604 642 680 718 756 794
Car door
1,2 1 0,8 0,6 0,4 0,2 0
400 420 440 460 480 500 520 540 560 580 600 620 640 660 680 700

General context

- Predictive visualisation – in interactive time – of complex materials for industry (CA^{2}O)
  - Coupling optical simulation and machine learning
    - How can rendering methods be combined with Deep Learning?
    - How can they be adapted for HPC architecture?

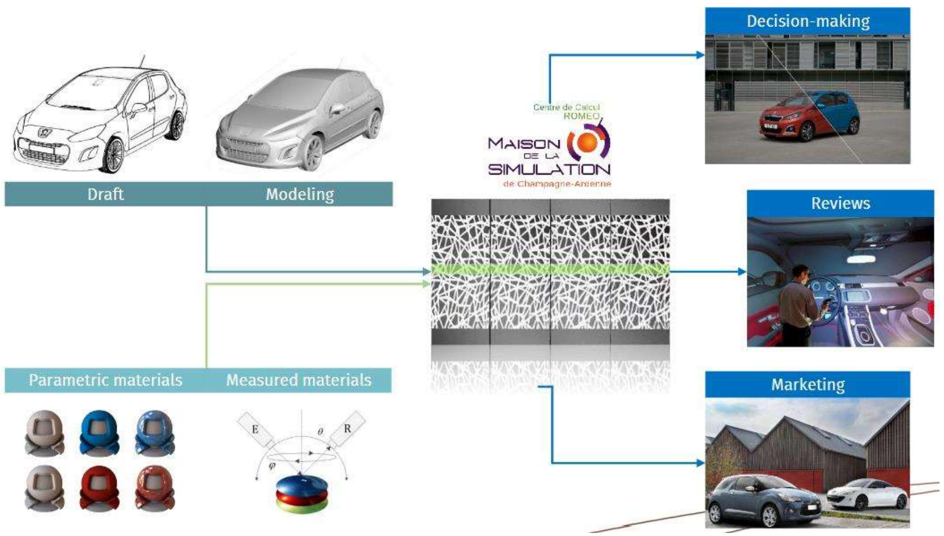

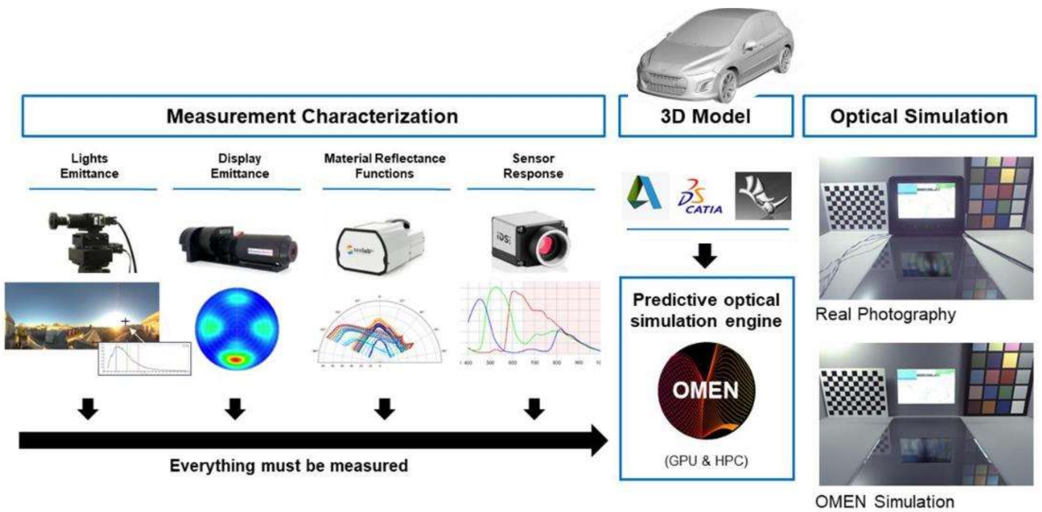

Monte Carlo Rendering

Rendering by UVR Predict Engine

- Rendering equation

$$L_o(x,\omega) = L_e(\ldots) + \int_{\Omega} f_r \cdot L_i(\ldots) \cdot \cos\theta \, d\vec{\omega}_i$$

- Complex analytic resolution
  - Recursive
  - High dimension

- Resolution based Monte Carlo approach

$$\langle F \rangle = \frac{1}{N}\left[\frac{f(X)}{p(X)}\right] \approx \int f(x)dx$$

- Approximate the solution from a number of sample

OUR GOAL

Conventional pipeline
LUCE's pipeline

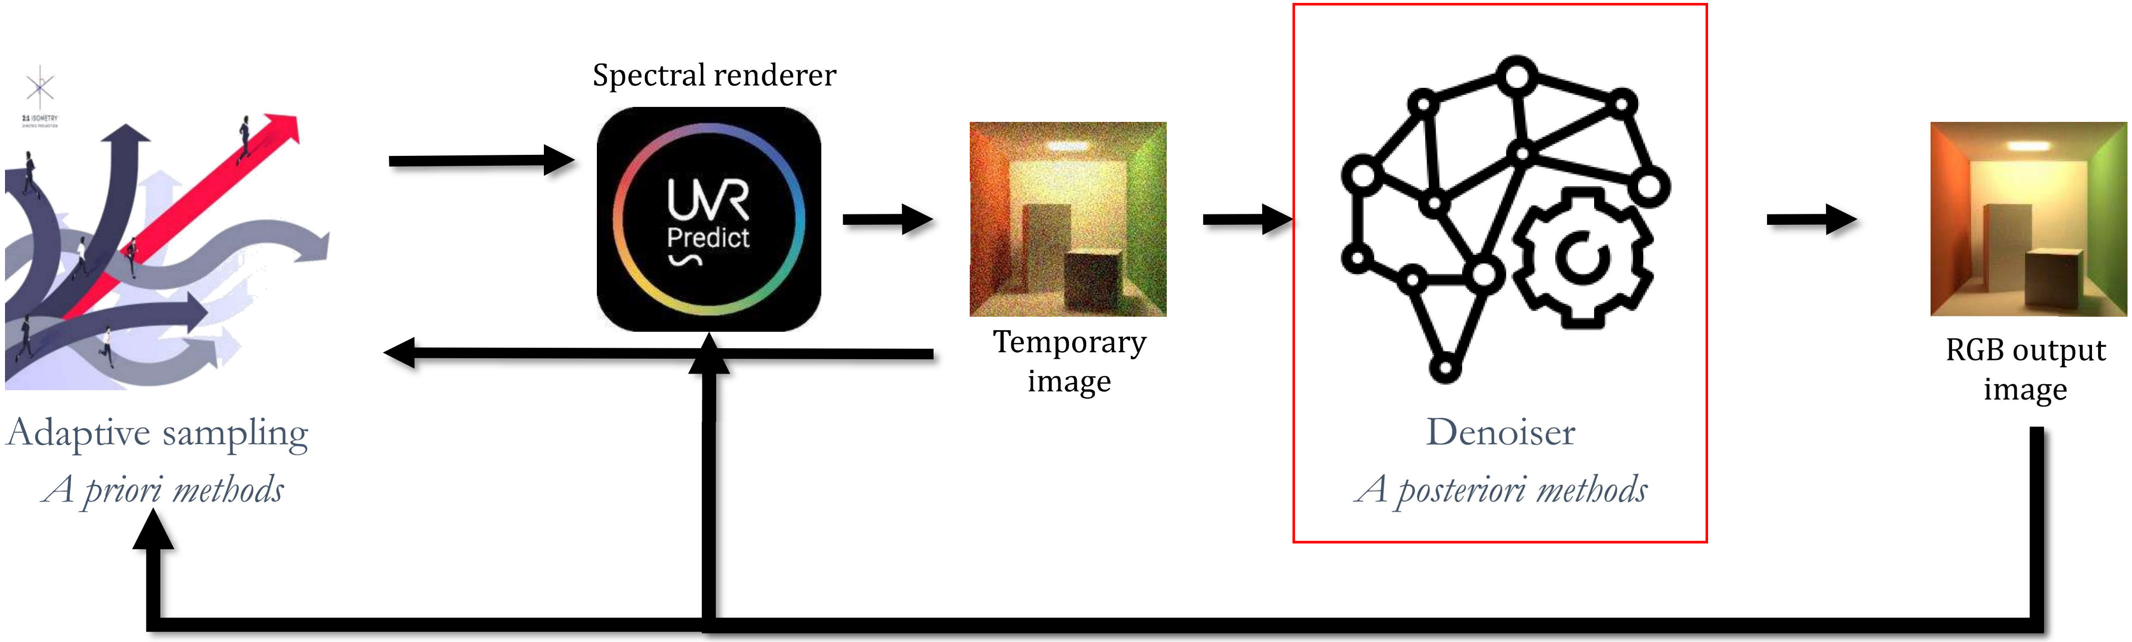

Available data from Spectral Monte Carlo Rendering

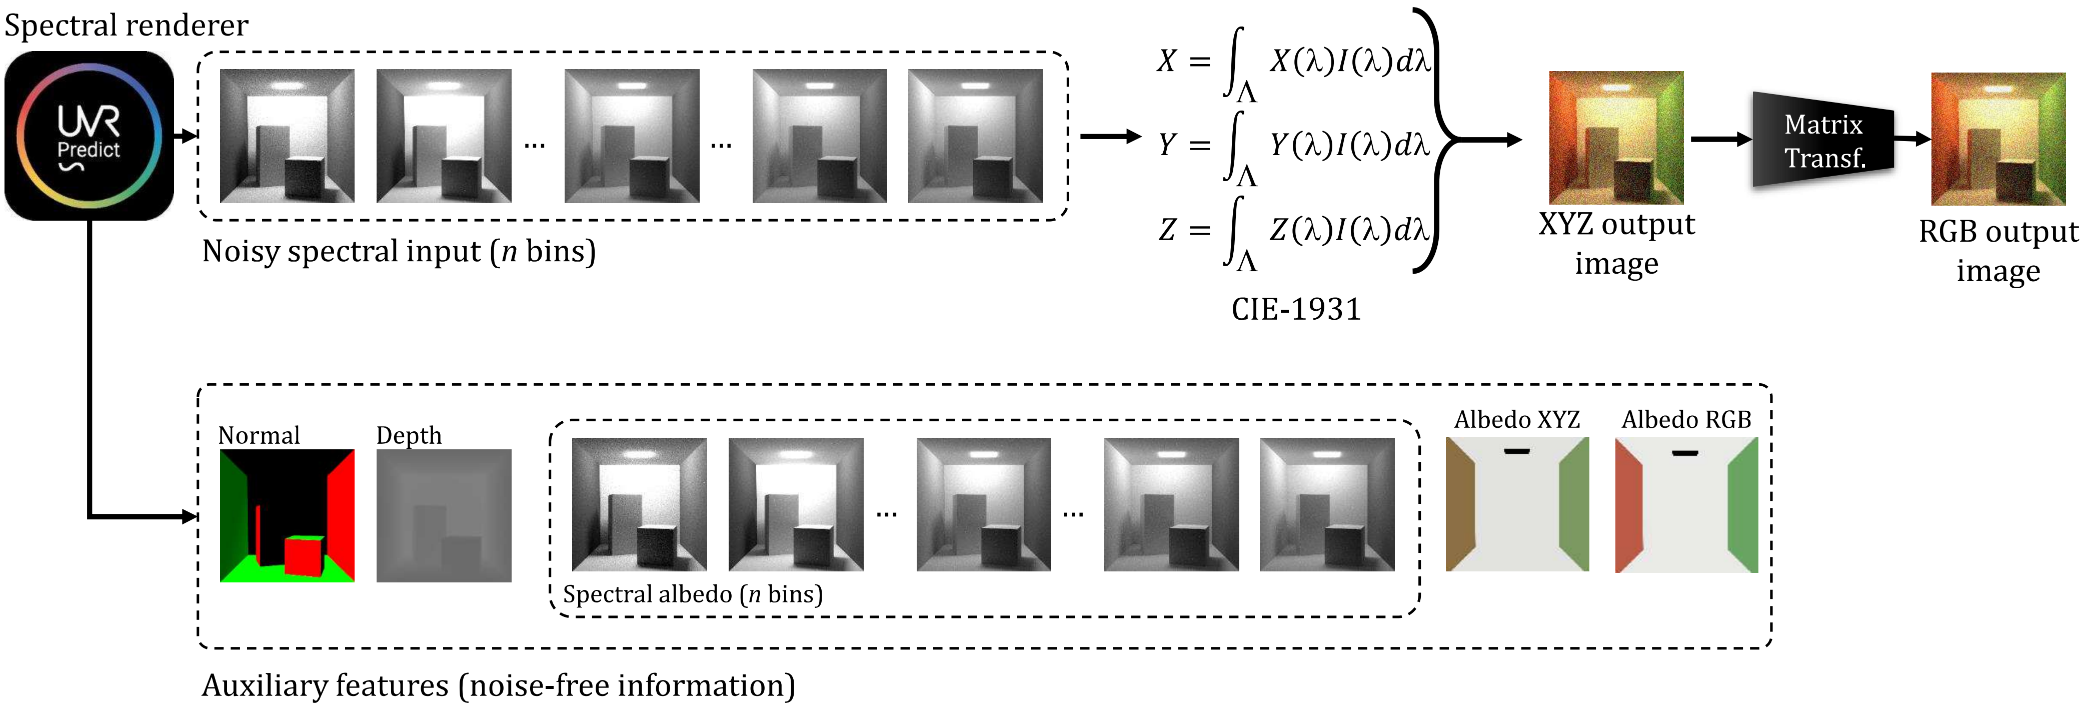

STATE OF THE ART - DENOISING

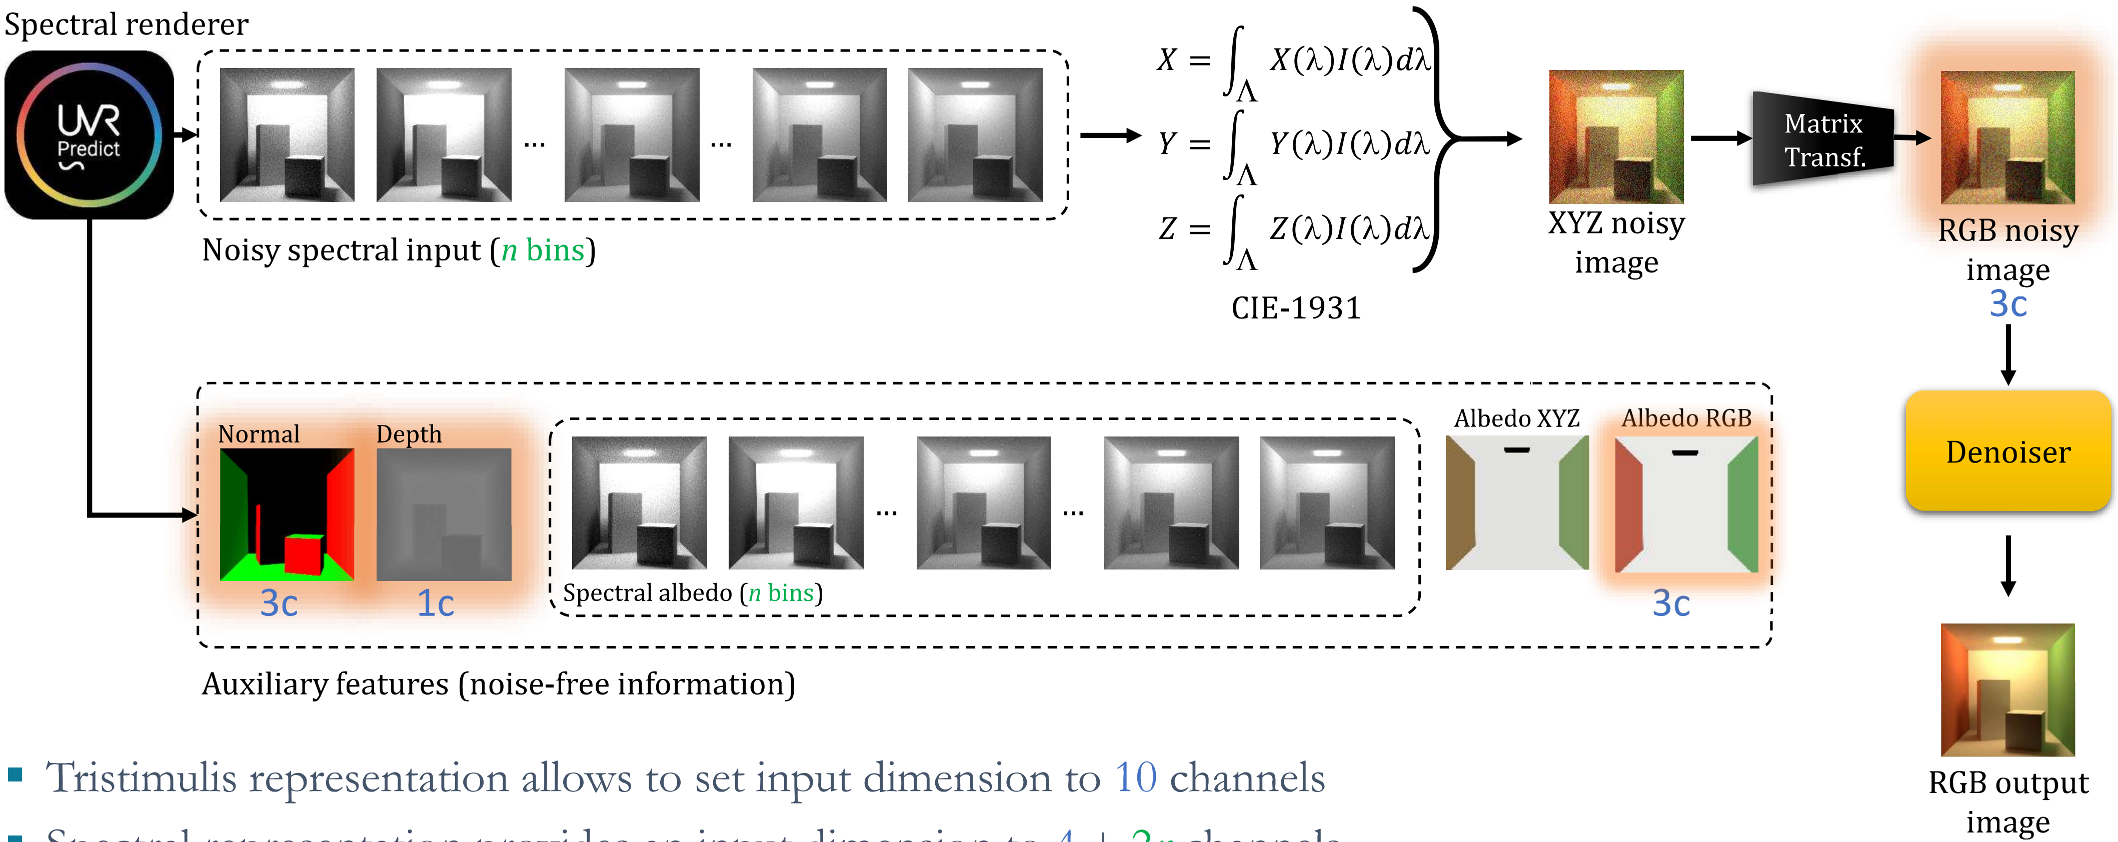

- Tristimulis representation allows to set input dimension to 10 channels
- Spectral representation provides an input dimension to $4 + 2n$ channels

STATE OF THE ART – OFF-THE-SHELL DENOISER

- Denoise the tristimulis representation (RGB, XYZ…)
  - NVIDIA OptiX denoiser
  - Intel Open Image Denoiser (IOID)
- Apply denoiser with a well sample rate
- Limits
  - Compress all spectrum information into 3 dimensions
  - Bring chromatic aberration
- Questions
  - Has off-the-shell denoiser train on spectral rendering?
  - Do spectral denoisers exist?

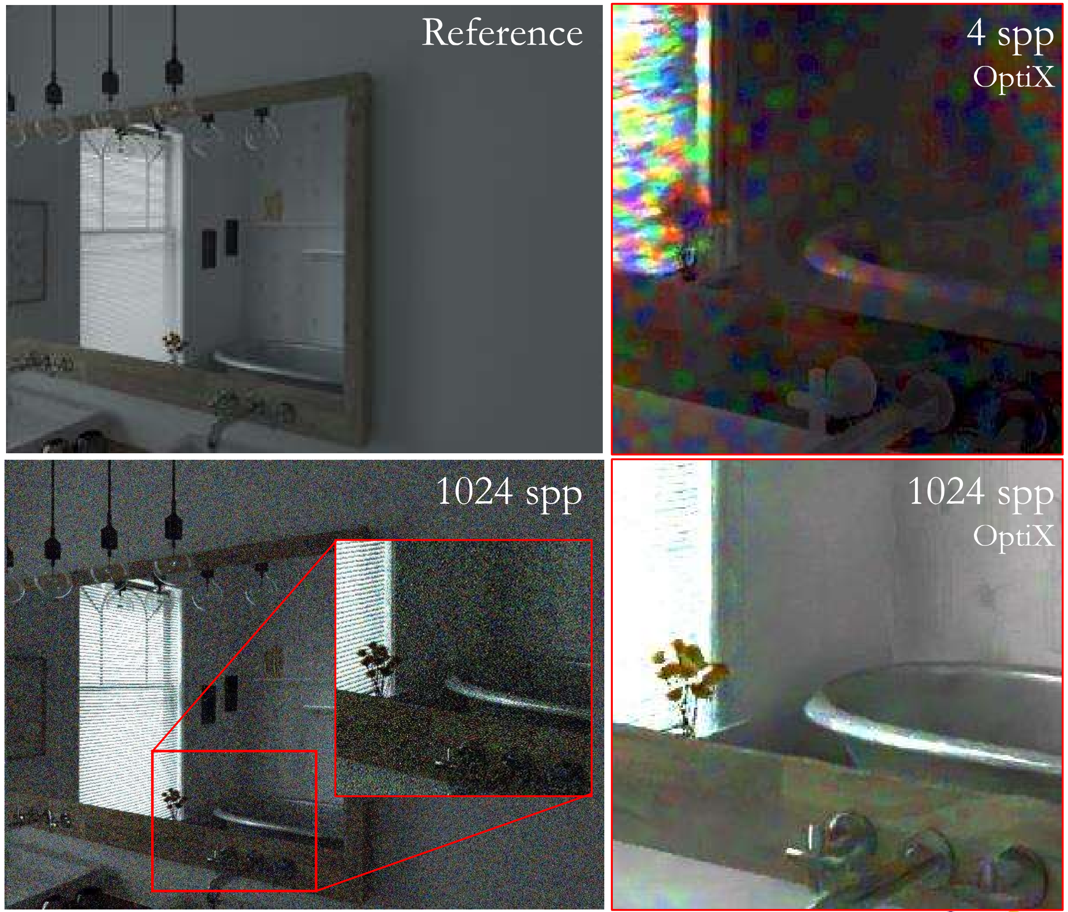

State of the art – Monte Carlo Denoising

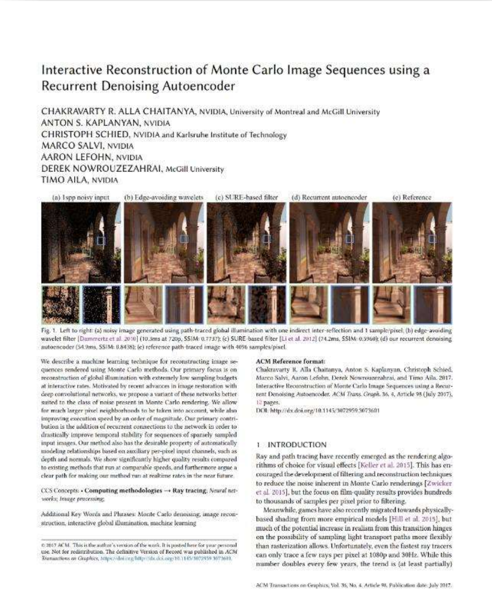

Interactive Reconstruction of Monte Carlo Image Sequences using a Recurrent Denoising Autoencoder

CHAKRAVARTY R. ALLA CHAITANYA, NVIDIA, University of Montreal and McGill University
ANTON S. KAPLANYAN, NVIDIA
CHRISTOPH SCHIED, NVIDIA and Karlsruhe Institute of Technology
MARCO SALVI, NVIDIA
AARON LEFOHN, NVIDIA
DEREK NOWROUZEZAHRAI, McGill University
TIMO AILA, NVIDIA

[Chaitanya et al. 2017]

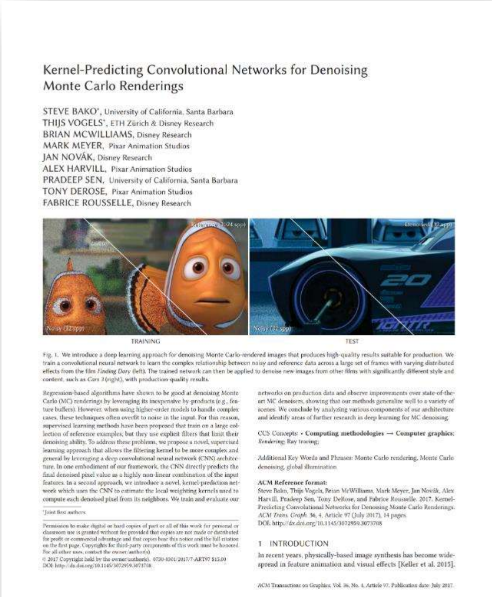

Kernel-Predicting Convolutional Networks for Denoising Monte Carlo Renderings

STEVE BAKO, University of California, Santa Barbara
THIJS VOGELS, ETH Zürich & Disney Research
BRIAN MCWILLIAMS, Disney Research
MARK MEYER, Pixar Animation Studios
JAN NOVÁK, Disney Research
ALEX HARVILL, Pixar Animation Studios
PRADEEP SEN, University of California, Santa Barbara
TONY DEROSE, Pixar Animation Studios
FABRICE ROUSSELLE, Disney Research

[Bako et al. 2017]

DEMC: A Deep Dual-Encoder Network for Denoising Monte Carlo Rendering

[Yang et al. 2019]

Deep Combiner for Independent and Correlated Pixel Estimates

JONGHEE BACK, Gwangju Institute of Science and Technology, South Korea
BINH-SON HUA, VinAI Research, Vietnam and VinUniversity, Vietnam
TOSHIYA HACHISUKA, The University of Tokyo, Japan
BOCHANG MOON, Gwangju Institute of Science and Technology, South Korea

[Back et al. 2020]

Direct prediction	Kernel prediction	Direct prediction	Kernel prediction
~1-4 spp*	~64-128 spp*	~1-4 spp*	independent of spp*
Loss details	Details reconstructed	Reduce artifact induce by auxilairy features	Improve diffuse details reconstruction
		Less of diffuse details lost	

spp* = 1 sample in RGB

LICIIS DiGIT
United Visual Researchers

DATA MANAGING

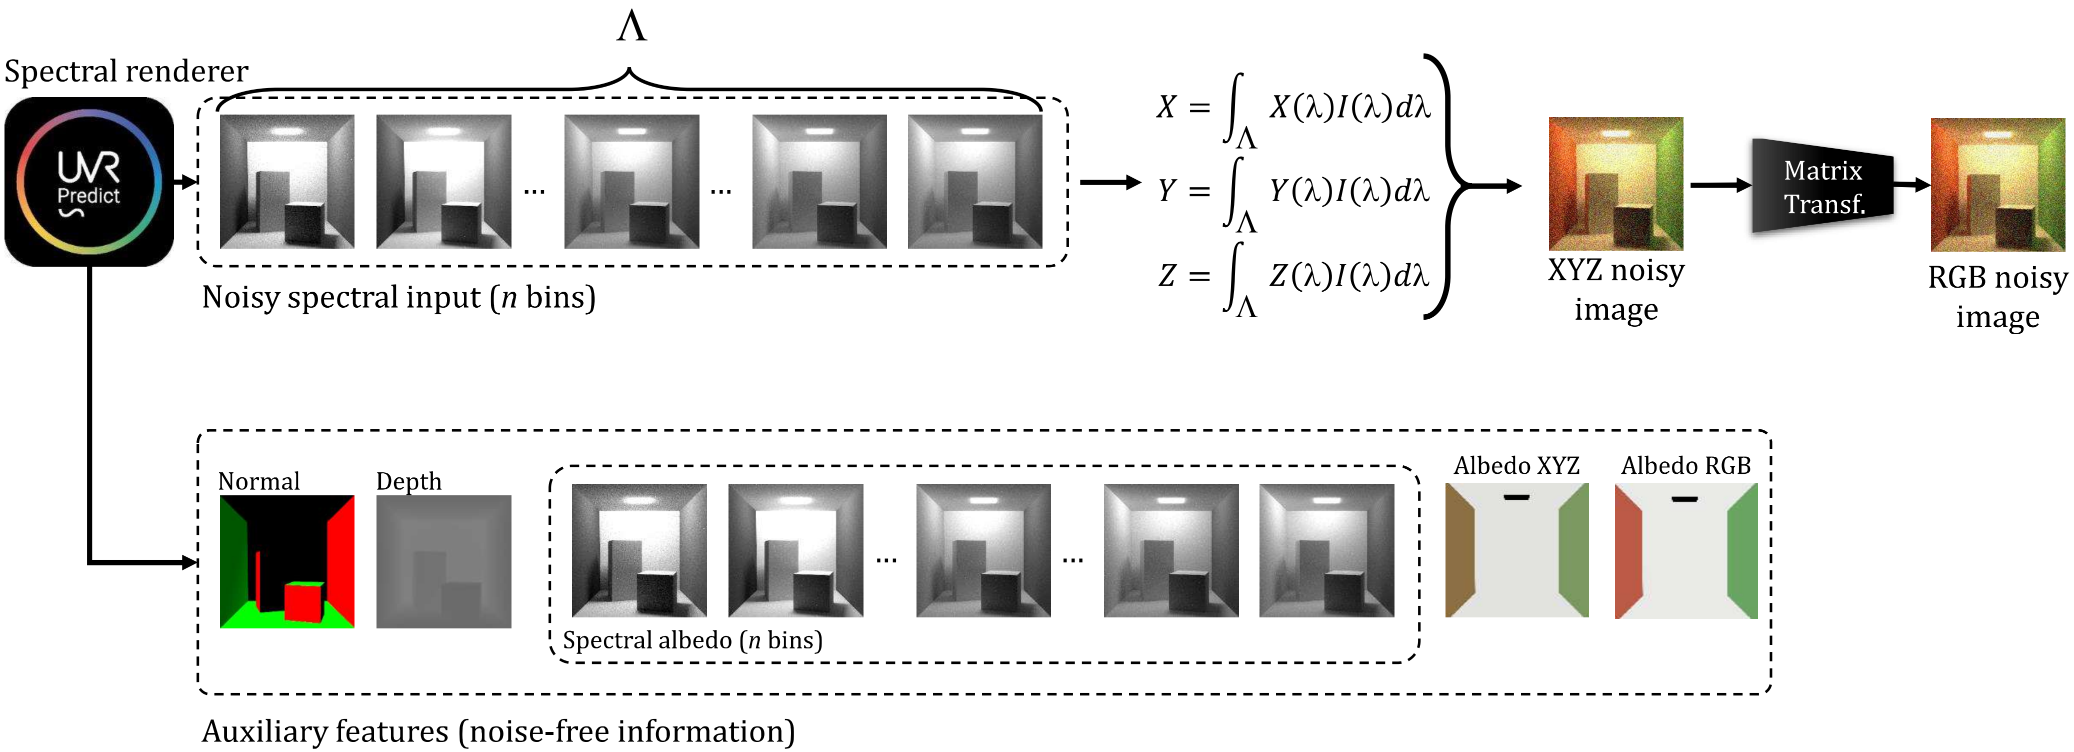

DATA MANAGING

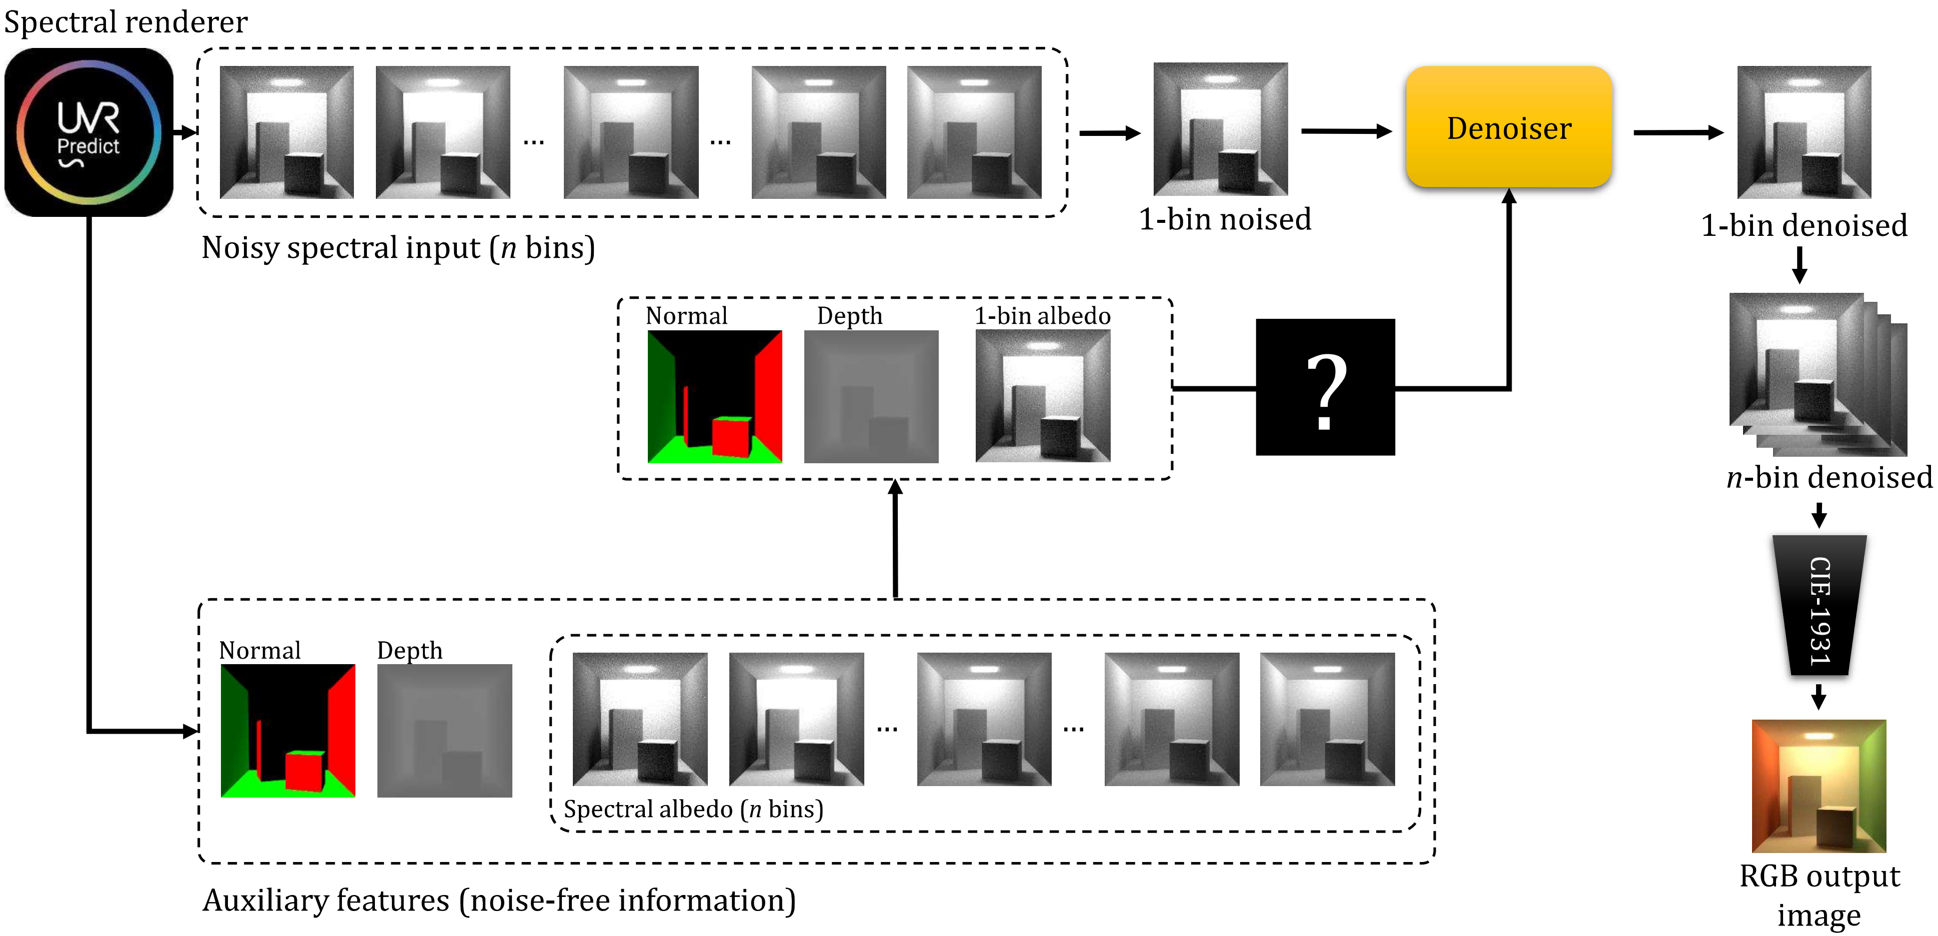

Spectral compression

- Aim to provide full light spectrum information
  - With fixed set input dimension
- Tri-chromatic representation
  - Set to 3 dimensions
  - Provide a displayable information
  - Change the nature of data representation
- PCA
  - No truncate spectral information
  - Compress without data lost
  - The 3 first PC represent 98.2% of initial information

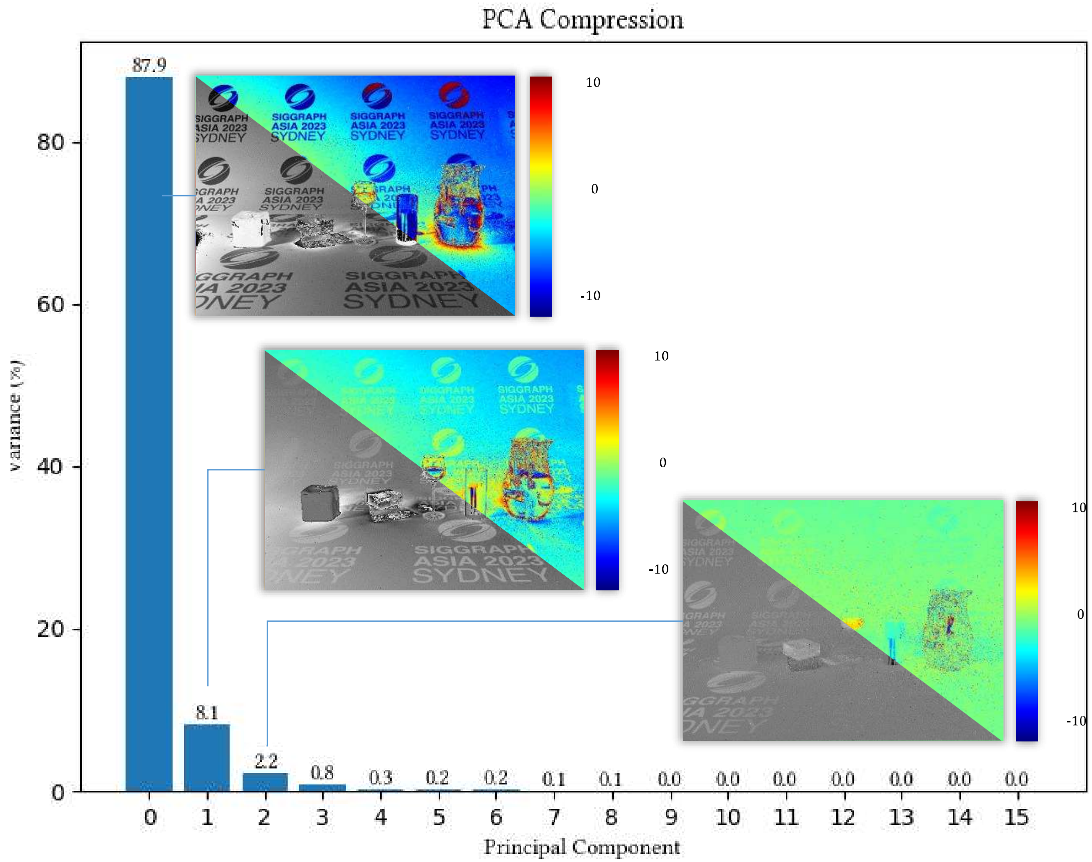

OUR CONTRIBUTION

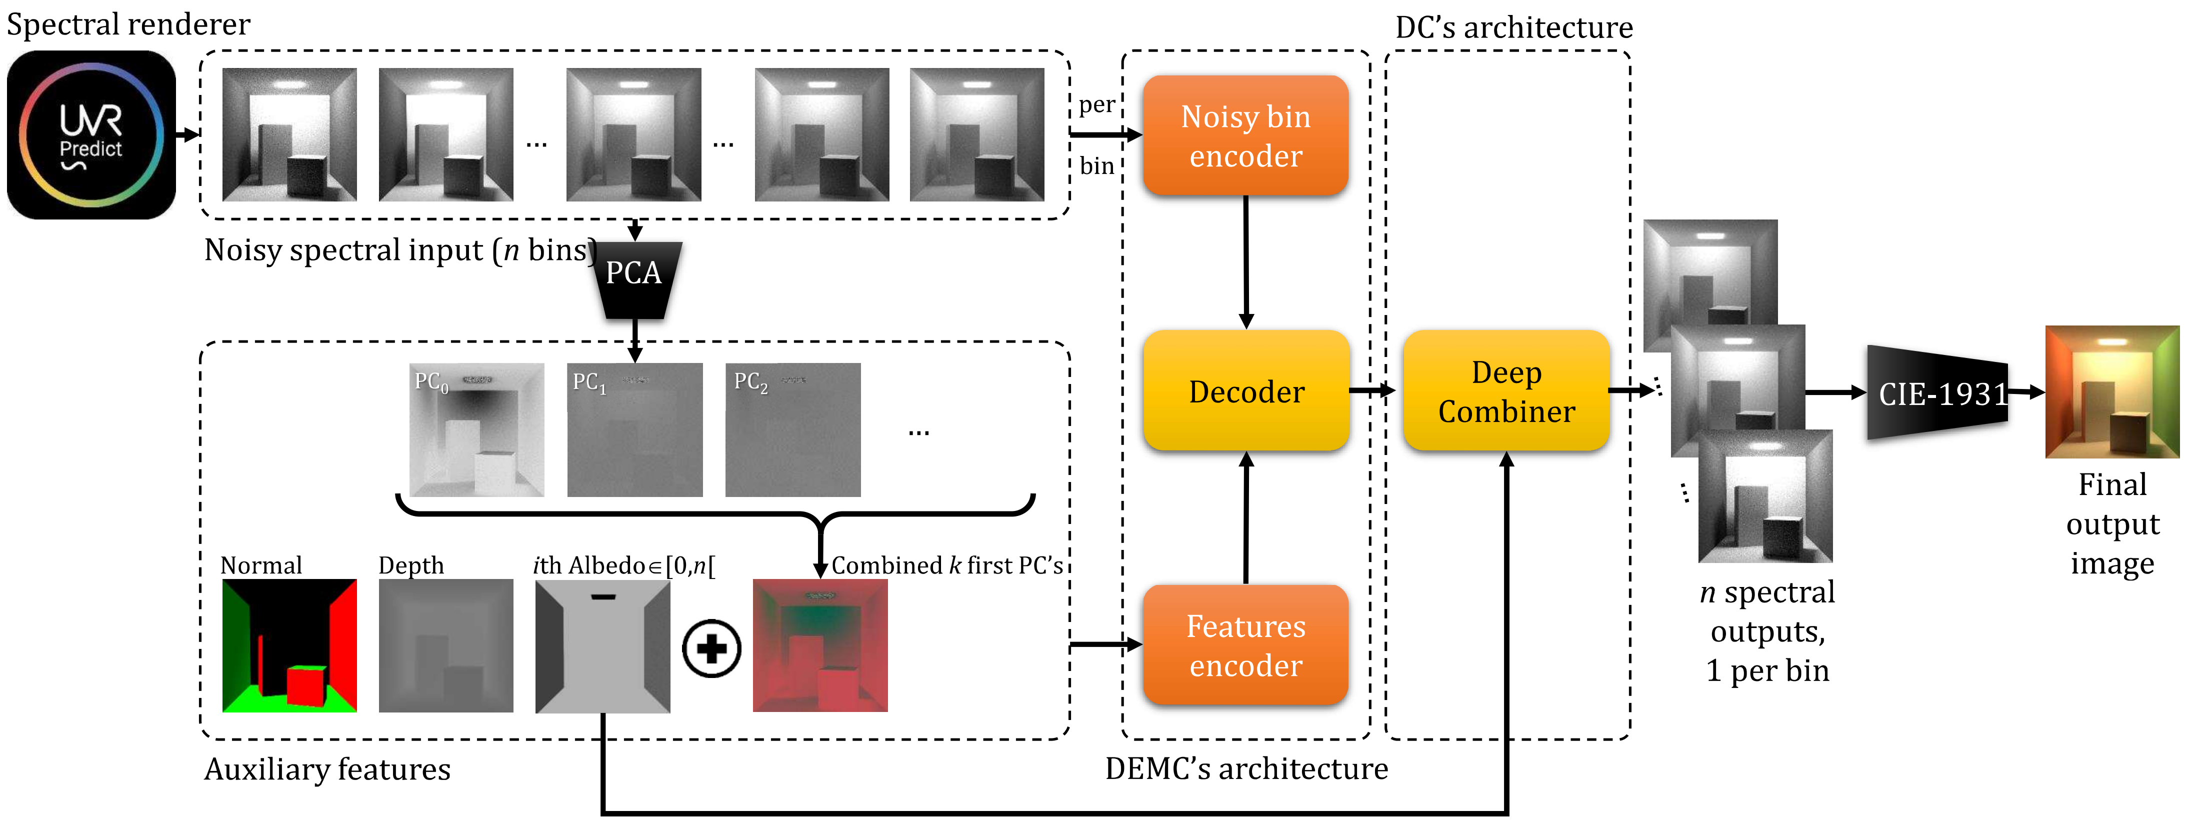

TRAINING INFORMATION

- Loss function: SMAPE
- Dataset
  - 5500 image peers (4554 for training, 946 for testing)
  - 23 scenes (22 points of views)
  - Resolution of image's crop 128×128
- Training parameters (for each network)
  - Epochs : 5000
  - Learning rate : 10^{-4}
  - Optimizer : Adam
- Training time: ~3 days
- 4 GPU Nvidia Tesla P100 (16 Go VRAM)

RESULTS

Noisy image

Reference image

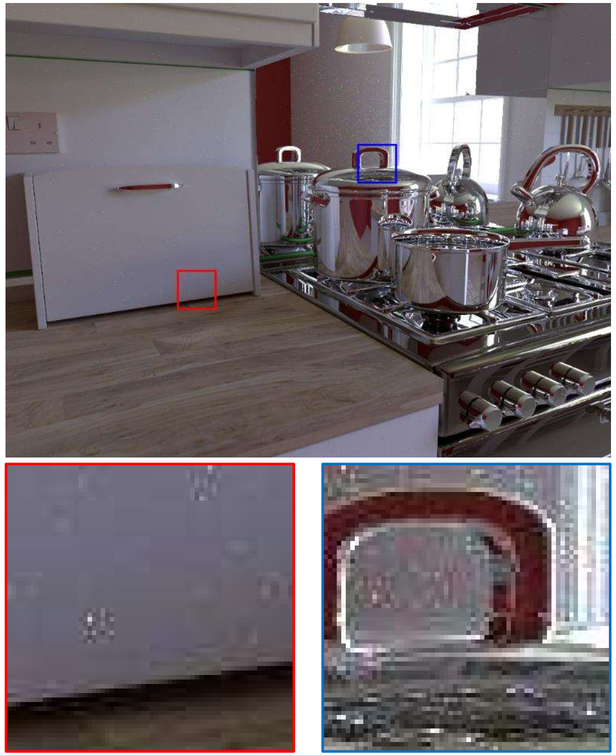

Denoised image with our method

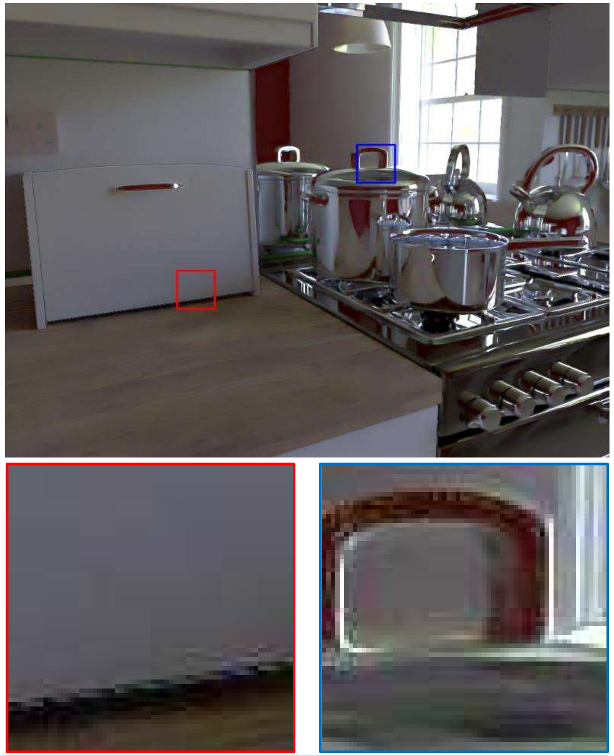

SPP:	1024	1 M	1024
Time:	89 s on 1 GPU	≈ 36 h on 4 GPUs	89.16 s on 1 GPU
RelMSE:	0.70	GT	0.0313

RESULTS

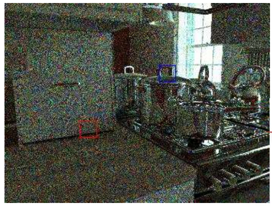

Noisy image
1024 SPP

Reference image
1M SPP

	Intel	OptiX	Ours

RelMSE:	0.176	0,205	0.0313

Conclusion & Future works

- Contributions
  - First spectral denoiser
    - based on the spectral bins processing
    - Tailoring input, auxiliary and output features to favorize spectral information
  - Out-perform off-the-shell denoiser (with RelMSE measure)
  - Submitted to Eurographics 2024

- Future works
  - Improve border reconstruction
  - Improve albedo computation to reduce artifacts

Thanks for your attention!