SIGGRAPH ASIA 2016 MACAO

Sponsored by acm

Image-space Control Variates for Rendering

Fabrice Rousselle[1] Wojciech Jarosz[2] Jan Novák[1]

[1]Disney Research [2]Dartmouth College

SA2016.SIGGRAPH.ORG

SPATIO-TEMPORAL COHERENCE

Original scene by Wig42, downloaded from blendswap.com

SPATIO-TEMPORAL COHERENCE

Original scene by Wig42, downloaded from blendswap.com

SPATIO-TEMPORAL COHERENCE

Original scene by Wig42, downloaded from blendswap.com

SPATIO-TEMPORAL COHERENCE

Original scene by Wig42, downloaded from blendswap.com

SPATIO-TEMPORAL COHERENCE

Original scene by Wig42, downloaded from blendswap.com

SPATIO-TEMPORAL COHERENCE

Original scene by Wig42, downloaded from blendswap.com

PREVIOUS WORKS

- Irradiance caching
  [Ward et al. 1988, …]
- Photon mapping
  [Jensen 1995, …]
- Image-space denoising
  [Rushmeier and Ward 1994, …]
- Many more…

PREVIOUS WORKS

Consistent scene editing [CSE], Günther and Grosch, EGSR 2015

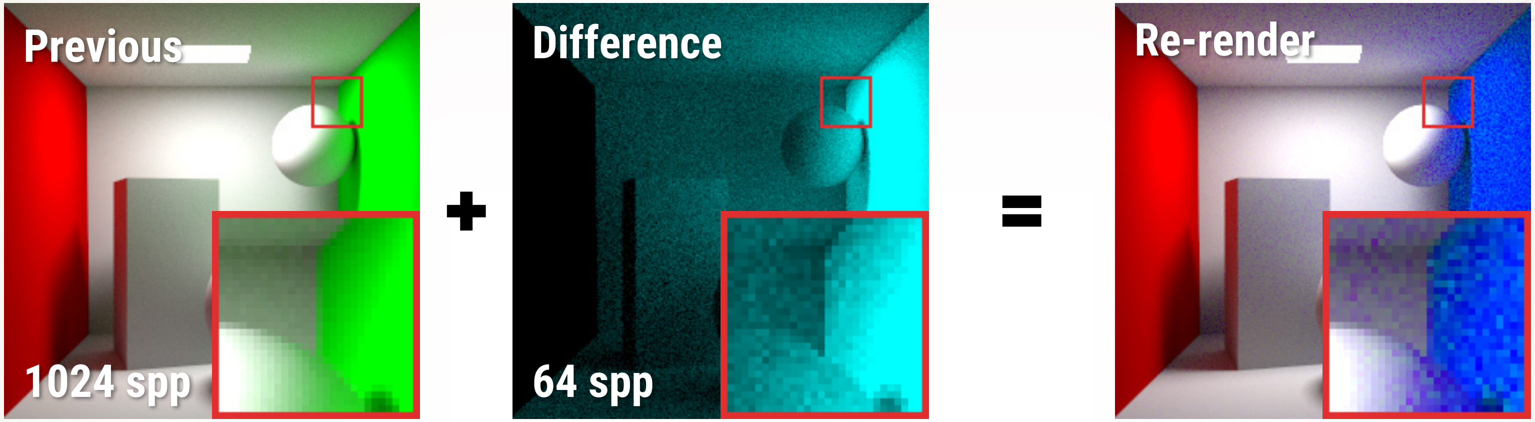

PREVIOUS WORKS

Consistent scene editing [CSE], Günther and Grosch, EGSR 2015

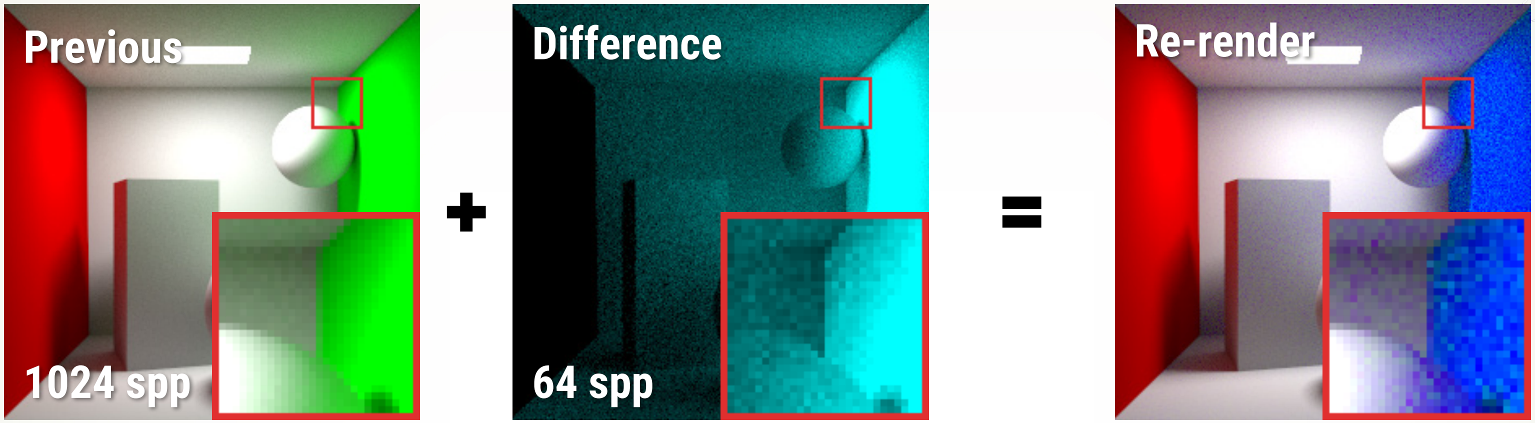

Gradient-domain path tracing [GDPT], Kettunen et al., SIGGRAPH 2015

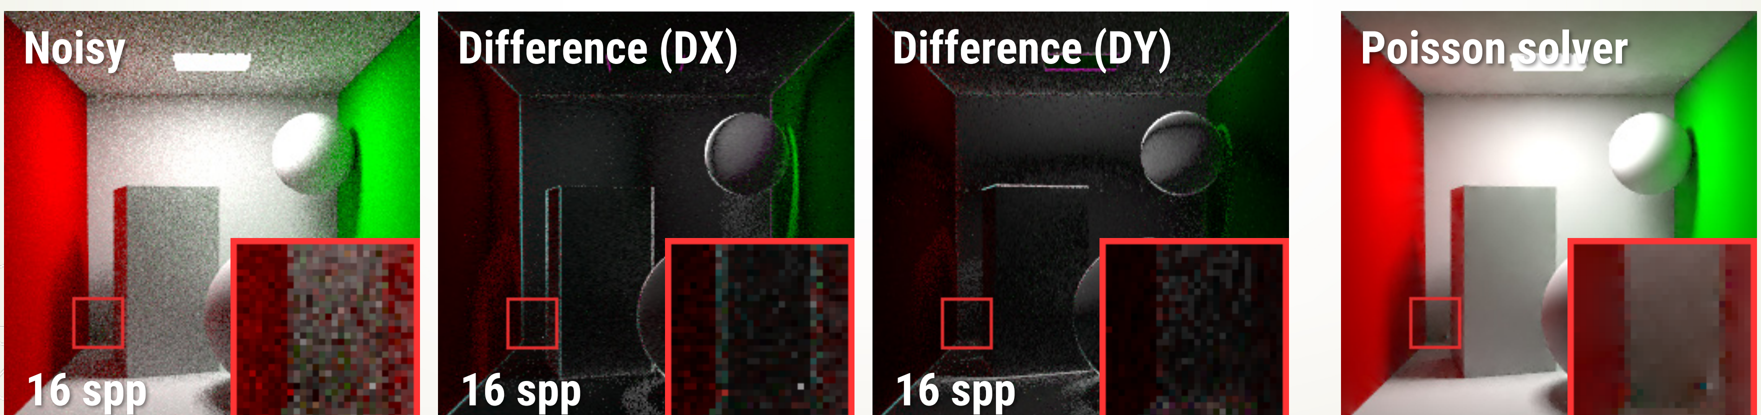

OUR WORK

Scene editing

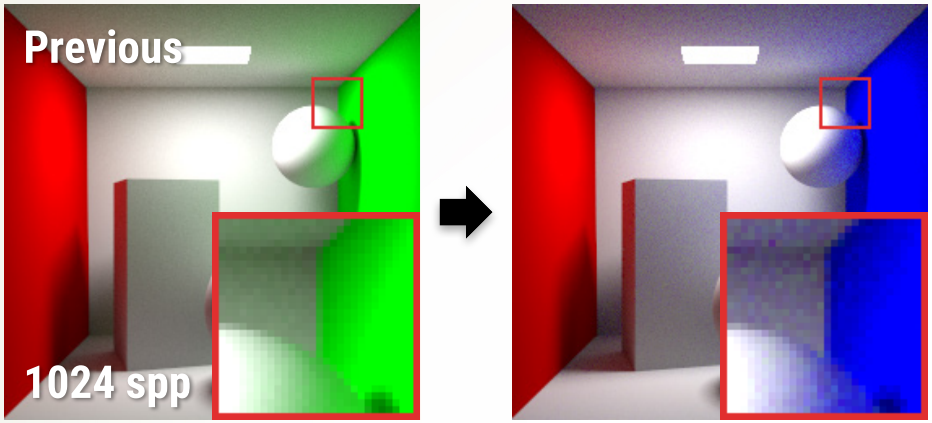

- Unified framework
- Control variates
- Provably optimal variance

Gradient-domain rendering

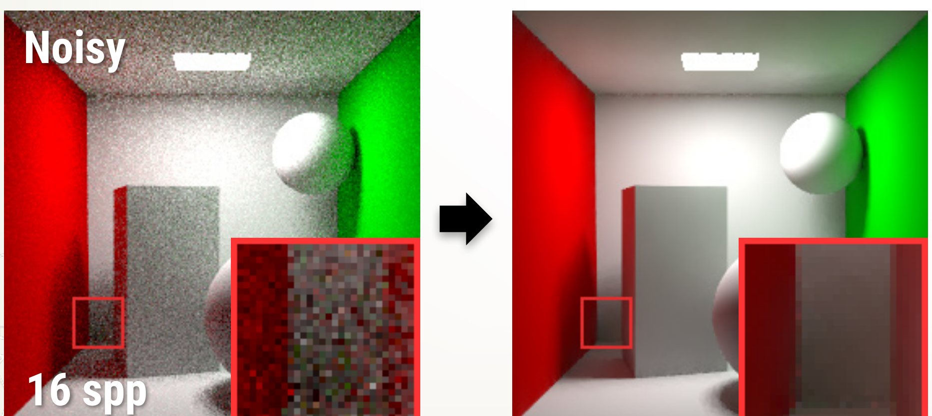

INTEGRATION WITH CONTROL VARIATES

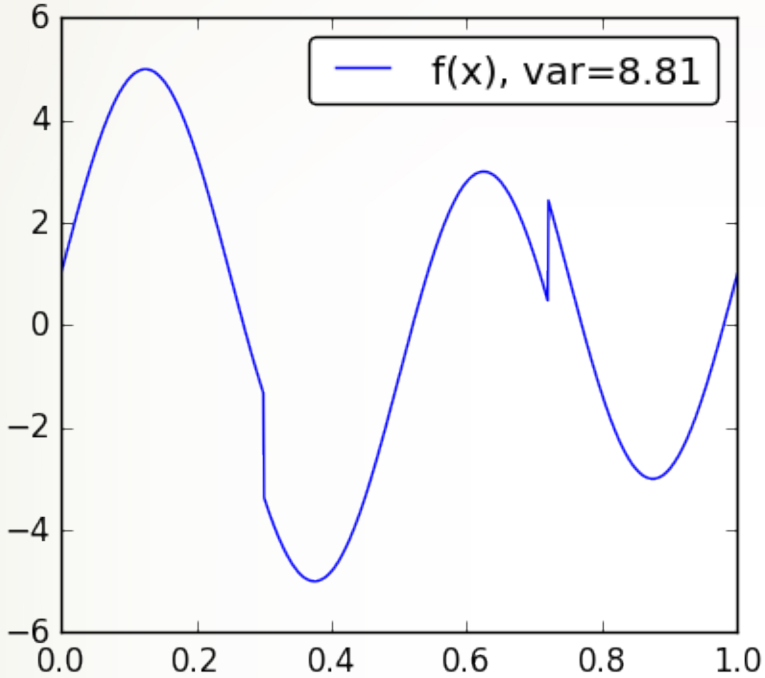

$$\int_{\Omega} f(x)\, dx$$

INTEGRATION WITH CONTROL VARIATES

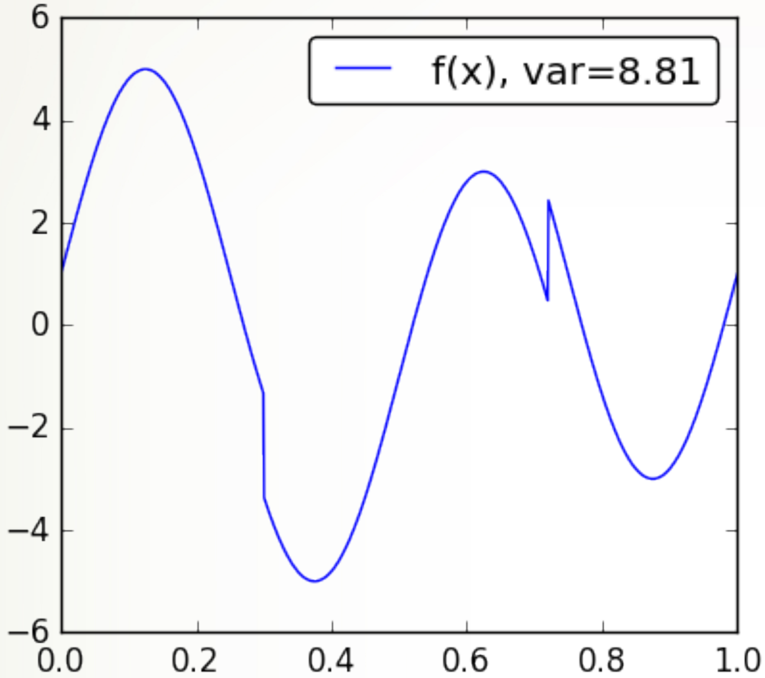

$$\int_\Omega f(x) - g(x)\, dx + G$$

Control variate with known integral

INTEGRATION WITH CONTROL VARIATES

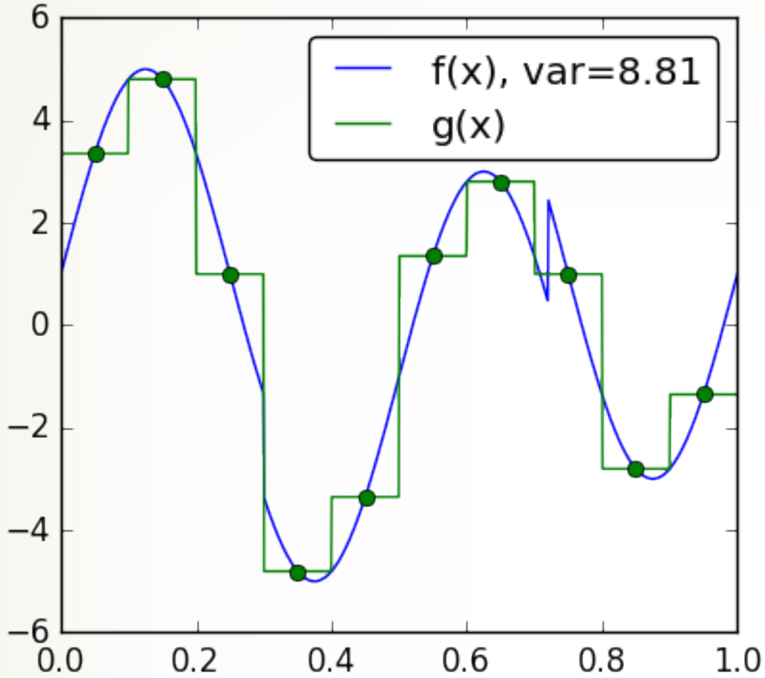

$$\int_{\Omega} f(x) - g(x)\, dx + G$$

INTEGRATION WITH CONTROL VARIATES

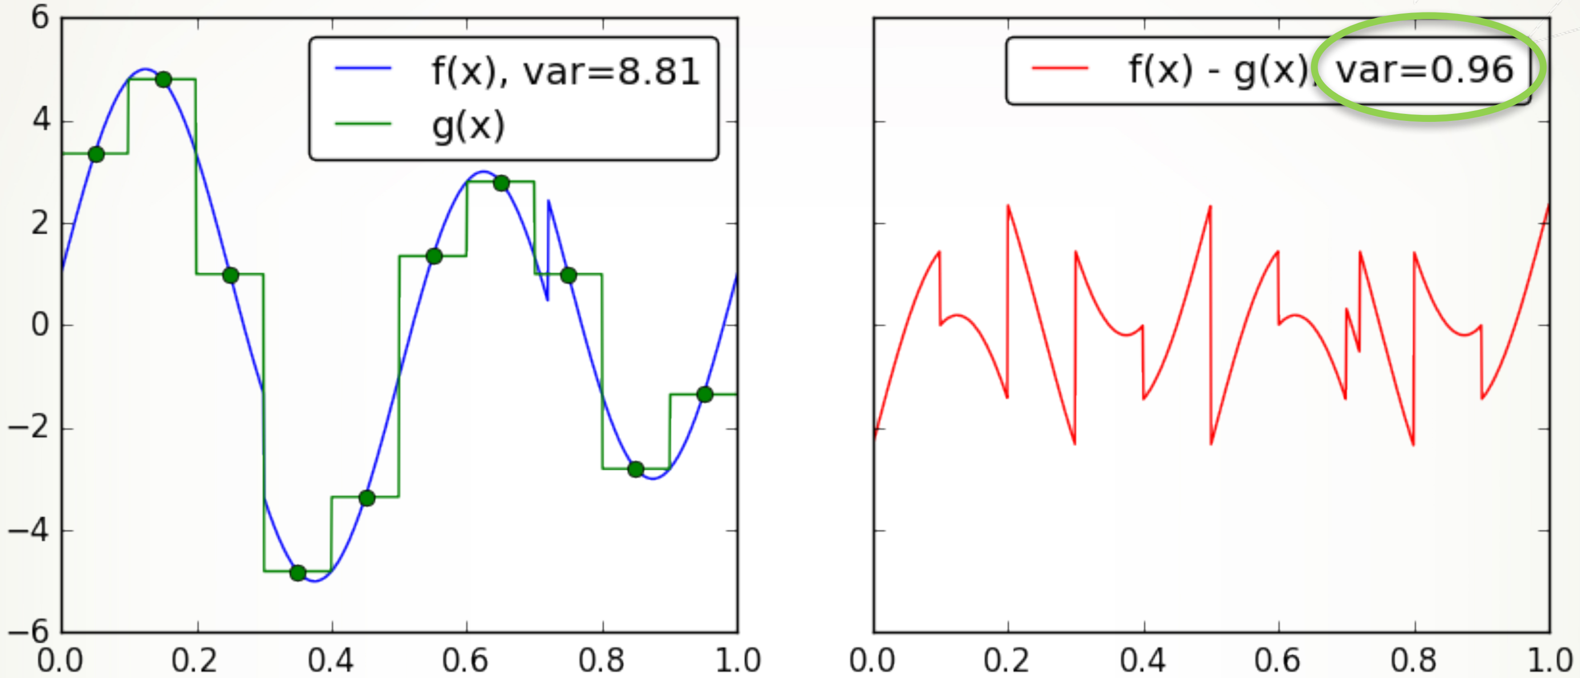

$$\int_{\Omega} f(x) - g(x)\, dx + G$$

INTEGRATION WITH CONTROL VARIATES

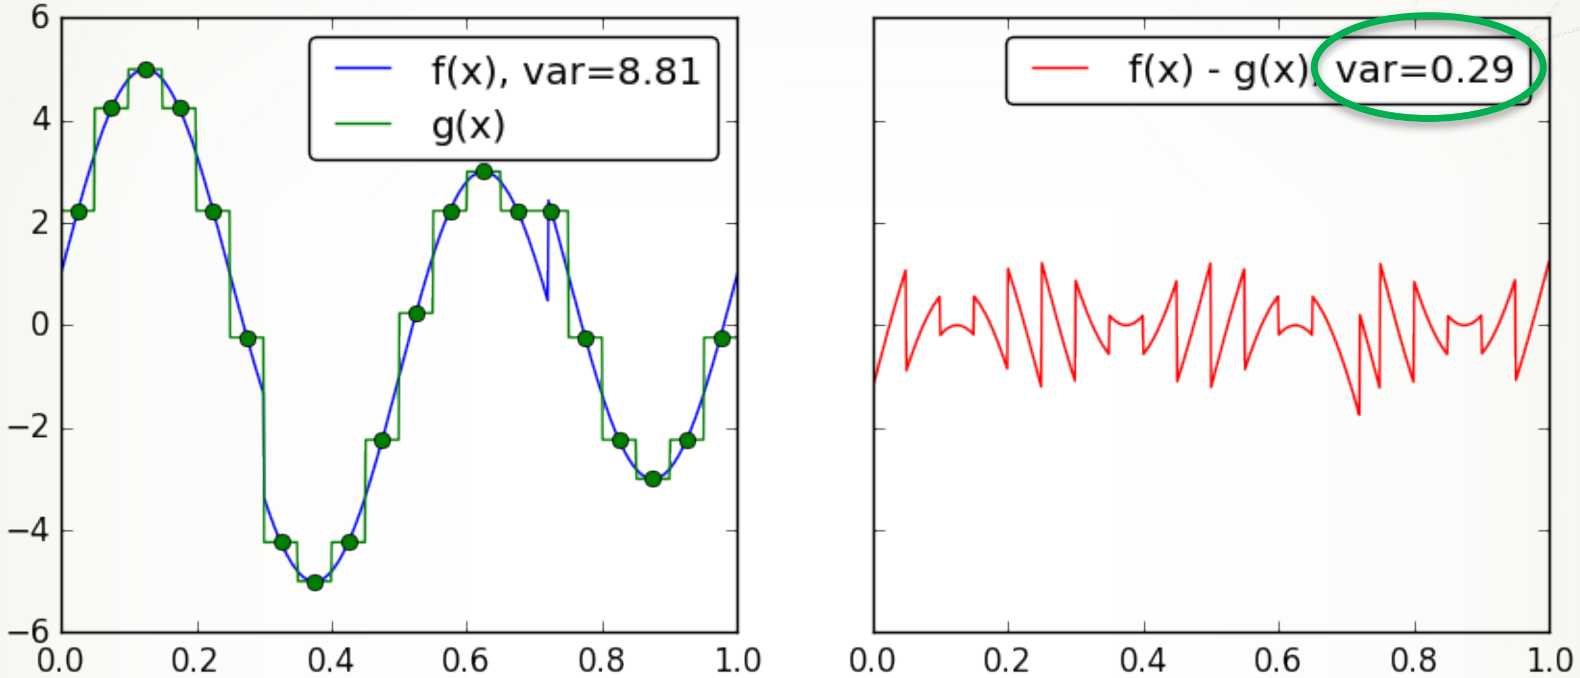

$$\int_{\Omega} f(x) - g(x)\, dx + G$$

INTEGRATION WITH CONTROL VARIATES

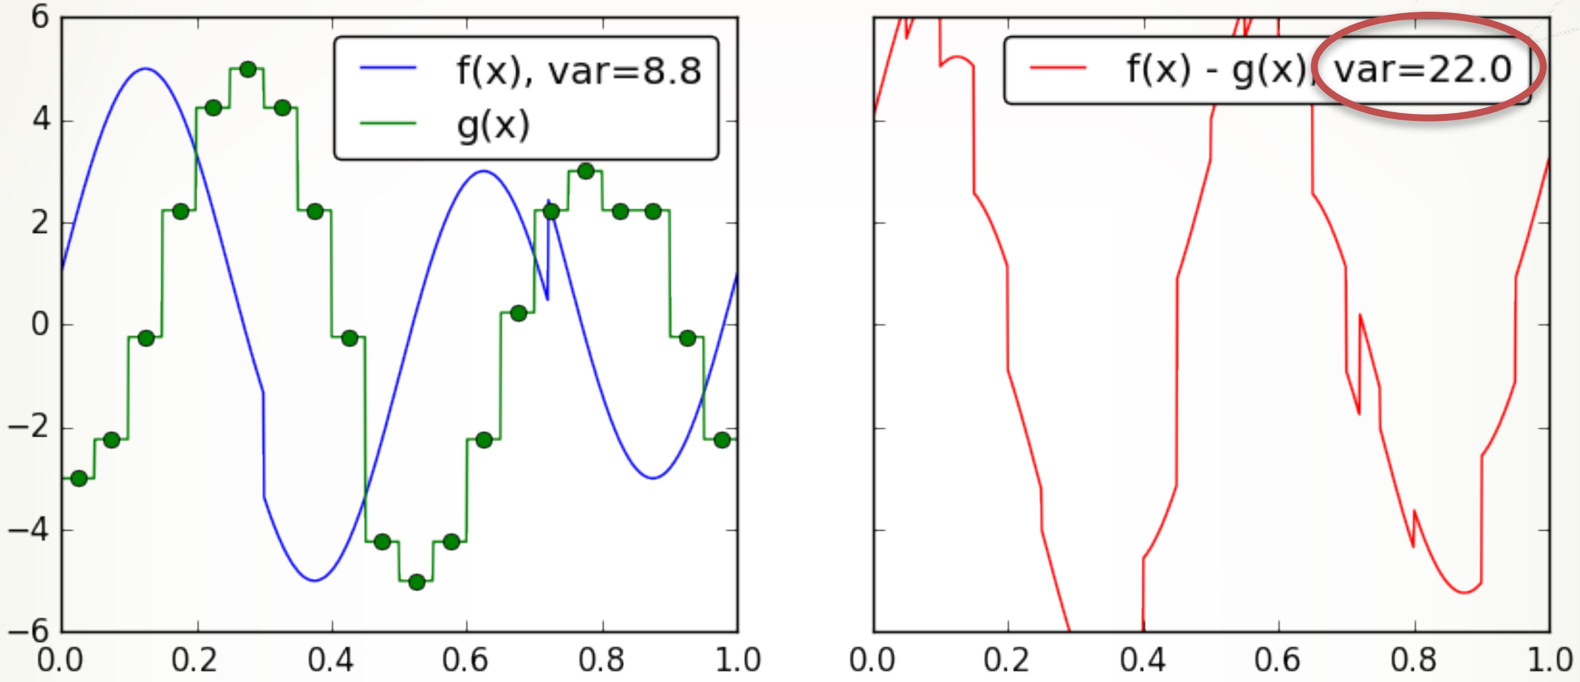

$$\int_{\Omega} f(x) - g(x)\, dx + G$$

SCENE EDITING

Previous

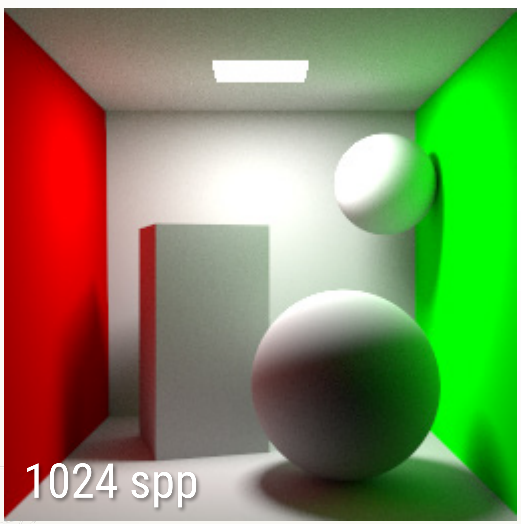

Re-render

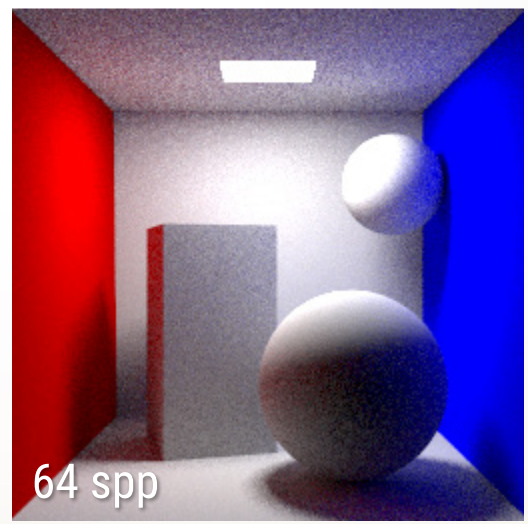

SCENE EDITING

Previous

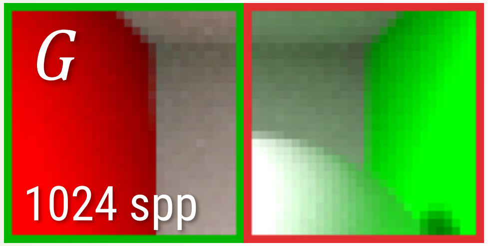

$$\int_{\Omega} f(x) - g(x)\, dx + G$$

Re-render

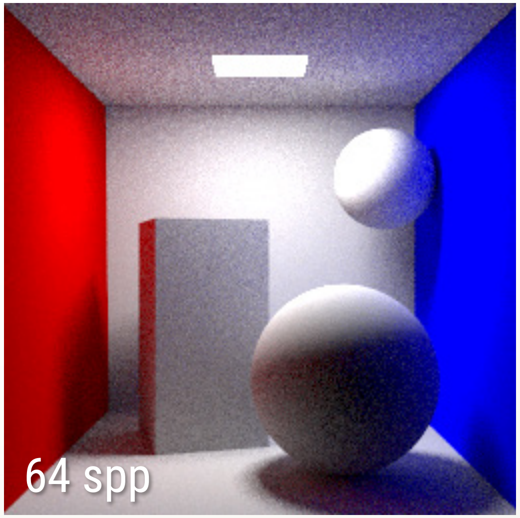

Previous

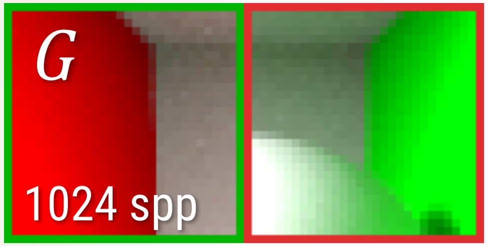

$$\int_{\Omega} f(x) - g(x)\,dx + G$$

Re-render

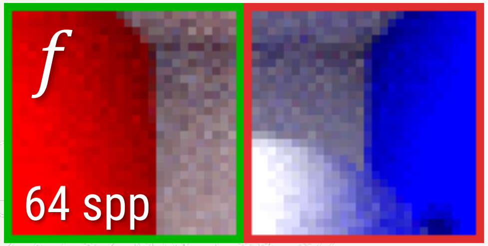

SCENE EDITING

Previous

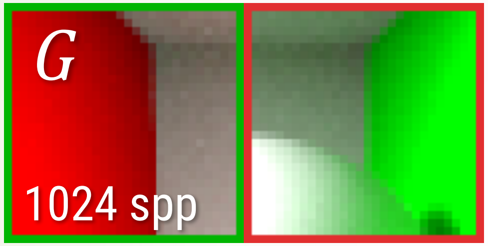

Re-render

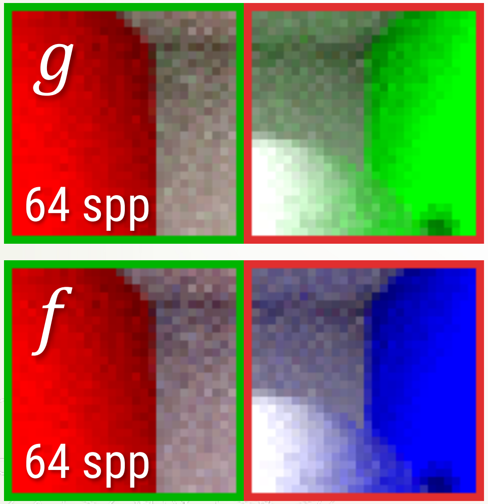

$$\int_{\Omega} f(x) - g(x)\, dx + G$$

Previous

Re-render *same random seed*

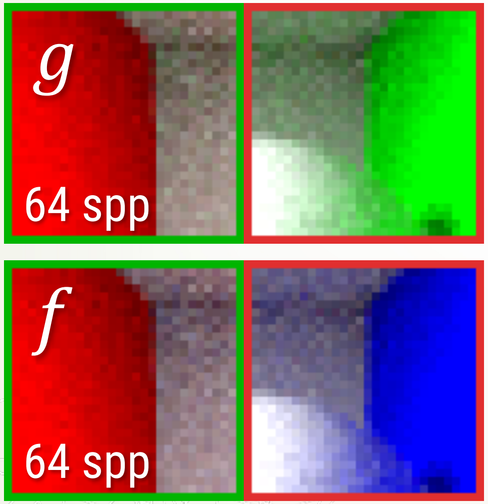

$$\int_{\Omega} f(\boxed{x}) - g(\boxed{x})\, dx + G$$

SCENE EDITING

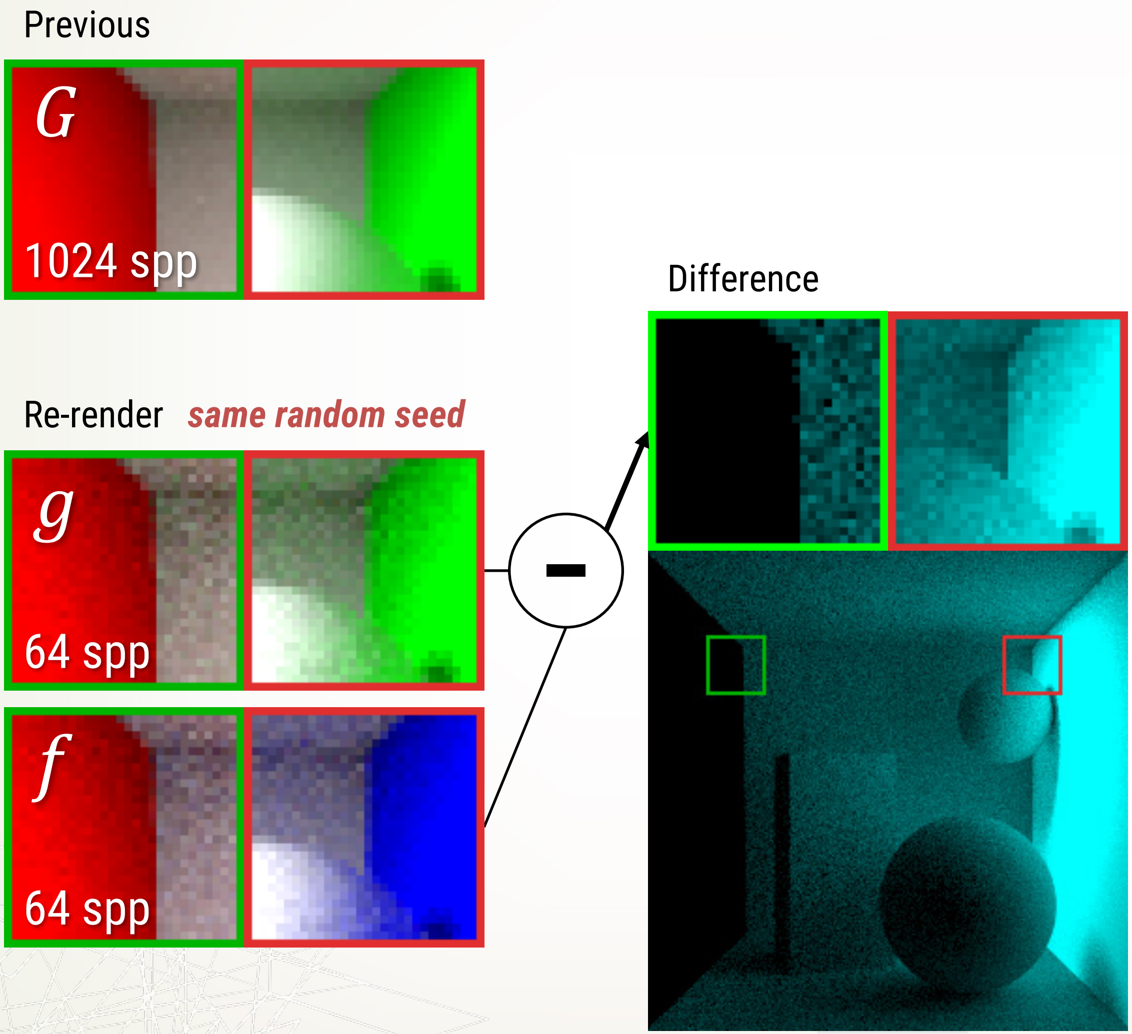

SCENE EDITING

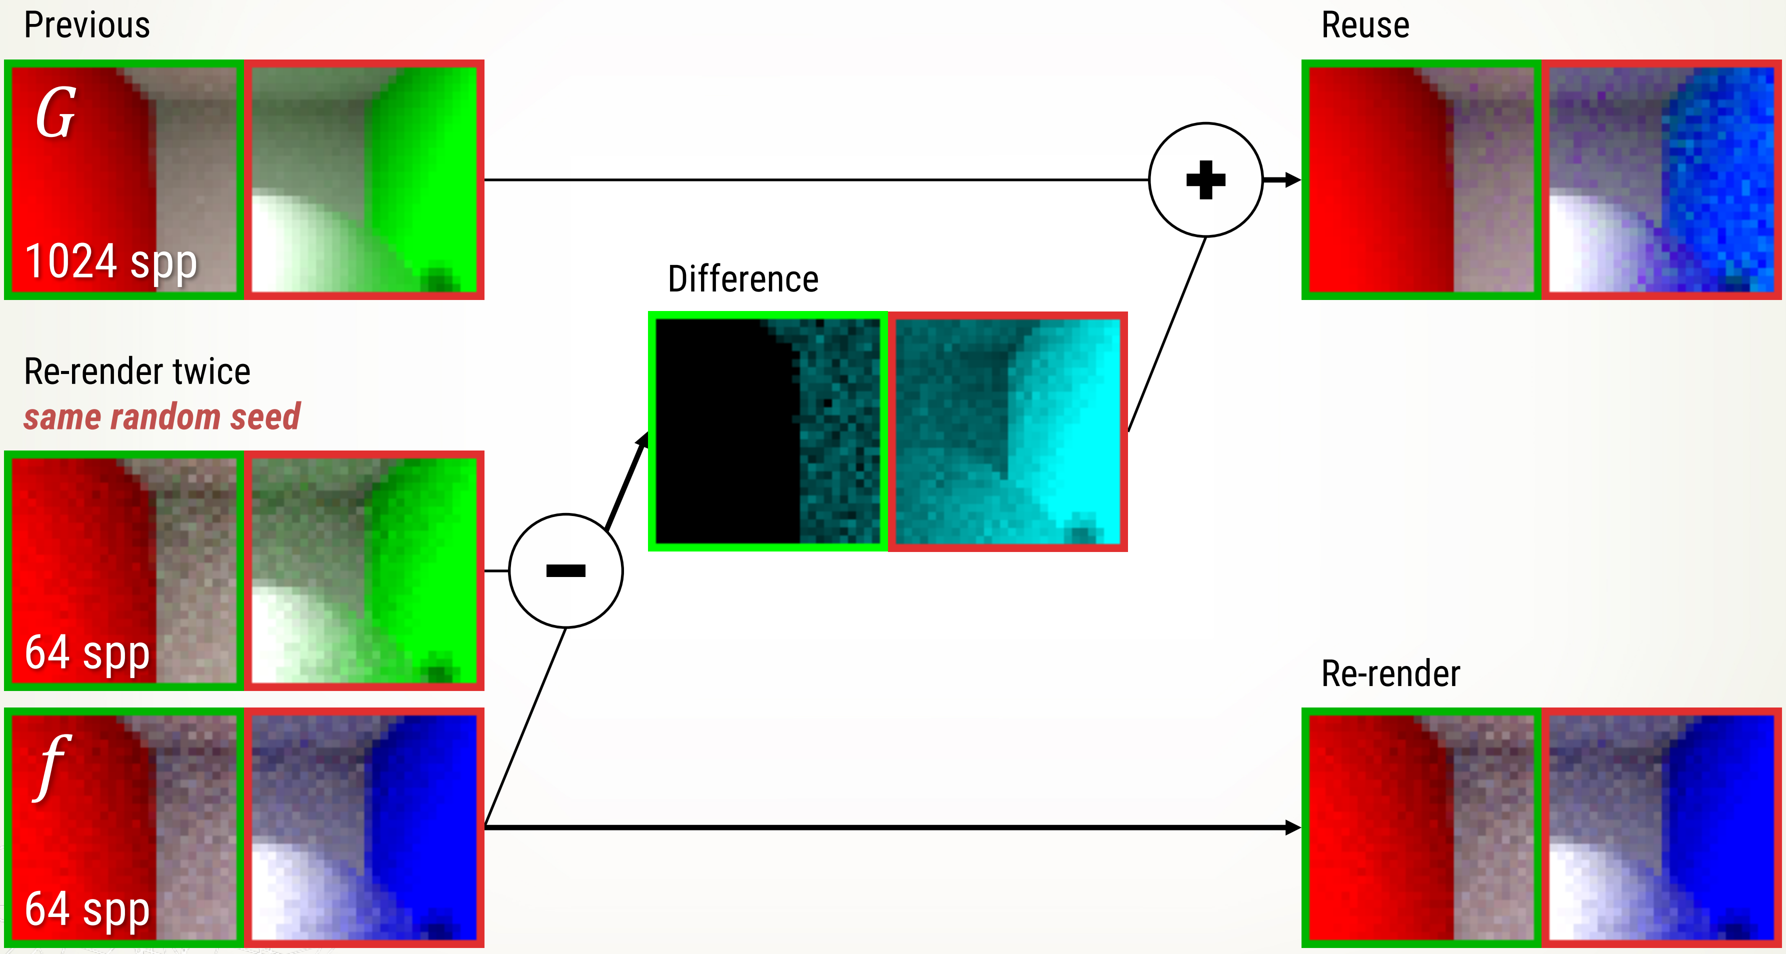

OPTIMAL COMBINATION OF ESTIMATORS

Reuse

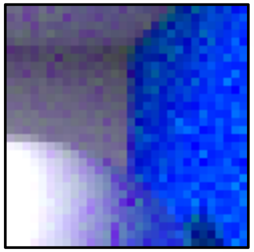

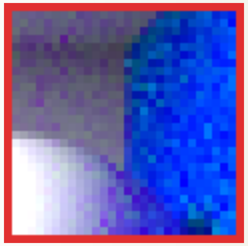

Re-render

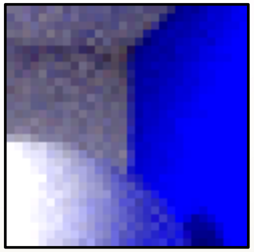

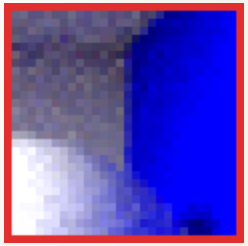

OPTIMAL COMBINATION OF ESTIMATORS

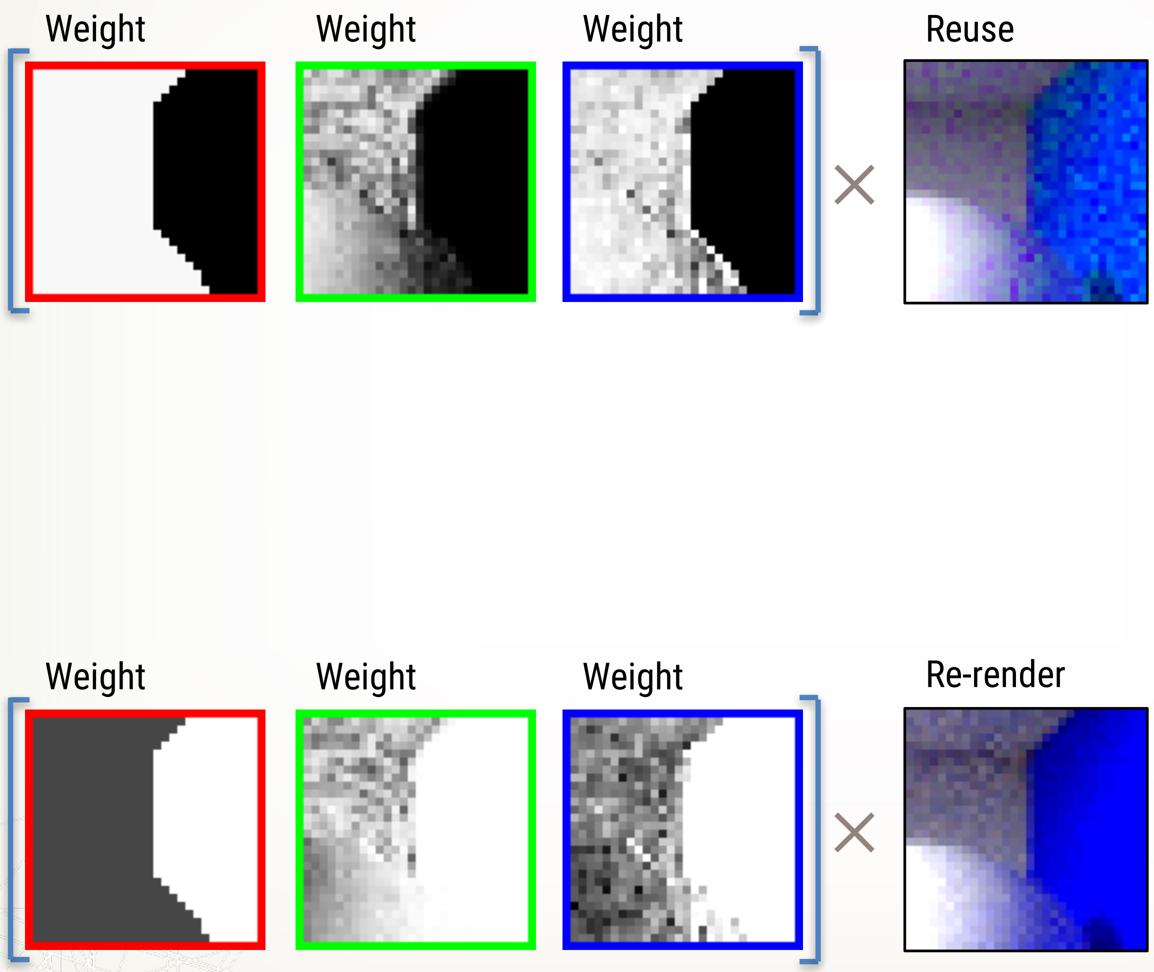

OPTIMAL COMBINATION OF ESTIMATORS

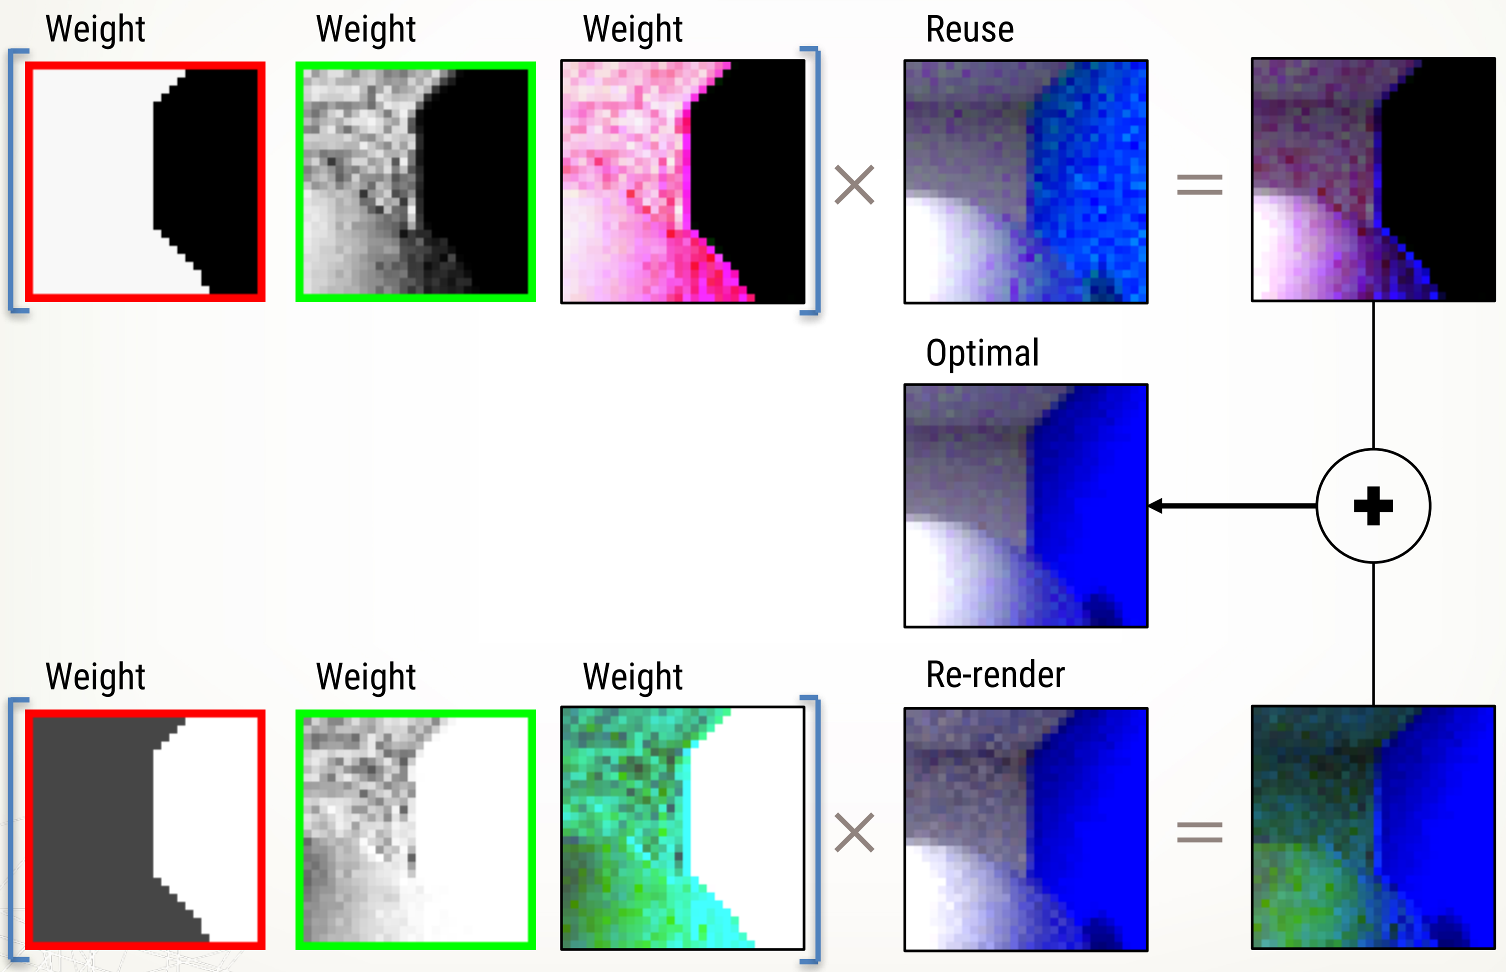

OPTIMAL COMBINATION OF ESTIMATORS

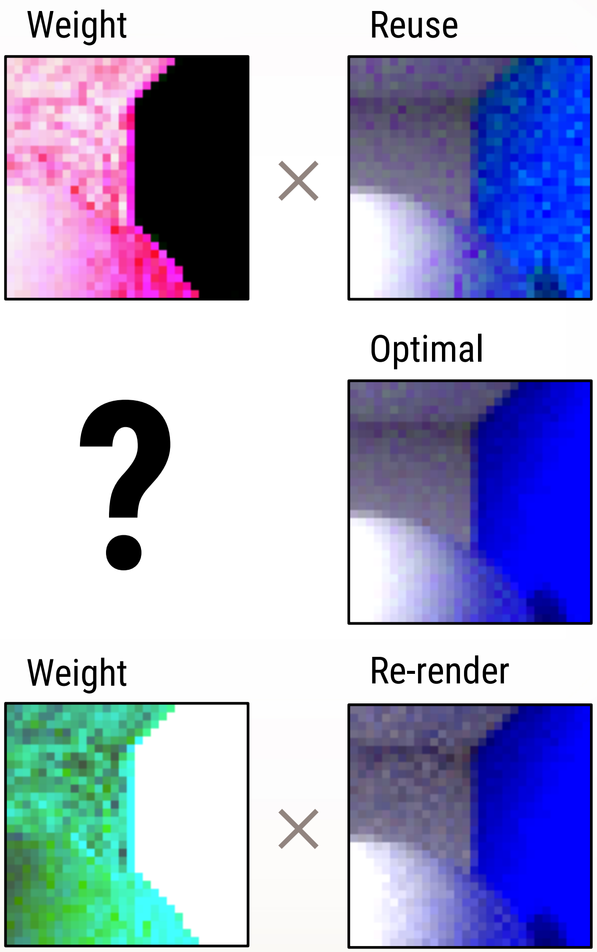

PREVIOUS WORK: SELECTION HEURISTIC [GG15]

Limitations

- Assumes hi-quality previous image
- User-defined threshold
- Selection discards data

Reuse

Use 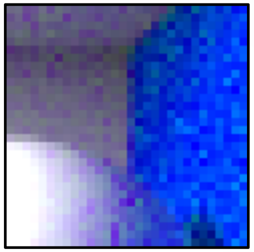 if

Difference

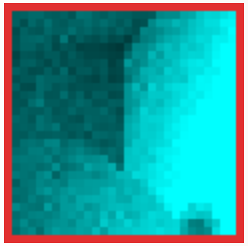 $< \tau$

Re-render

Use 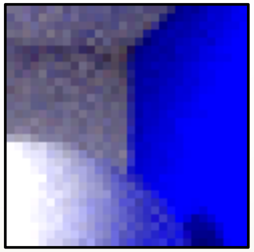 otherwise

OPTIMAL COMBINATION OF ESTIMATORS

Independent

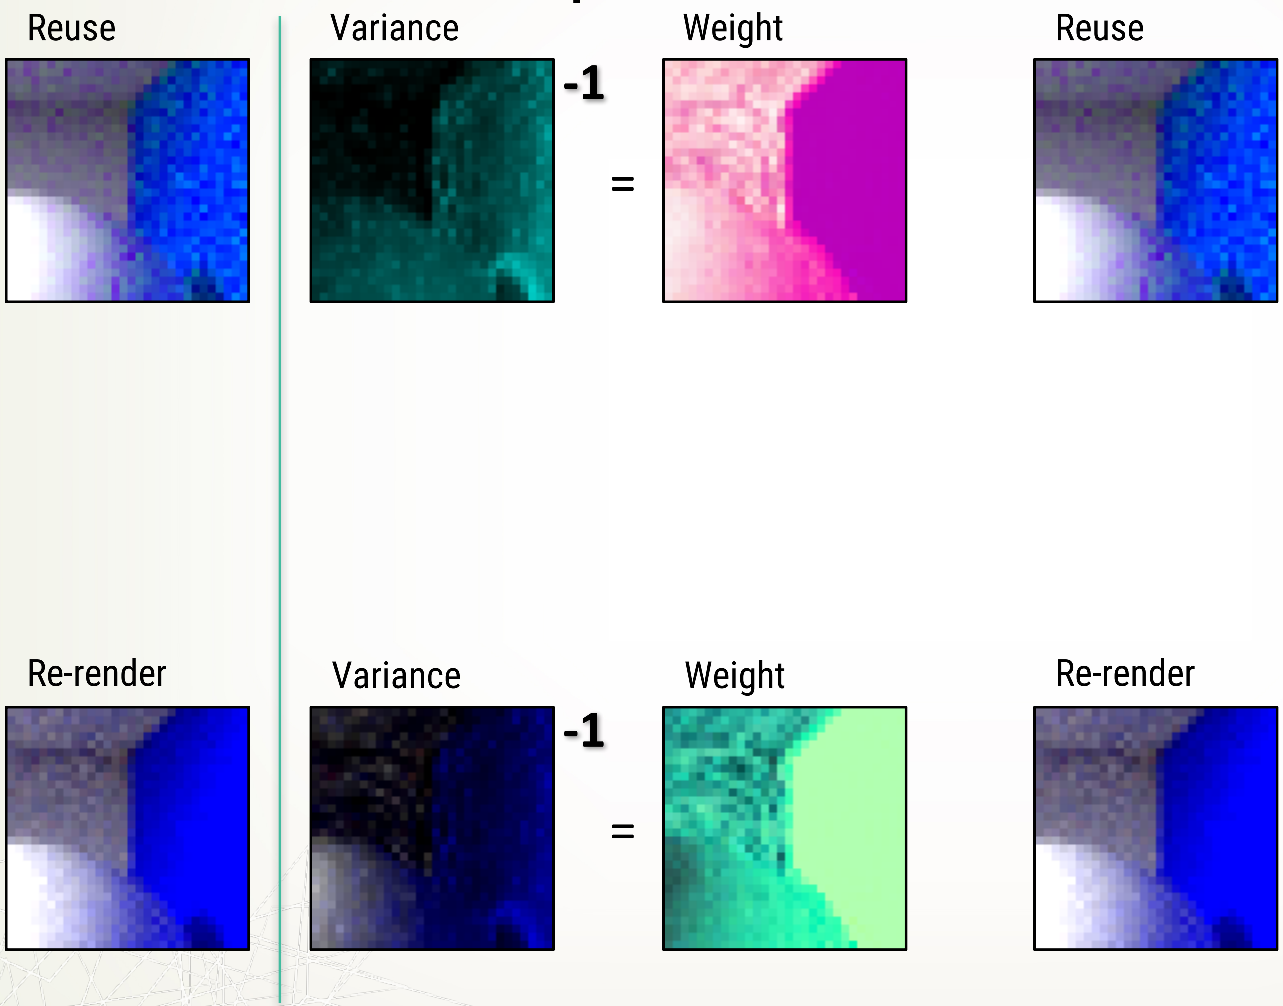

OPTIMAL COMBINATION OF ESTIMATORS

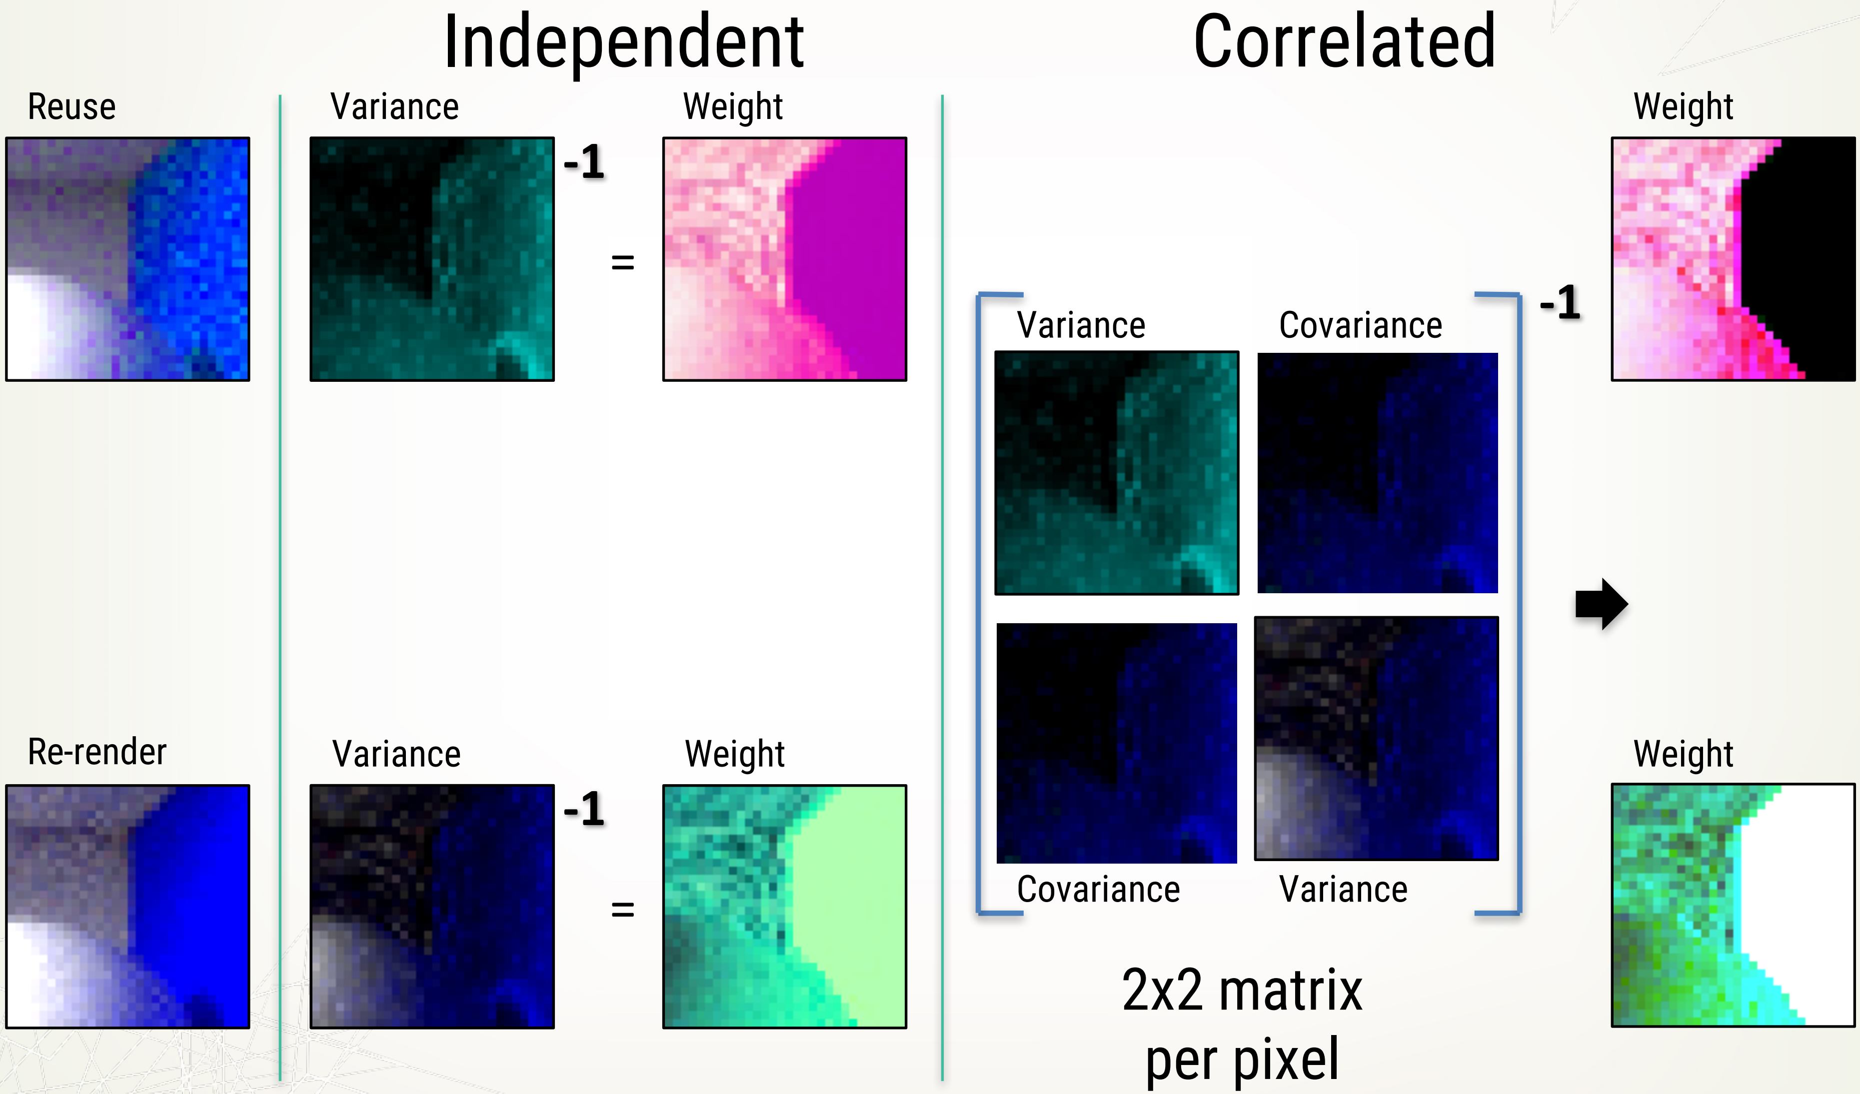

COMPARISON TO PREVIOUS WORK [GG15] – relative MSE

Previous

1024 spp – ***HQ***

Re-render

64 spp – 0.363

COMPARISON TO PREVIOUS WORK [GG15] – relative MSE

Previous

1024 spp – ***HQ***

Re-render

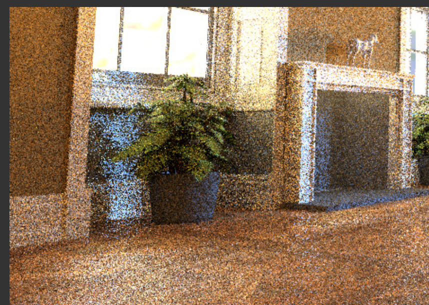

64 spp – 0.363

[GG15]

0.230

Ours

0.084

COMPARISON TO PREVIOUS WORK [GG15] – relative MSE

Previous | Re-render | [GG15] | Ours


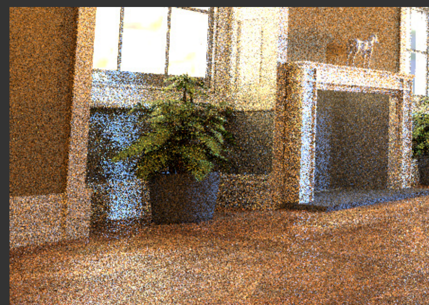

1024 spp – ***HQ*** | 64 spp – 0.363 | 0.230 | 0.084


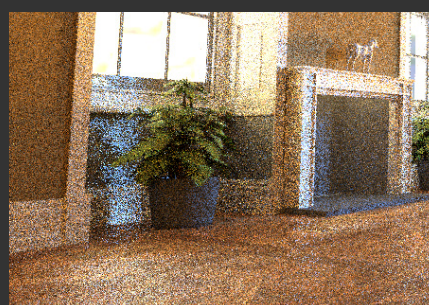
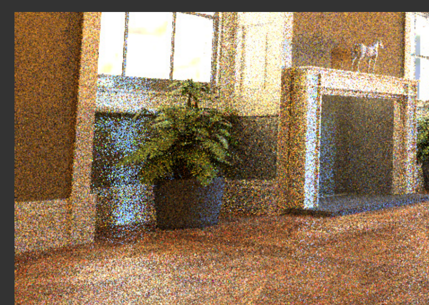
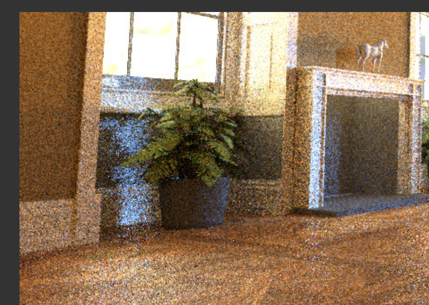

64 spp – ***LQ*** | 64 spp – 0.363 | 0.384 | 0.153

GRADIENT-DOMAIN PATH TRACING

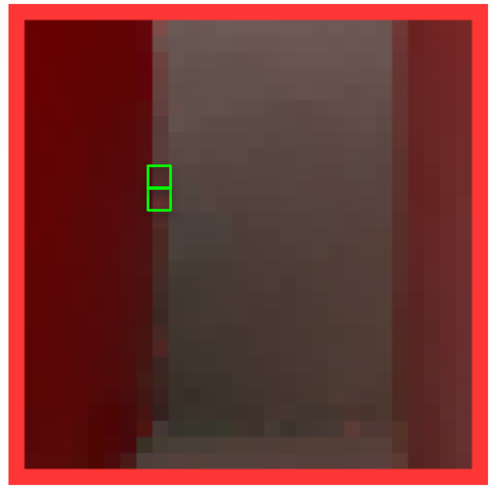

Throughput

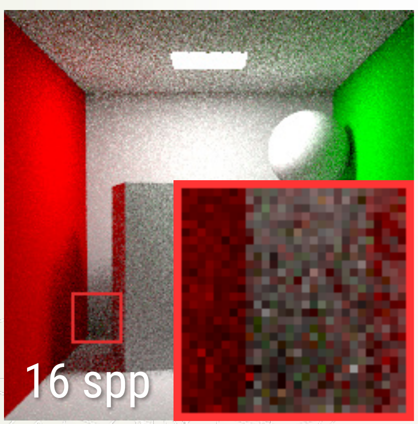

Horizontal gradient

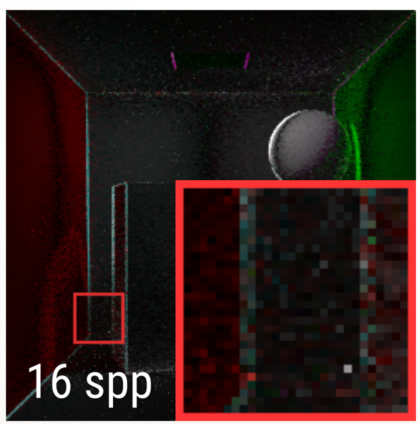

Vertical gradient

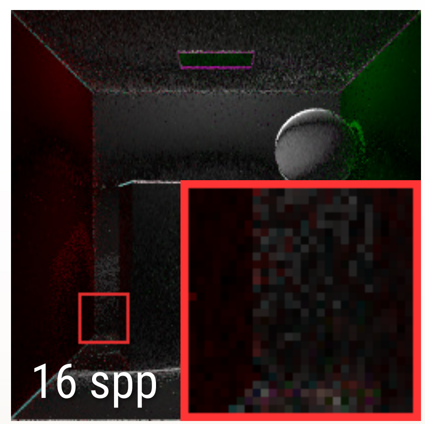

PREVIOUS WORK: POISSON RECONSTRUCTION

Throughput

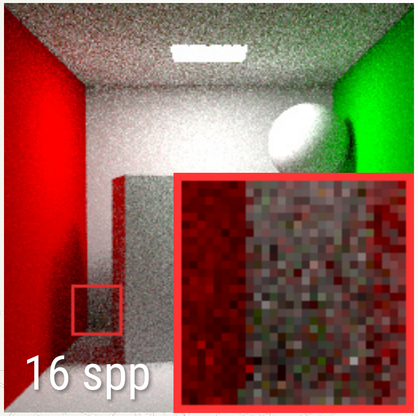

Horizontal gradient

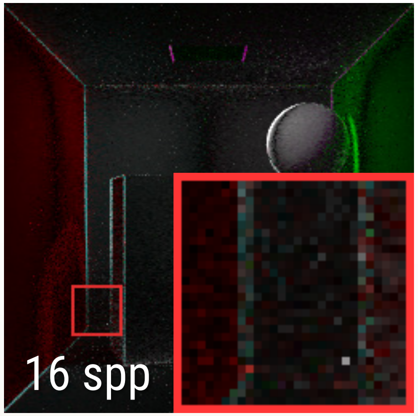

Vertical gradient

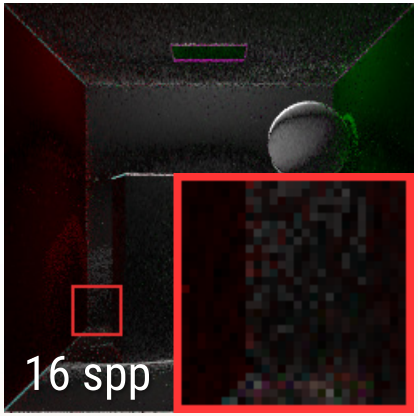

Poisson solver (L2)

Disney Research

PREVIOUS WORK: POISSON RECONSTRUCTION

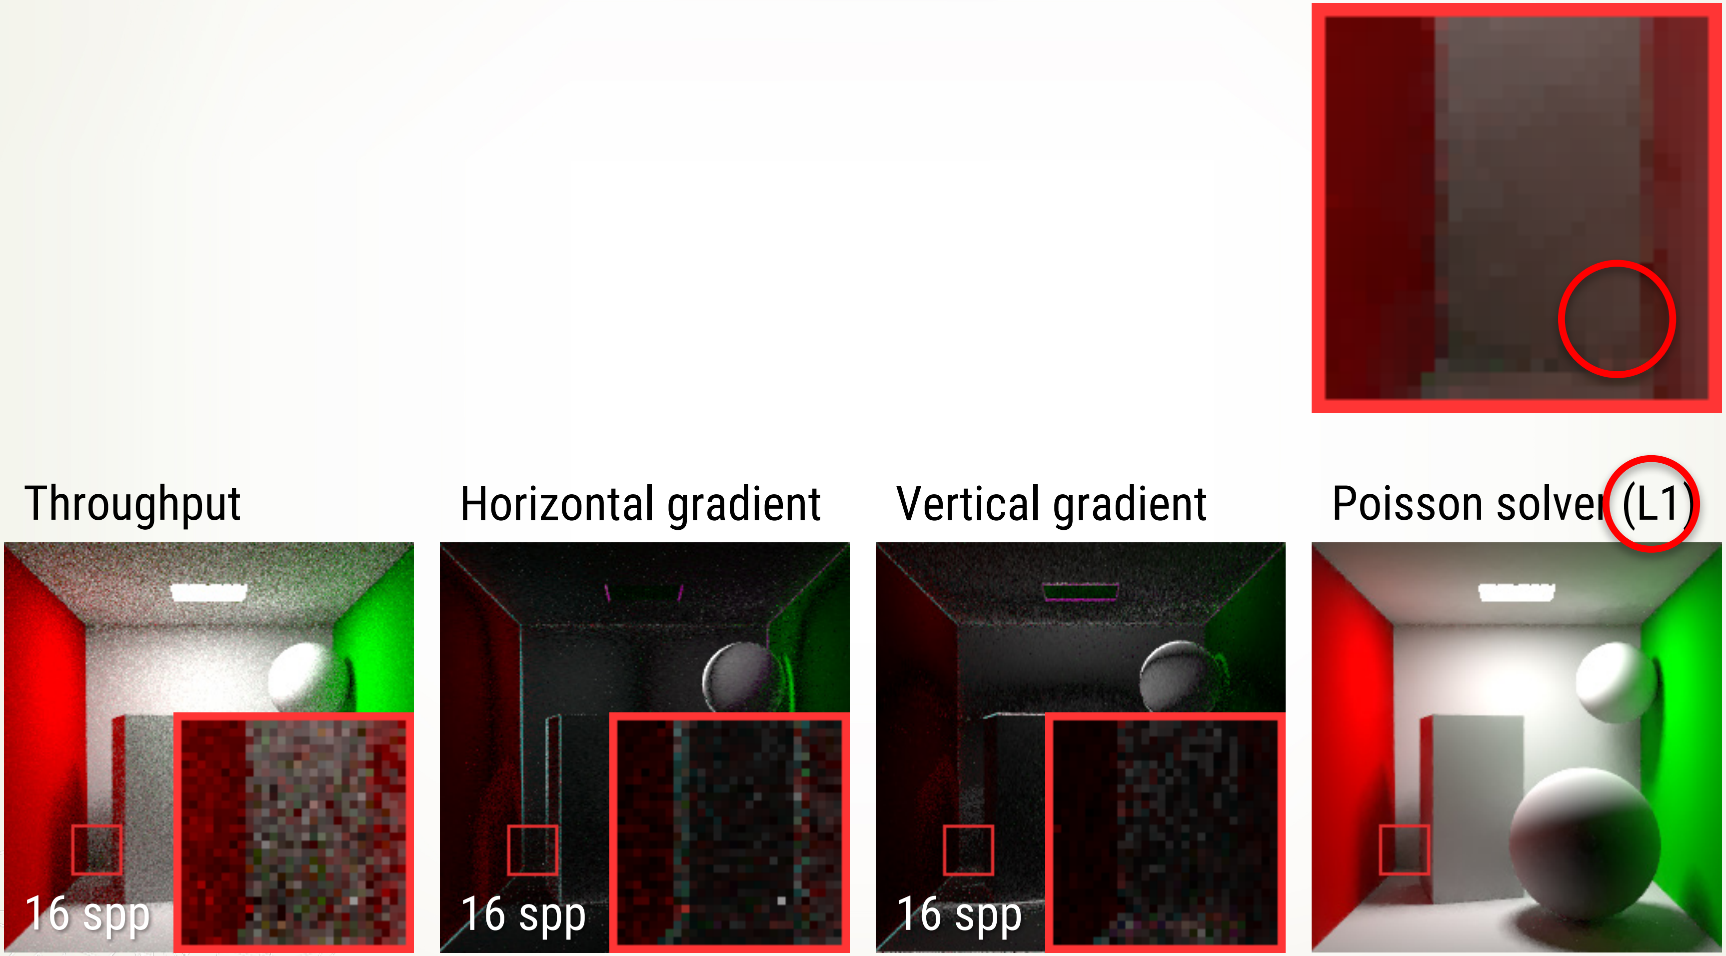

EXTENSION TO MULTIPLE ESTIMATORS: GDR

Throughput

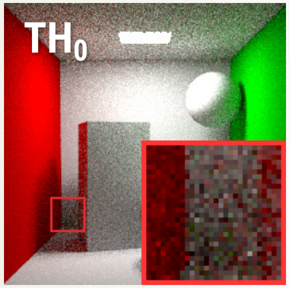

Hori. gradient

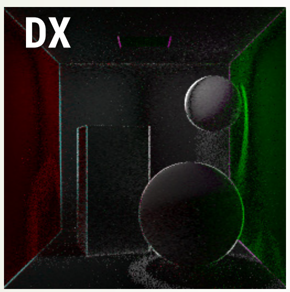

Vert. gradient

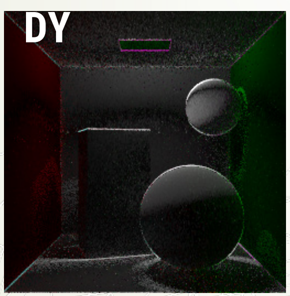

Control Variate

None

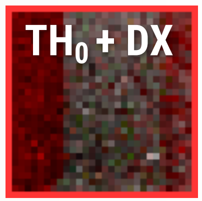

Left neighbor

Right neighbor

Bottom neighbor

Top neighbor

EXTENSION TO MULTIPLE ESTIMATORS: GDR

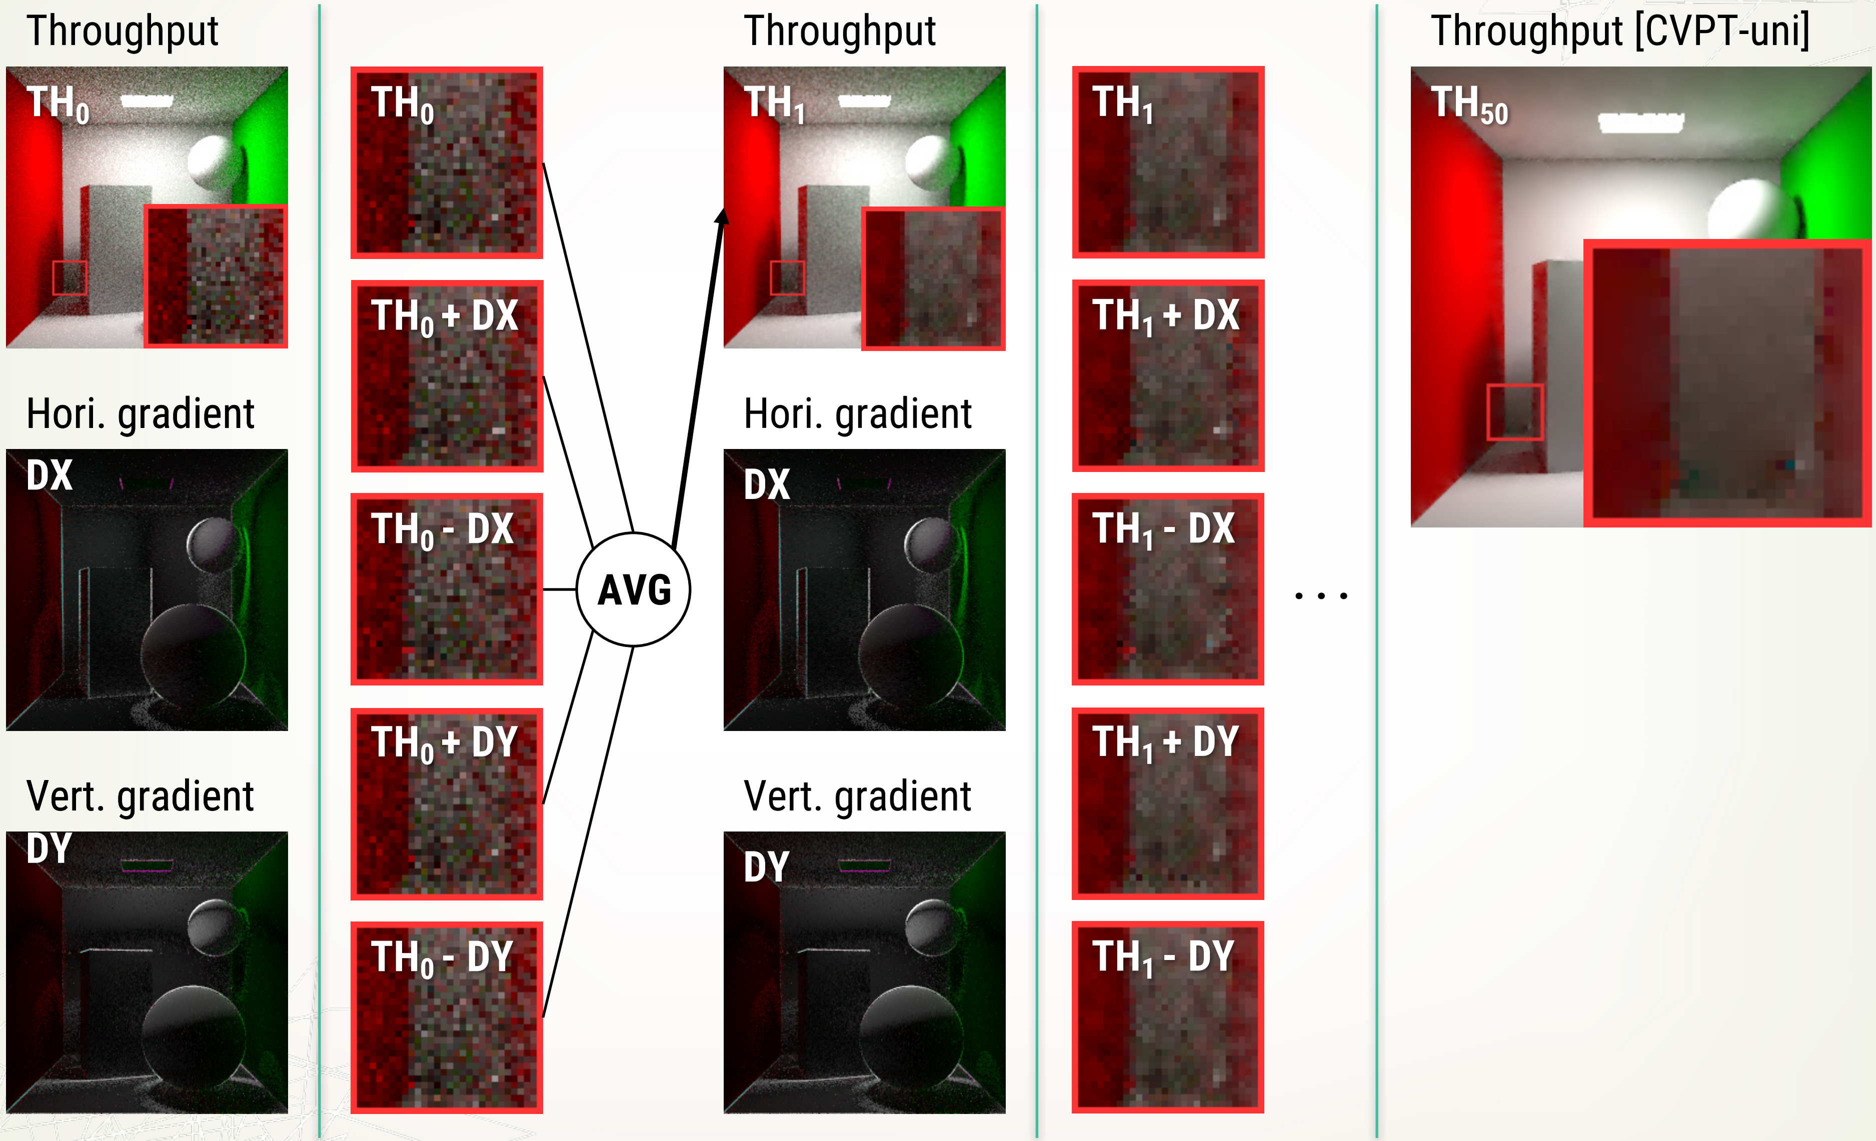

EXTENSION TO MULTIPLE ESTIMATORS: GDR

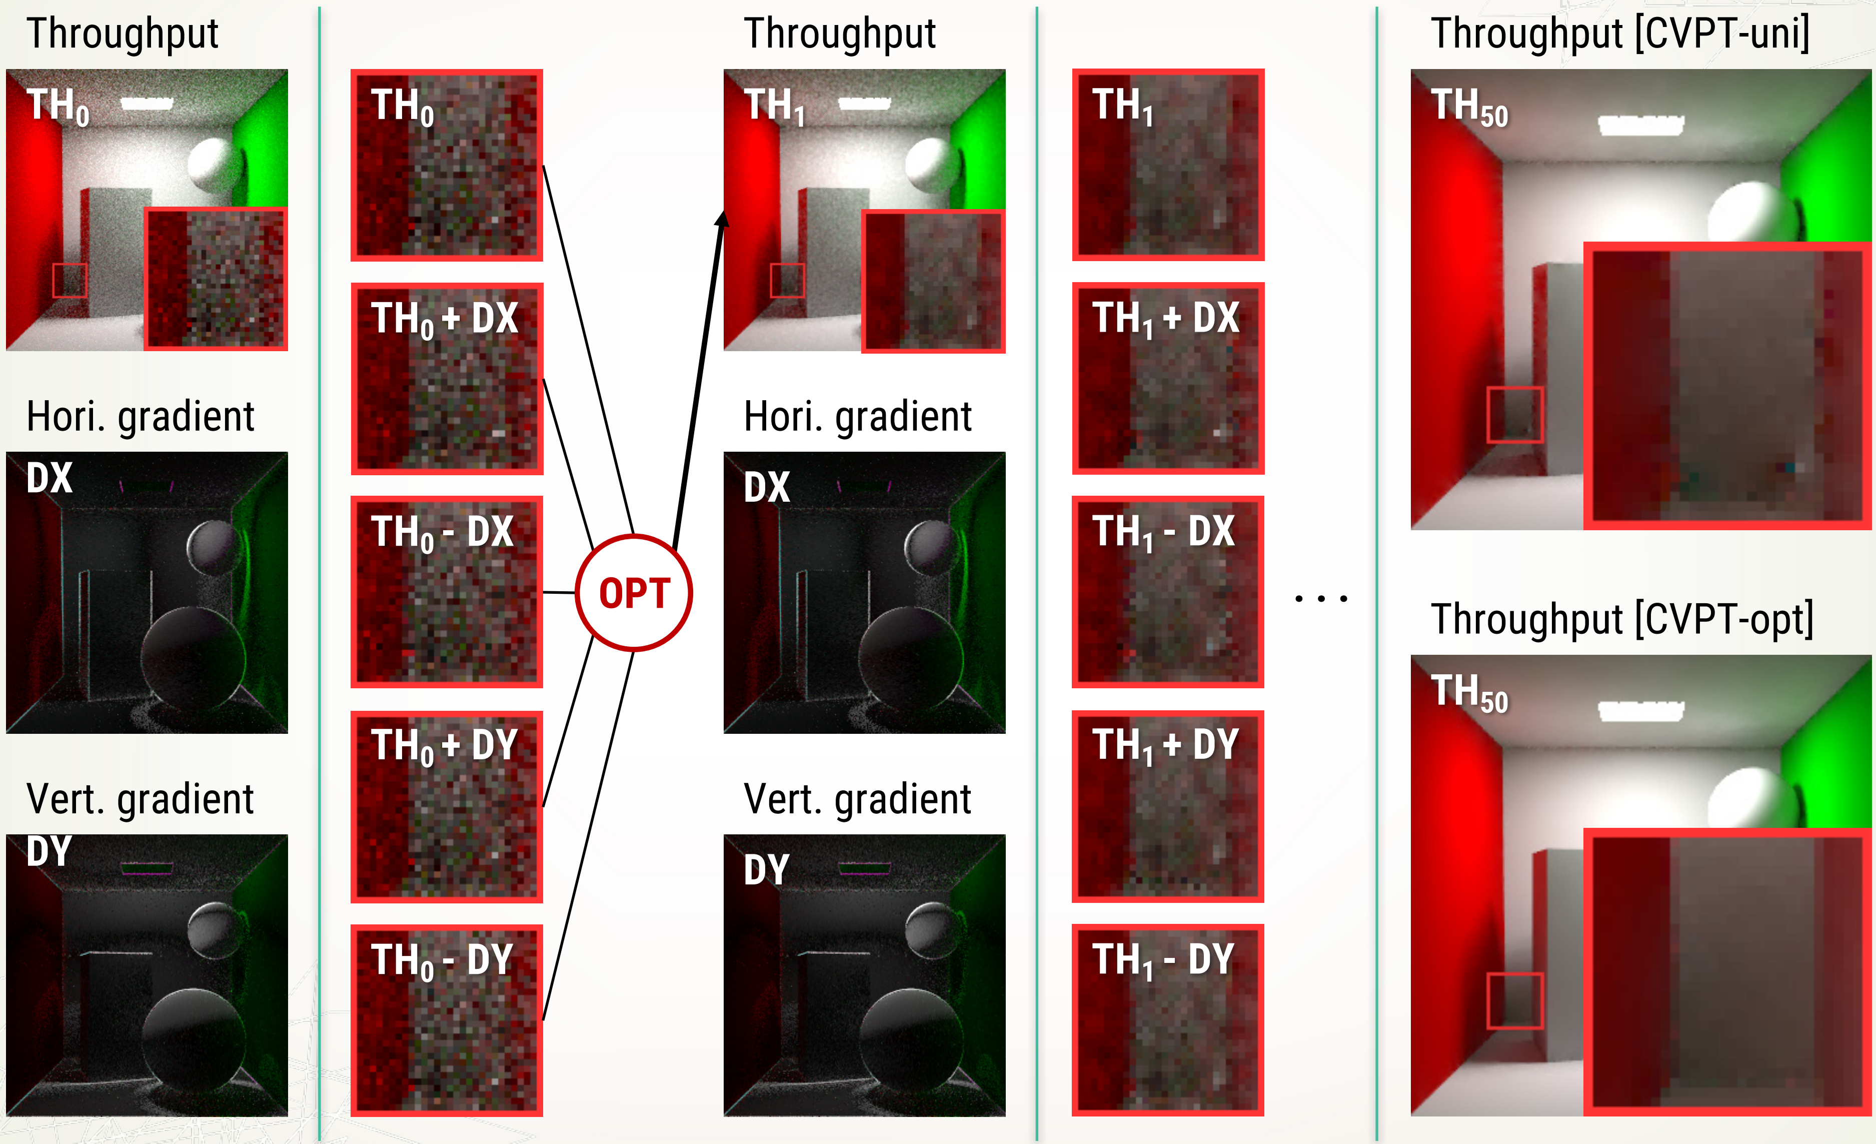

RELATION TO POISSON SOLVER

Iterated reconstruction (Ours)

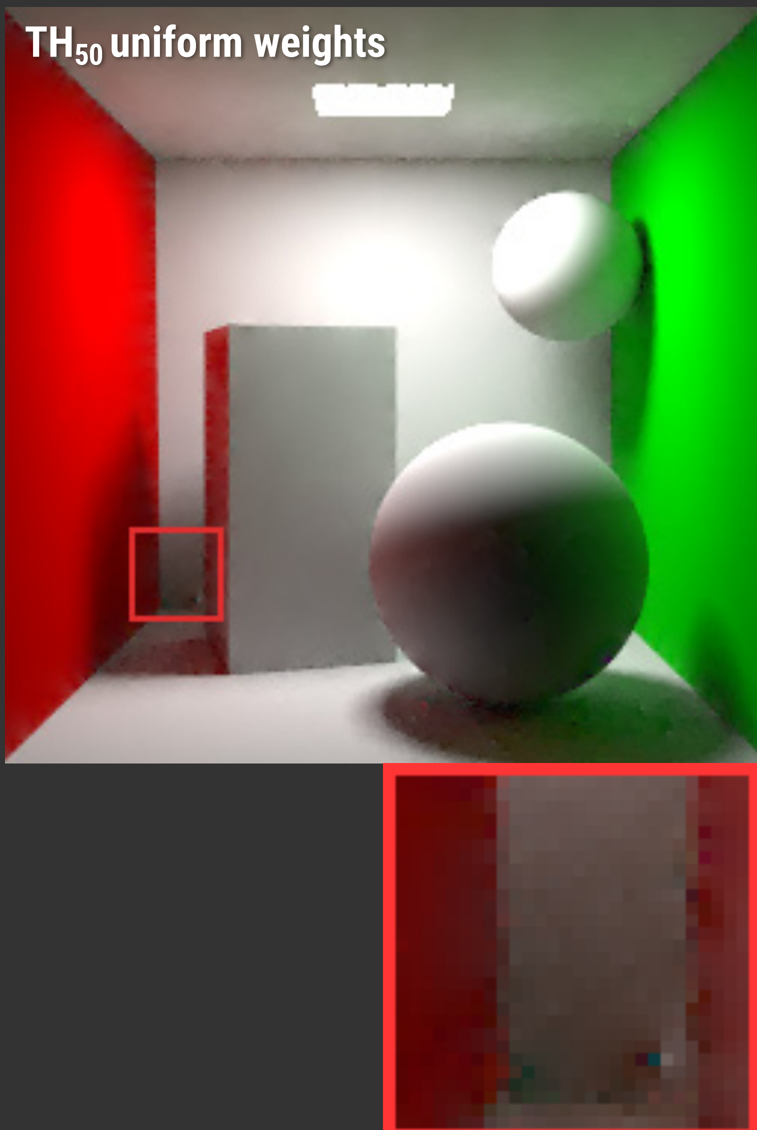

≈

Poisson Solver [GDPT]

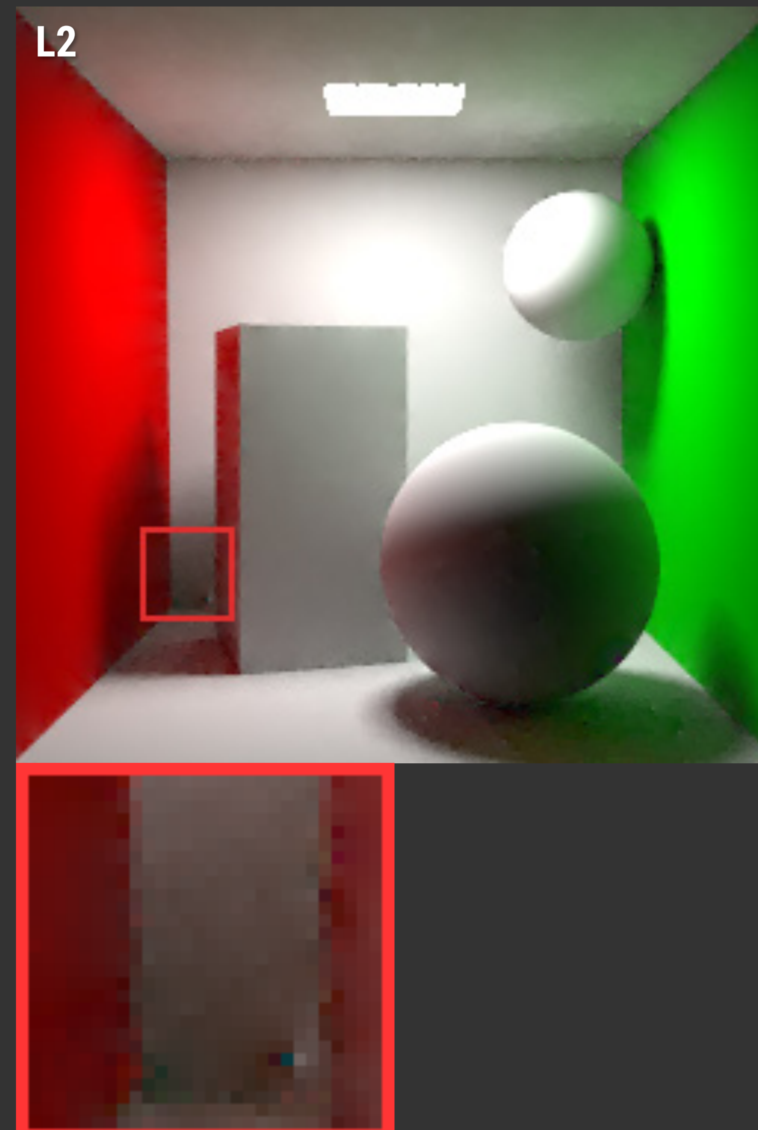

COMPARISON TO PREVIOUS WORK [GDPT] – relative MSE

Sponza scene – 16 spp

Input	GDPT-L2	Ours CVPT-uni	GDPT-L1	Ours CVPT-wgt	Reference
0.0282	0.0034	0.0034	0.0032	0.0018	

COMPARISON TO PREVIOUS WORK [GDPT] – relative MSE

Bookshelf scene – standard path tracing 0.0590

COMPARISON TO PREVIOUS WORK [GDPT] – relative MSE

Bookshelf scene – L1 Poisson reconstruction 0.0094

COMPARISON TO PREVIOUS WORK [GDPT] – relative MSE

Bookshelf scene – Ours with optimized weights 0.0044

MORE IN THE PAPER: UNBIASED RECONSTRUCTIONS

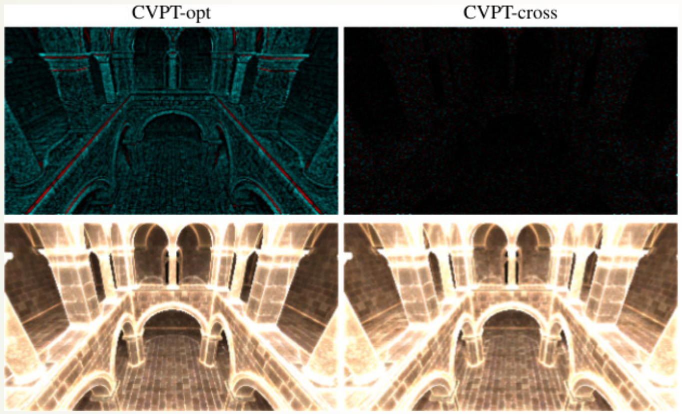

Re-rendering Test Suite

The following scenes were rendered with the PBRTv3 renderer. For each scene, we show multiple images: the control image (Control), the standard path tracer output (PT), the result of using integration with a fully weighted control variate (CVPT), the result of using our optimized weights (CVPT-opt), and the result using our optimized weights based on independently estimated statistics (CVPT-cross). For our optimized weights, we first filter the statistics using an NL-Means filter, and therefore provide results without this prefilter step (CVPT-opt-raw, CVPT-cross-raw). Lastly, we show results using the heuristic of Günther and Grosch [2015] to pick between the PT and CVPT estimates on a per pixel basis. That heuristic uses a threshold parameter and we show results for two settings of this parameter (including the setting recommended in the original publication, 0.1). For all results, we indicate underneath whether the estimate is biased or not.

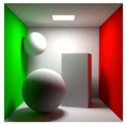
Luminance Test

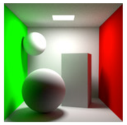
Color Test

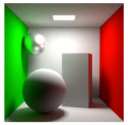
Roughness Test


Horse Room Test

Gradient-domain Reconstruction Test Suite

The following scenes were rendered using the public implementation of the Gradient-domain Path Tracing method of Kettunen et al. [2015], which uses Wenzel Jakob's Mitsuba renderer. We reconstruct the scenes with an array of techniques, some of which operate in the gradient domain (GDPT-L2, GDPT-L1, GDPT-L1L2, GDPT-WL2, CVPT-uni, CVPT-opt, CVPT-cross), and some of which don't (PT, RDFC, NFOR). Techniques with the prefix GDPT all use a screened Poisson solver, whereas techniques with the prefix CVPT use an iterated reconstruction using control variates. Since computing the gradients incurs a ~2.5x rendering time overhead, we increased the sampling rate by a factor of 2.5 for the PT, RDFC, and NFOR results. Please see our paper for a description of the CVPT-uni, CVPT-opt, CVPT-cross, and GDPT-WL2 reconstructions. The GDPT-L2 and GDPT-L1 reconstructions were proposed by Kettunen et al. [2015], the GDPT-L1L2 reconstruction by Manzi et al. [2016], the RDFC reconstruction by Rousselle et al. [2013], and the NFOR reconstruction by Bitterli et al. [2016]. The GDPT-L1L2 reconstruction is essentially a weighted screened Poisson solver where only the gradient constraints are reweighted, hence the name L1L2; in their original publication, Manzi et al. regularize the Poisson solver with constraints based on auxiliary buffers, but in our comparison we use a standard screened Poisson solver.


Sponza


Veach Door


Bathroom


Kitchen

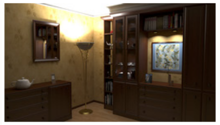
Bookshelf

CONCLUSION

- Integration with control variates…
  - simple yet powerful tool
  - relatively unexplored in rendering
  - should be used with optimal weighting scheme

- Future work
  - animations, stereo rendering, light fields
  - better sampling of the difference buffer

RELATION TO POISSON SOLVER

Iterated reconstruction (Ours)

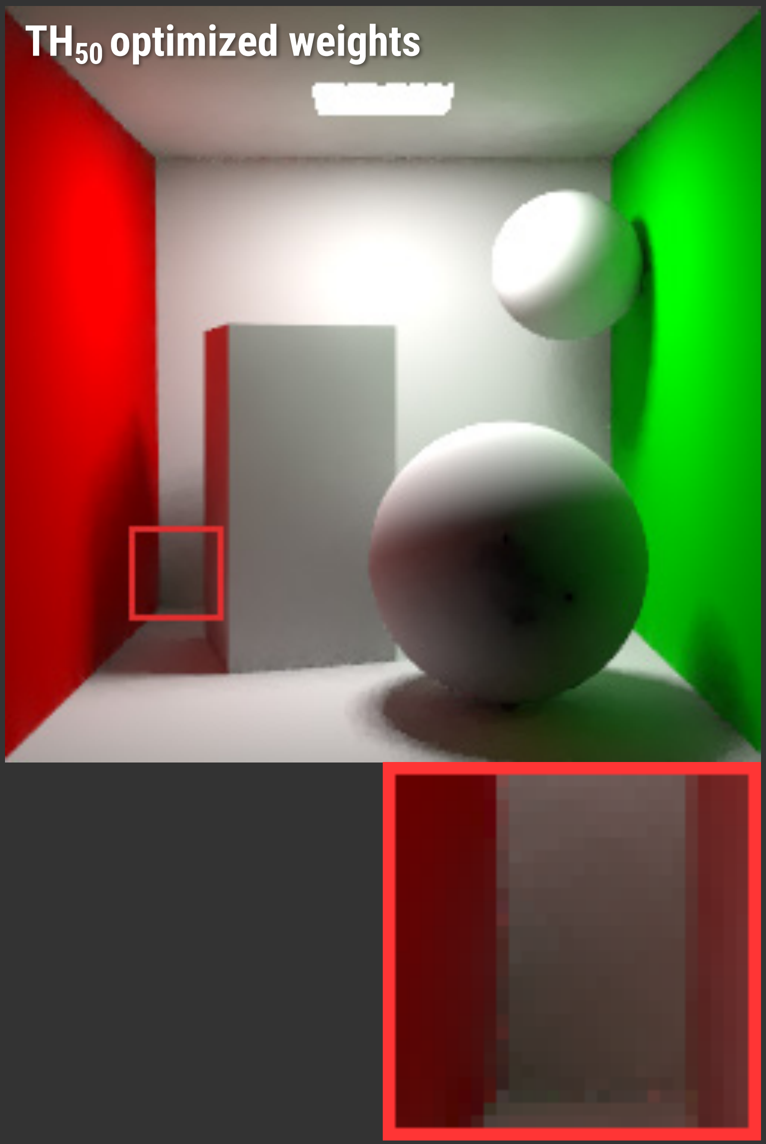

Poisson Solver [GDPT]

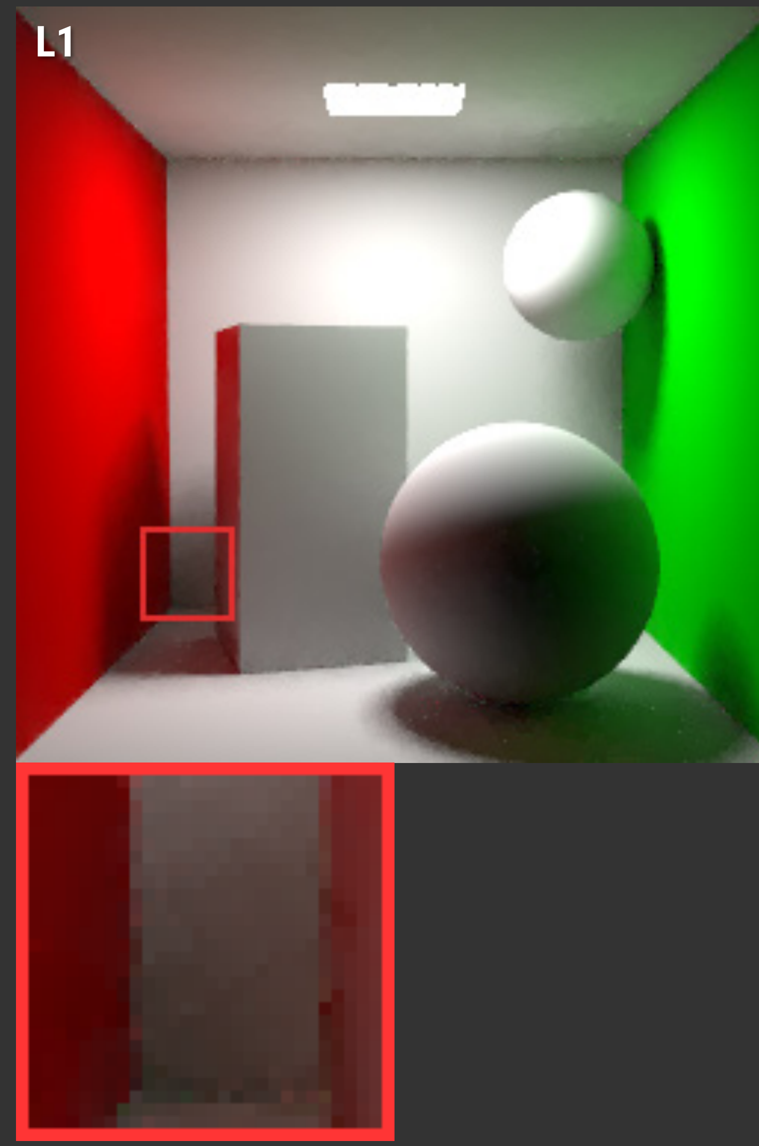

INTEGRATION WITH CONTROL VARIATES

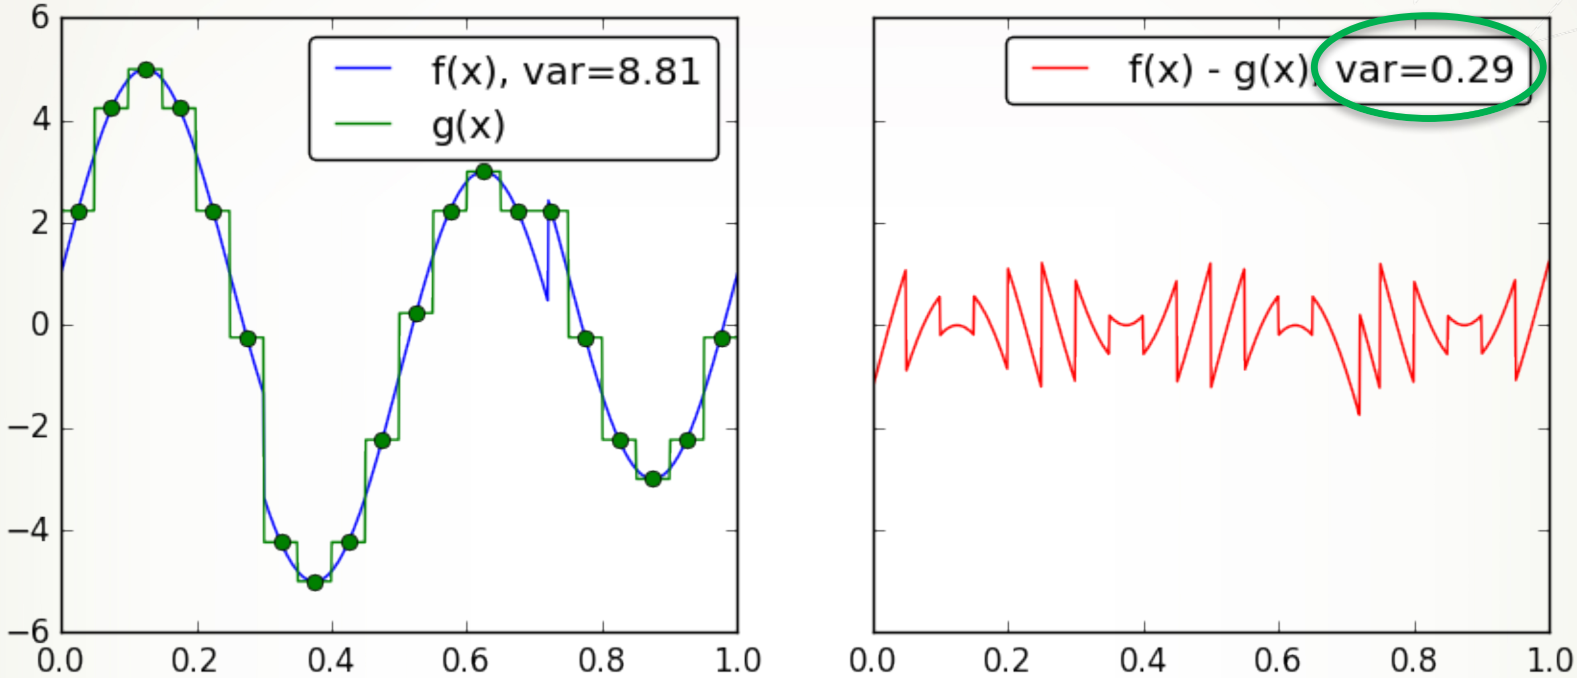

$$\int_{\Omega} f(x) - g(x)\, dx + G$$

INTEGRATION WITH CONTROL VARIATES

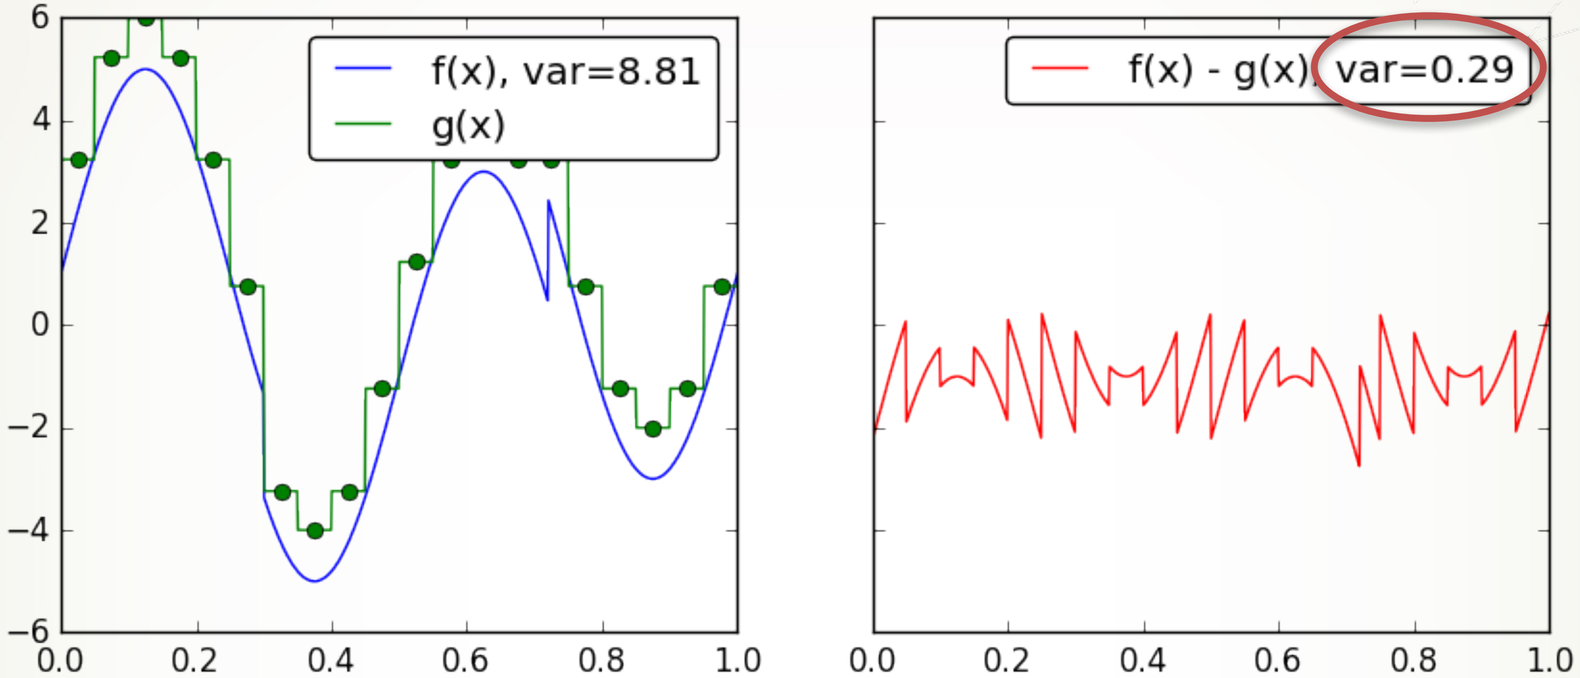

$$\int_{\Omega} f(x) - g(x)\, dx + G$$

COMPARISON TO PREVIOUS WORK [GDPT] – relative MSE

Bookshelf scene – 256 spp

Input	GDPT-L2	Ours CVPT-uni	GDPT-L1	Ours CVPT-wgt	Reference

Input	GDPT-L2	Ours CVPT-uni	GDPT-L1	Ours CVPT-wgt	Reference
0.206	0.0365	0.0362	0.0314	0.0129	

COMPARISON TO PREVIOUS WORK [GDPT] – relative MSE

Bathroom scene – 256 spp

Input	GDPT-L2	Ours CVPT-uni	GDPT-L1	Ours CVPT-wgt	Reference

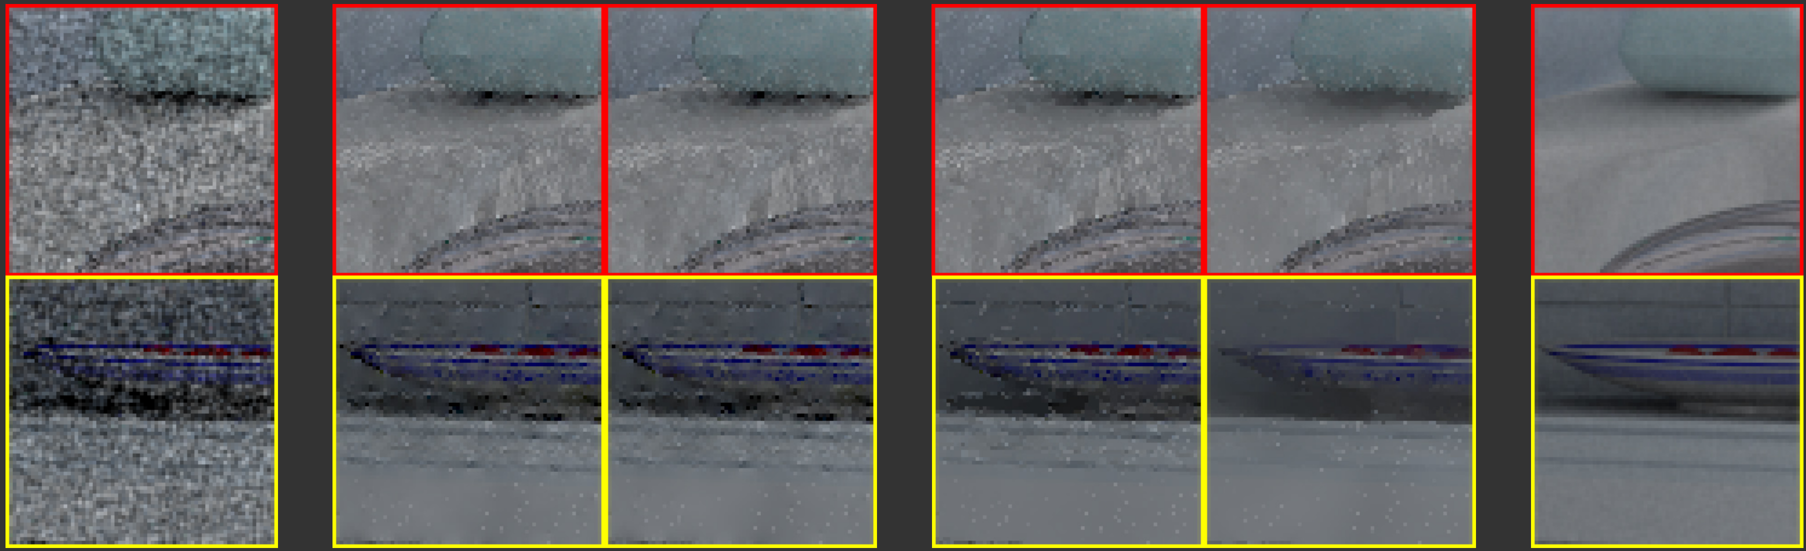

0.0616	0.0098	0.0099	0.0088	0.0070	

- Irradiance caching

A ray tracing solution for diffuse interreflection
Ward et al., SIGGRAPH 1988

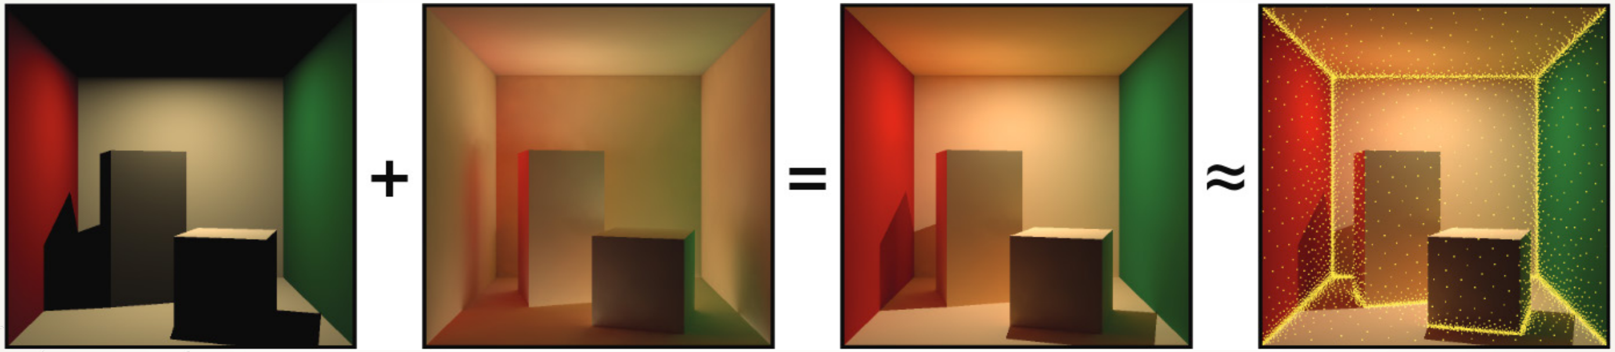

Image by Wojciech Jarosz

LEVERAGING COHERENCE: PREVIOUS WORK

- Irradiance caching
- Photon mapping

Global illumination using photon maps
Jensen, EGSR 1996

Image by Henrik Wann Jensen

LEVERAGING COHERENCE: PREVIOUS WORK

- Irradiance caching
- Photon mapping
- Image-space denoising

Energy-preserving non-linear filters
Rushmeier and Ward, SIGGRAPH 1994

Scene by Guillermo M. Lean Llaguno
Image by Benedikt Bitterli