FlexISP: A Flexible Camera Image Processing Framework

Dr. Kari Pulli

VP of Computational Imaging

https://light.co

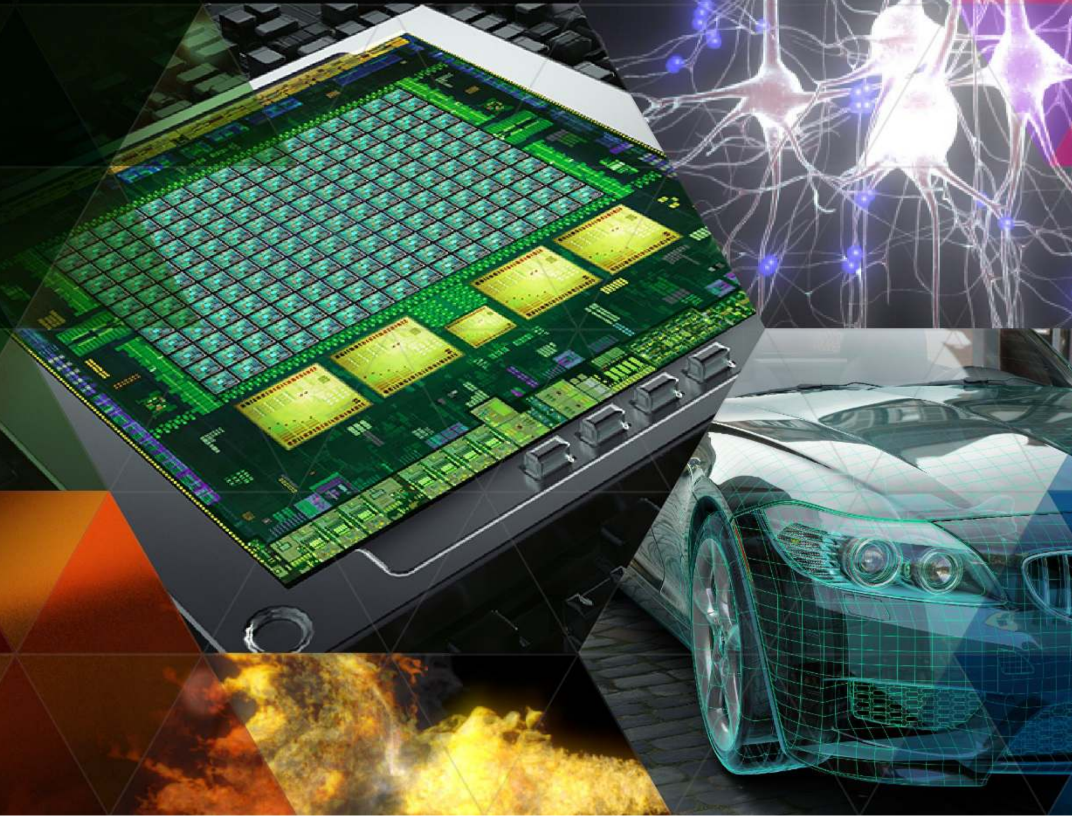

FlexISP: A Flexible Camera Image Processing Framework

Felix Heide[1,2] Markus Steinberger[3] Yun-Ta Tsai[1]
Nasa Rouf[1,2] Dawid Pająk[1] Dikpal Reddy[1]
Orazio Gallo[1] Jing Liu[4] Wolfgang Heidrich[5]
Karen Egiazarian[1,6] Jan Kautz[1] Kari Pulli[1]

[1] NVIDIA
[2] UBC
[3] TU Graz
[4] UCSC
[5] KAUST
[6] TUT

SIGGRAPH ASIA 2014

IMAGE PROCESSING PIPELINE

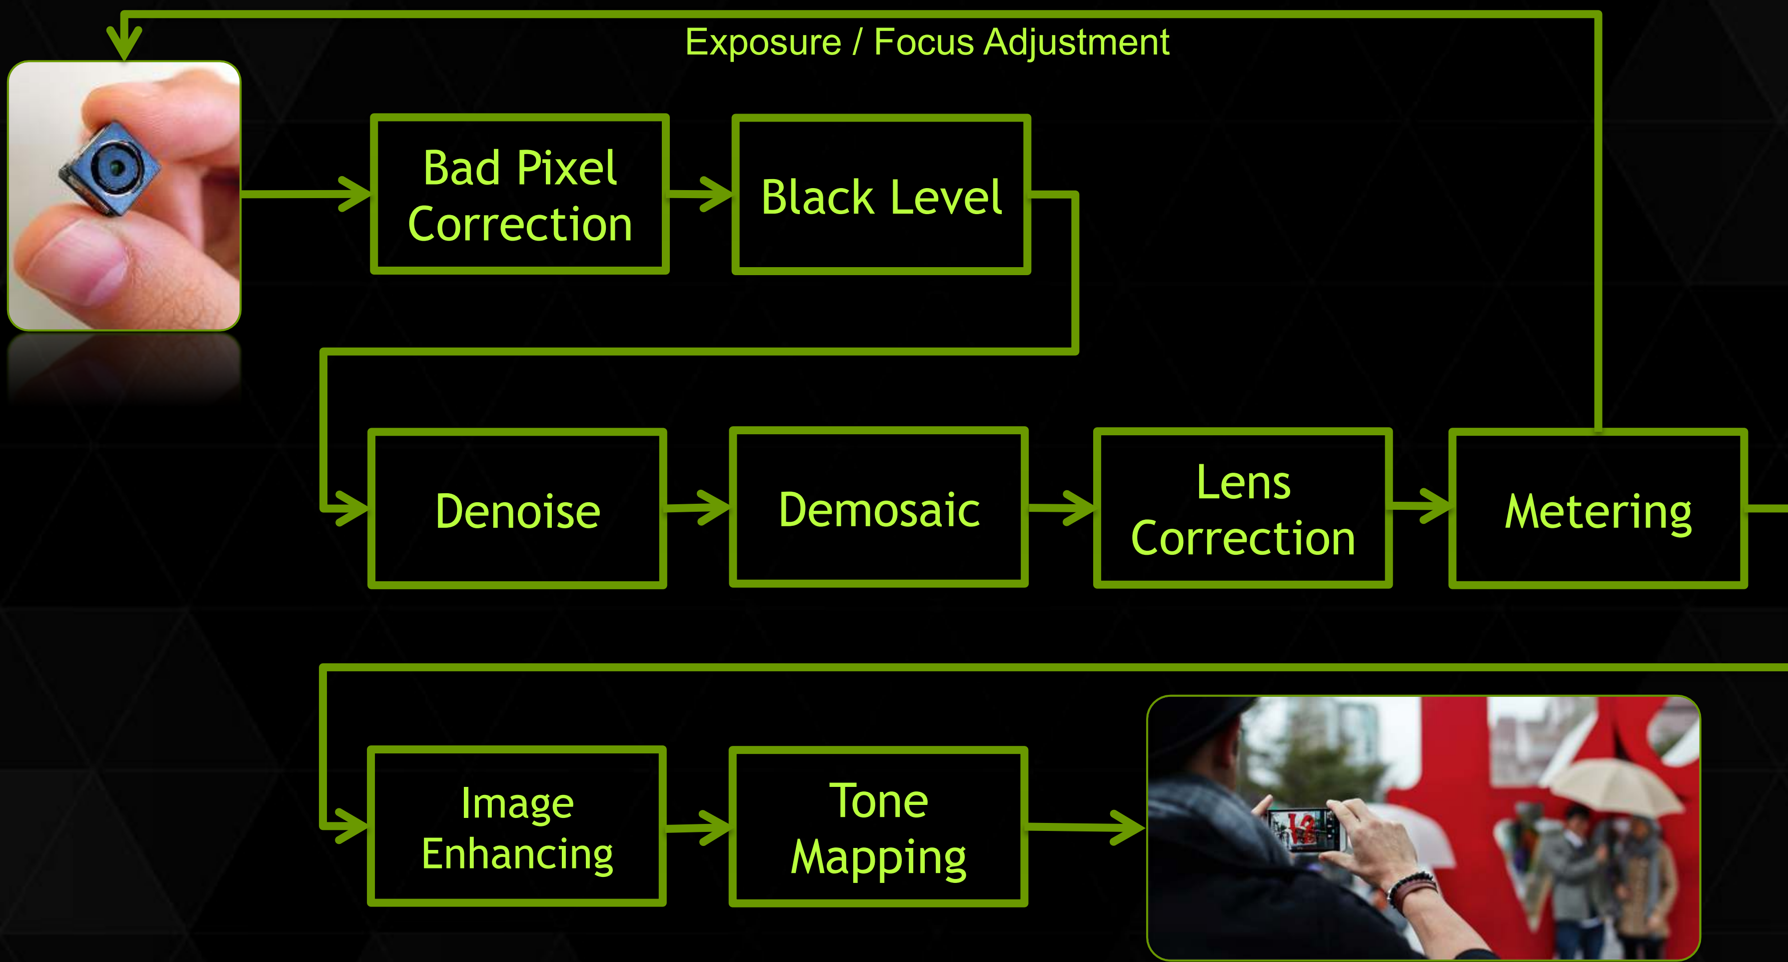

IMAGE PROCESSING PIPELINE

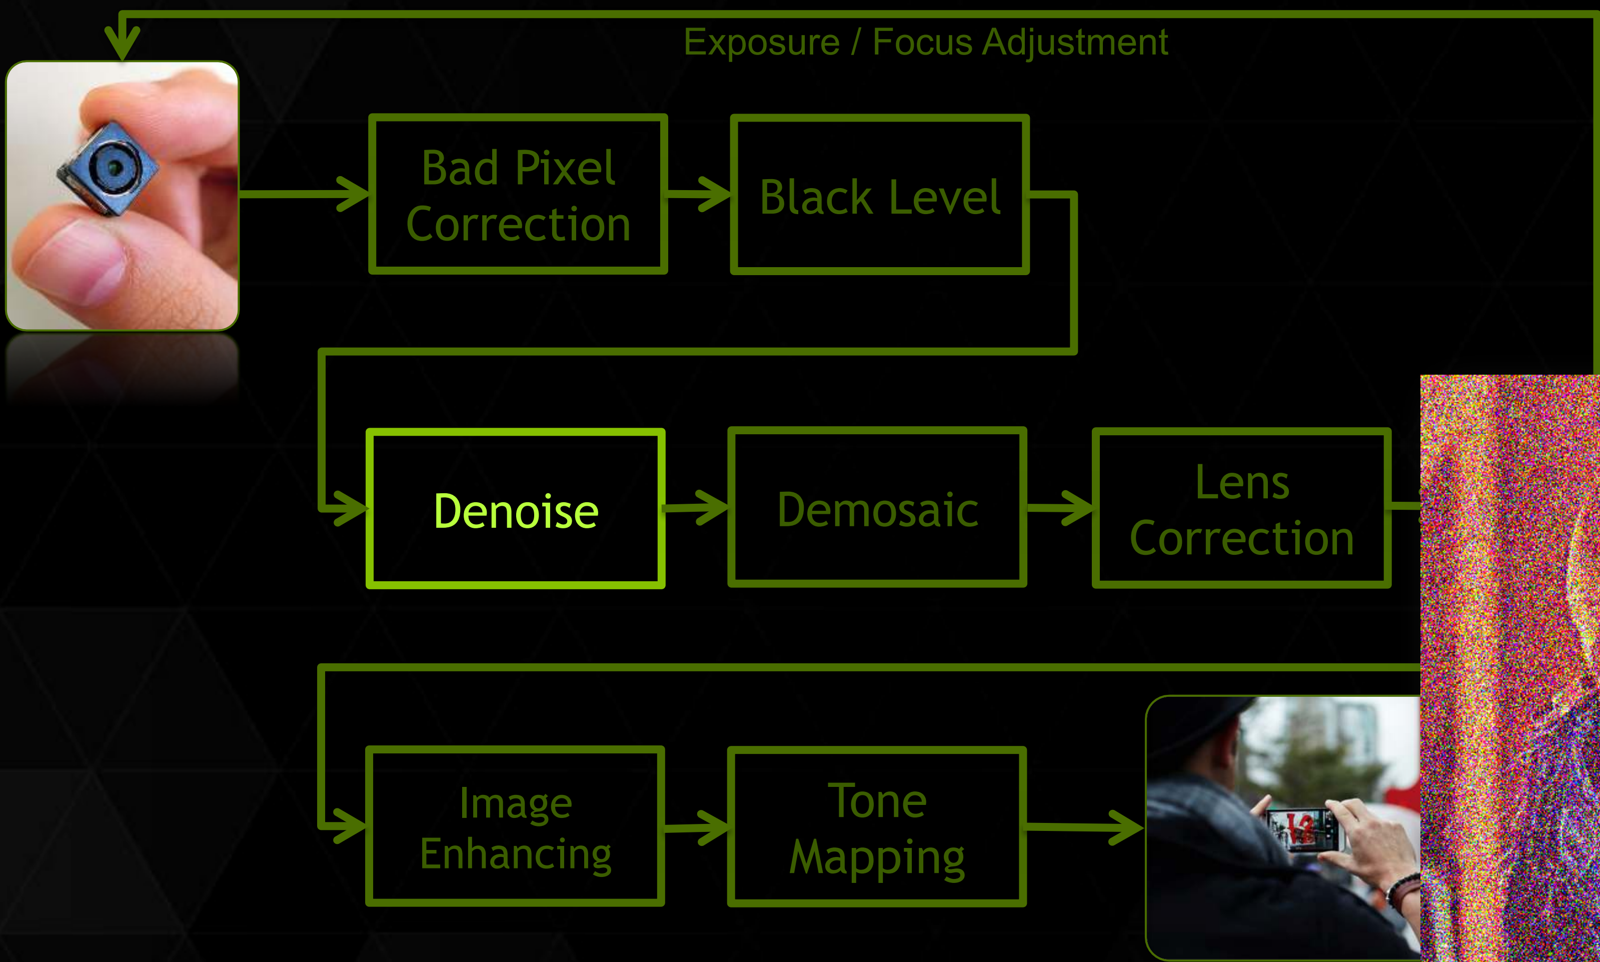

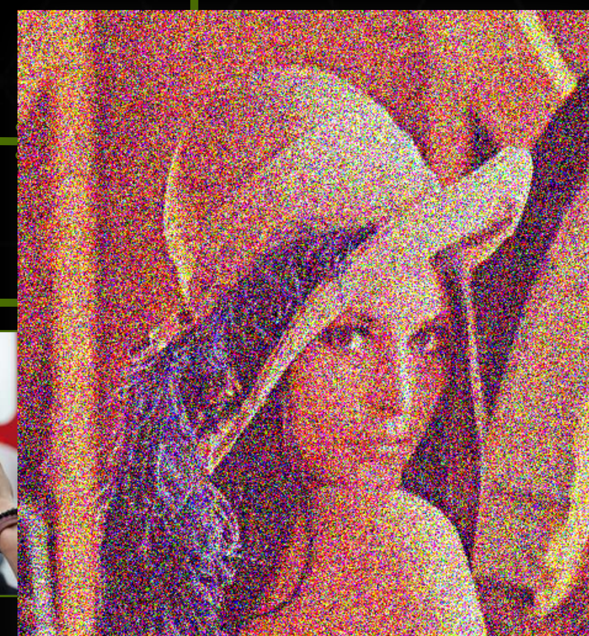

IMAGE PROCESSING PIPELINE

Exposure / Focus Adjustment

Bad Pixel Correction

Black Level

Denoise

Demosaic

Lens Correction

Image Enhancing

Tone Mapping

IMAGE PROCESSING PIPELINE

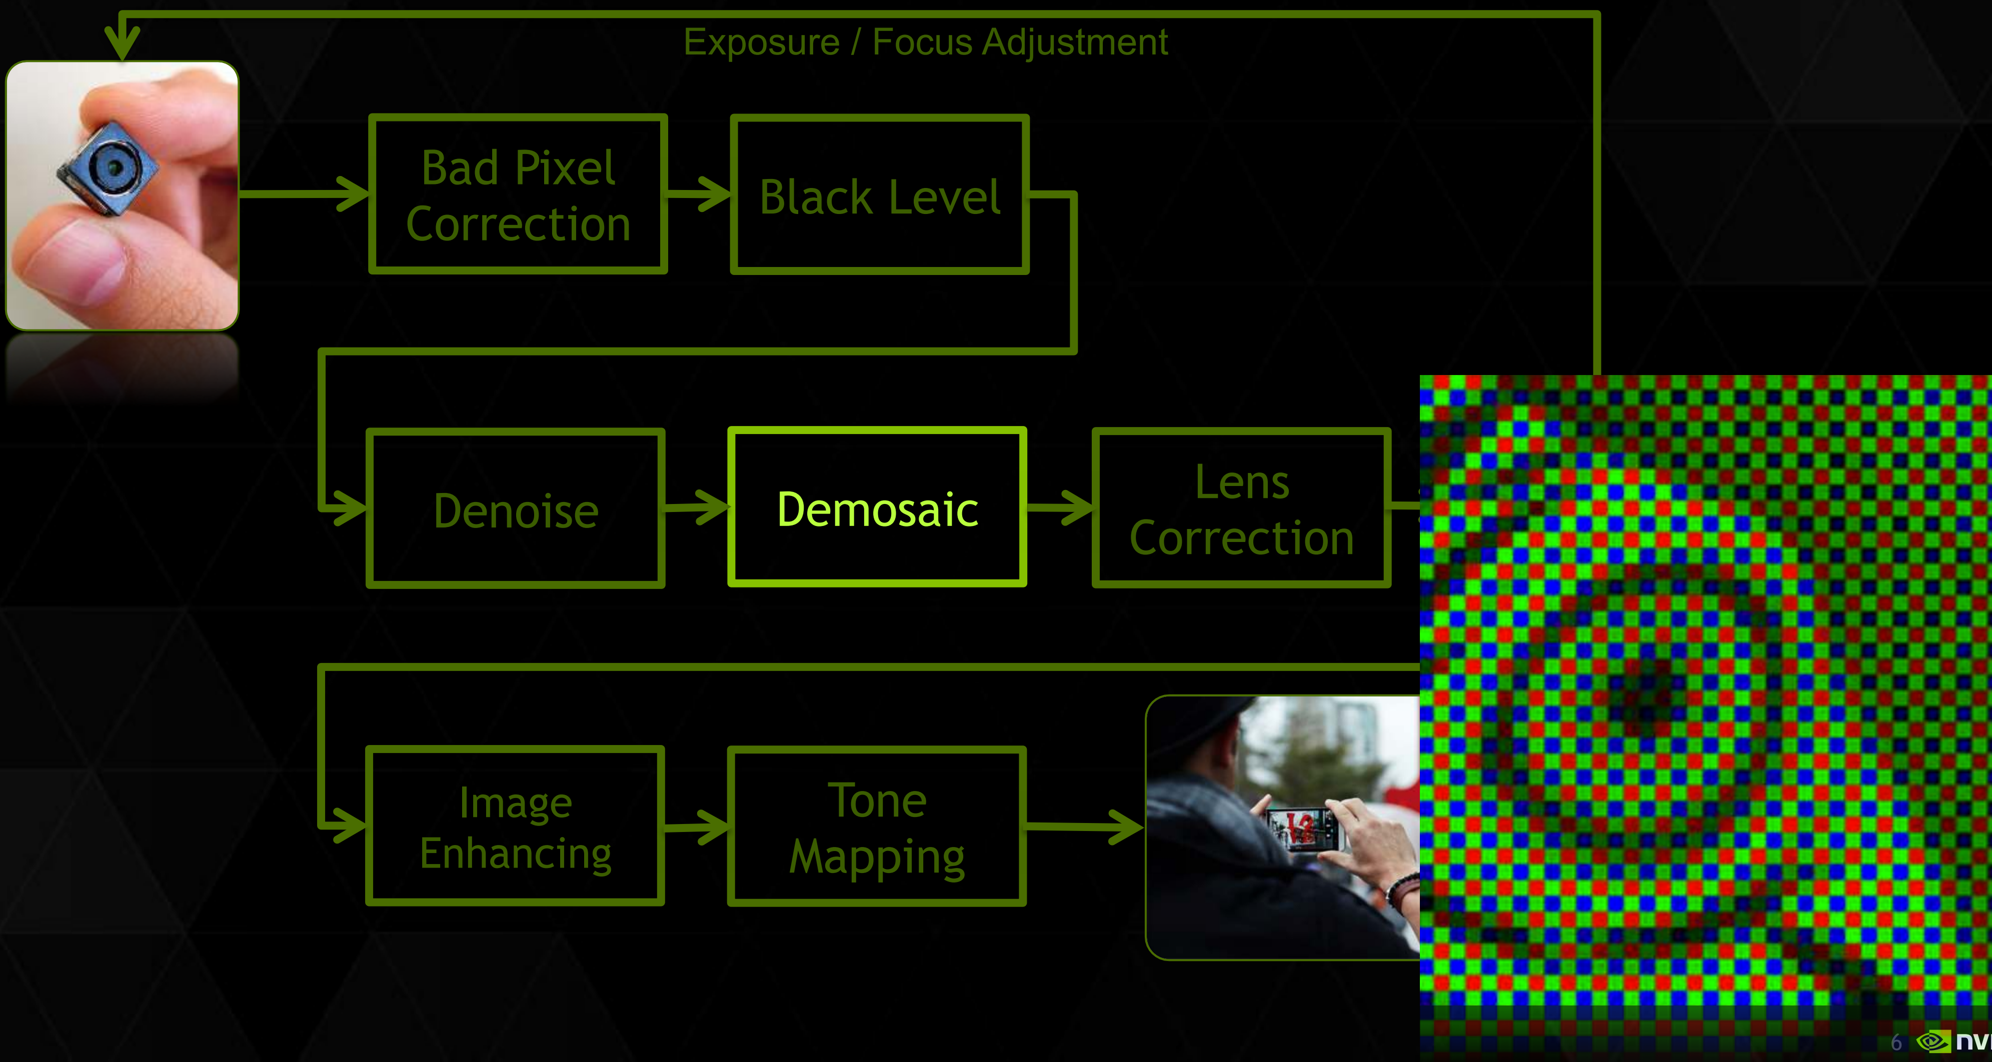

IMAGE PROCESSING PIPELINE

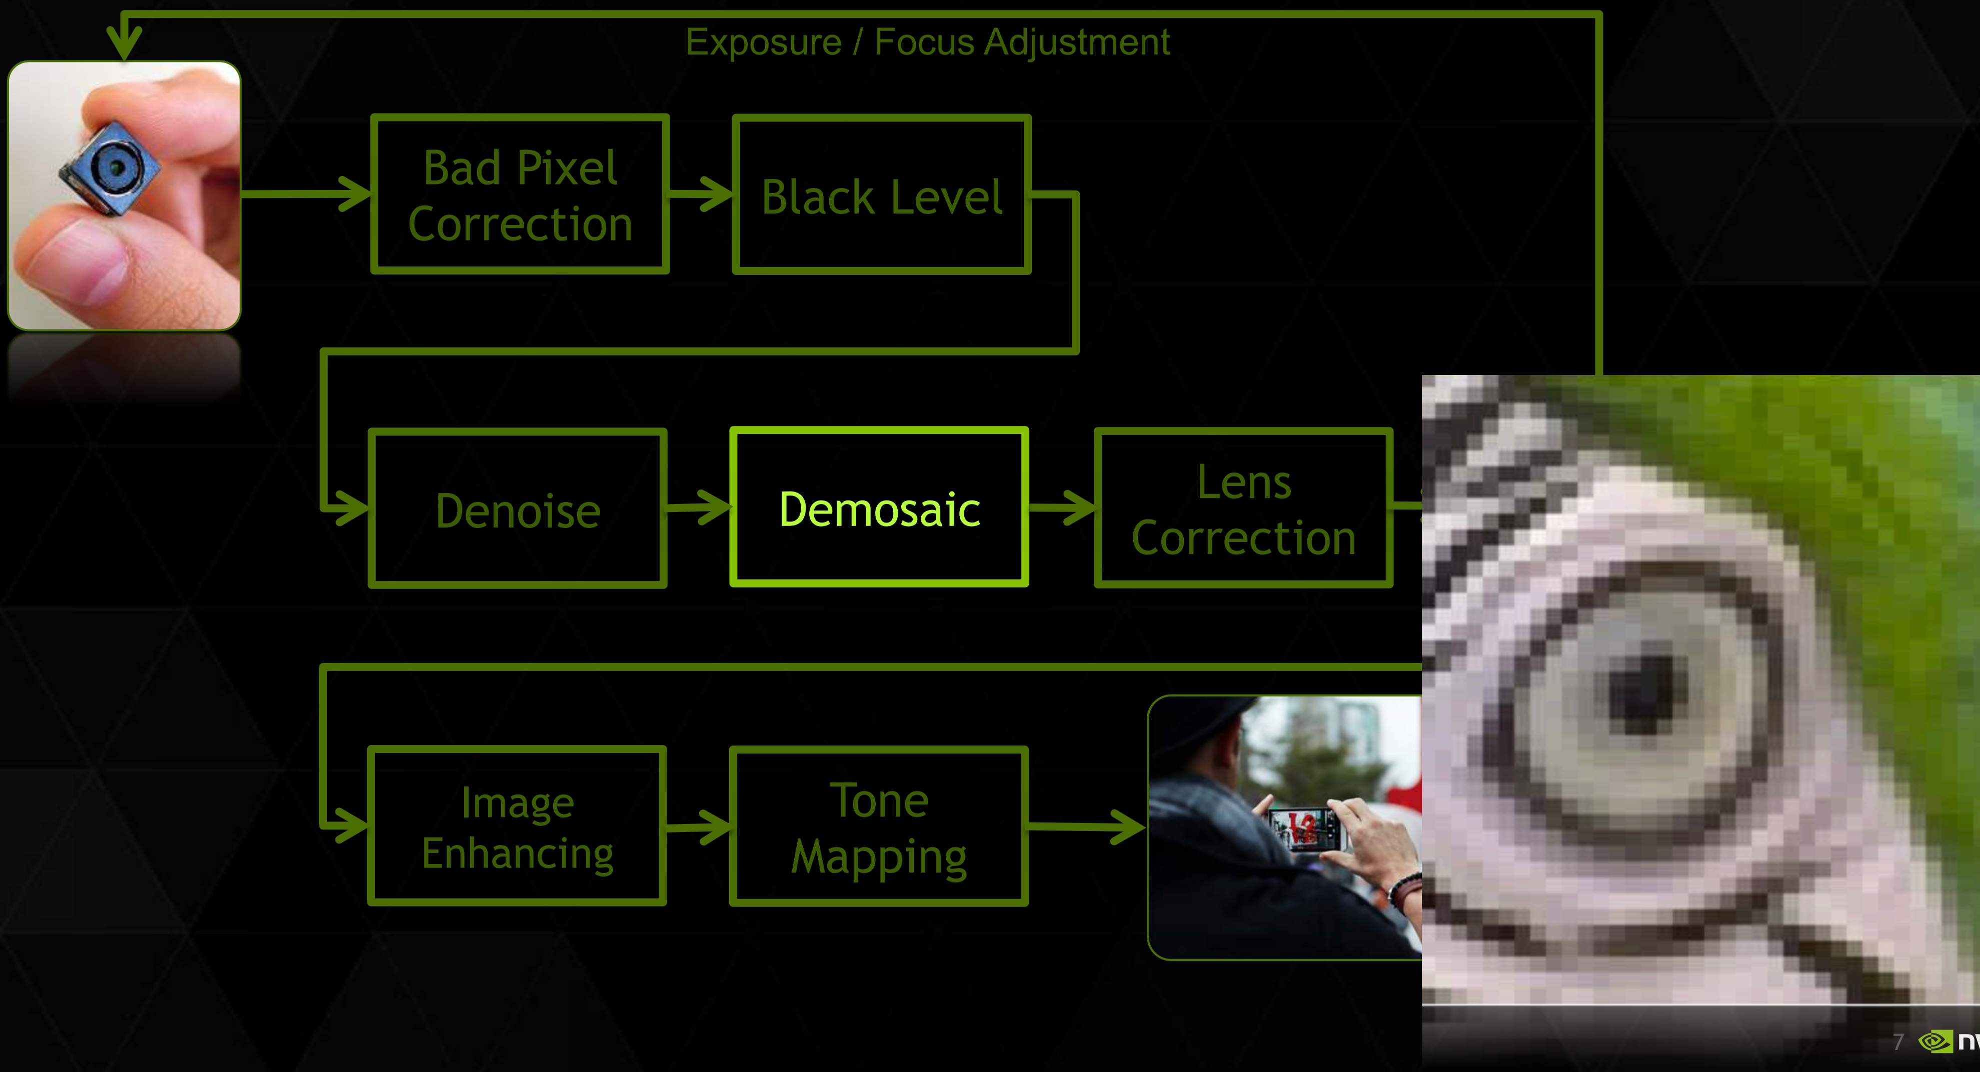

IMAGE PROCESSING PIPELINE

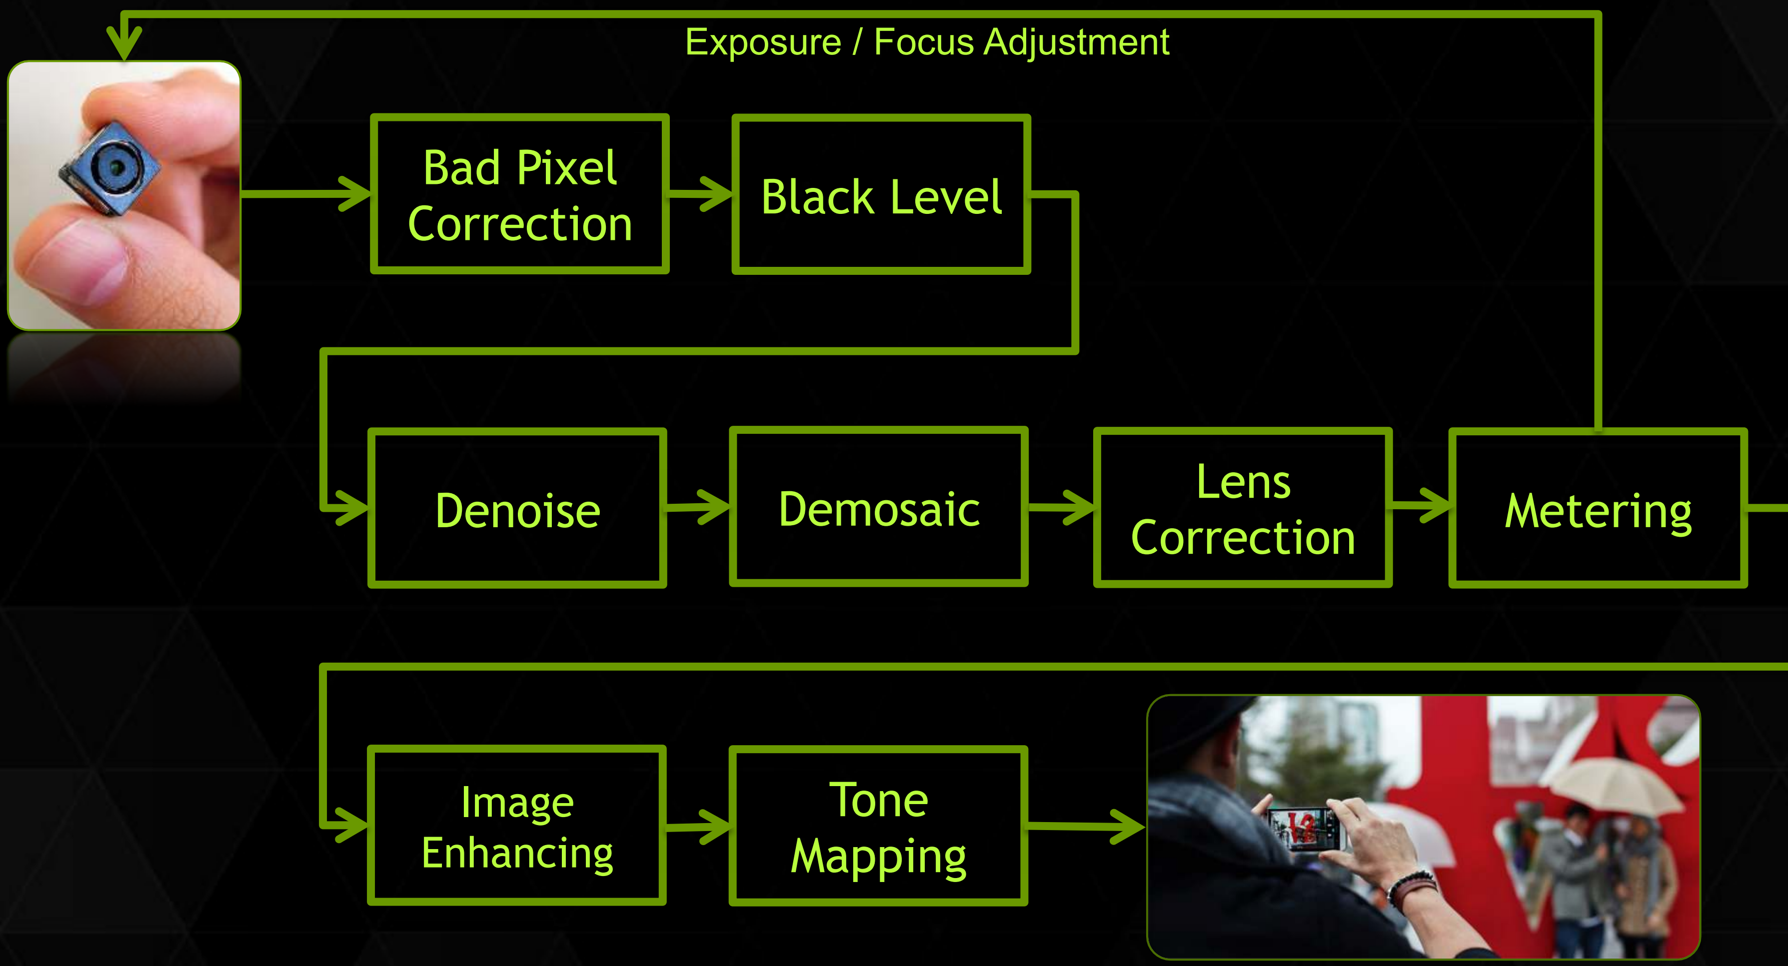

IMAGE PROCESSING PIPELINE

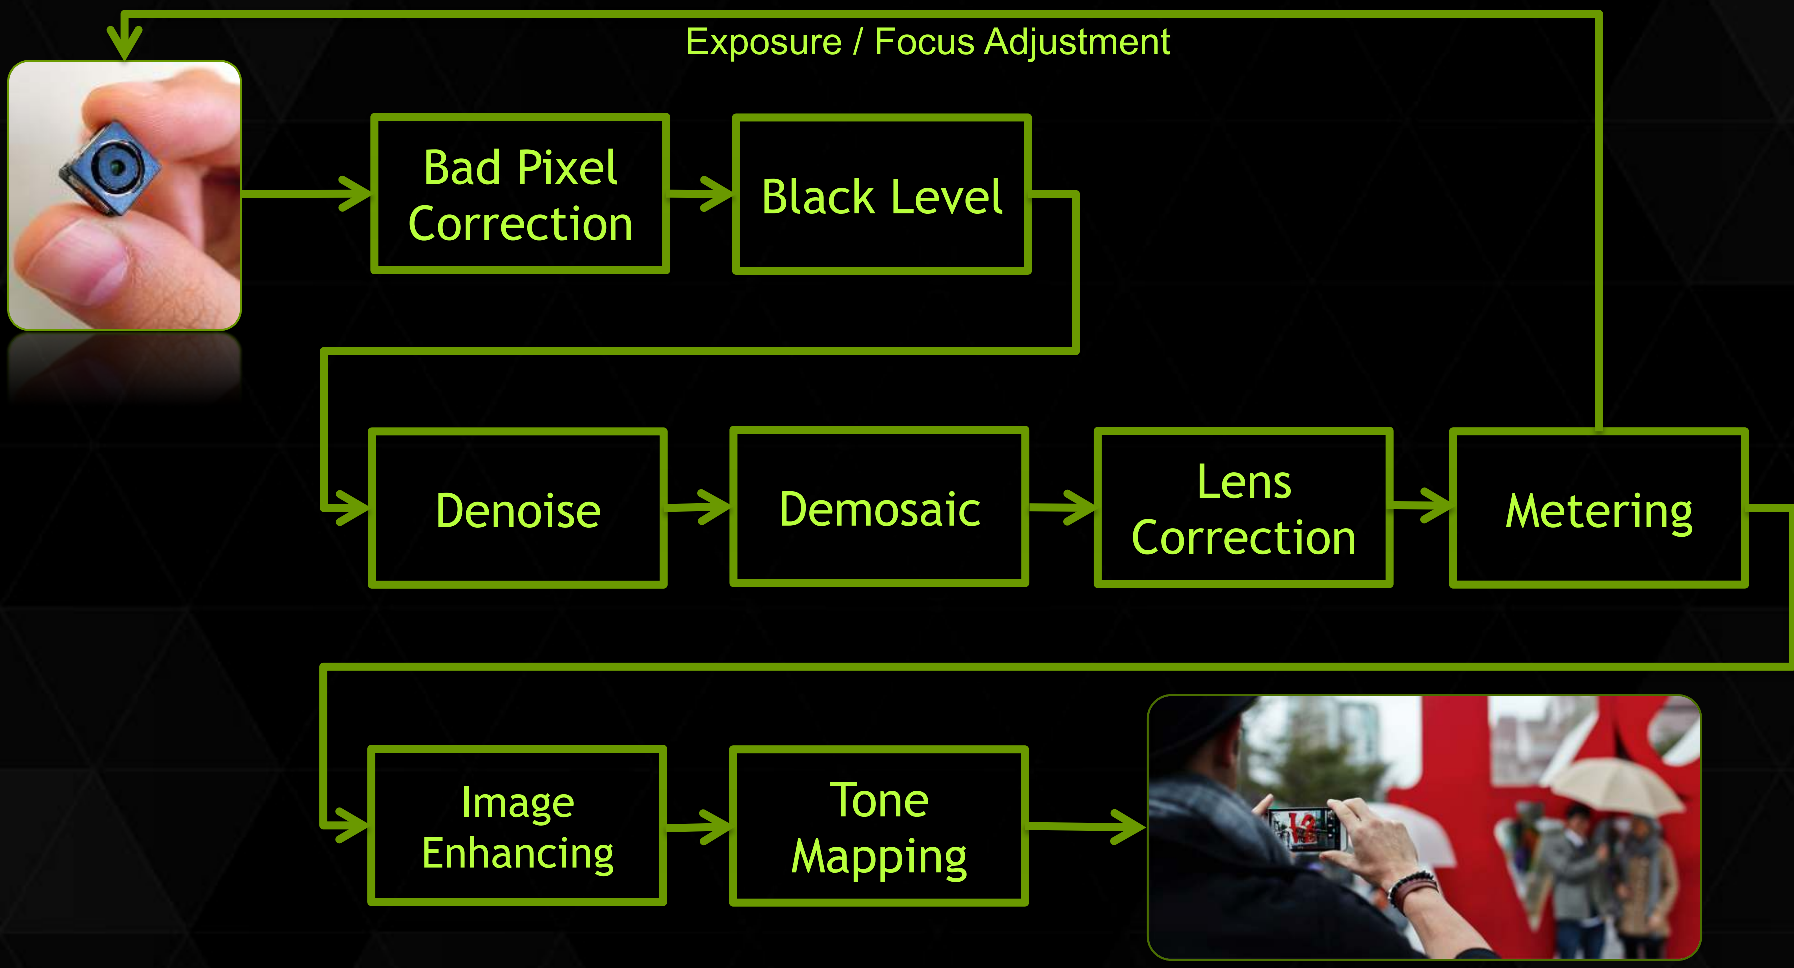

IMAGE PROCESSING PIPELINE

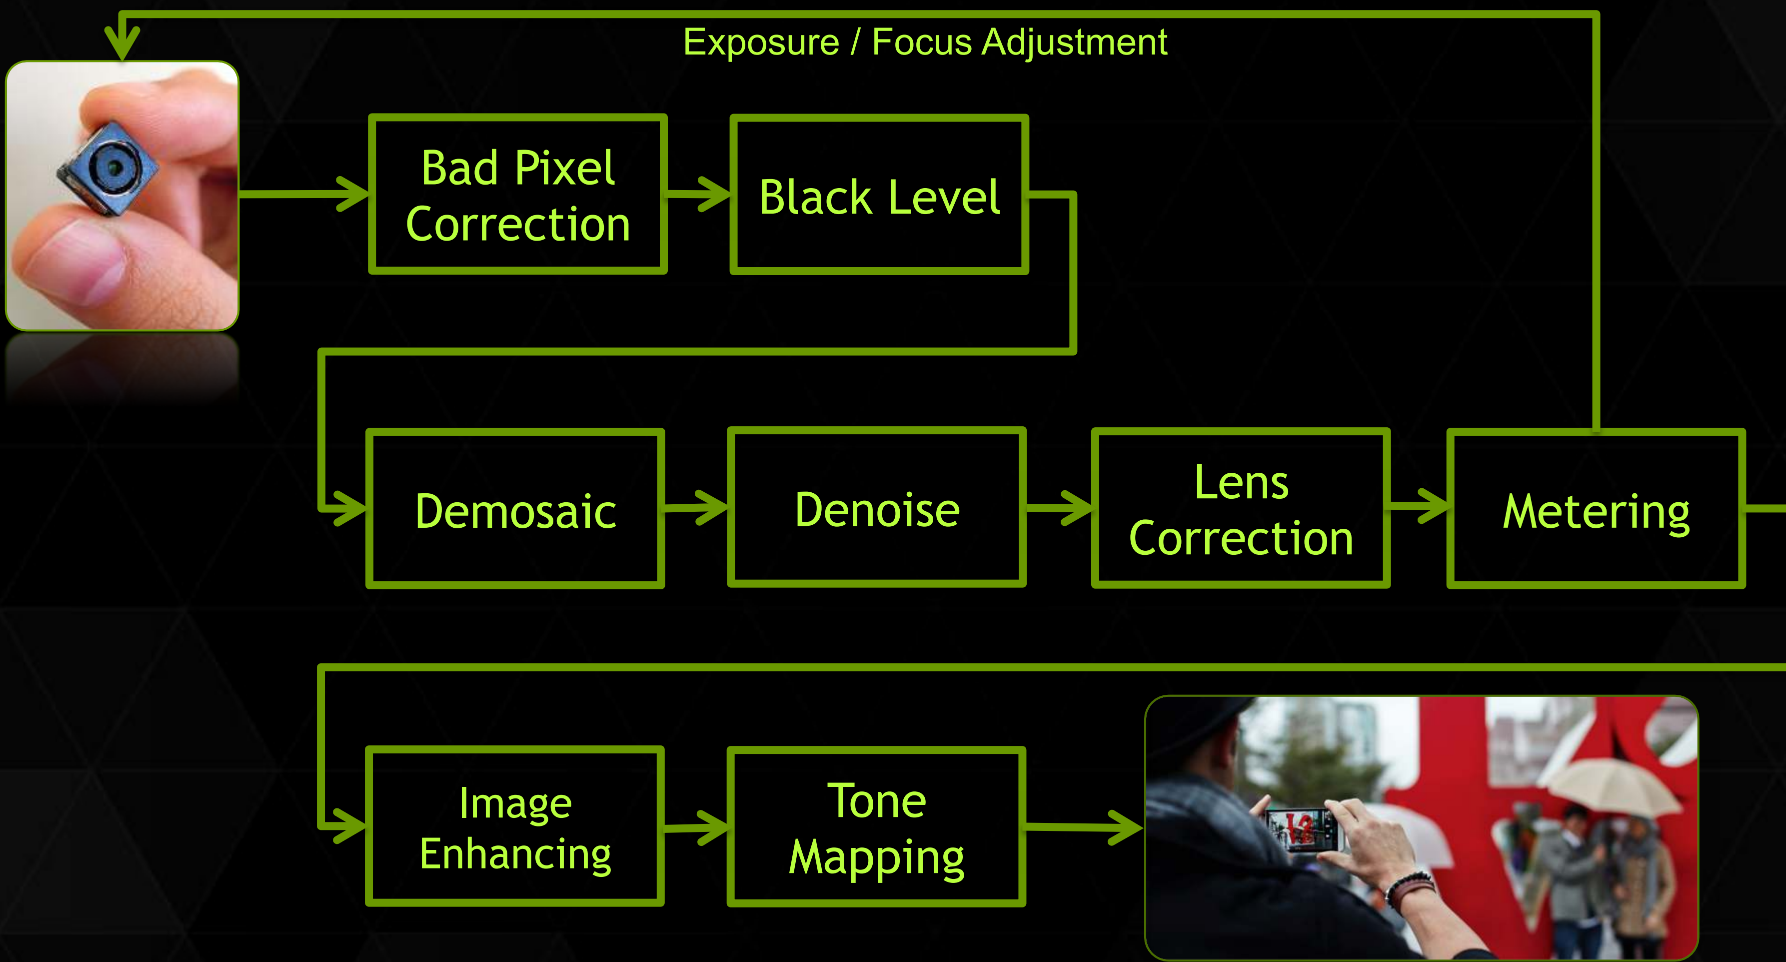

IMAGE PROCESSING PIPELINE

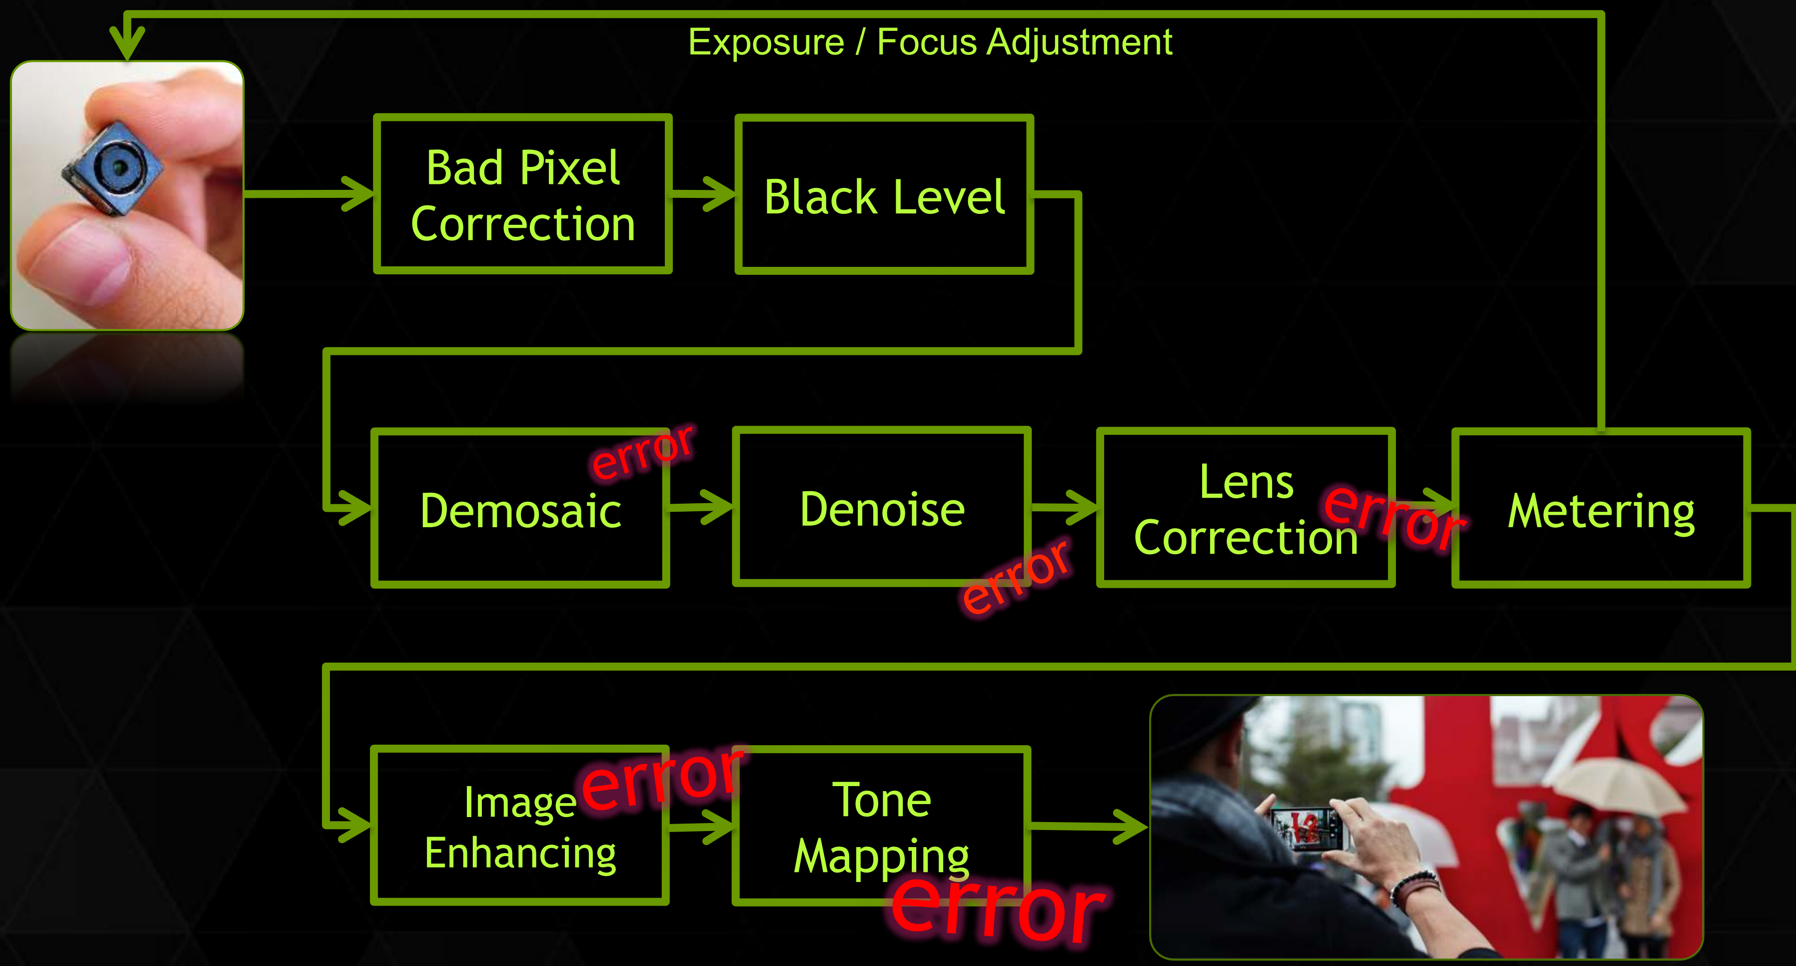

IMAGE PROCESSING PIPELINE

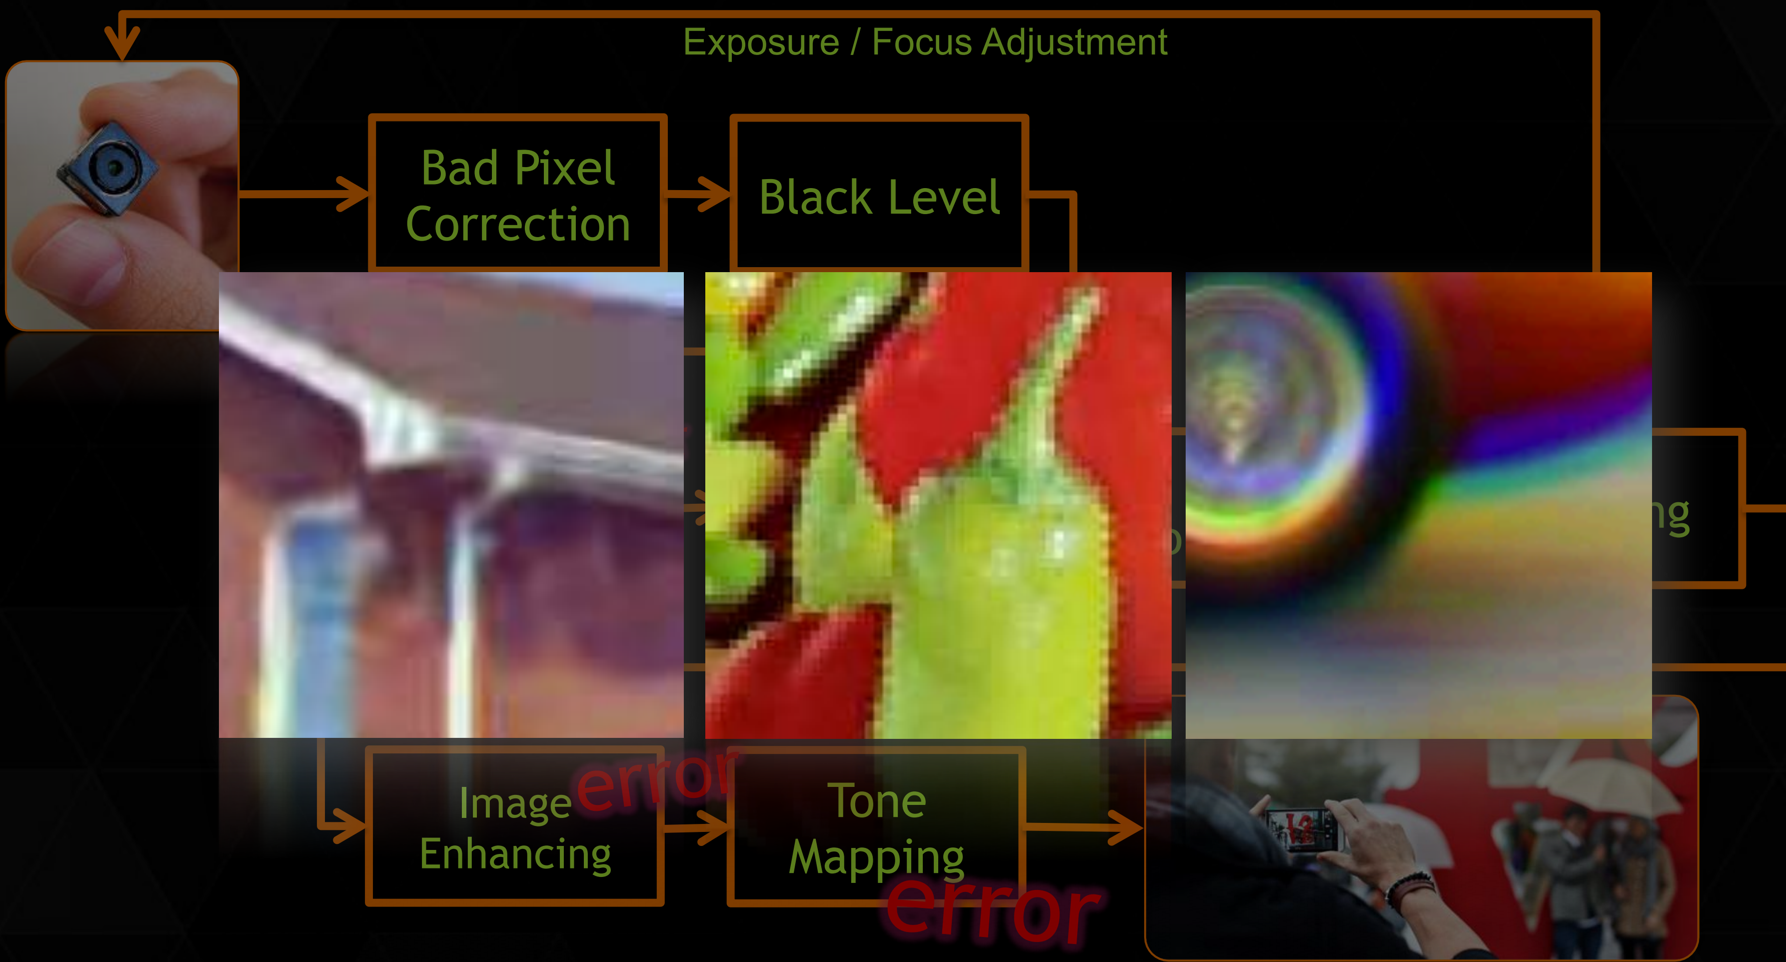

A BETTER IMAGE PROCESSOR

NO PIPELINE

FLEXIBILITY

PERFORMANCE

FLEXISP

Bayer | Burst | iHDR | Color Camera Arra

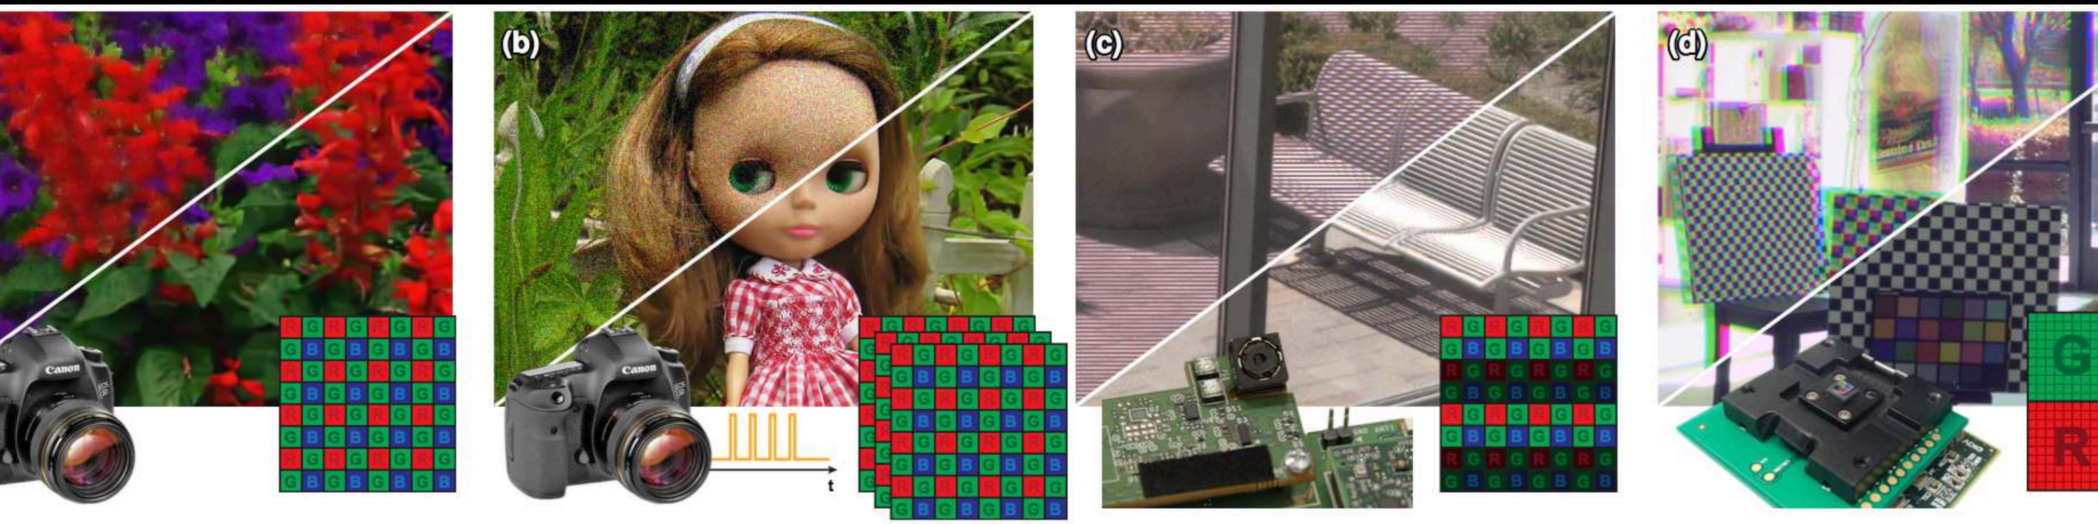

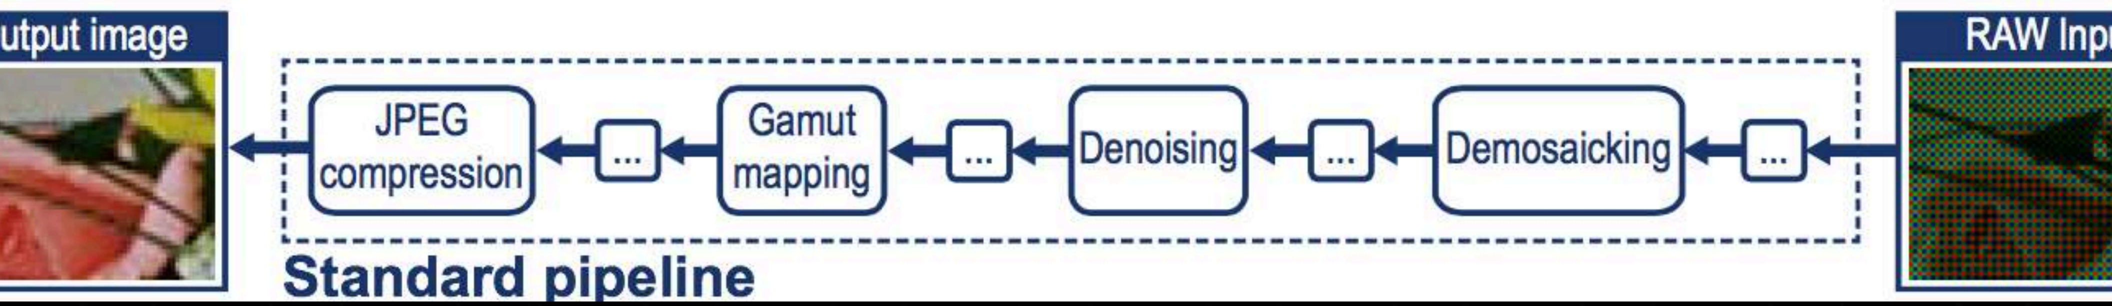
utput image
JPEG compression
...
Gamut mapping
...
Denoising
...
Demosaicking
...
RAW Inp
Standard pipeline

JOINT OPTIMIZATION

Lens Blur

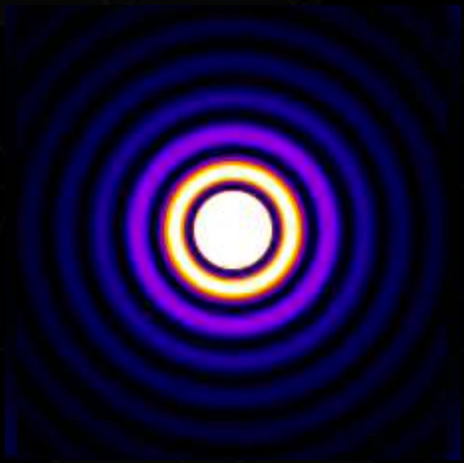

Bayer Patterns

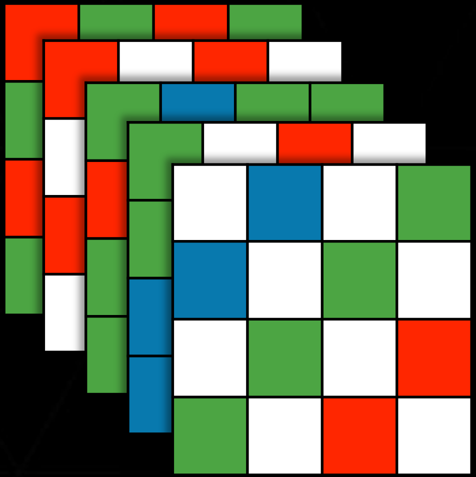

Noise

Motion / View

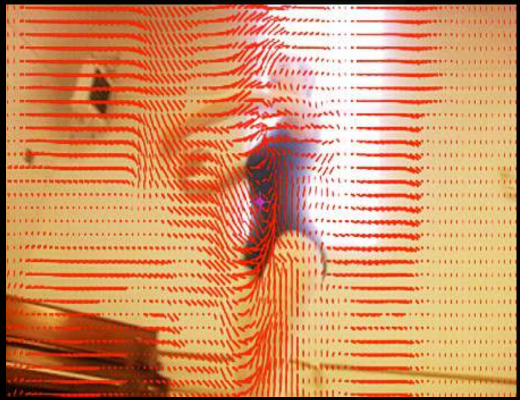

primal (solution image)

dual (constraints)

$$\min_{\mathbf{x},\mathbf{y}} \frac{1}{2} \|\mathbf{z} - \mathbf{A}\mathbf{x}\|_2^2 + F(\mathbf{y}), \quad \text{subject to } \boxed{\mathbf{K}\mathbf{x} = \mathbf{y}}$$

image formation model

measurement

Algorithm 1 – First order primal-dual

Initialization: $\gamma\tau\|\mathbf{K}\|^2 < 1$, $\theta \in [0,1]$, $(\mathbf{x}^0, \mathbf{y}^0)$, $\boxed{\bar{\mathbf{x}}^0 = \mathbf{x}^0}$.

Repeat until convergence:

$$\boxed{\text{penalty:} \quad \mathbf{y}^{k+1} = \mathrm{prox}_{\gamma F^*}\left(\mathbf{y}^k + \gamma \mathbf{K}\bar{\mathbf{x}}^k\right)}$$

$$\boxed{\text{data fidelity:} \quad \mathbf{x}^{k+1} = \mathrm{prox}_{\tau G}\left(\mathbf{x}^k - \tau \mathbf{K}^T \mathbf{y}^{k+1}\right)}$$

$$\boxed{\text{extrapolation:} \quad \bar{\mathbf{x}}^{k+1} = \mathbf{x}^{k+1} + \theta\left(\mathbf{x}^{k+1} - \mathbf{x}^k\right)}$$

primal (solution image)

dual (constraints)

$$\min_{\mathbf{x},\mathbf{y}} \frac{1}{2} \|\mathbf{z} - \mathbf{A}\mathbf{x}\|_2^2 + F(\mathbf{y}), \quad \text{subject to} \quad \mathbf{K}\mathbf{x} = \mathbf{y}$$

measurement

image formation model

Algorithm 1 – First order primal-dual

Initialization: $\gamma\tau\|\mathbf{K}\|^2 < 1$, $\theta \in [0,1]$, $(\mathbf{x}^0, \mathbf{y}^0)$, $\bar{\mathbf{x}}^0 = \mathbf{x}^0$.

Repeat until convergence:

penalty:	$\mathbf{y}^{k+1} = \text{prox}_{\gamma F^*}\left(\mathbf{y}^k + \gamma \mathbf{K}\bar{\mathbf{x}}^k\right)$
data fidelity:	$\mathbf{x}^{k+1} = \text{prox}_{\tau G}\left(\mathbf{x}^k - \tau \mathbf{K}^T \mathbf{y}^{k+1}\right)$
extrapolation:	$\bar{\mathbf{x}}^{k+1} = \mathbf{x}^{k+1} + \theta\left(\mathbf{x}^{k+1} - \mathbf{x}^k\right)$

$$\text{prox}_{\tau H}(\mathbf{v}) = \underset{\mathbf{u}}{\arg\min}\left(H(\mathbf{u}) + \frac{1}{2\tau}\|\mathbf{u} - \mathbf{v}\|_2^2\right)$$

3.3.1 Data Fidelity Operator

The data fidelity (primal) step in Alg. 1 updates the current estimate $\mathbf{x}$ of the latent image. Loosely speaking, this operator takes a gradient descent step of the data term while remaining close to the argument of the operator. From Eqs. 2 and 4 we see that

$$\mathbf{x} = \text{prox}_{\tau G}(\mathbf{v}) = \arg\min_{\mathbf{u}} \left(\frac{1}{2} \|\mathbf{z} - \mathbf{A}\mathbf{u}\|_2^2 + \frac{1}{2\tau} \|\mathbf{u} - \mathbf{v}\|_2^2 \right). \quad (6)$$

This is a simple least-squares problem, whose solution $\mathbf{x}$ is obtained as the solution to the following linear system:

$$\left(\tau \mathbf{A}^T \mathbf{A} + \mathbb{I}\right) \mathbf{x} = \left(\tau \mathbf{A}^T \mathbf{z} + \mathbf{v}\right). \quad (7)$$

There are many ways to solve this linear system, but as the matrix is symmetric and positive semi-definite, we can use the Conjugate Gradient (CG) algorithm. The advantage of CG in our setting is that the image formation model $\mathbf{A}$ and its transpose can be represented algorithmically, without explicitly generating the system matrix.

IMAGE FORMATION

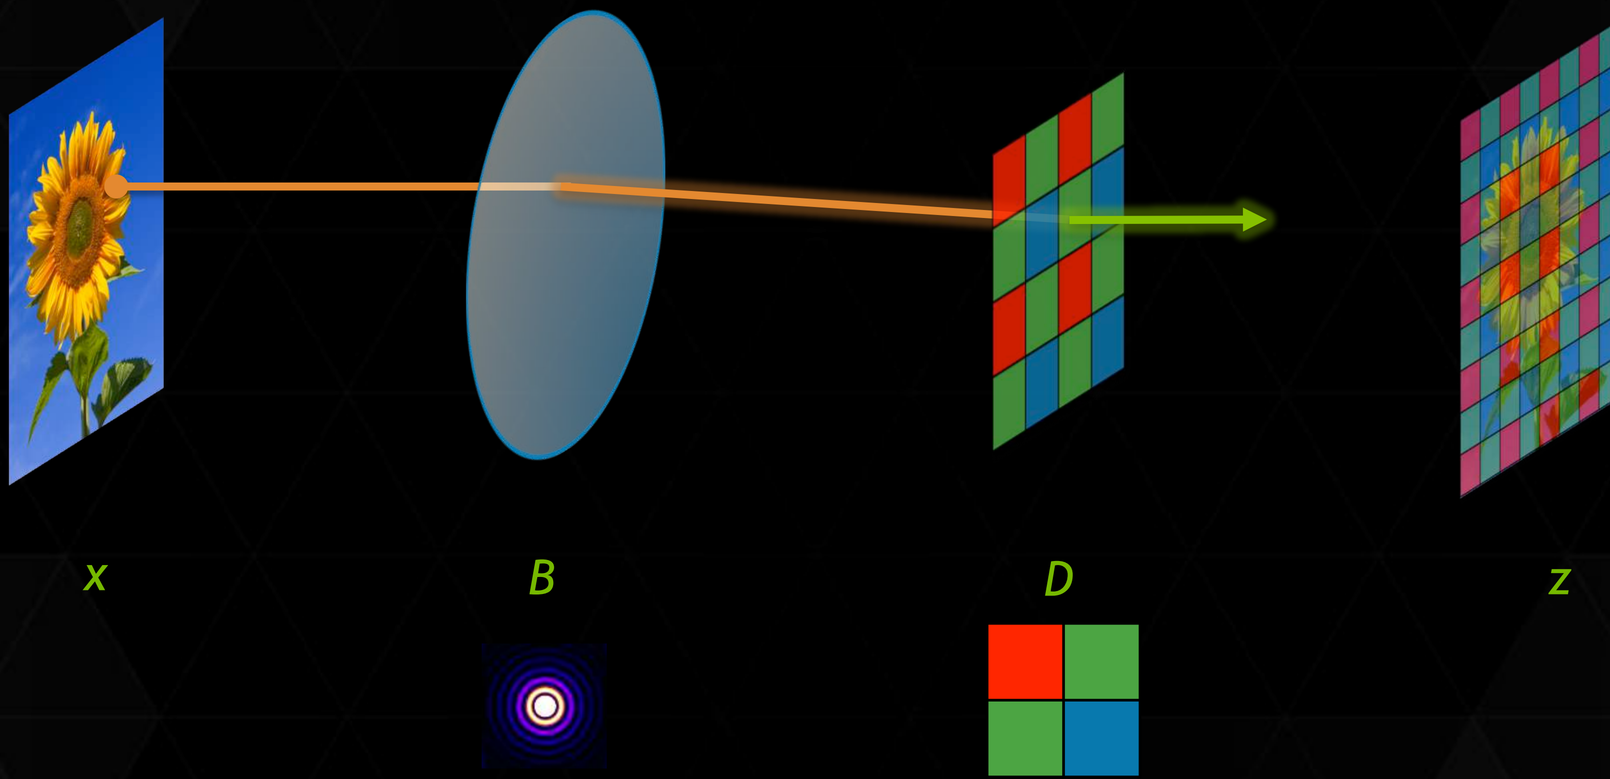

IMAGE FORMATION

$$z = DBx$$

x

B

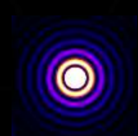

D

z

IMAGE FORMATION

$$z = A\,x$$

D *B*

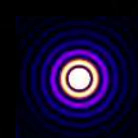

LINEAR SYSTEM

ill posed!

$$argmin_x \underbrace{|z - Ax|_2}_{\text{errors w.r.t. observation}}$$

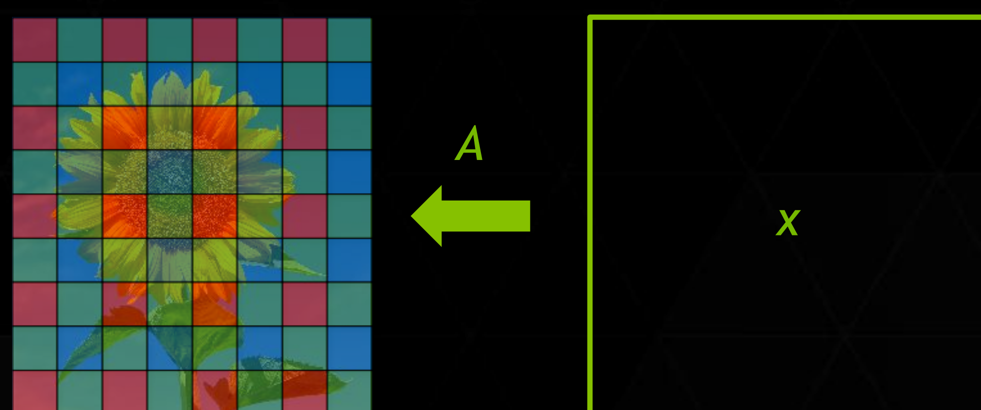

LINEAR SYSTEM

$$argmin_x \; |z - Ax|_2 + \underbrace{\lambda(x)}_{\text{regularization}}$$

TOTAL VARIATION

(Rudin et al. '92)

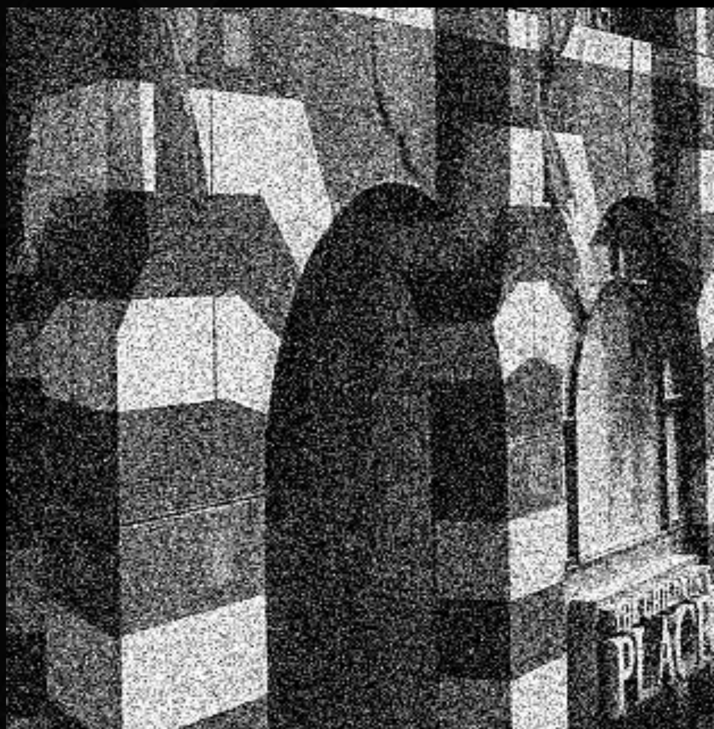

$$\mathrm{TV} = \sum \|\nabla X\|_1$$

CROSS-CHANNEL

(Heide et al. ‘13)

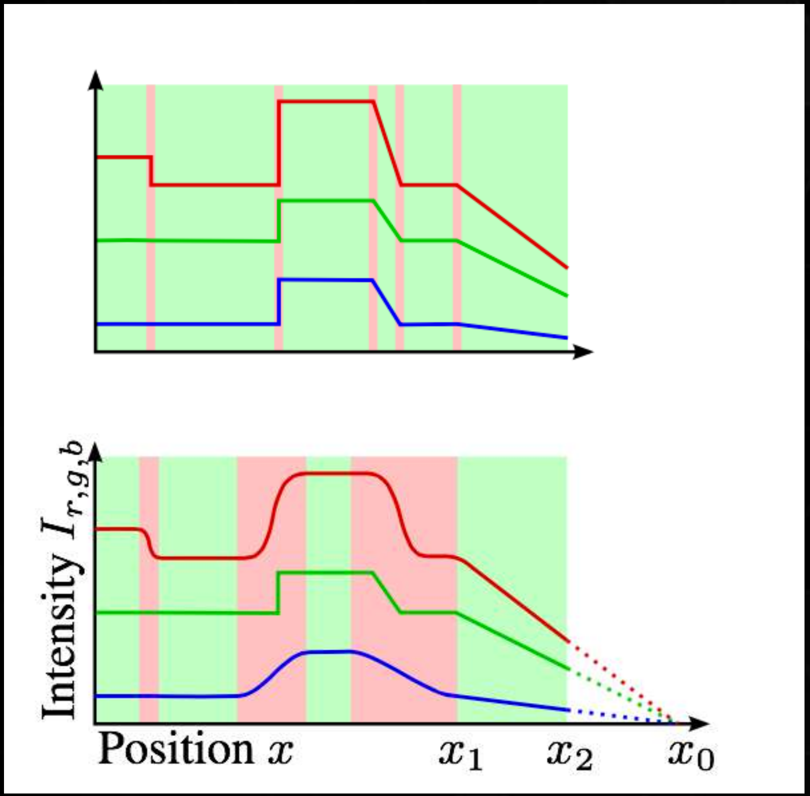

$$\frac{\nabla X_r}{X_r} \approx \frac{\nabla X_g}{X_g} \Rightarrow \nabla X_r \cdot X_g - \nabla X_g \cdot X_r \approx 0$$

COLLABORATIVE DENOISING

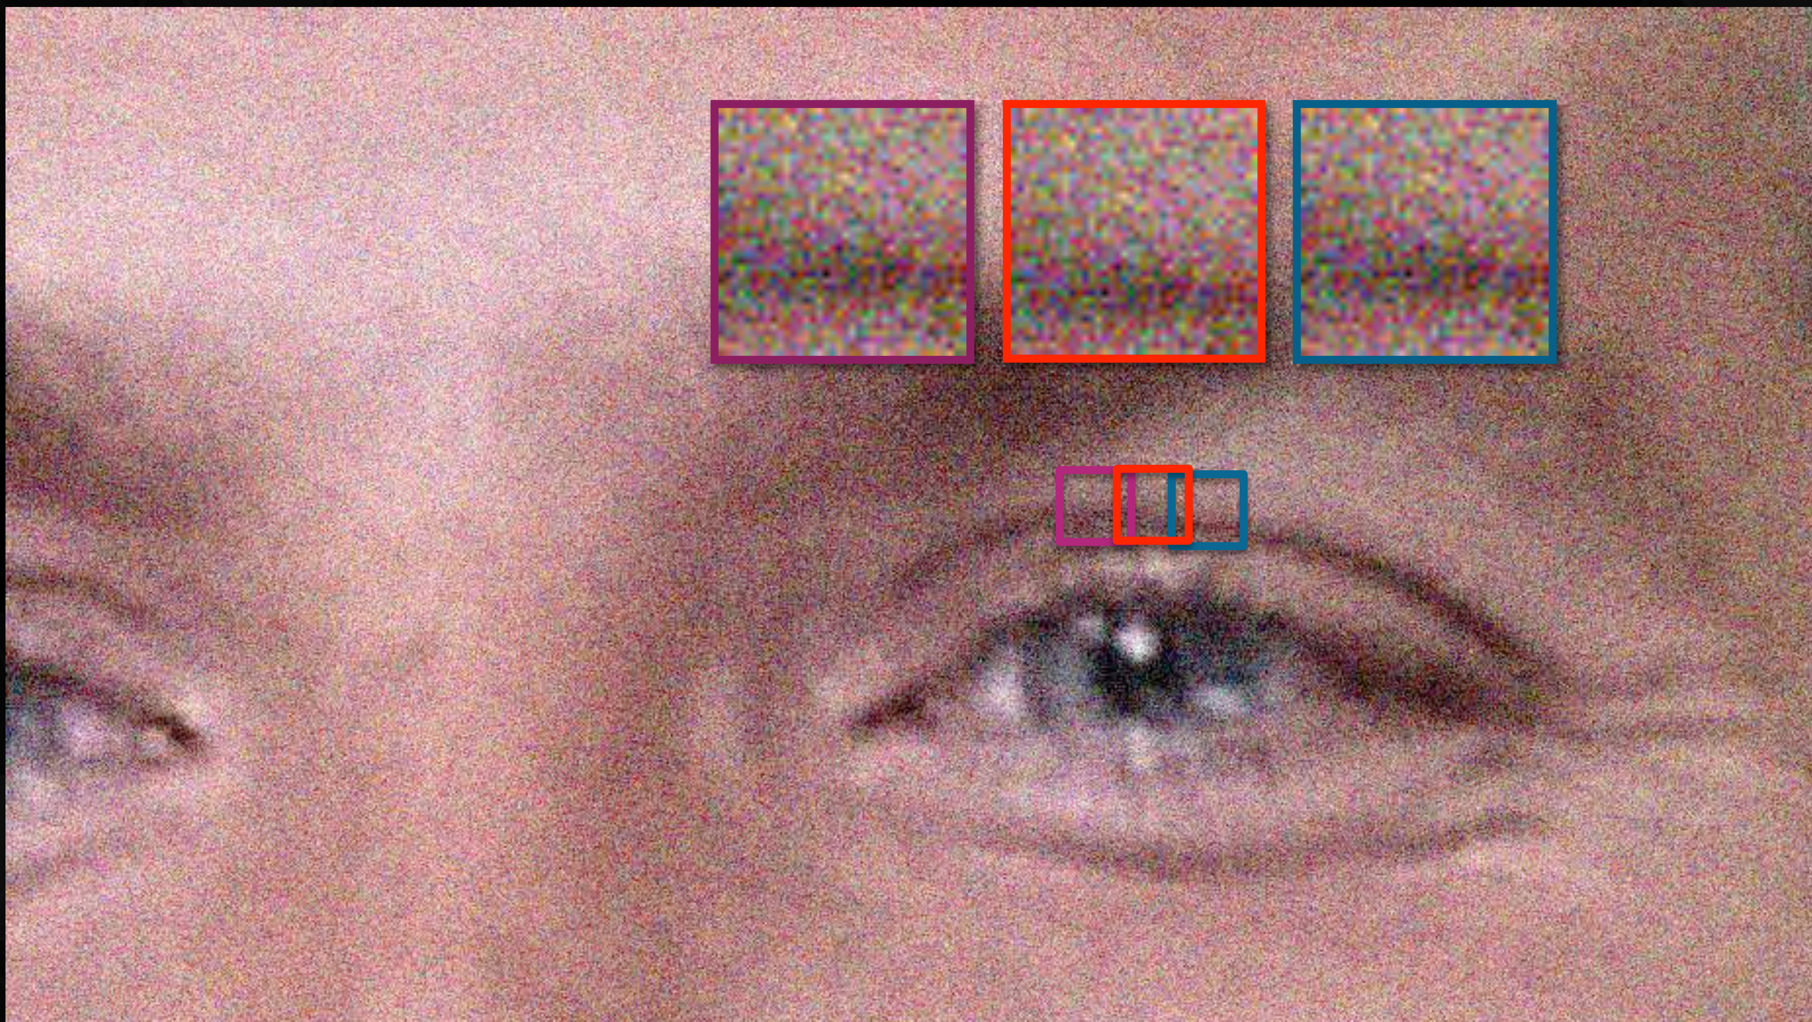

COLLABORATIVE DENOISING

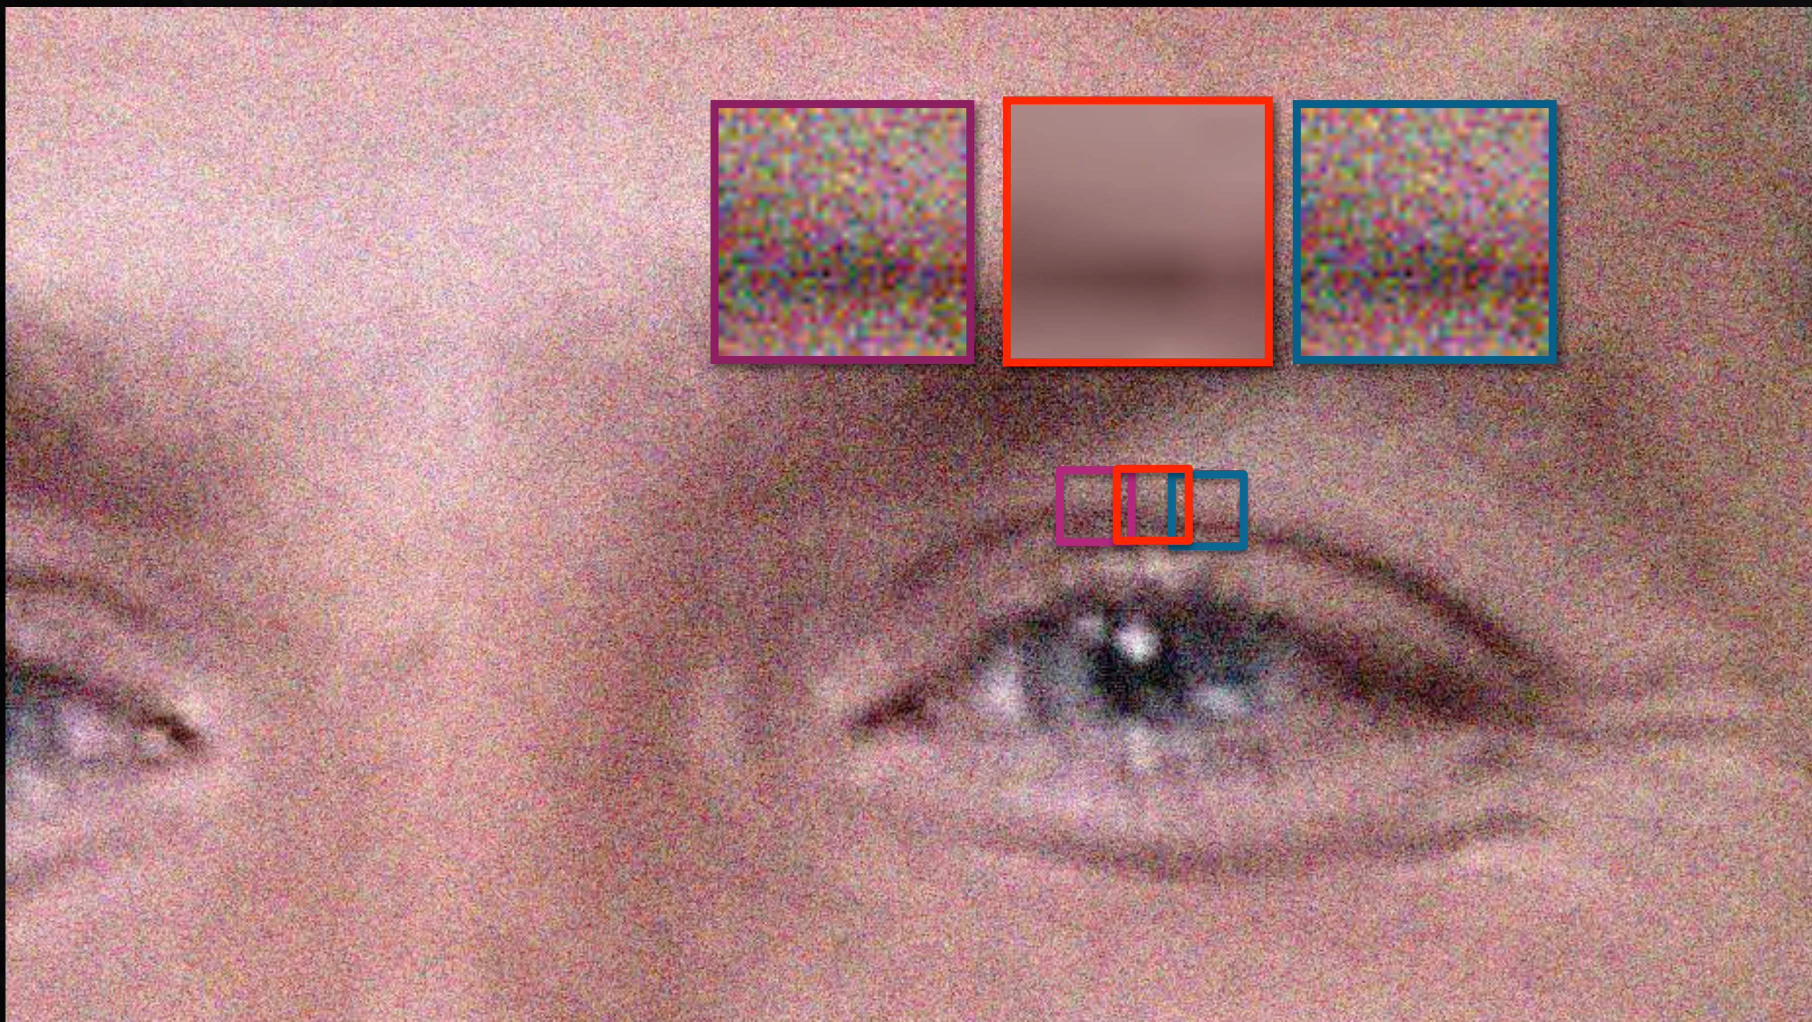

LINEAR SYSTEM

$$\operatorname{argmin}_x \, |z - Ax|_2 + \lambda(x)$$

λ

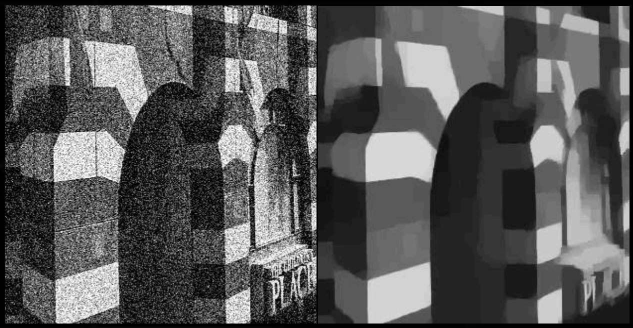

Total variation

+

Cross-channel
(Heide et al. '13)

+

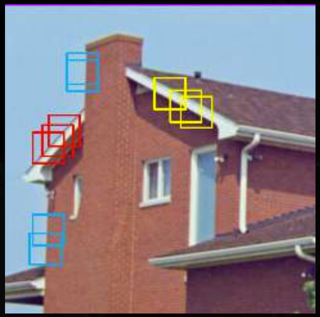

Self-similarity
denoising

WEIGHTS

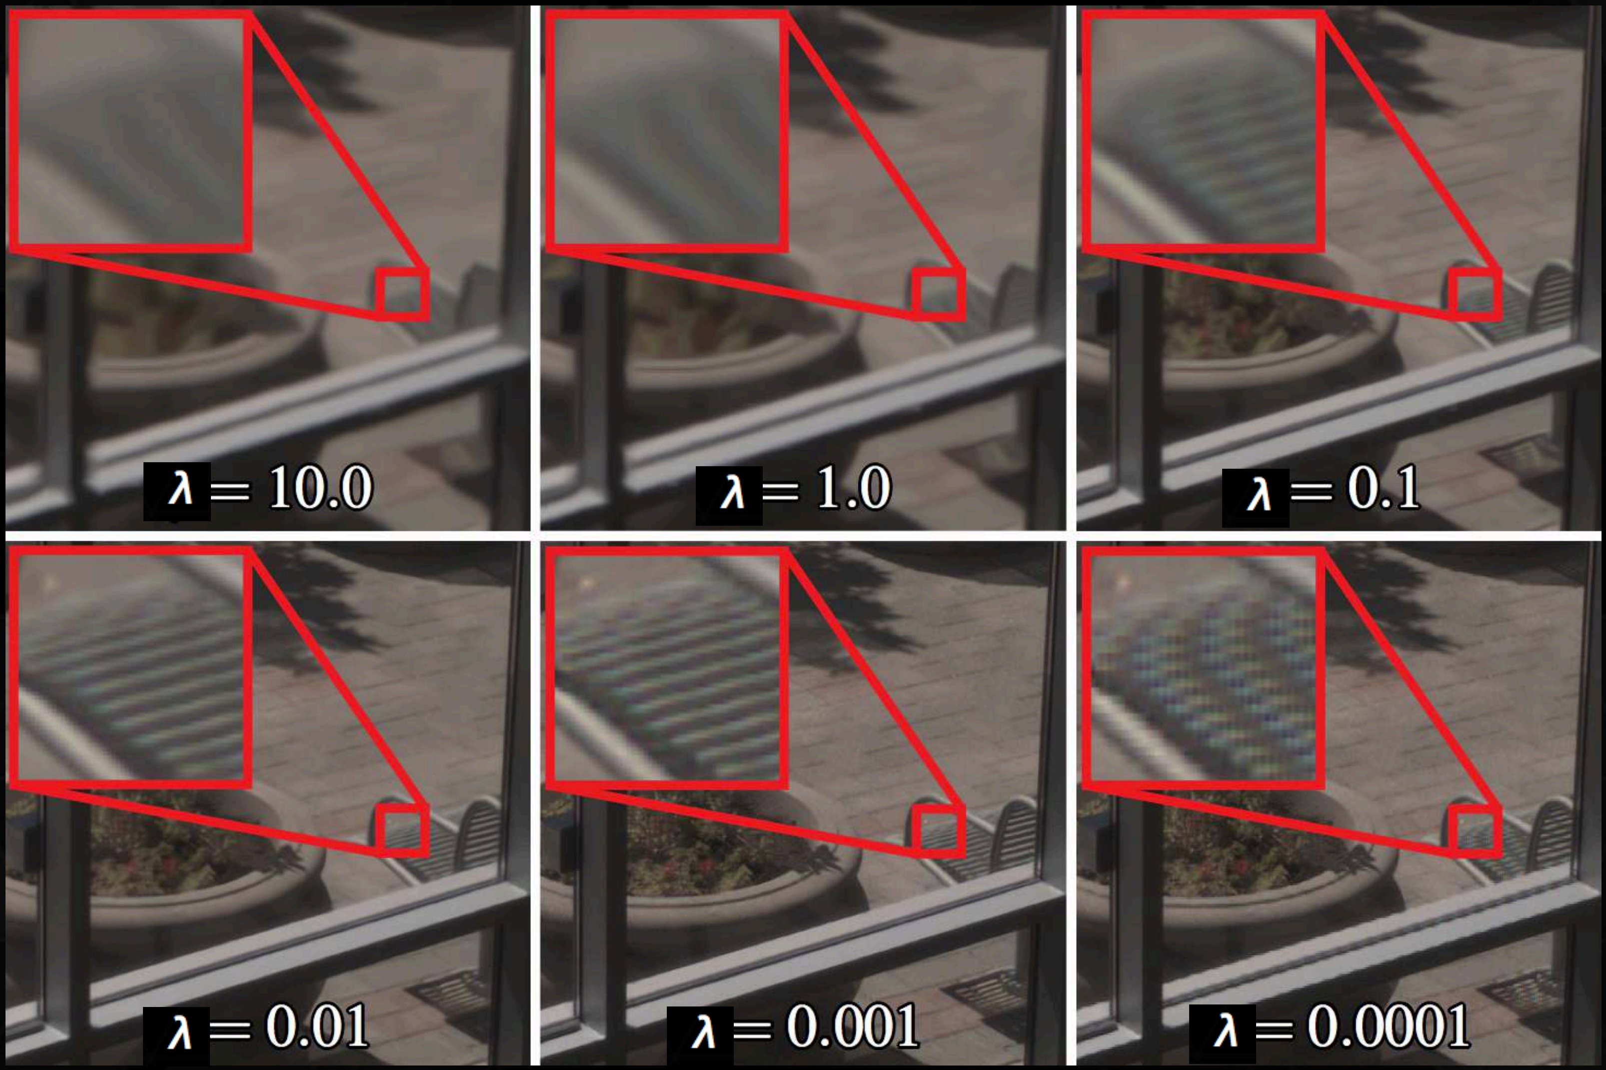

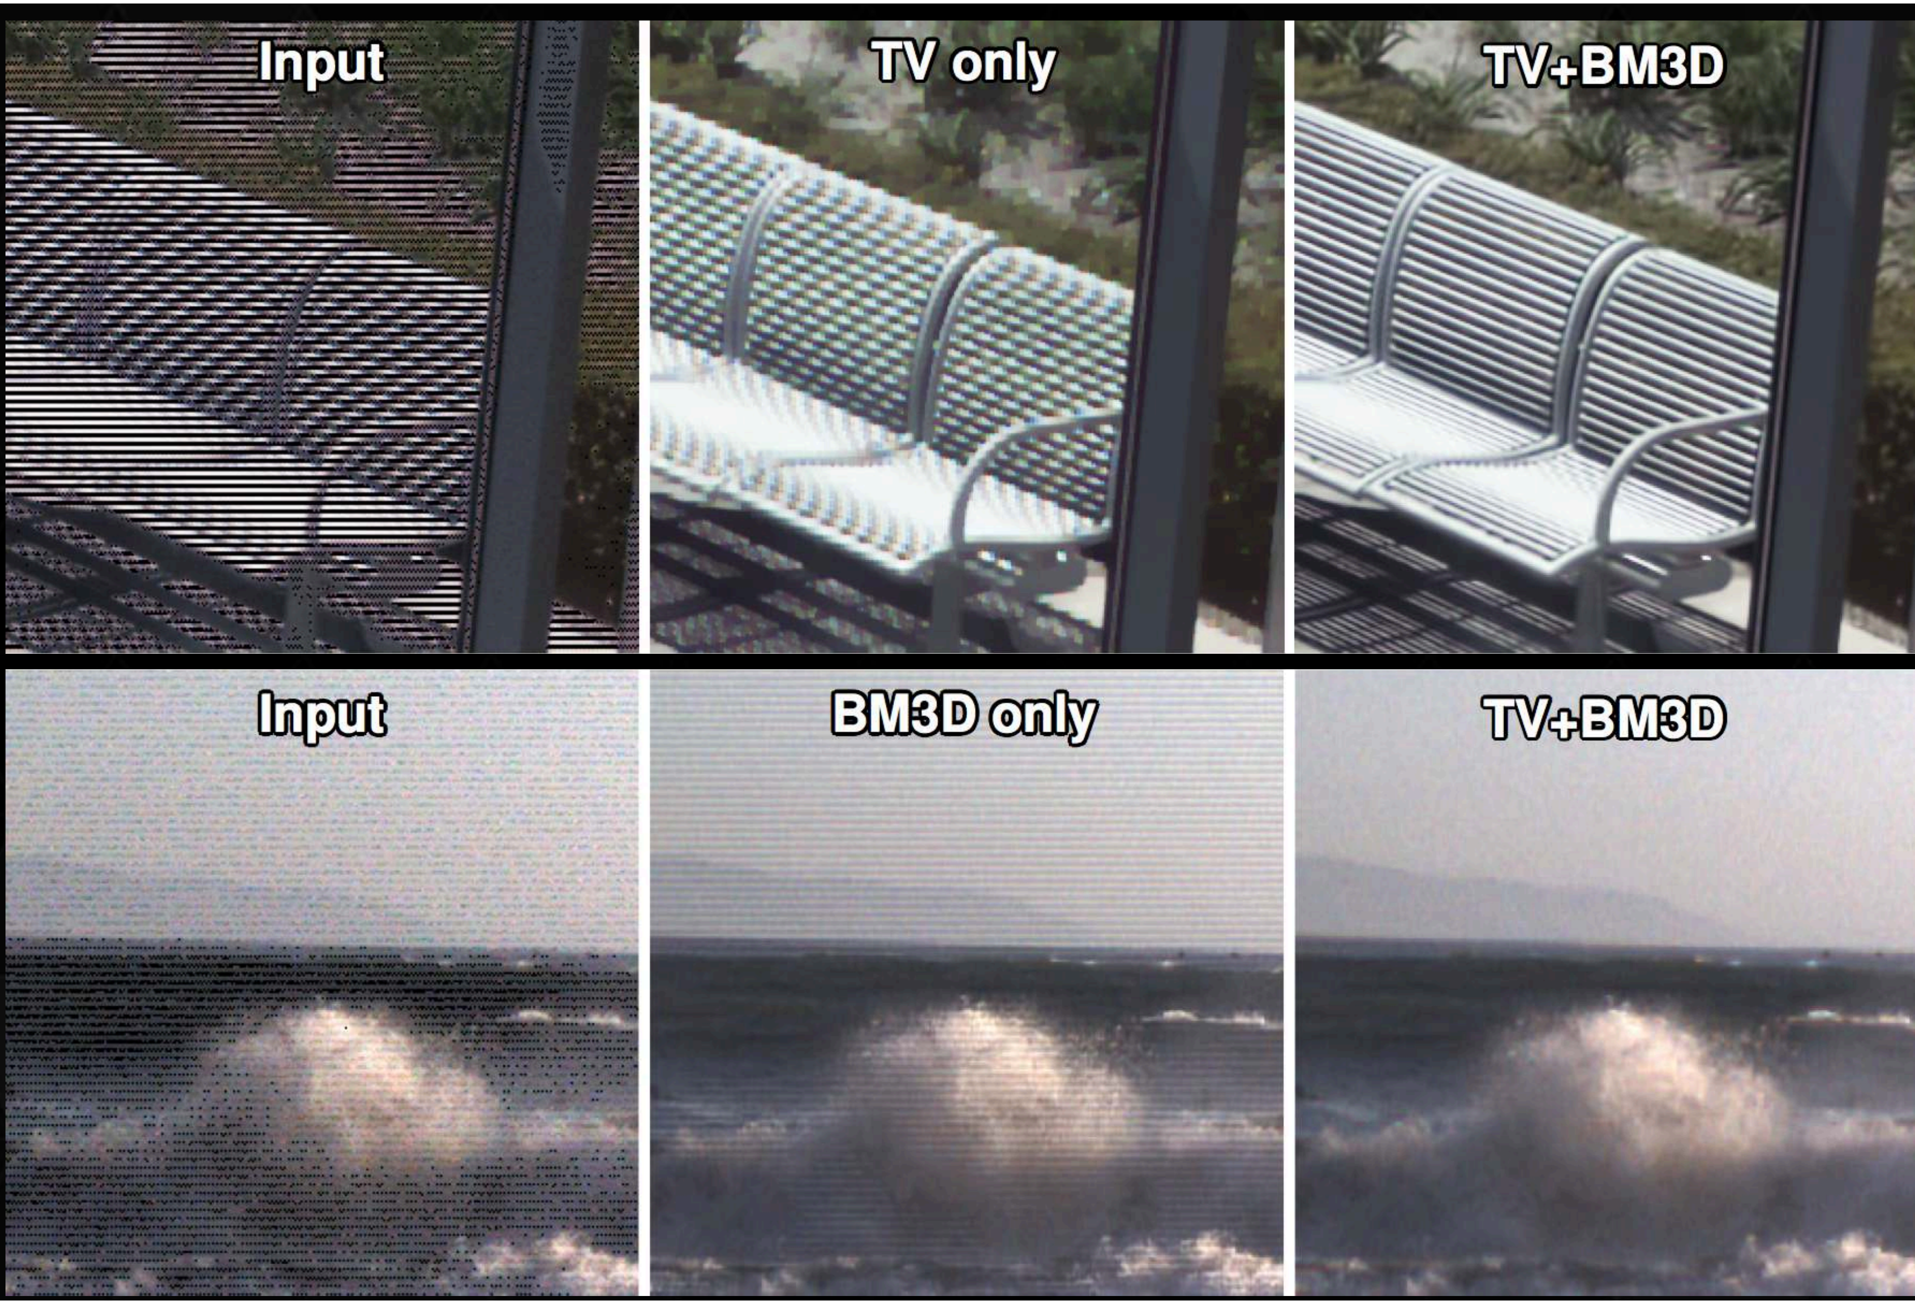
Input
TV only
TV+BM3D
Input
BM3D only
TV+BM3D

MORE ON PRIORS…

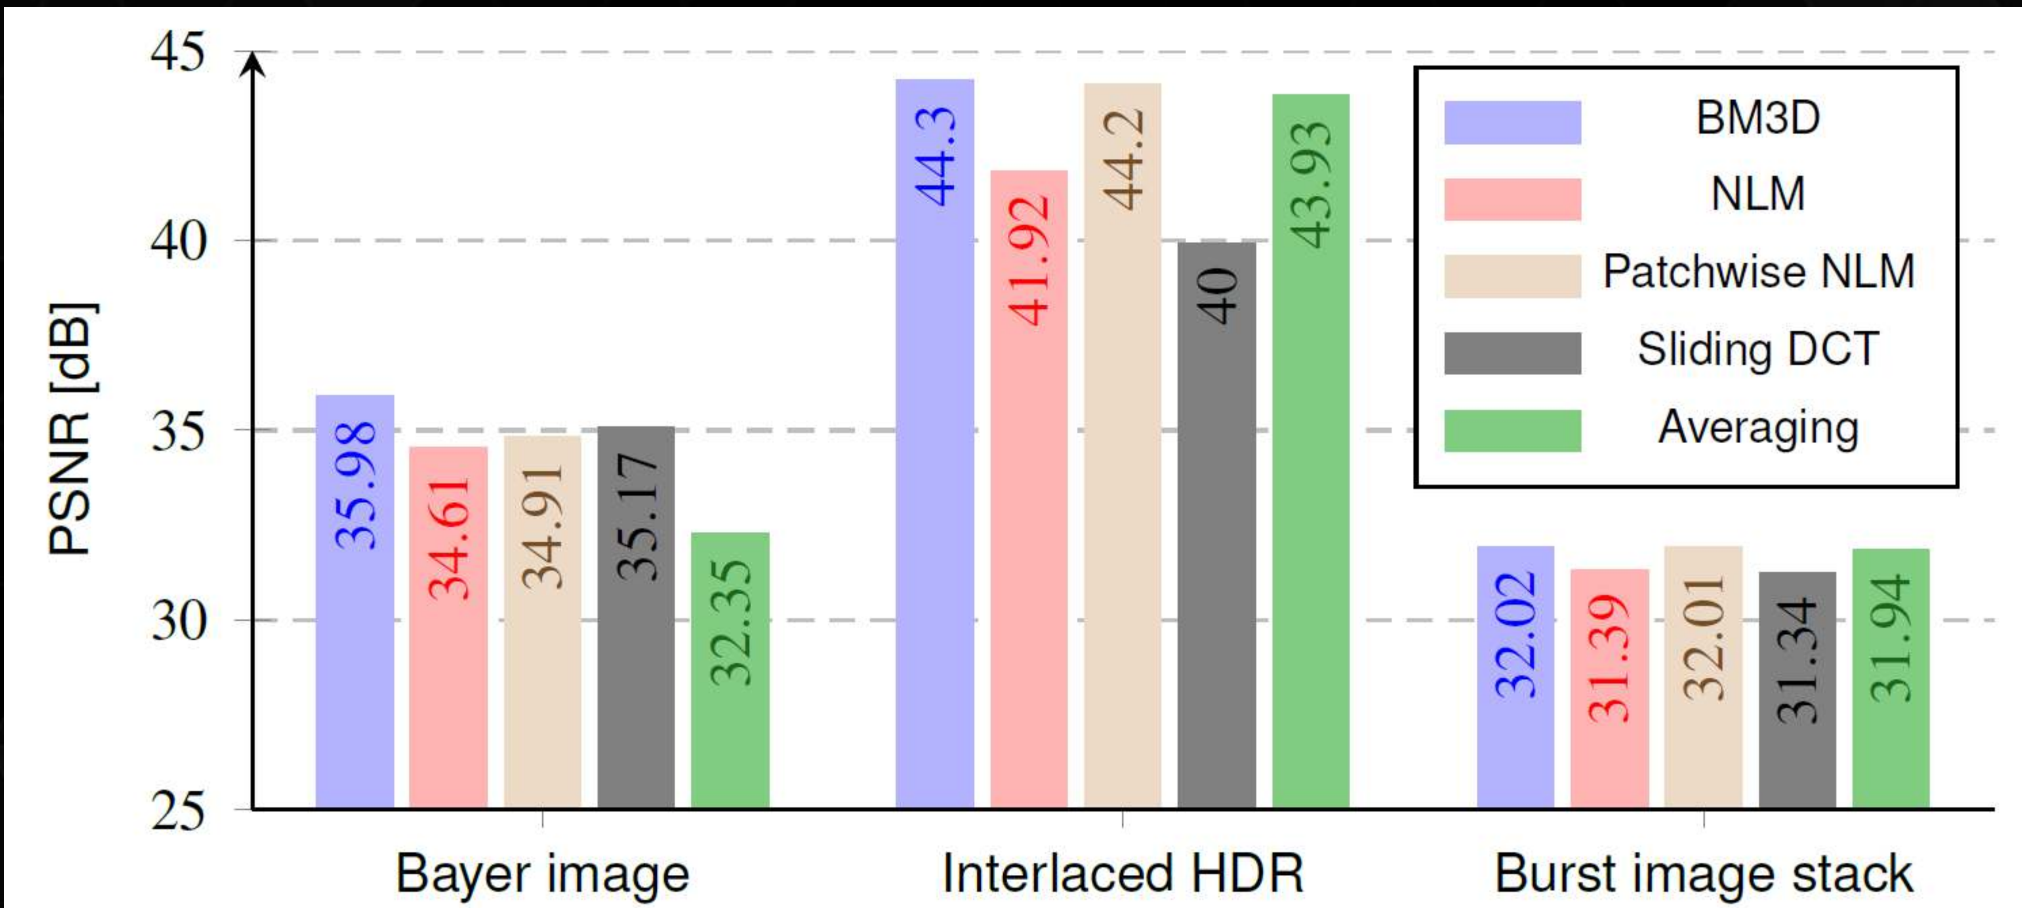

CONVERGENCE (DEMOSAICKING)

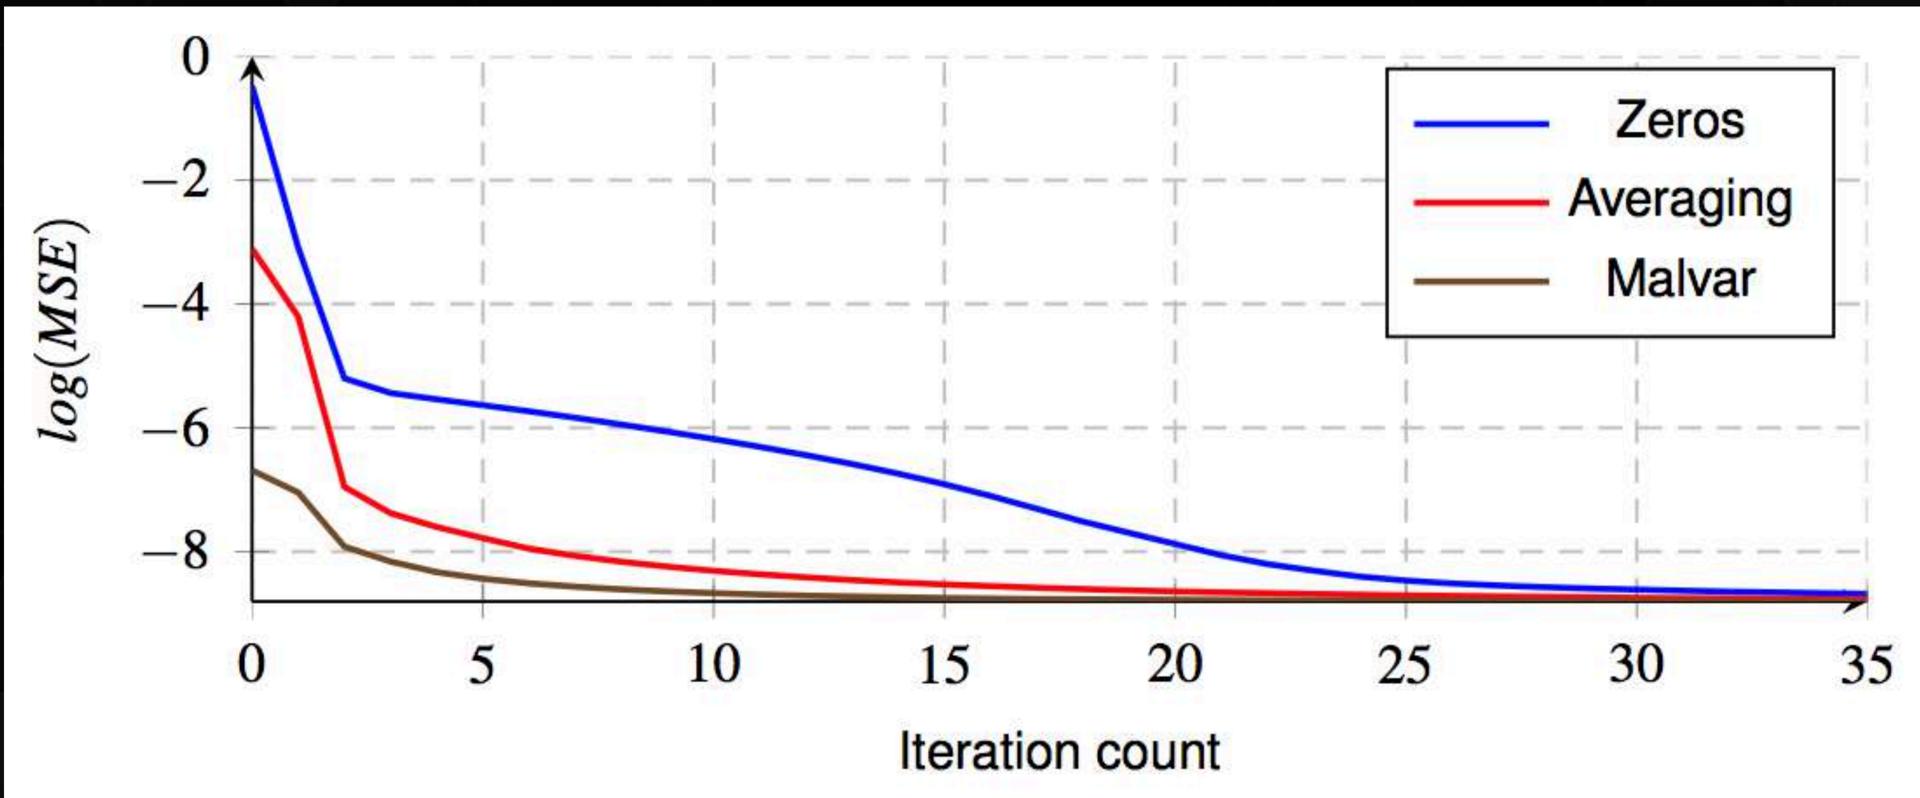

Results

CAMERA SYSTEMS

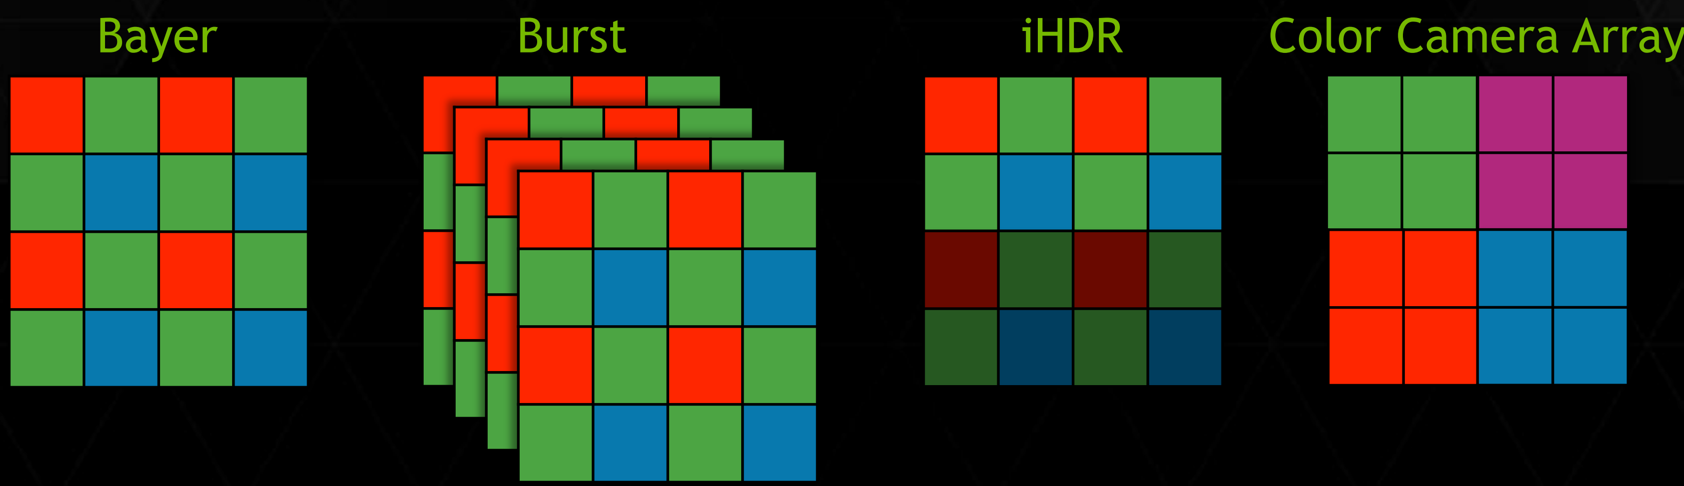

CAMERA SYSTEMS

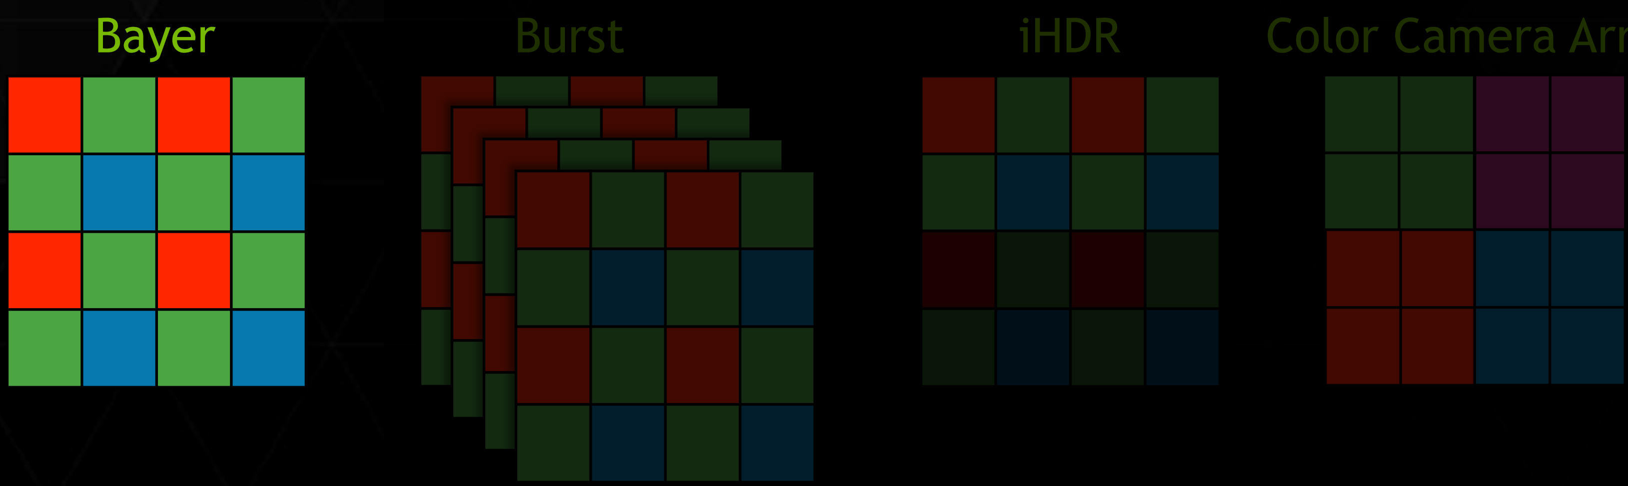

CAMERA SYSTEMS

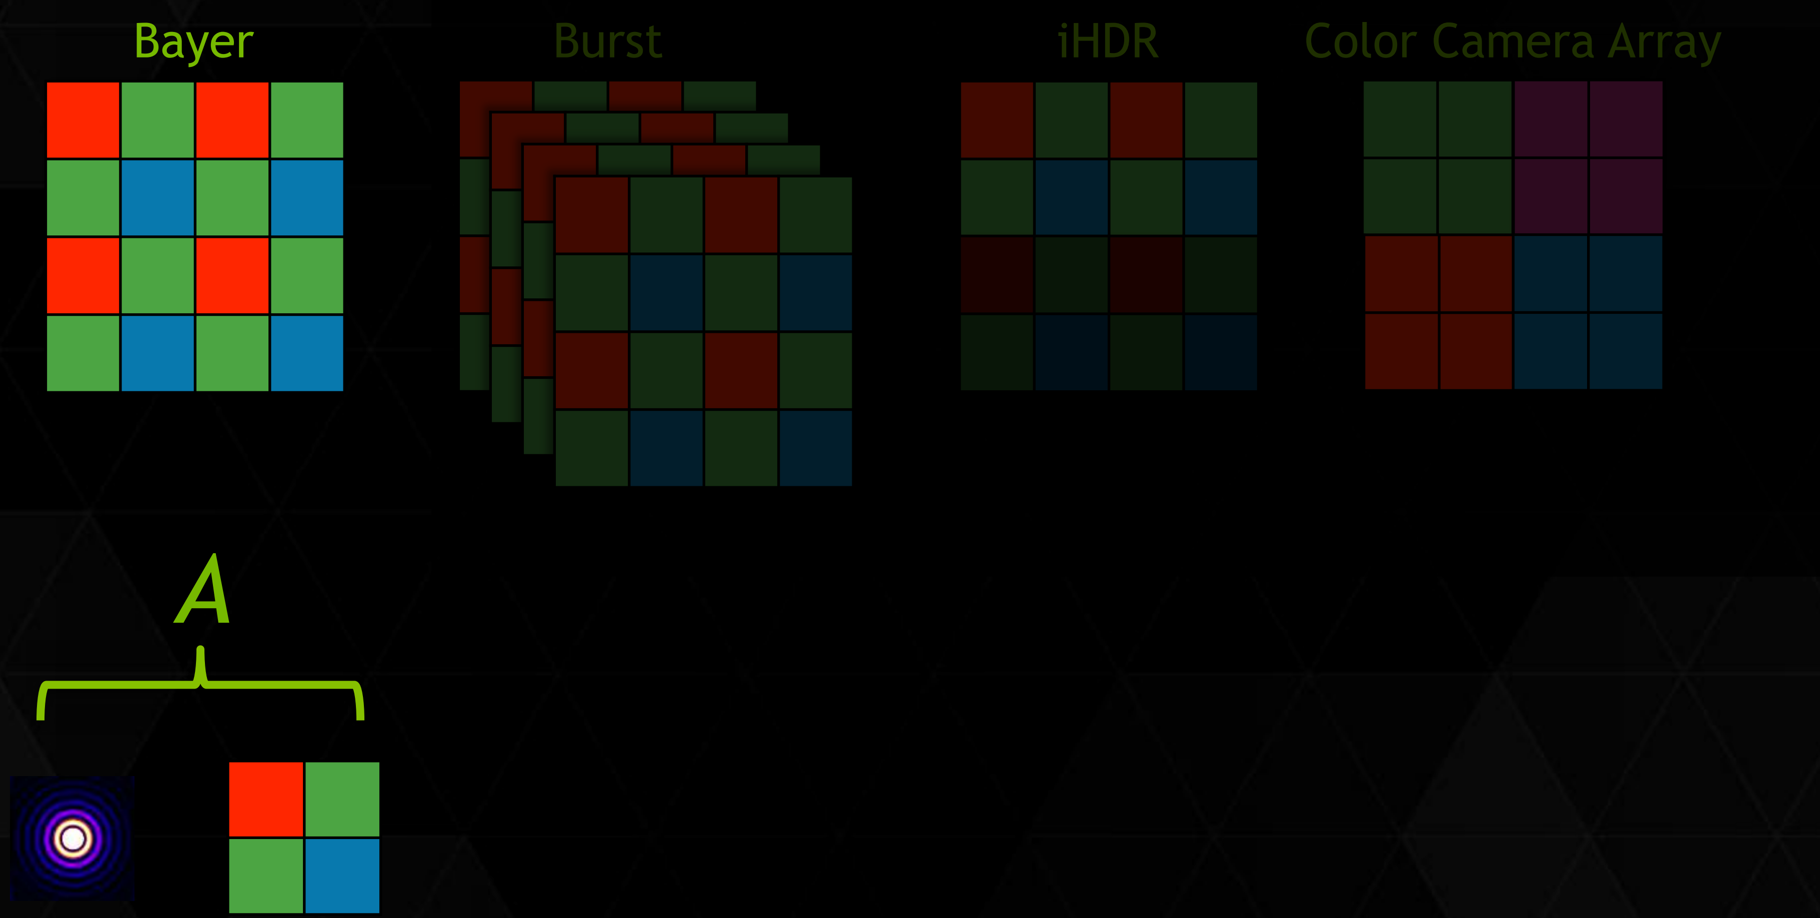

obe (31.64 dB)


obe (31.64 dB)
LDI-NAT (34.79 dB)

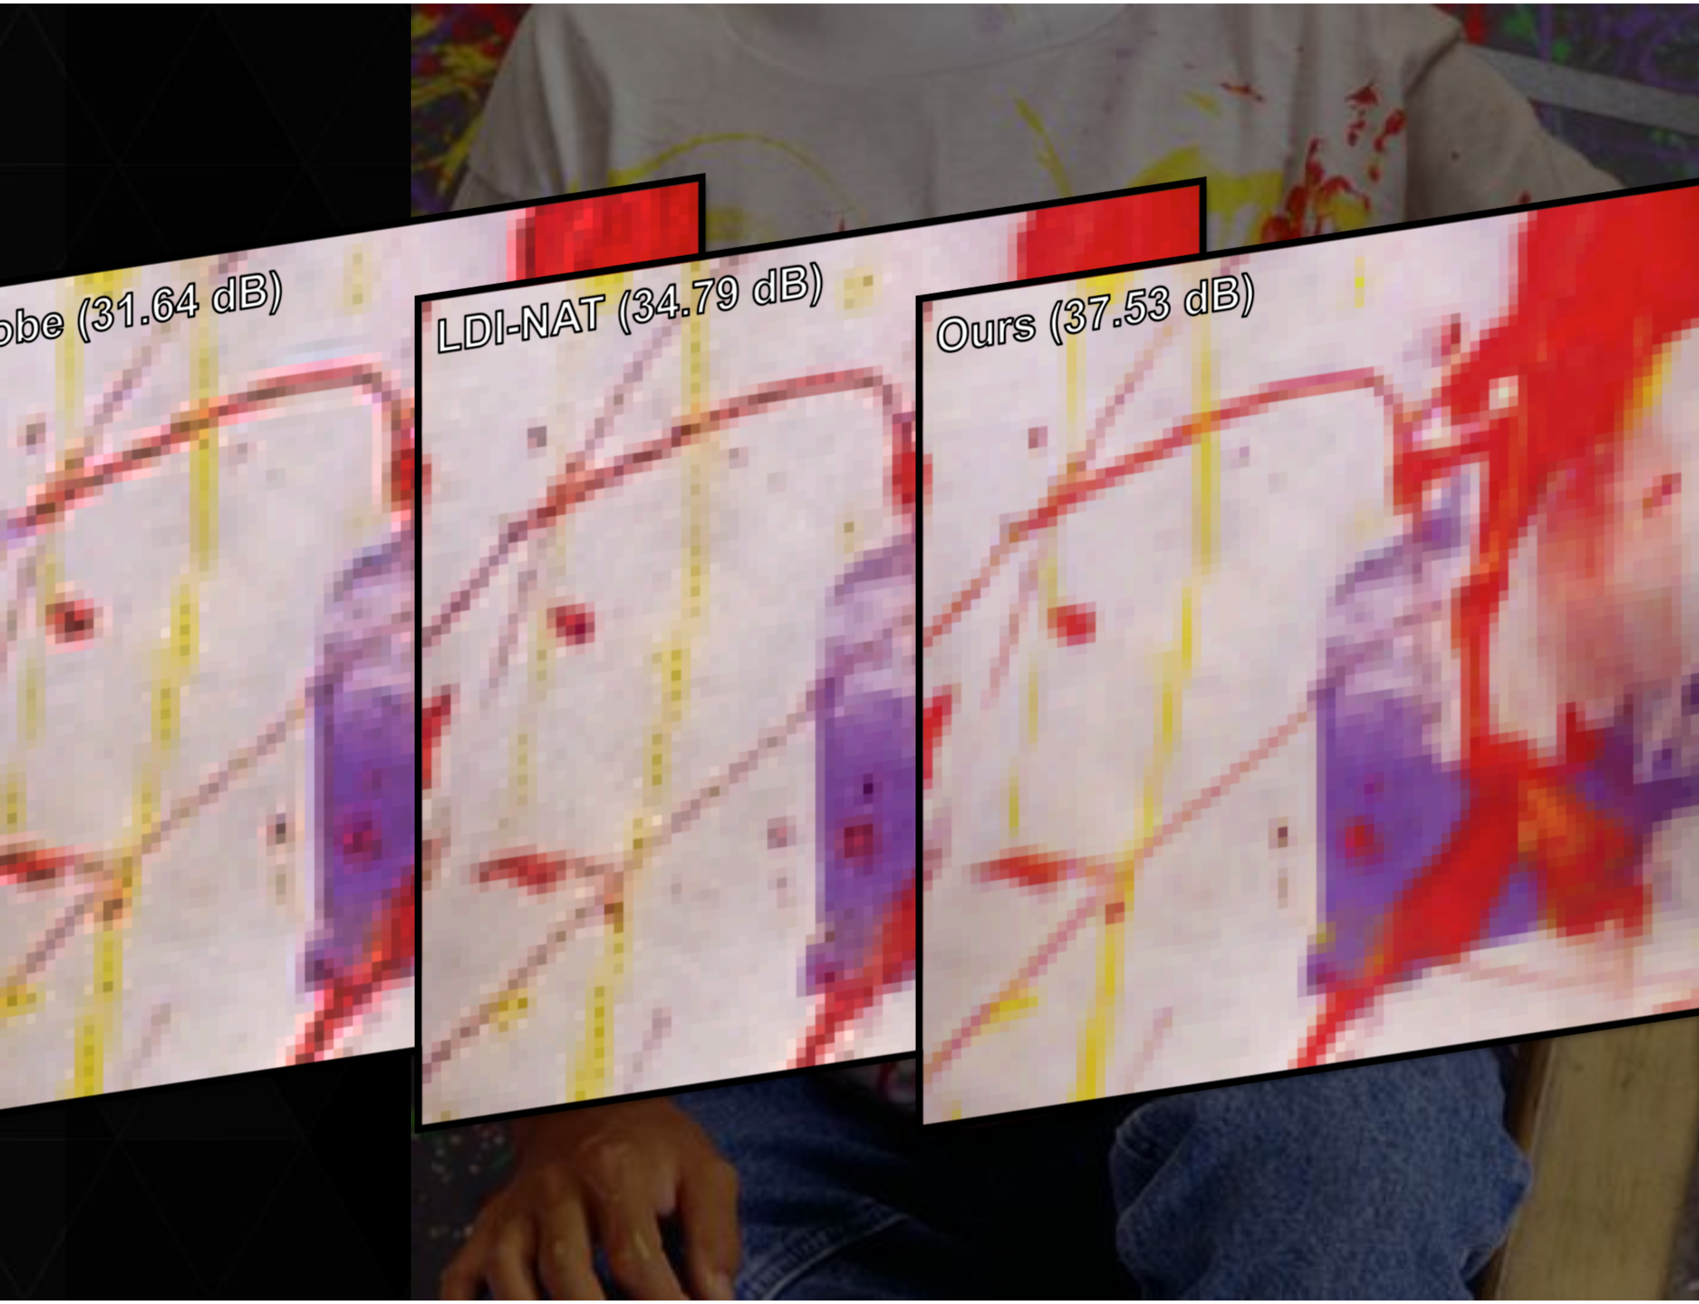
obe (31.64 dB)
LDI-NAT (34.79 dB)
Ours (37.53 dB)

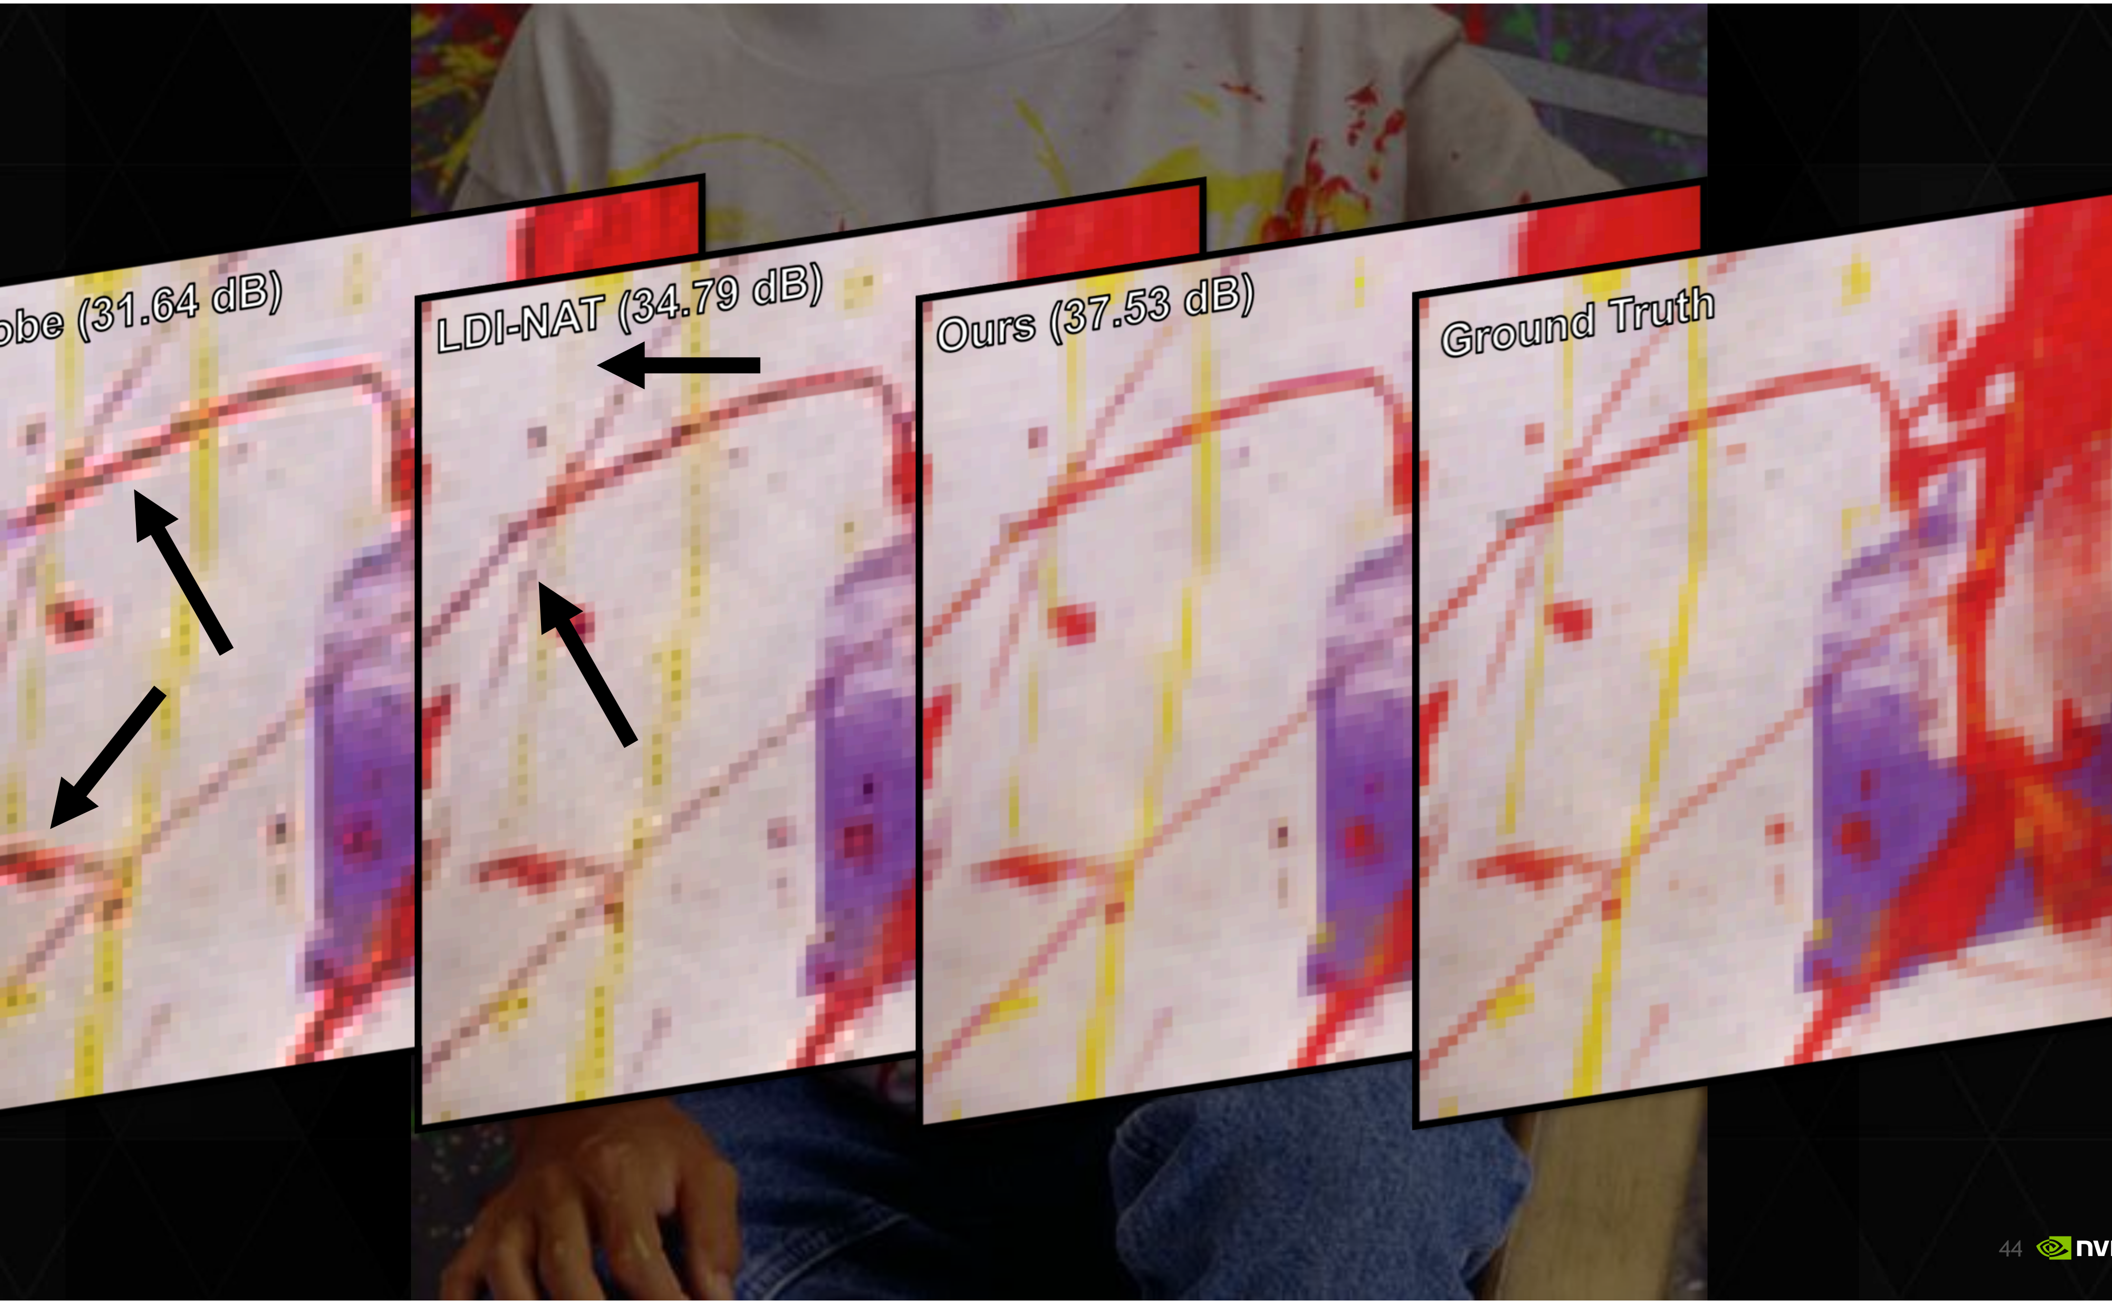

obe (31.64 dB)
LDI-NAT (34.79 dB)
Ours (37.53 dB)
Ground Truth

CAMERA SYSTEMS

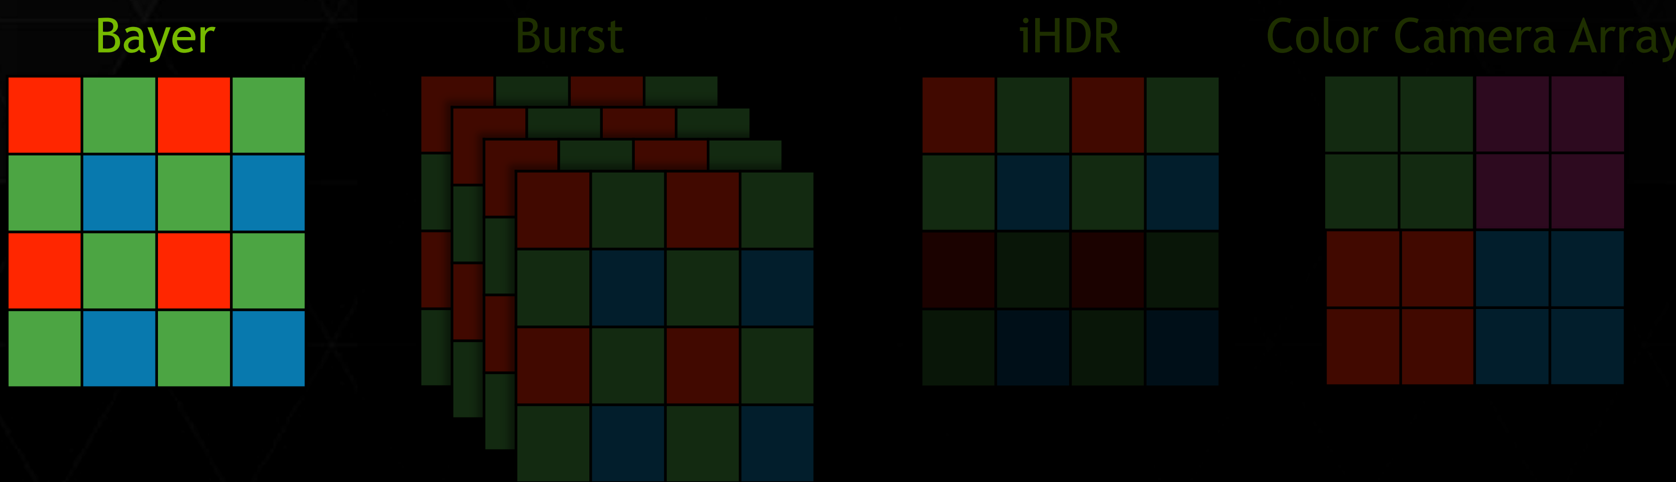

CAMERA SYSTEMS

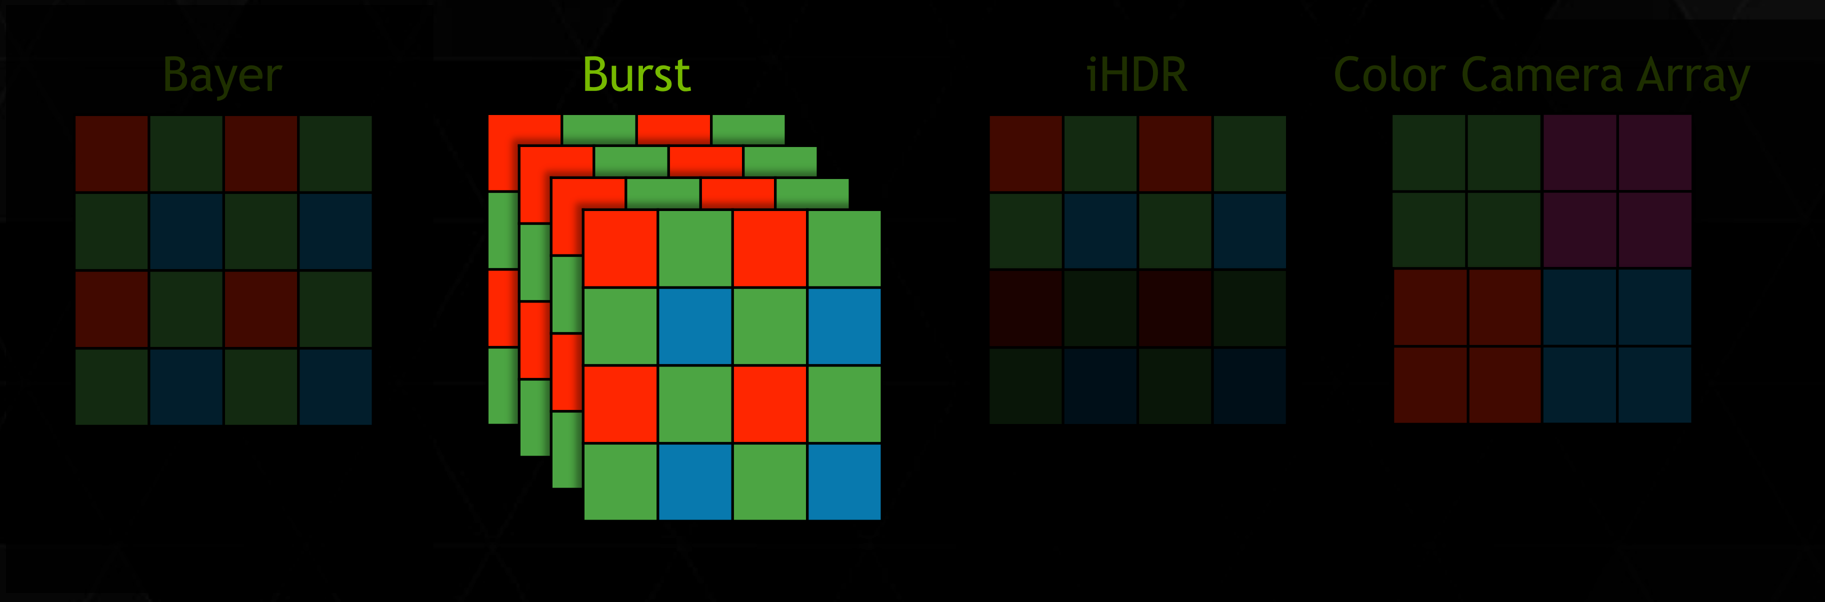

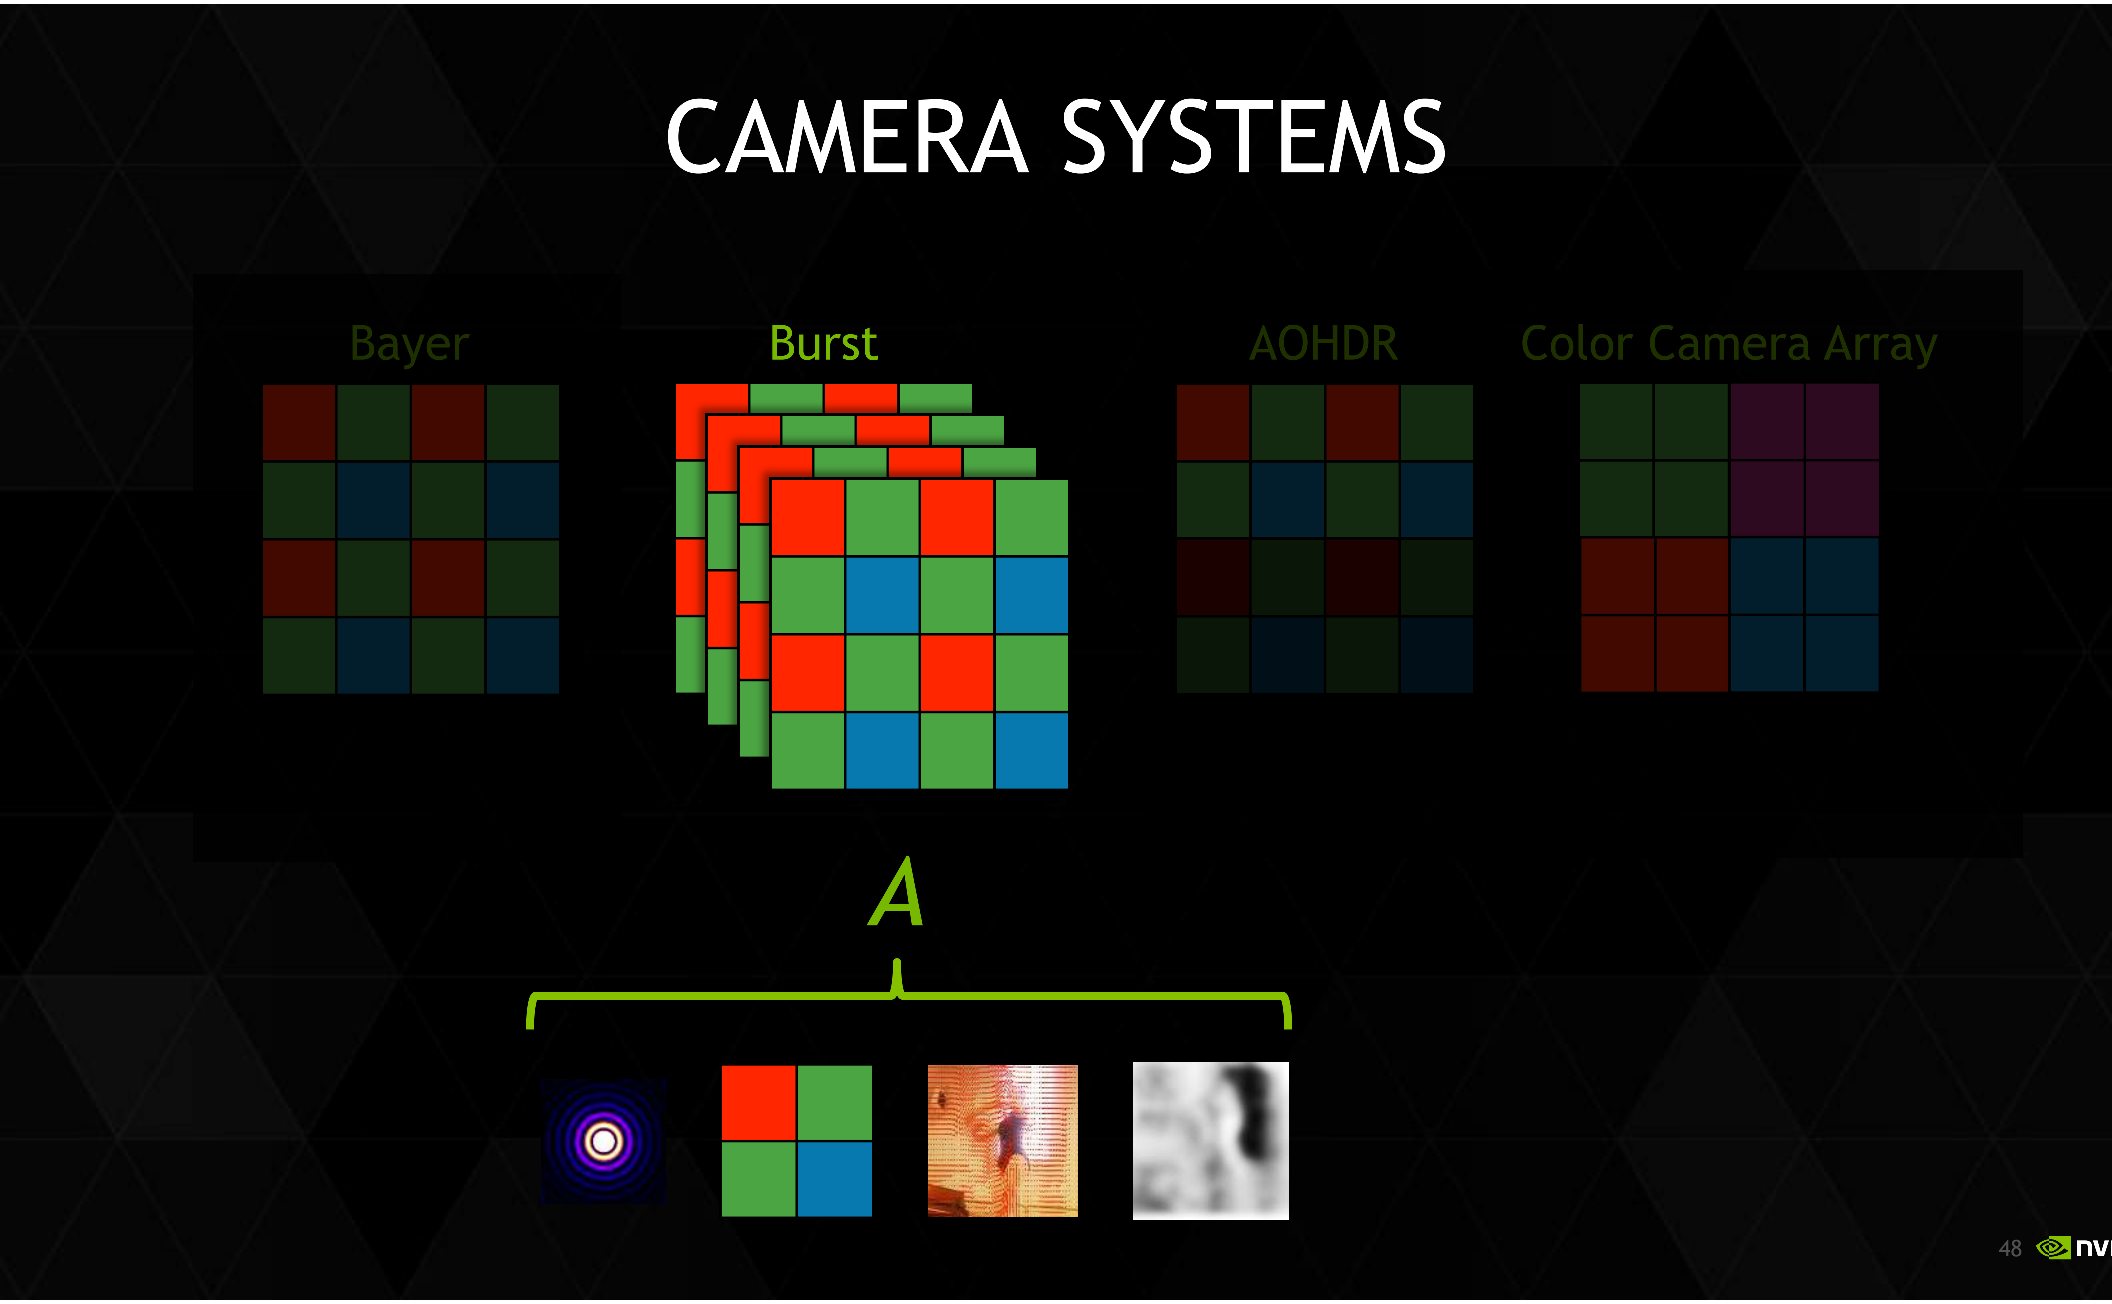
CAMERA SYSTEMS
Bayer
Burst
AOHDR
Color Camera Array
A
48


Canon EOS 650D / ISO 12800
INTERIOR
SEMI-GLOSS ENAMEL
BEHR
PREMIUM PLUS
ULTRA PURE WHITE
ZINSSER
SATIN
QUICK-15
FAST-DRY
CLEAR FINISH AND SEALER
INTERIOR
FLAT
BEHR
PREMIUM PLUS
CEILING PAINT
PRIMER
SEALER
Stain Kill


Canon EOS 650D / ISO 12800
BEHR
PREMIUM PLUS
ULTRA PURE WHITE
ZINSSER
SATIN
QUICK-15
BEHR
PREMIUM PLUS
CEILING PAINT

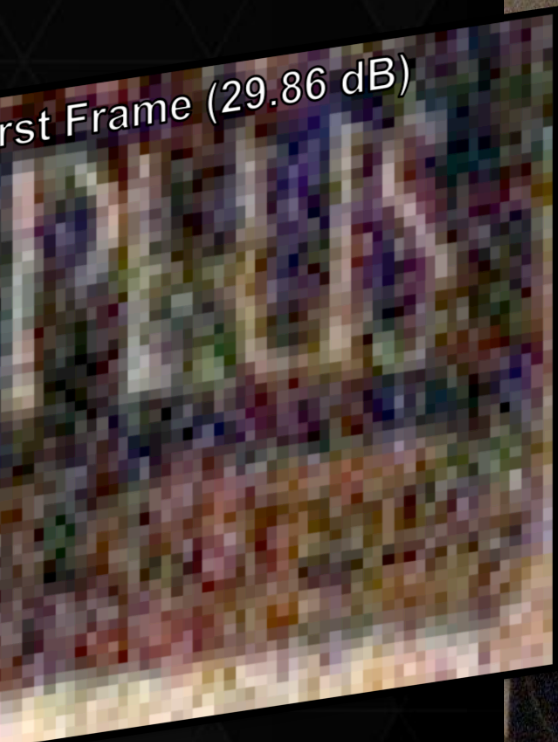
rst Frame (29.86 dB)

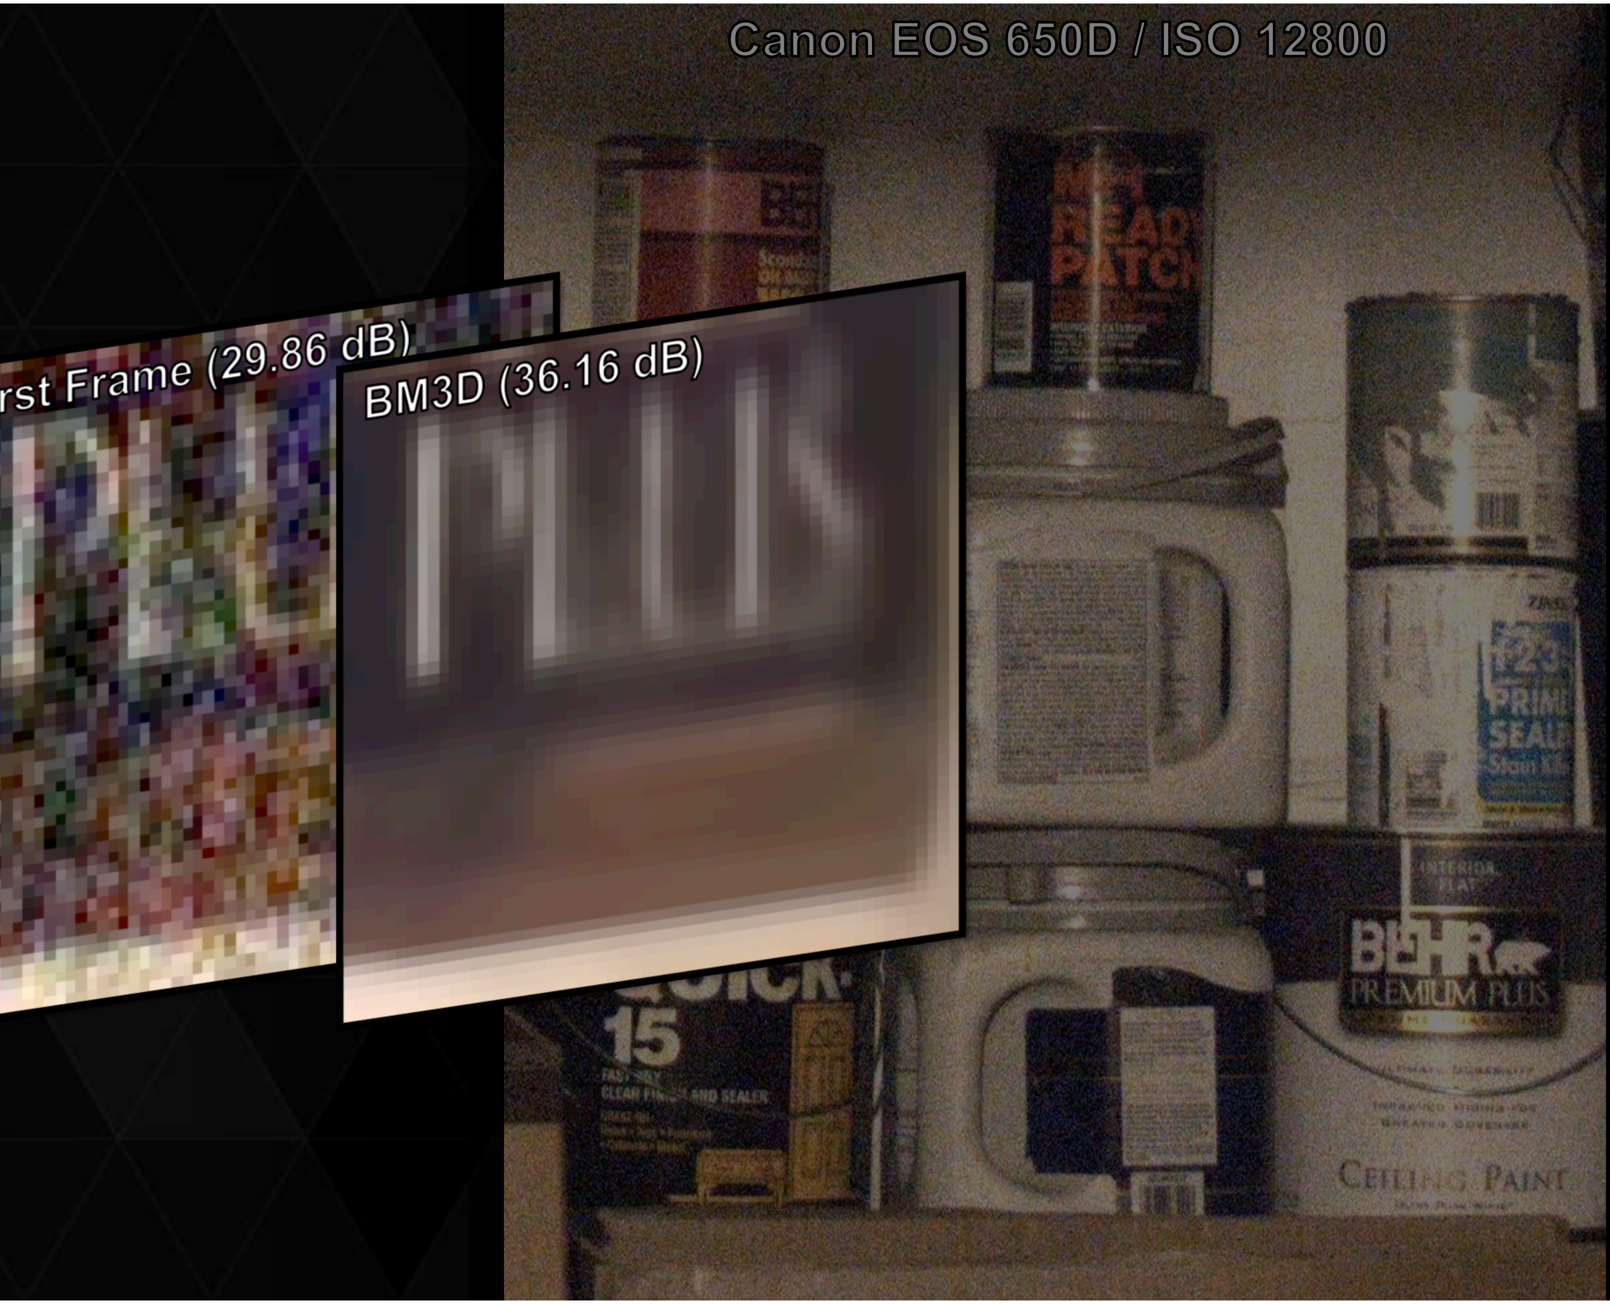
Canon EOS 650D / ISO 12800
rst Frame (29.86 dB)
BM3D (36.16 dB)


Canon EOS 650D / ISO 12800
rst Frame (29.86 dB)
BM3D (36.16 dB)
VBM3D (37.81 dB)
PLUS
BEHR
PREMIUM PLUS
CEILING PAINT


Canon EOS 650D / ISO 12800
rst Frame (29.86 dB)
BM3D (36.16 dB)
VBM3D (37.81 dB)
Ours (39.67 dB)


Canon EOS 650D / ISO 12800
rst Frame (29.86 dB)
BM3D (36.16 dB)
VBM3D (37.81 dB)
Ours (39.67 dB)
Ground Truth

CAMERA SYSTEMS

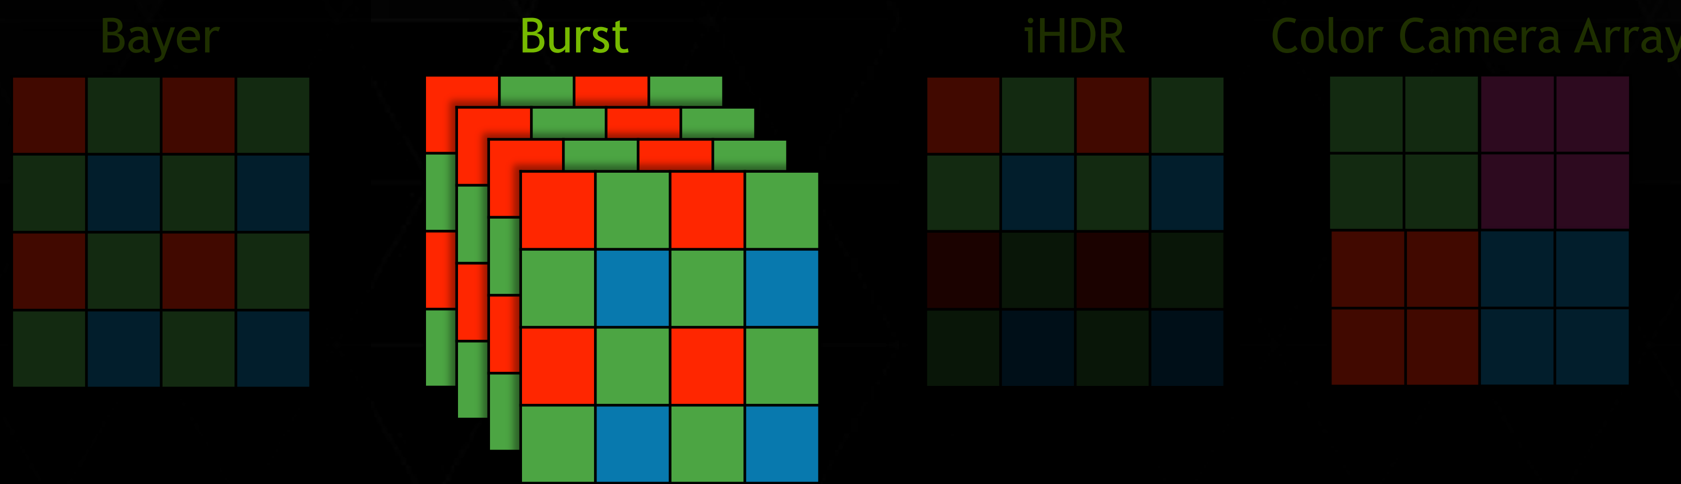

CAMERA SYSTEMS

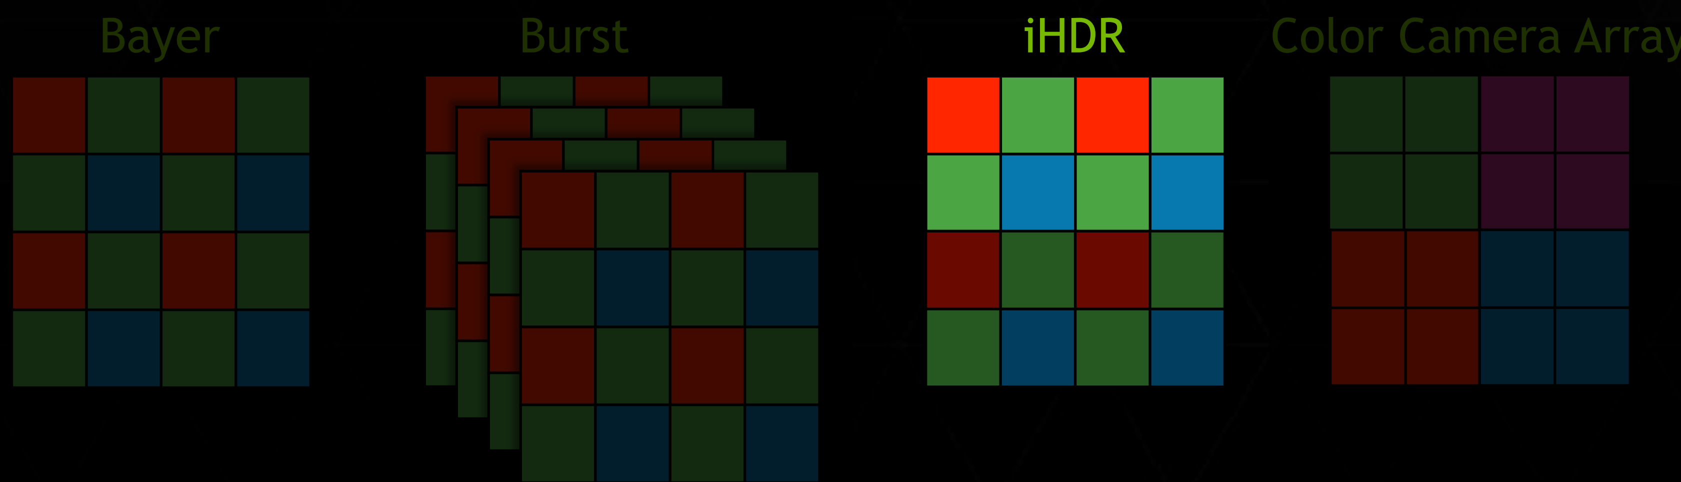

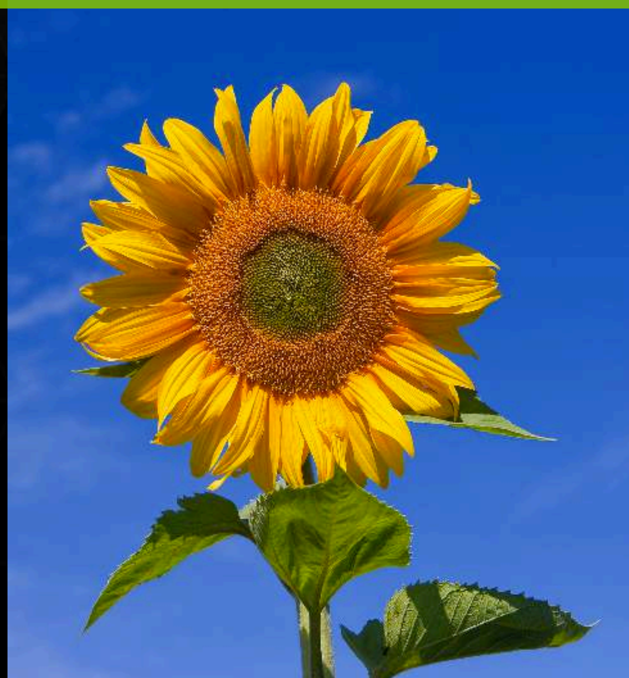

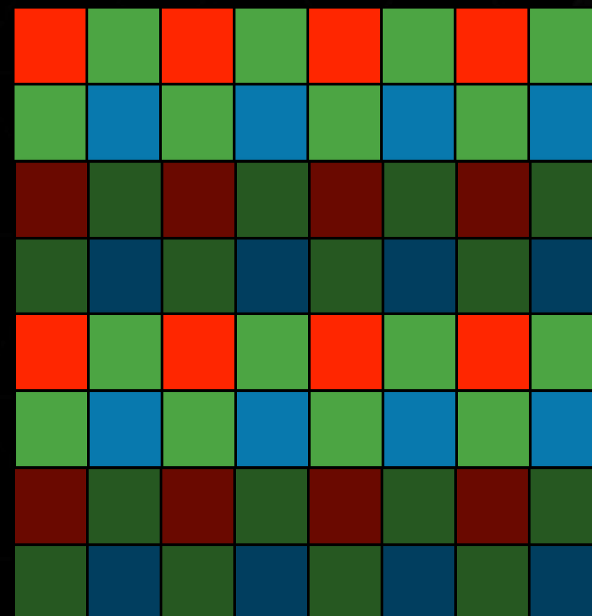

Long Exposur

Short Exposur

Long Exposur

Short Exposur

Short Exposure

Long Exposure

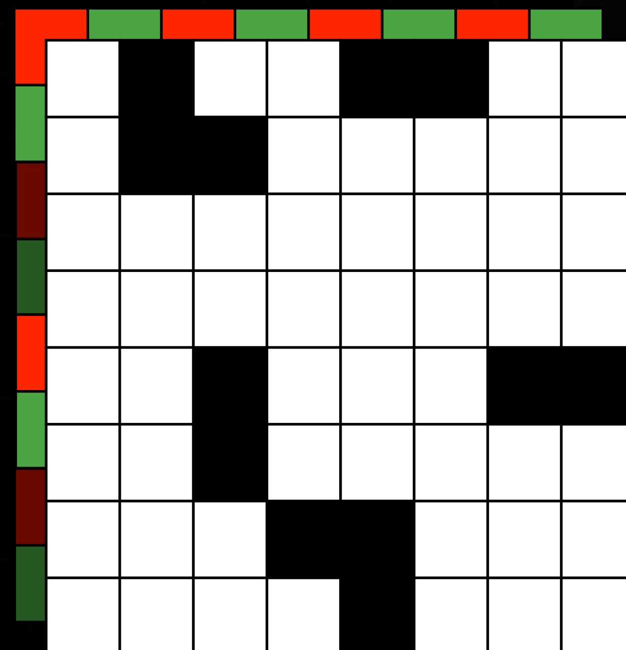
Long Exposur
Short Exposur
Binary Mask
Long Exposur
Short Exposur

CAMERA SYSTEMS

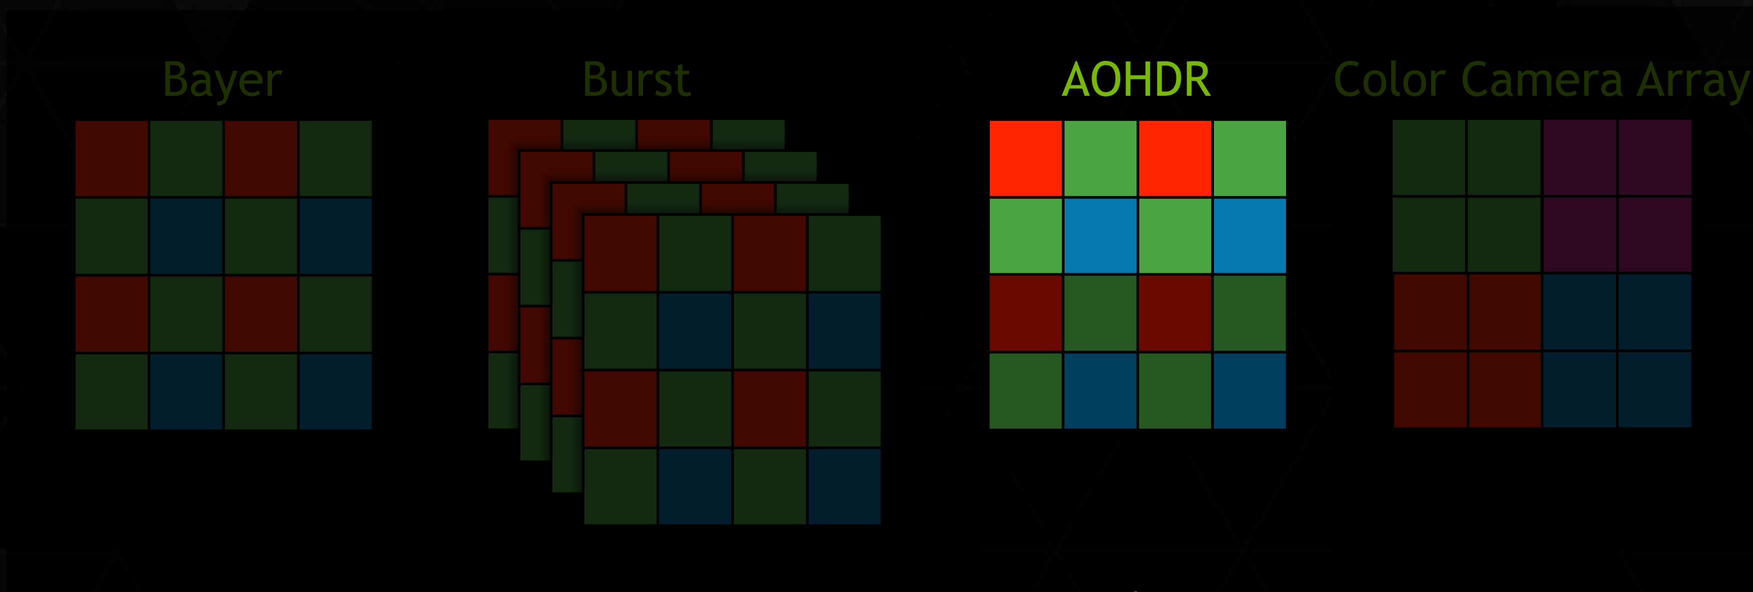

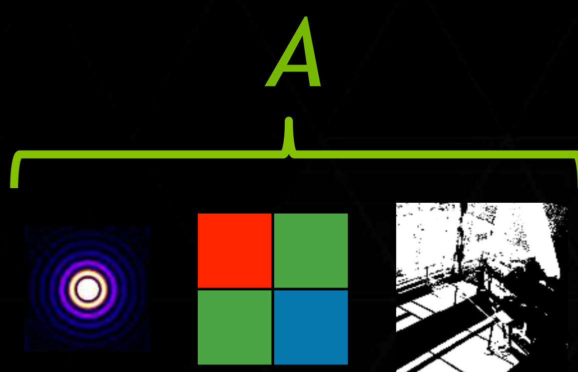

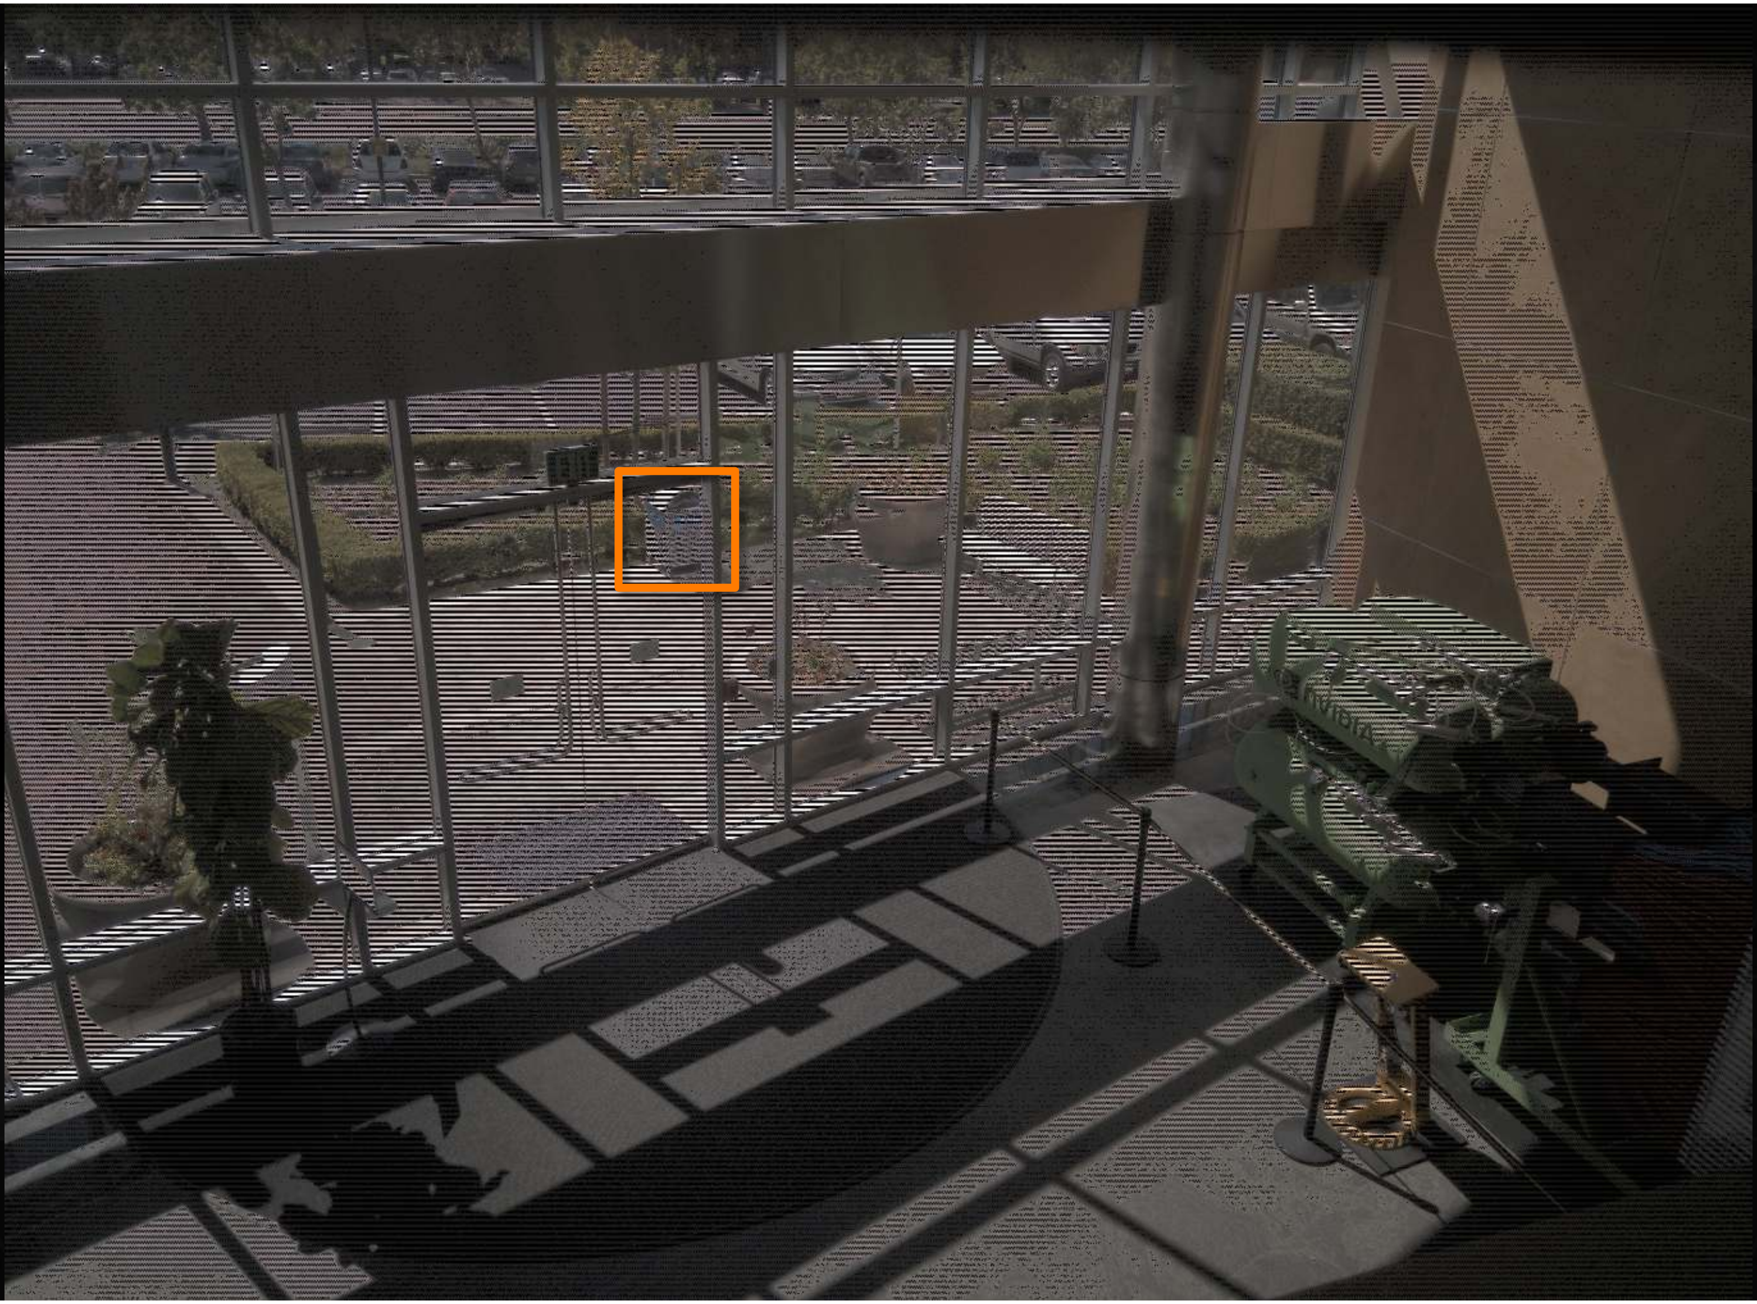


parse Observation
Magic Lantern
SHIELD Tablet
Ours

CAMERA SYSTEMS

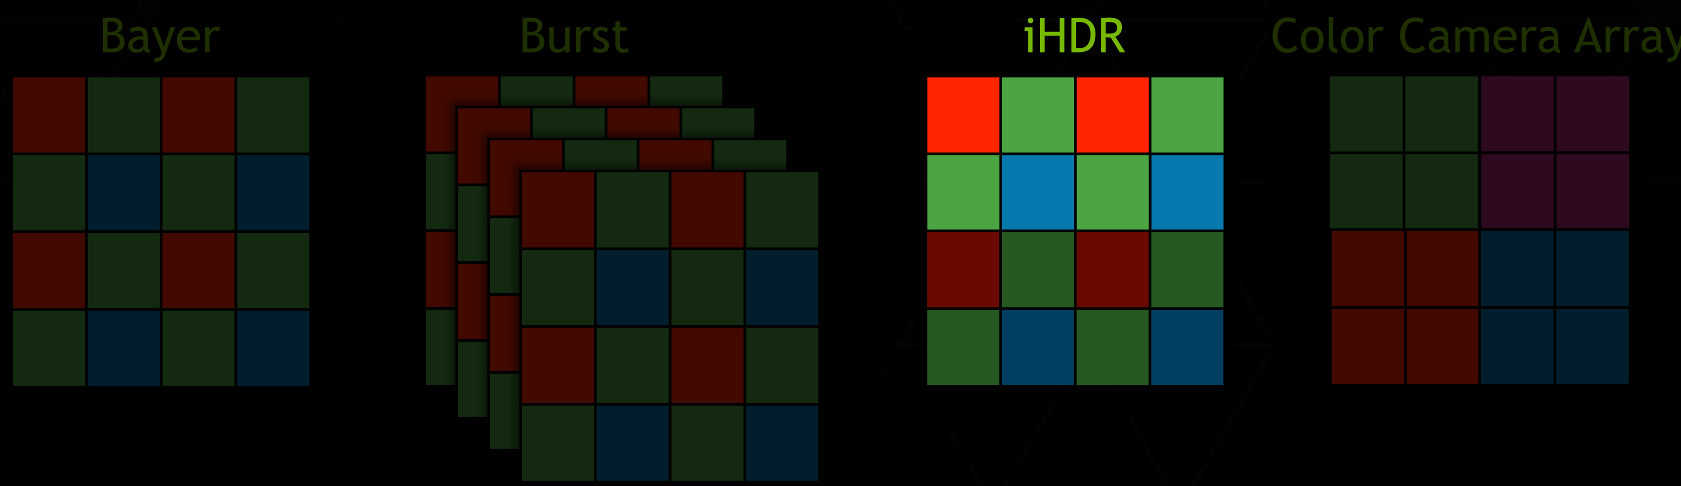

CAMERA SYSTEMS

Bayer

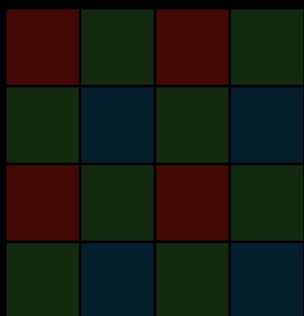

Burst

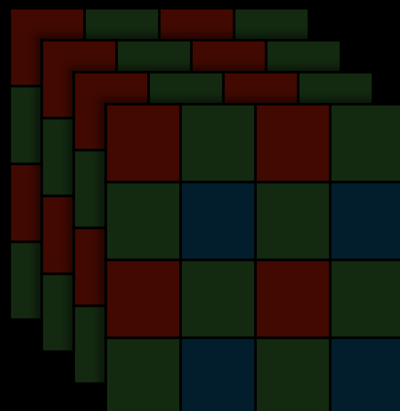

iHDR

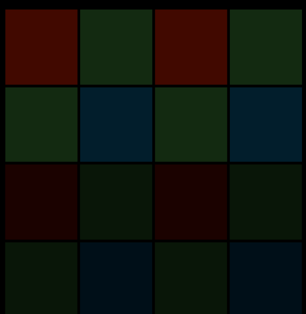

Color Camera Array

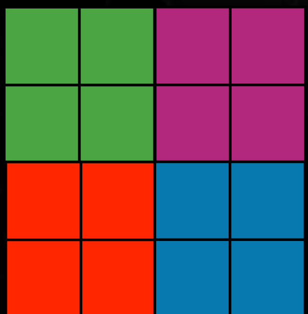

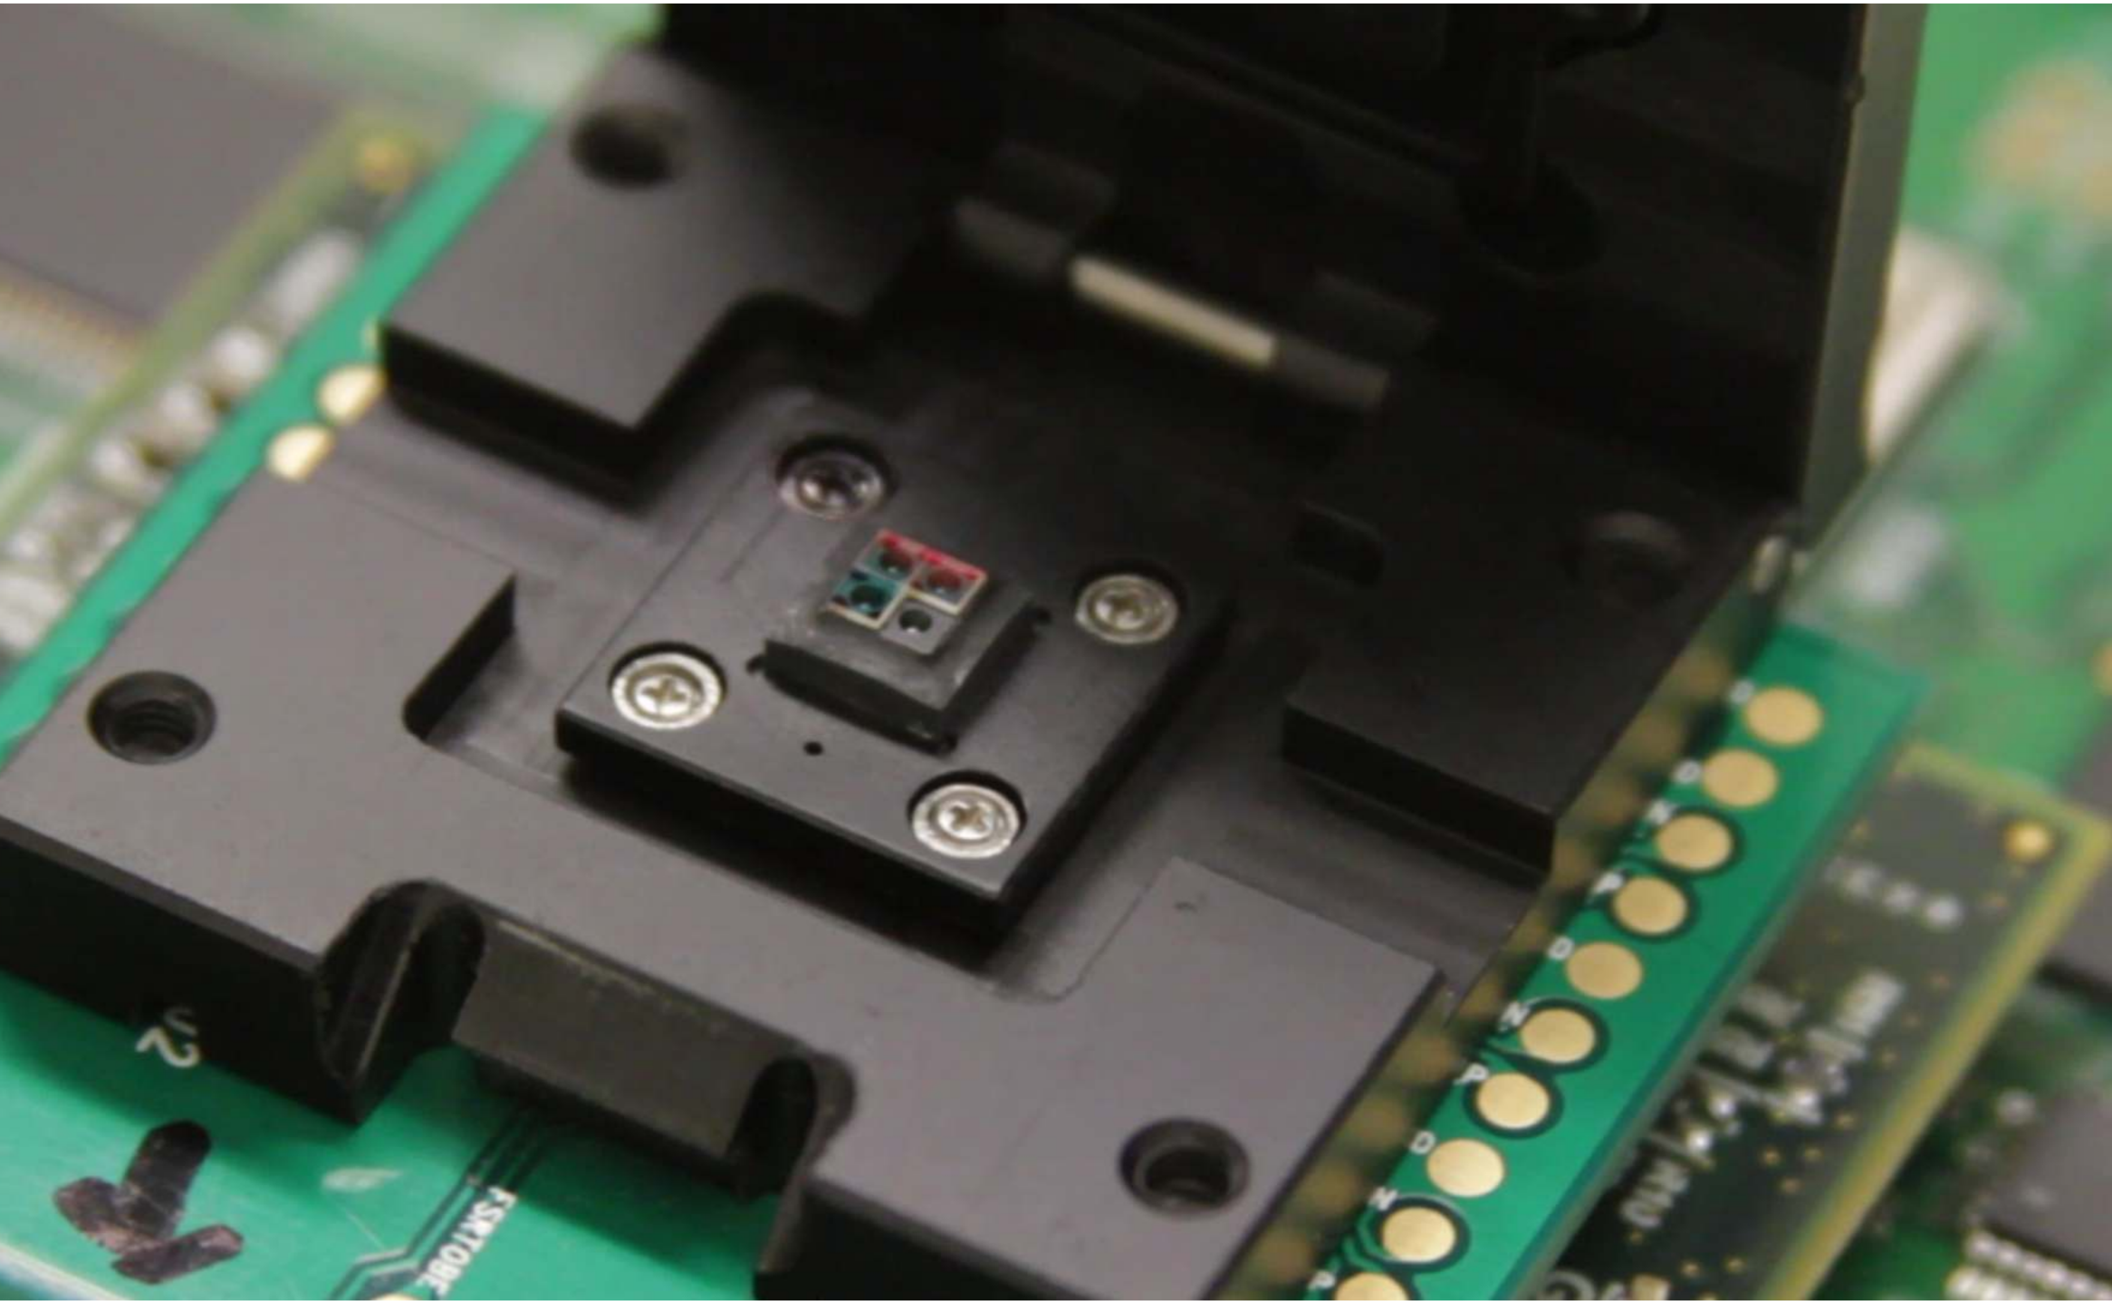


TAZO
ARANCIATA

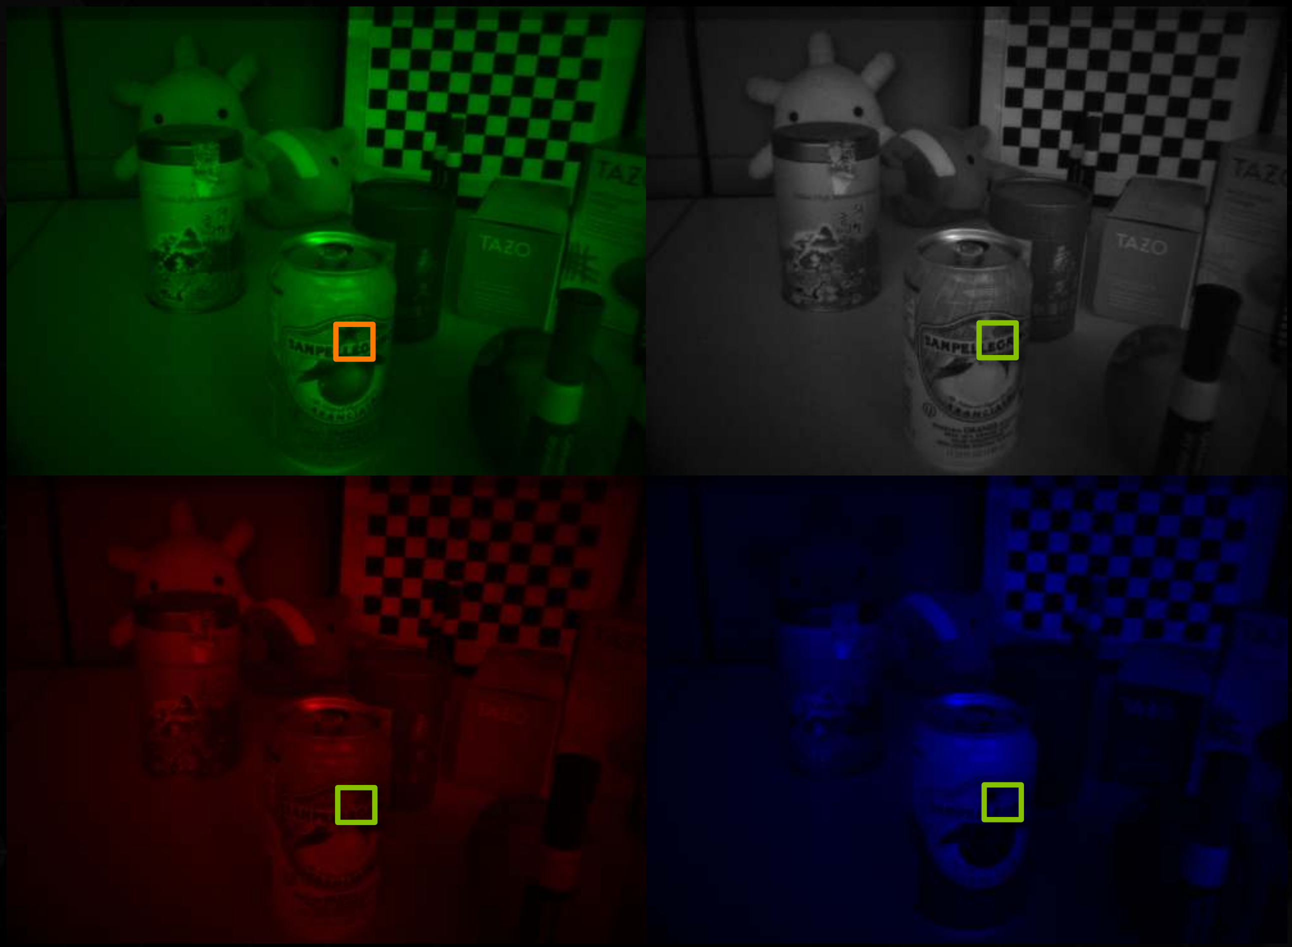
TAZO

CAMERA SYSTEMS

Bayer

Burst

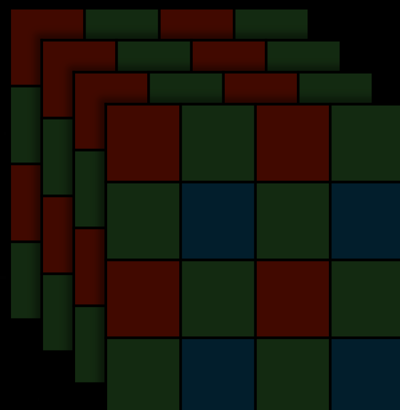

iHDR

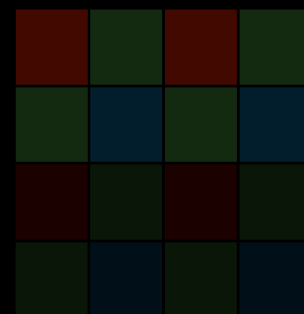

Color Camera Array

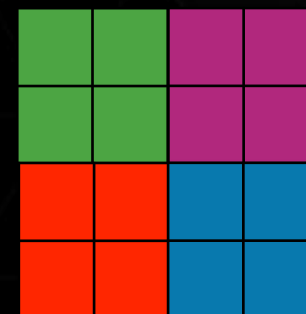

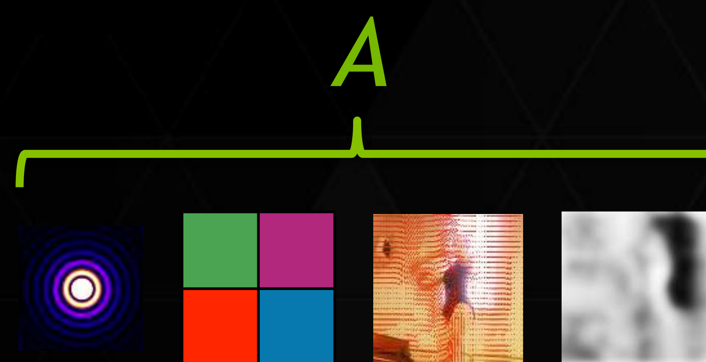


EXIT

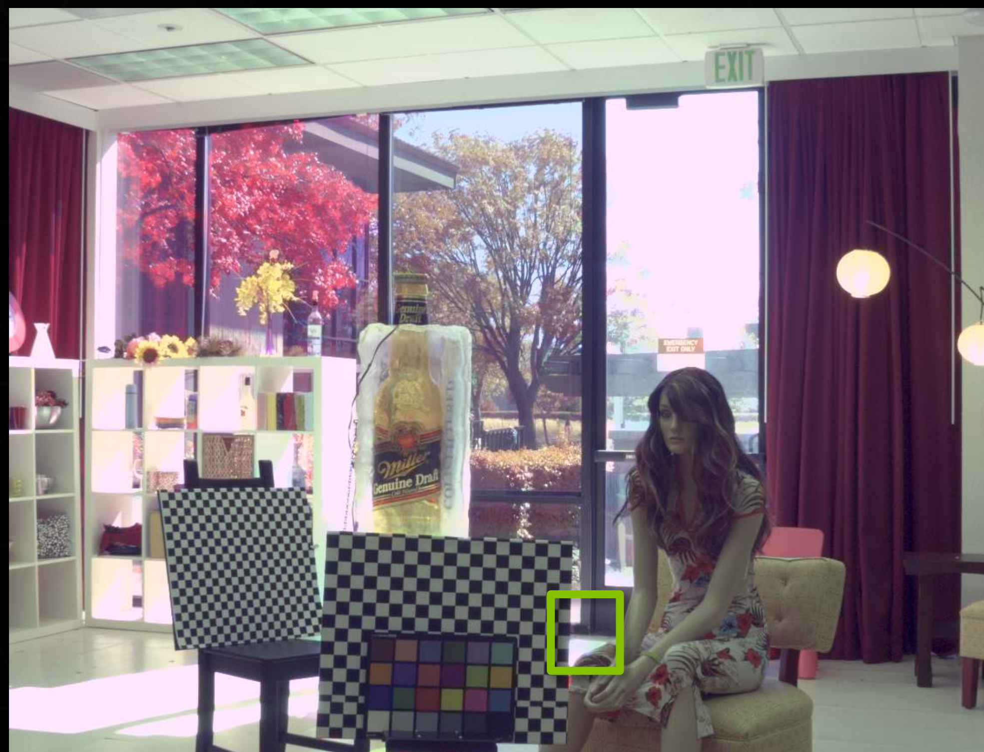
EXIT
Miller
Genuine Draft

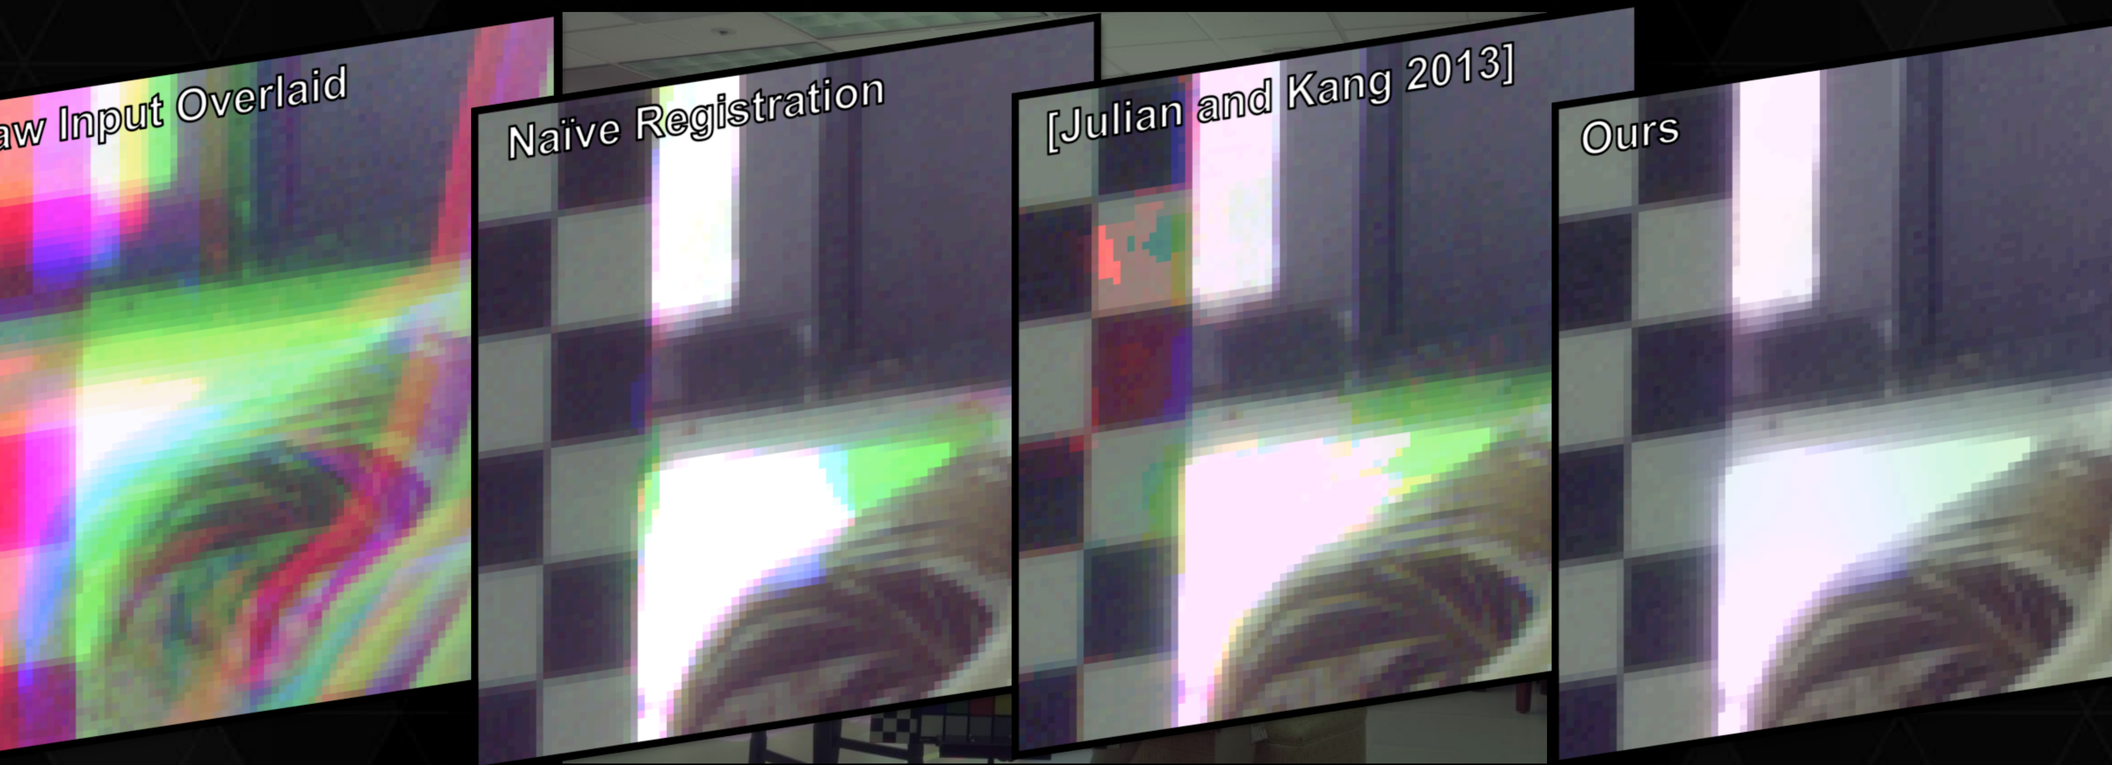
aw Input Overlaid
Naïve Registration
[Julian and Kang 2013]
Ours

PERFORMANCE

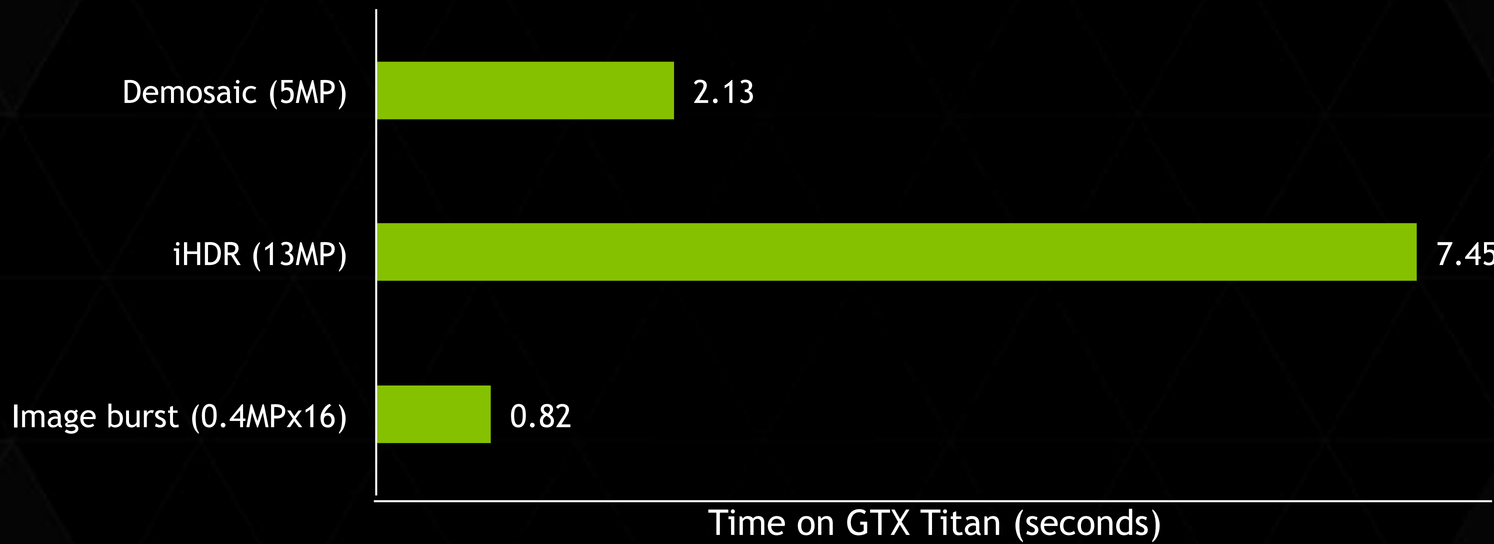


Noisy Capture
Reconstruction
Live Low-light Burst Capture

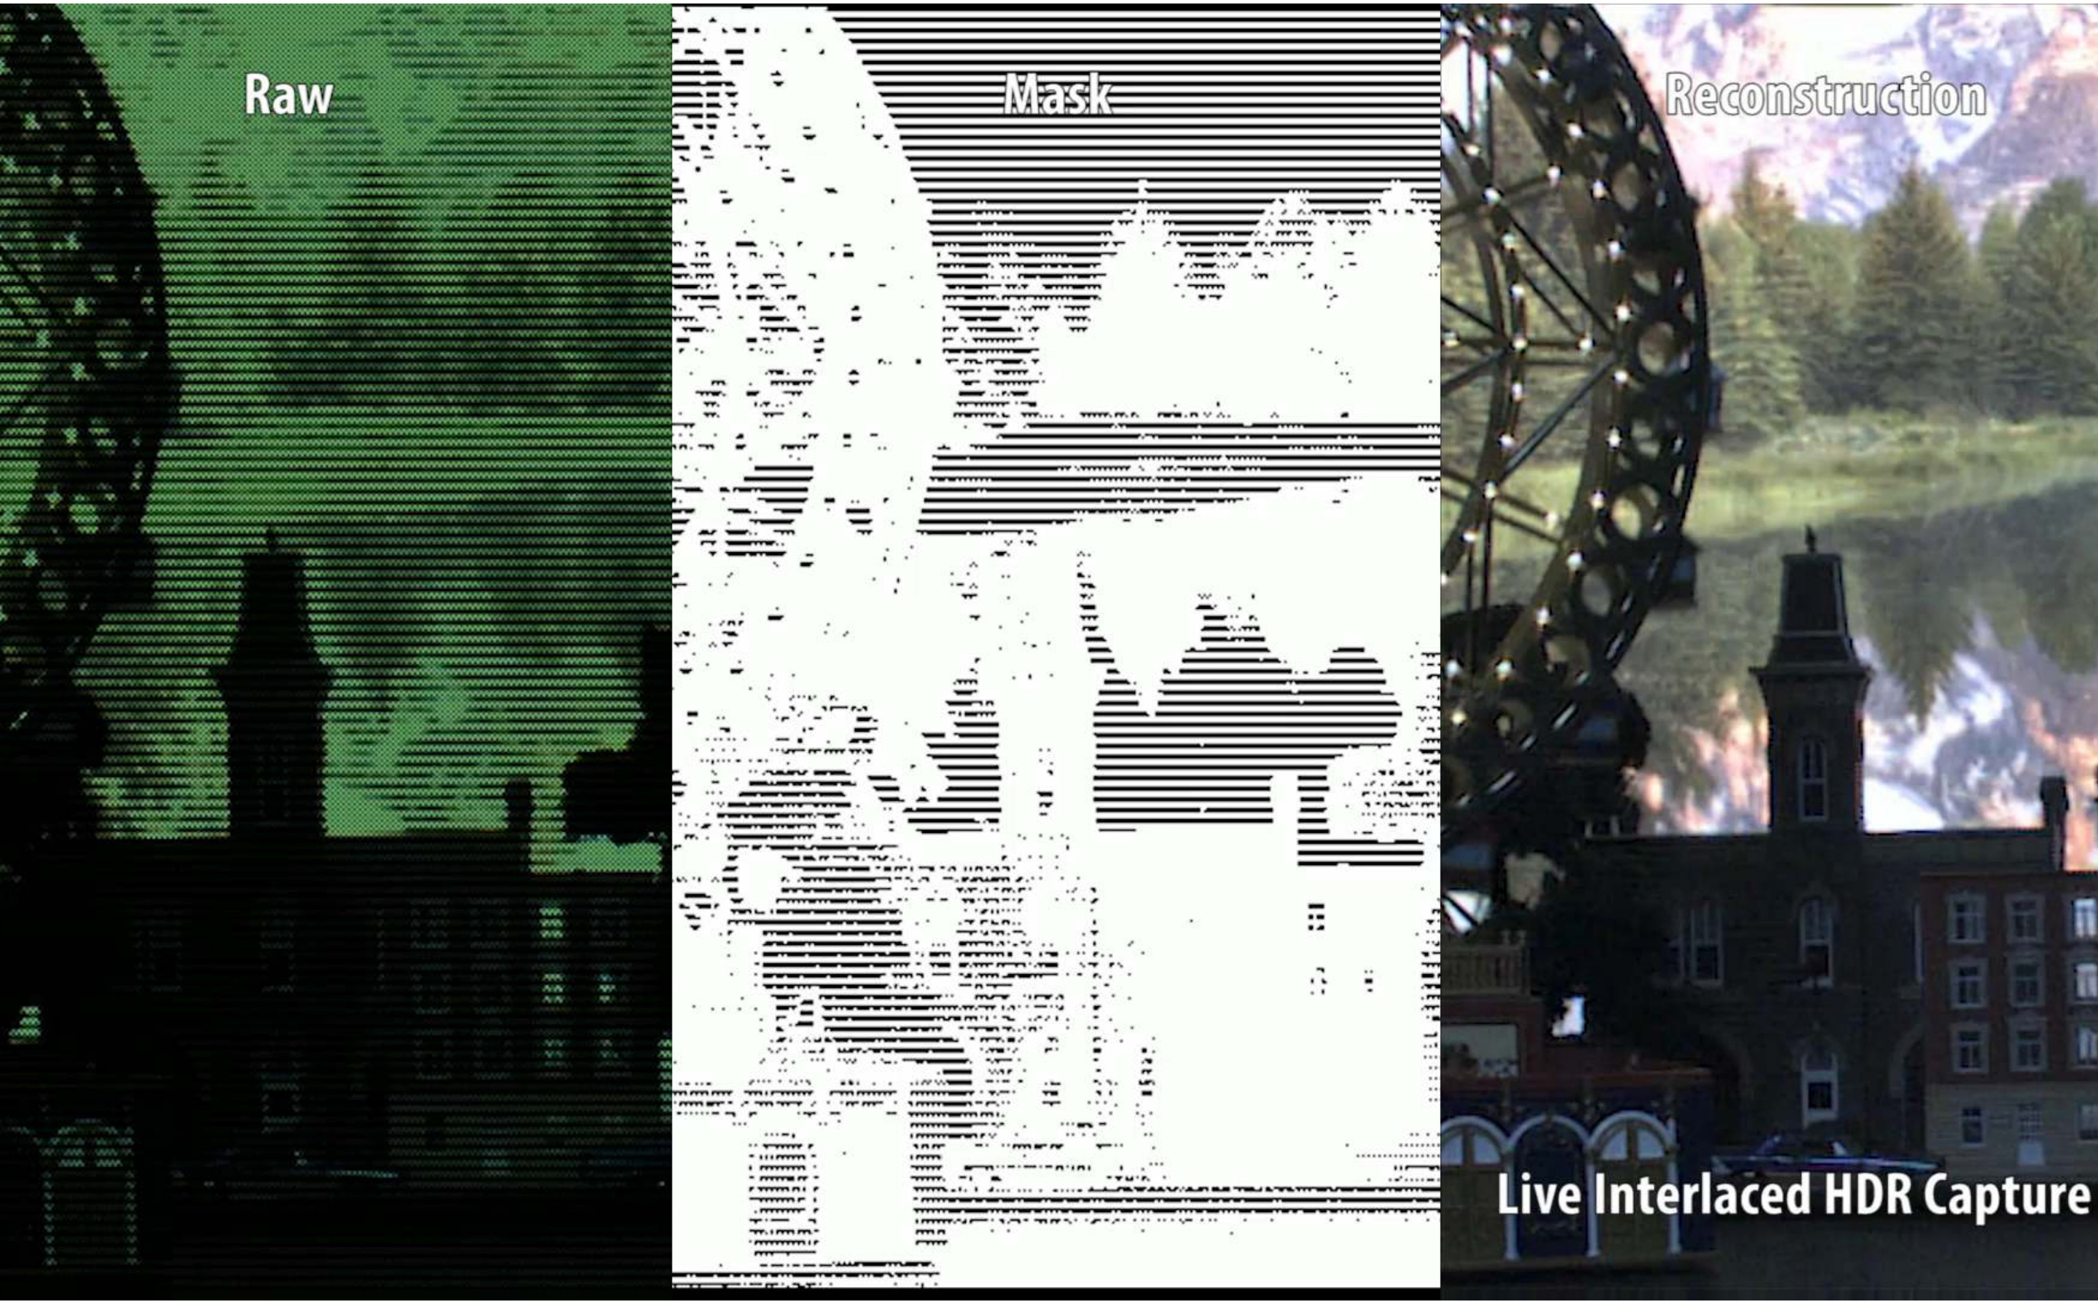
Raw
Mask
Reconstruction
Live Interlaced HDR Capture

PERFORMANCE

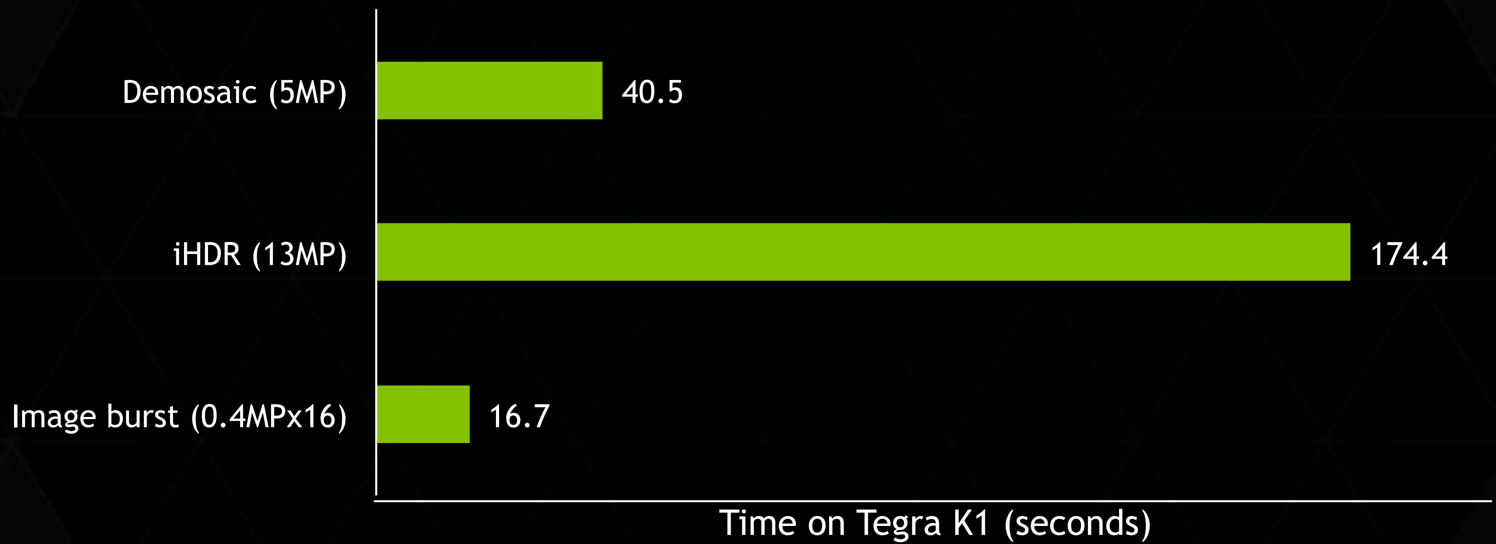

Conclusion

Bayer Burst iHDR Color Camera Array

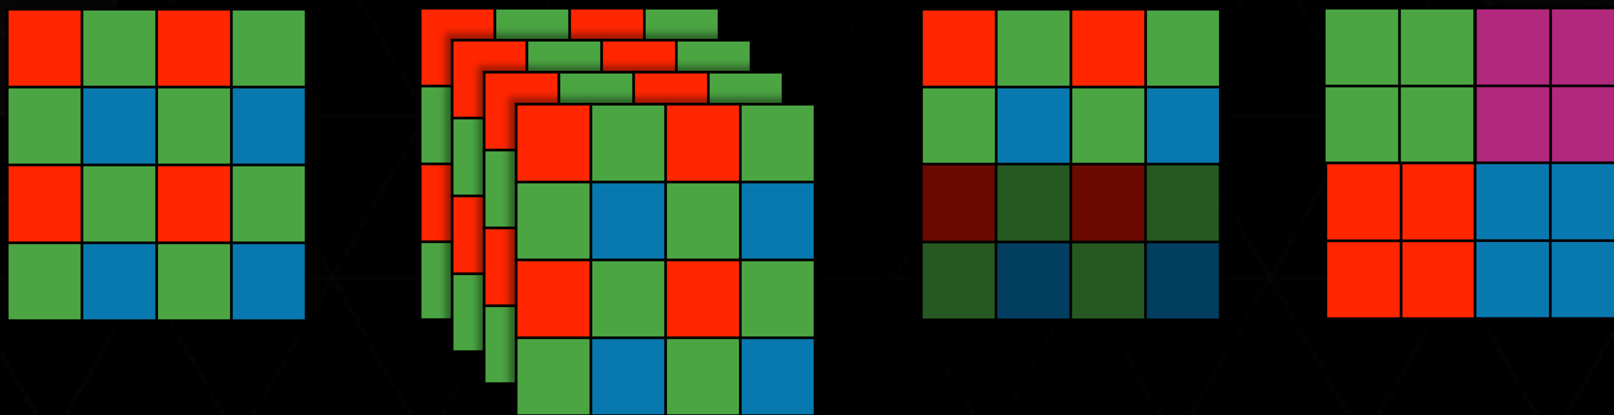

+ others such as:

deconvolution, super-resolution, JPEG compression, and de-blocking

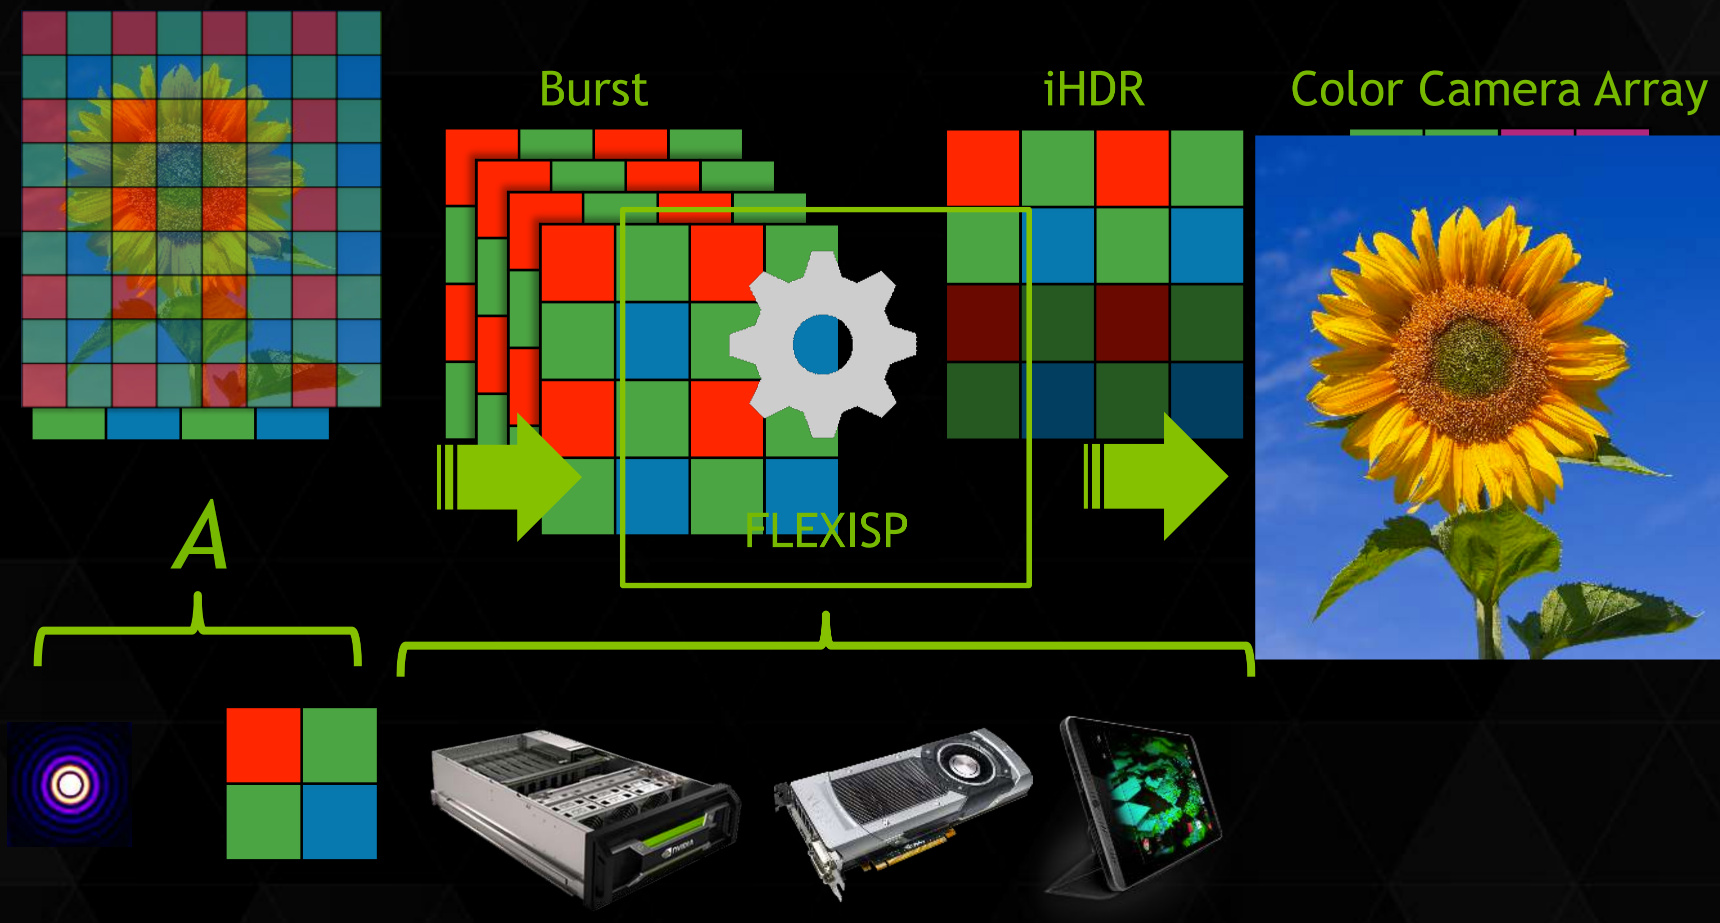
Burst
iHDR
Color Camera Array
A
FLEXISP