Hyperprior bayesian approach for inverse problems in imaging.

Application to single shot HDR.

Julie Delon
with Cecilia Aguerrebere, Andrés Almansa, Yann Gousseau and Pablo Musé

UNIVERSITÉ PARIS DESCARTES
LABORATOIRE MAP5

Teaser : what is High Dynamic Range Imaging (HDR) ?

Capture a scene containing a large range of intensity levels...

Teaser : what is High Dynamic Range Imaging (HDR) ?

... using a standard digital camera.

Limited dynamic range of the camera → loss of details in bright and/or dark areas.

Teaser : what is High Dynamic Range Imaging (HDR) ?

... using a standard digital camera.

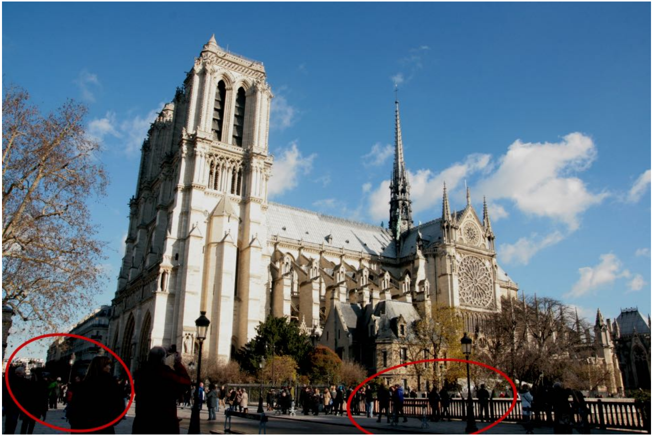

Limited dynamic range of the camera → loss of details in bright and/or dark areas.

Teaser : what is High Dynamic Range Imaging (HDR) ?

... using a standard digital camera.

Limited dynamic range of the camera → loss of details in bright and/or dark areas.

Teaser : what is High Dynamic Range Imaging (HDR) ?

... using a standard digital camera.

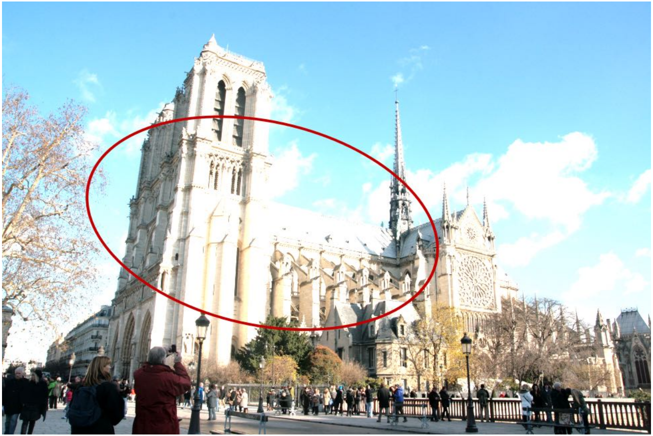

Limited dynamic range of the camera → loss of details in bright and/or dark areas.

Teaser : what is High Dynamic Range Imaging (HDR) ?

... using a standard digital camera.

Limited dynamic range of the camera $\rightarrow$ loss of details in bright and/or dark areas.

Teaser : what is High Dynamic Range Imaging (HDR) ?

Limited dynamic range of the camera $\rightarrow$ loss of details in bright and/or dark areas.

Teaser : what is High Dynamic Range Imaging (HDR) ?

Limited dynamic range of the camera $\rightarrow$ loss of details in bright and/or dark areas.

Motivation : High Dynamic Range Imaging (HDR)

Usual approach for HDR image generation : **fusion of mutiple exposures.**

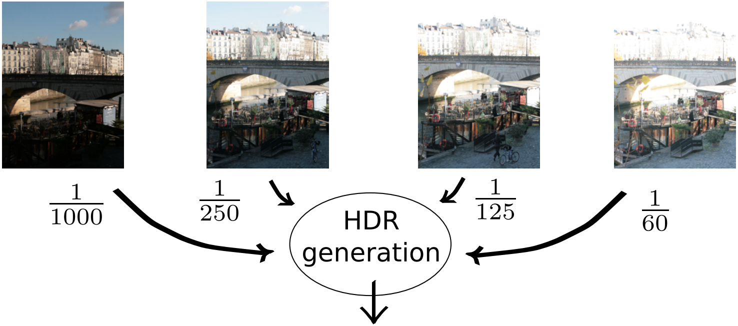

Irradiance Map
(number of photons reaching each pixel per unit time)

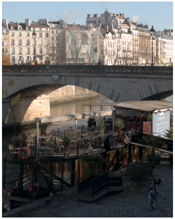

Challenges of HDR imaging in dynamic scenes

moving objects

noise

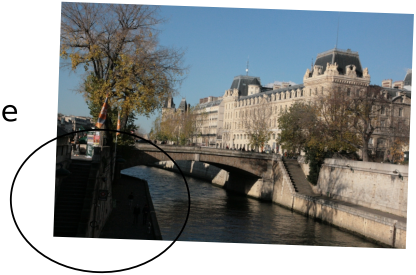

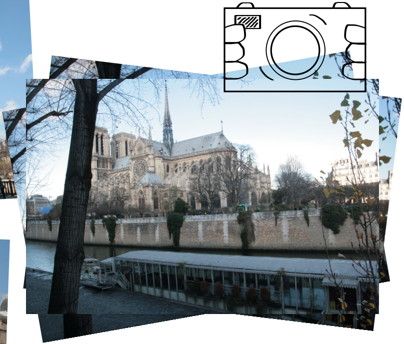

camera motion

Challenges of HDR imaging in dynamic scenes

camera + object motion

ghosting effect

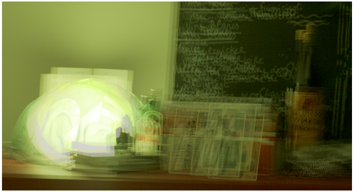

Would it be possible to create a HDR image from a single shot ?

Would it be possible to create a HDR image from a single shot ?

First, let's focus on a very generic inverse problem...

A generic inverse problem

Original image

A generic inverse problem

Noise

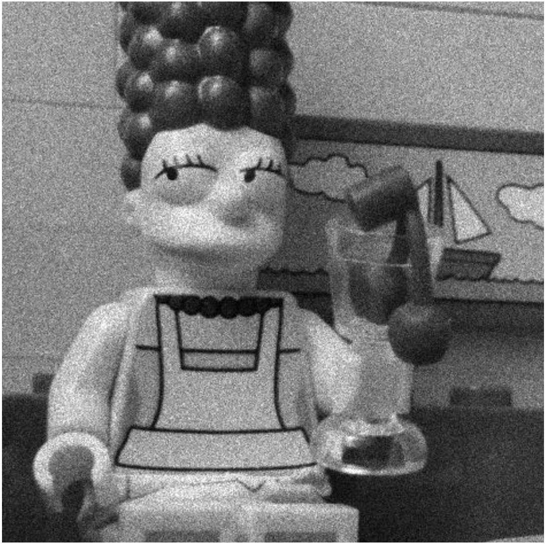

A generic inverse problem

Missing pixels

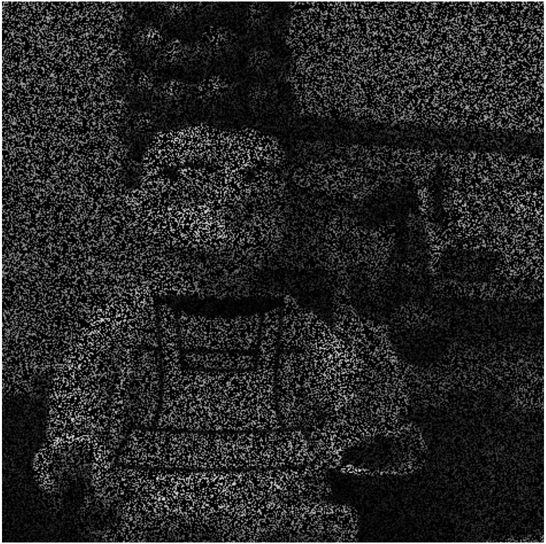

Inverse Problem

Degradation model

$$\tilde{u} = Au + n$$

- u reference image
- A is a diagonal operator
- Additive noise n may depend on u:
  Exemple RAW data (shot noise and readout noise)

$$n(x) \sim \mathcal{N}(0, \alpha(x)u(x) + \beta(x))$$

Inverse Problem

Degradation model

$$\tilde{u} = Au + n$$

Notation for patches

Z_i $= p_i(\tilde{u})$ (degraded patch of size $d = f \times f$ centered at i)

C_i $= p_i(u)$ (unknown reference patch)

N_i $= p_i(n)$ (additive noise patch)

D_i restriction of A to $p_i(u)$

Degradation model for a patch centered at pixel i

$$\mathbf{Z_i} = \mathbf{D_i C_i} + \mathbf{N_i}$$

Patch degradation Model

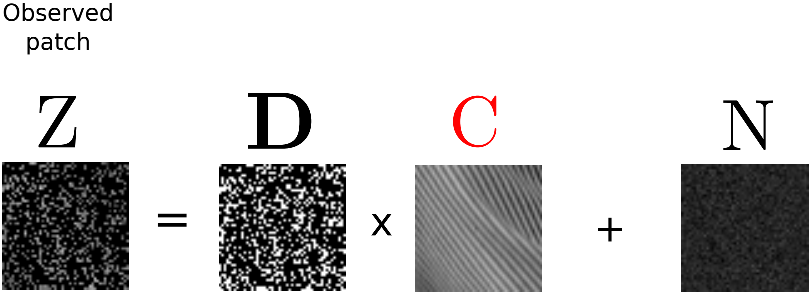

Patch degradation Model

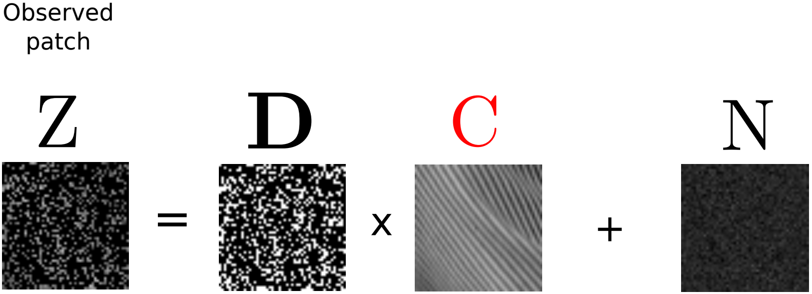

Assumptions :

- D is known
- $N \sim \mathcal{N}(0, \Sigma_N)$, eventually depends on C but $\mathrm{Cov}(N, C) = 0$

Patch degradation Model

Observed patch

Patch we seek to estimate

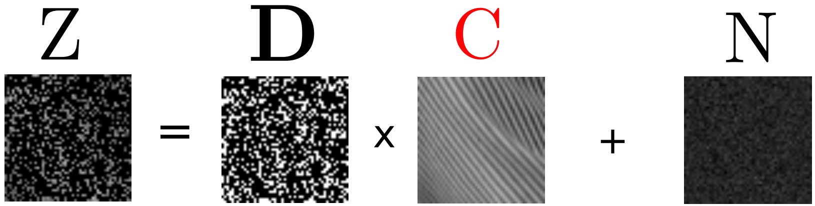

Gaussian prior for patches

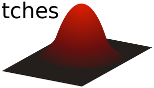

Assumptions :

- D is known
- $N \sim \mathcal{N}(0, \Sigma_N)$, eventually depends on C but $\mathrm{Cov}(N, C) = 0$
- $C \sim \mathcal{N}(\mu, \Sigma)$ with μ and Σ unknown

How to set Gaussian prior parameters μ and Σ ?

1. **Classical choice : MLE [NL-Bayes - Lebrun et al. 2013]**
   Set of similar patches $Z_1, \ldots, Z_M$, such that all the (unknown) C_i follow the same law $\mathcal{N}(\mu, \Sigma)$.

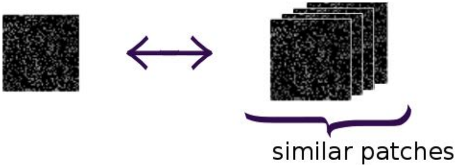

$$\widehat{\mu} = \frac{1}{M} \sum_{i=1}^{M} Z_i \quad \text{and} \quad \widehat{\Sigma} = \frac{1}{M-1} \sum_{i=1}^{M} [Z_i - \widehat{\mu}][Z_i - \widehat{\mu}]^T$$

How to set Gaussian prior parameters μ and Σ ?

1. **Classical choice : MLE [NL-Bayes - Lebrun et al. 2013]**
   Set of similar patches $Z_1, \ldots, Z_M$, such that all the (unknown) C_i follow the same law $\mathcal{N}(\mu, \Sigma)$.

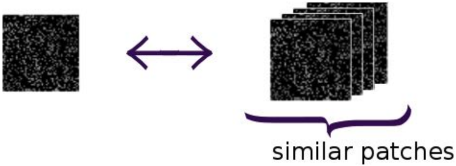

$$\widehat{\mu} = \frac{1}{M}\sum_{i=1}^{M} Z_i \quad \text{and} \quad \widehat{\Sigma} = \frac{1}{M-1}\sum_{i=1}^{M}[Z_i - \widehat{\mu}][Z_i - \widehat{\mu}]^T$$

Not reliable when pixels are missing !

How to set Gaussian prior parameters μ and Σ ?

1. **Classical choice : MLE [NL-Bayes - Lebrun et al. 2013]**
   Set of similar patches $Z_1, \ldots, Z_M$, such that all the (unknown) C_i follow the same law $\mathcal{N}(\mu, \Sigma)$.

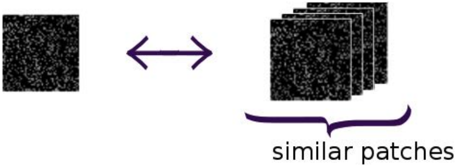

$$\widehat{\mu} = \frac{1}{M} \sum_{i=1}^{M} Z_i \quad \text{and} \quad \widehat{\Sigma} = \frac{1}{M-1} \sum_{i=1}^{M} [Z_i - \widehat{\mu}][Z_i - \widehat{\mu}]^T$$

Not reliable when pixels are missing !

2. **Gaussian Mixture Model prior on patches [EPLL - Zoran and Weiss 2011 ; PLE - Yu et al. 2012]**

How to set Gaussian prior parameters μ and Σ ?

MAP with an hyperprior on (μ, Σ)

$$\underset{\{C_i\},\mu,\Sigma}{\operatorname{argmax}} \; p(\{C_i\}_i, \mu, \Sigma \mid \{Z_i\}_i) =$$

$$\underset{\{C_i\},\mu,\Sigma}{\operatorname{argmax}} \; p(\{Z_i\} \mid \{C_i\}, \mu, \Sigma) \,.\, p(\{C_i\} \mid \mu, \Sigma) \,.\, p(\mu, \Sigma).$$

Rappel

- $Z_i \mid C_i, \mu_i, \Sigma_i \;\sim\; \mathcal{N}(D_i C_i, \Sigma_{N_i})$
- $C_i \mid \mu_i, \Sigma_i \;\sim\; \mathcal{N}(\mu, \Sigma)$
- (μ, Σ) ?

Inclusion of hyperprior information compensates for missing pixels.

Hyperprior on (μ, Σ)

Conjugate prior for a **multivariate normal distribution**

- Normal prior on the mean (conditionnal on the covariance)

$$\mu \mid \Sigma \sim \mathcal{N}(\mu_0, \Sigma/\kappa) \propto |\Sigma|^{-\frac{1}{2}} \exp\left(-\frac{\kappa}{2}(\mu - \mu_0)^T \Sigma^{-1}(\mu - \mu_0)\right)$$

- inverse Wishart prior on the covariance matrix

$$\Sigma \sim \mathcal{IW}(\nu\Sigma_0, \nu) \propto |\Sigma|^{-\frac{\nu+d+1}{2}} \exp\left(-\frac{1}{2}\mathrm{trace}[\nu\Sigma_0\Sigma^{-1}]\right)$$

Hyperprior on (μ, Λ) with $\Lambda = \Sigma^{-1}$ (precision matrix)

Conjugate prior for a **multivariate normal distribution**

- Normal prior on the mean (conditionnal on the covariance)

$$\mu \mid \Lambda \sim \mathcal{N}(\mu_0, \Lambda^{-1}/\kappa) \propto |\Lambda|^{\frac{1}{2}} \exp\left(-\frac{\kappa}{2}(\mu - \mu_0)^T \Lambda (\mu - \mu_0)\right)$$

- Wishart prior on the inverse covariance matrix

$$\Lambda \sim \mathcal{W}((\nu\Sigma_0)^{-1}, \nu) \propto |\Lambda|^{\frac{\nu-d-1}{2}} \exp\left(-\frac{1}{2}\mathrm{trace}[\nu\Sigma_0\Lambda]\right)$$

Minimization with respect to $\{C_i\}_i$

- $\{Z_i\}$ set of similar patches;
- μ, Λ fixed.

$$\begin{aligned}
&\underset{\{C_i\}}{\text{argmax}} \;\; p(\{Z_i\} \mid \{C_i\}) \,.\, p(\{C_i\} \mid \mu, \Lambda) \,.\, p(\mu, \Lambda) \\
&= \underset{\{C_i\}}{\text{argmax}} \;\; \Pi_{i=1}^{M} \left(p(Z_i \mid C_i) \,.\, p(C_i \mid \mu, \Lambda) \right) \\
&= \underset{\{C_i\}}{\text{argmax}} \;\; \Pi_{i=1}^{M} \left(g_{0,\Sigma_{N_i}}(Z_i - D_i C_i) \,.\, g_{\mu,\Lambda^{-1}}(C_i) \right) \qquad .
\end{aligned}$$

Solution given by Wiener estimator for each i separately

$$\widehat{C}_i = \underbrace{\underbrace{\Lambda^{-1} D_i^T}_{\mathbb{E}(C_i Z_i^T)} \underbrace{(D_i \Lambda^{-1} D_i^T + \Sigma_{N_i})}_{\mathbb{E}(Z_i Z_i^T)}{}^{-1}}_{W_i} (Z_i - D_i \mu) + \mu$$

Minimization with respect to μ, Λ

- $\{C_i\}_i$ fixed.

$$\underset{\mu,\Lambda}{\operatorname{argmax}} \; \underbrace{p(\{Z_i\} \mid \{C_i\}, \mu, \Lambda)}_{\text{HYP : independent of } \mu,\Lambda} \; . \; p(\{C_i\} \mid \mu, \Lambda) \; . \; p(\mu, \Lambda)$$

$$\simeq \underset{\mu,\Lambda}{\operatorname{argmax}} \; p(\{C_i\} \mid \mu, \Lambda) \; . \; p(\mu, \Lambda)$$

$$= \underset{\mu,\Lambda}{\operatorname{argmax}} \; \Pi_{i=1}^{M} \, g_{\mu,\Lambda^{-1}}(C_i) \; g_{\mu_0,\Lambda^{-1}/\kappa}(\mu) \; w_{\Lambda_0/\nu,\nu}(\Lambda).$$

Explicit solution

$$\begin{cases} \widehat{\mu} = \frac{M\overline{C}+\kappa\mu_0}{M+\kappa} \\ \widehat{\Lambda}^{-1} = \frac{\nu\Sigma_0+\kappa(\widehat{\mu}-\mu_0)(\widehat{\mu}-\mu_0)^T+\sum_{i=1}^{M}(\widehat{C}_i-\widehat{\mu})(\widehat{C}_i-\widehat{\mu})^T}{\nu+M-d} \end{cases}$$

Loop in (μ, Λ)

In the previous formula, $\hat{C}_i$, $\hat{\mu}$ and $\hat{\Lambda}$ depend on each other. Replacing $\hat{C}_i$ by its expression in (μ, Λ) and reinjecting this in the formula of $(\hat{\mu}, \hat{\Lambda})$, we get

$$\widehat{\mu} = \left(\kappa Id + \sum_{j=1}^{M} W_j D_j \right)^{-1} \left(\sum_{j=1}^{M} W_j Z_j + \kappa \mu_0 \right)$$

$$\begin{aligned} (\nu + M - d)\, \widehat{\Lambda}^{-1} &= \sum_{j=1}^{M} (W_j (Z_j - D_j \mu))(W_j (Z_j - D_j \mu))^T \\ &+ \kappa(\mu - \mu_0)(\mu - \mu_0)^T + \nu \Sigma_0 \end{aligned}$$

with $W_j = \Lambda^{-1} D_j^T (D_j \Lambda^{-1} D_j^T + \Sigma_{N_j})^{-1}$.

Algorithm

Initialization : compute Oracle image C_{oracle}

For k = 1 : maxit

- For each Patch Z
  1. Find patches similar to Z in C_{oracle}
  2. Compute μ_0 and Σ_0 from this set of similar patches in C_{oracle}
  3. Compute first $\widehat{\mu}$, $\widehat{\Sigma}$ with a small loop and then $\hat{C}$.
- Restore image from restored patches and update C_{oracle} = restored image.

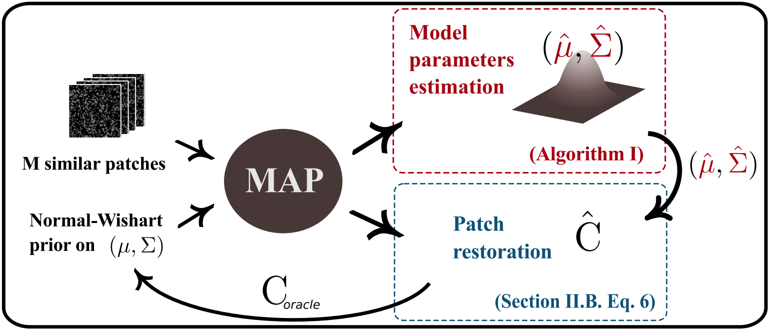

Initialization

From PLE [Yu et al., 2012]:

K predefined models : (K-1) edges with different orientations + DCT for isotropic patterns

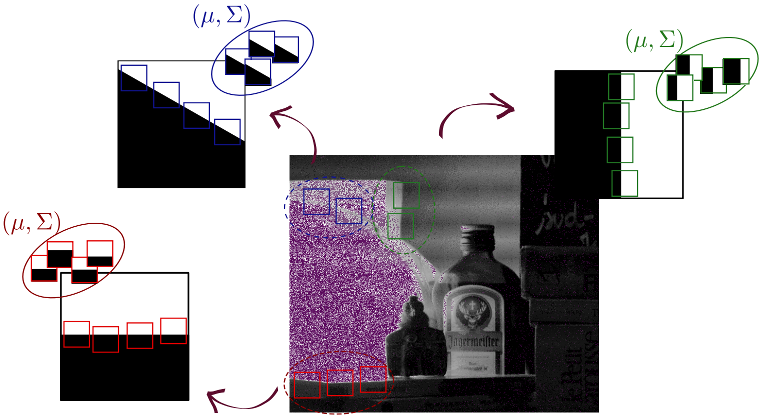

Results

(a) Ground-truth (b) HBE (**30.01 dB**) (c) PLE (26.78 dB)

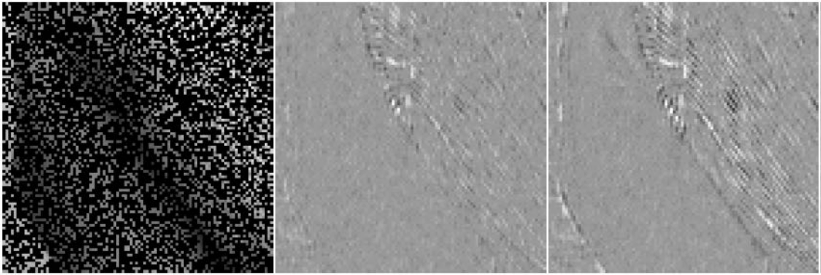

Synthetic data, 70% missing pixels.

Results

(f) Ground-truth (g) HBE (**30.20 dB**) (h) PLE (27.89 dB)

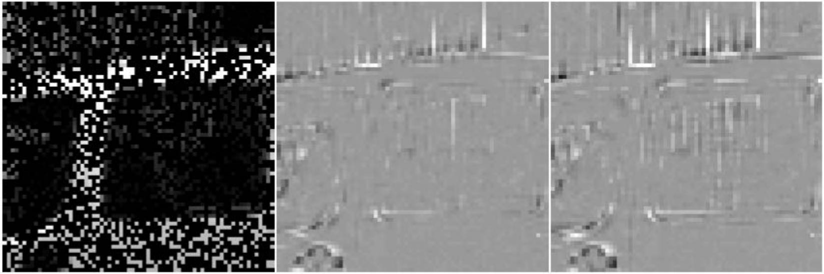

Synthetic data, 70% missing pixels.

Results

(a) Ground-truth (b) HBE (**26.20 dB**) (c) PLE (24.76 dB)

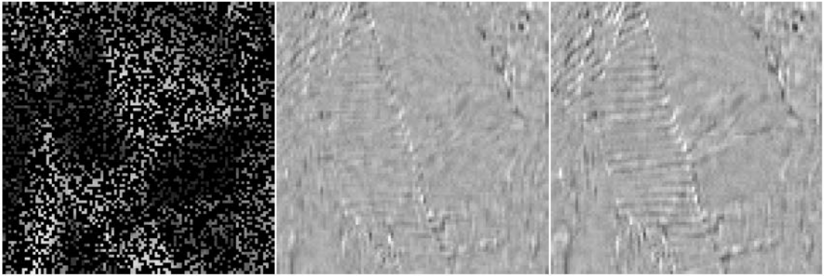

Synthetic data, 70% missing pixels, gaussian noise $\sigma = 10$.

Results

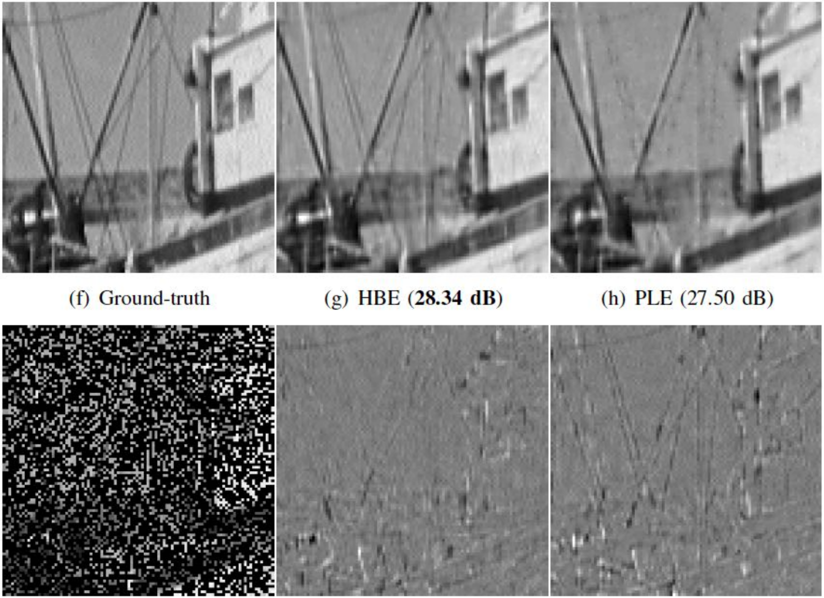

(f) Ground-truth (g) HBE (**28.34 dB**) (h) PLE (27.50 dB)

Synthetic data, 70% missing pixels, gaussian noise $\sigma = 10$.

Results

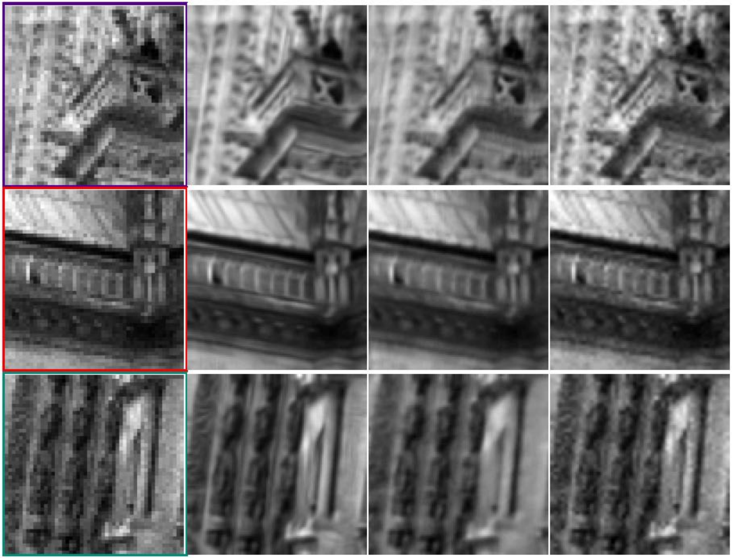

Zoom on Real data. Left to right : Input low-resolution image, HBE, PLE, bicubic.

Now, how can we do HDR imaging from a single shot ?

Spatially Varying pixel Exposures (SVE)
[Nayar and Mitsunaga, 2000]

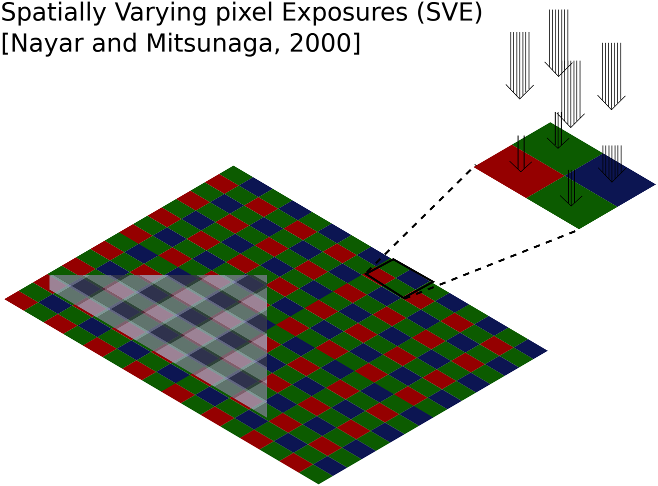

1 image = N exposures

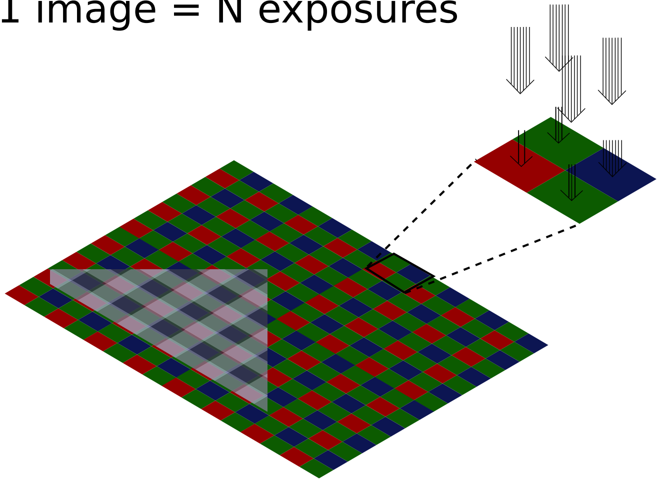

SVE Single-image HDR

- ✓ No need for image alignment.
- ✓ No need for motion detection.
- ✓ No ghosting problems.
- ✓ No large saturated regions to fill.

- × Resolution loss : unknown pixels to be restored (over and under exposed pixels).
- × Noise.
- × Need to modify the standard camera.
  - Alternative without camera modification [Hirakawa and Simon, 2011].

SVE: Regular or Random?

Random pattern to avoid aliasing [Schöberl et al., 2012]

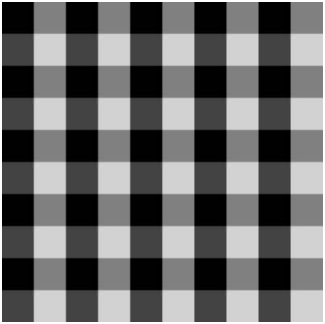

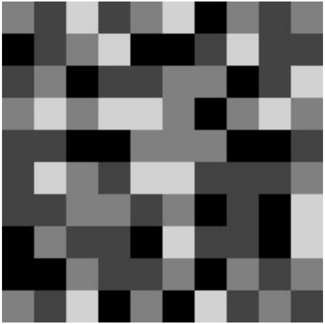

Inverse Problem for HDR

Degradation model for a patch centered at pixel i

$$Z_i = D_i C_i + N_i$$

- D_i is a diagonal operator
  - $D_{ii} = 0 \quad \Rightarrow$ over- or under-exposed pixel (ignored)
  - $D_{ii} = 1 \quad \Rightarrow$ well-exposed pixel (kept)
- C_i irradiance at pixel i (reference image)
- Noise model for RAW data (shot noise and readout noise)

$$N_i \sim \mathcal{N}(0, \Sigma_{N_i})$$

with diagonal covariance matrix Σ_{N_i} such that

$$(\Sigma_{N_i})_k = \alpha_k C_k + \beta_k,$$

with α and β known.

Results HDR - Synthetic data

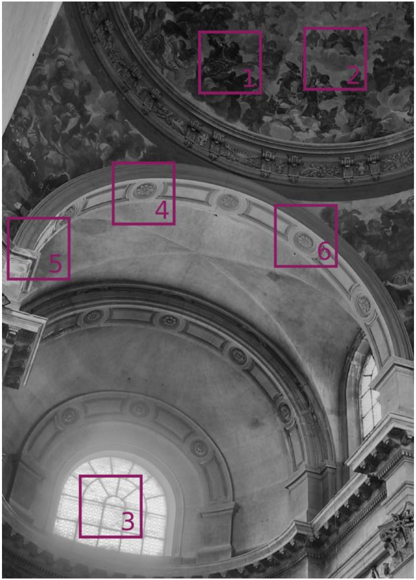

Results HDR - Synthetic data

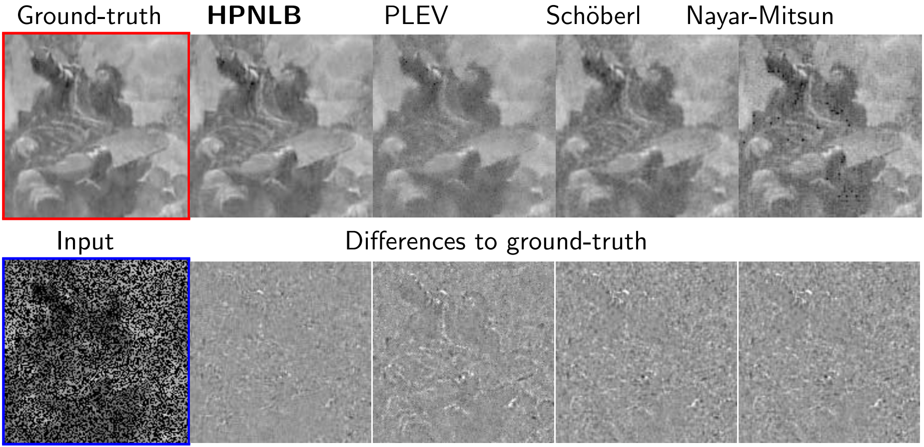

PSNR:	**33.1dB**	29.7dB	30.4dB	29.4dB

Results HDR - Synthetic data

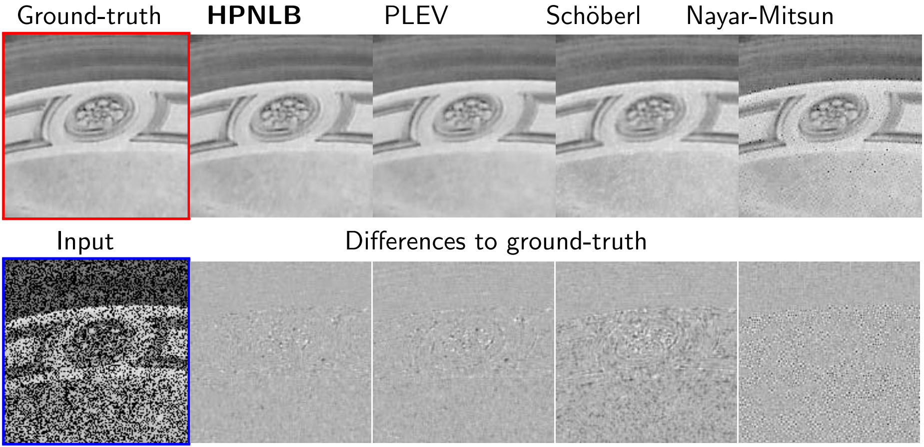

PSNR:	**35.1dB**	34.0dB	30.0dB	28.5dB

Real data: experimental protocol

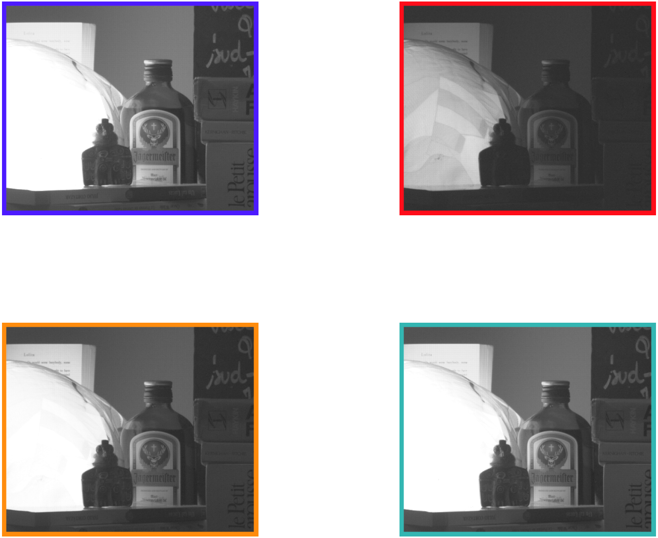

Real data: experimental protocol

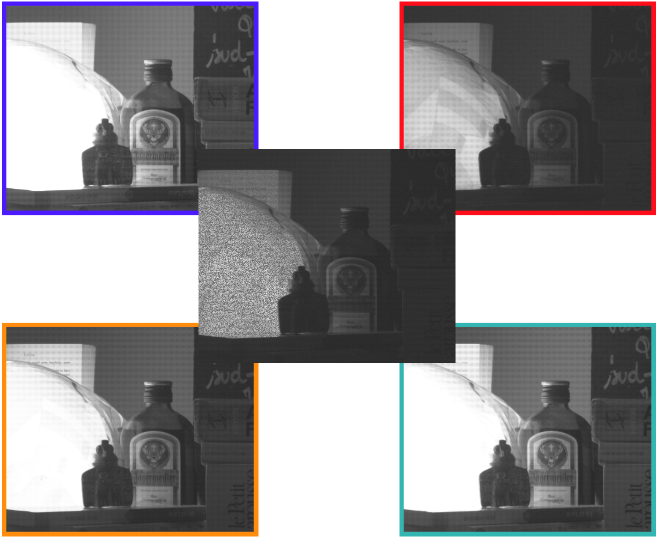

Real data: experimental protocol

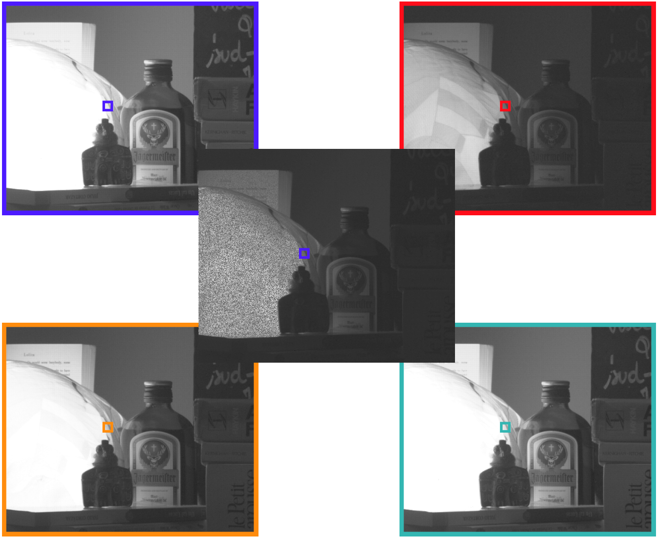

Real data: experimental protocol

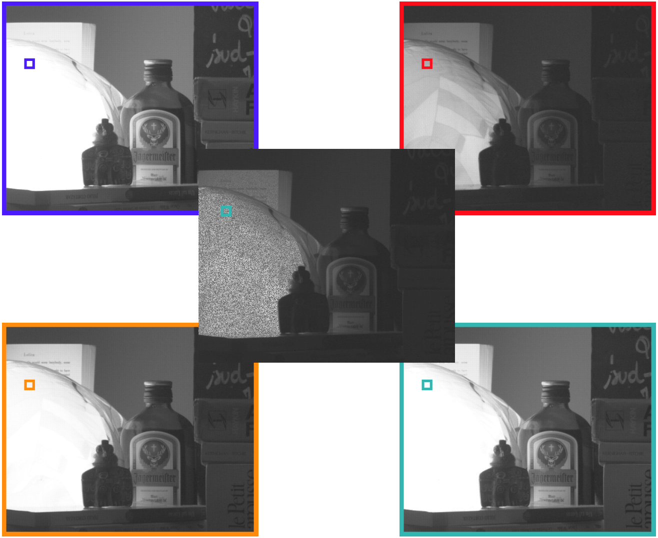

Real data: experimental protocol

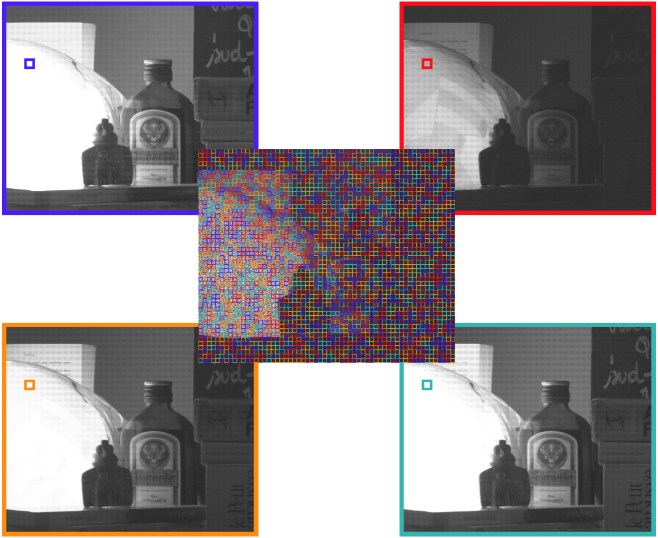

Results HDR - Real data

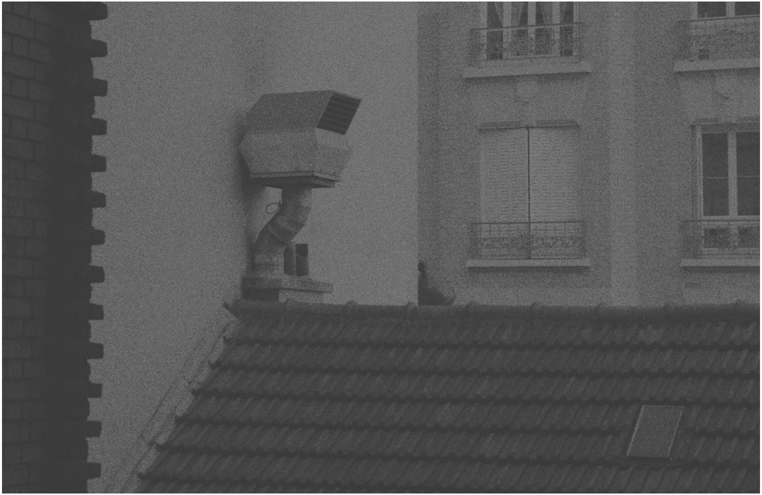

Results HDR - Real data

Results HDR - Real data

Results HDR - Real data

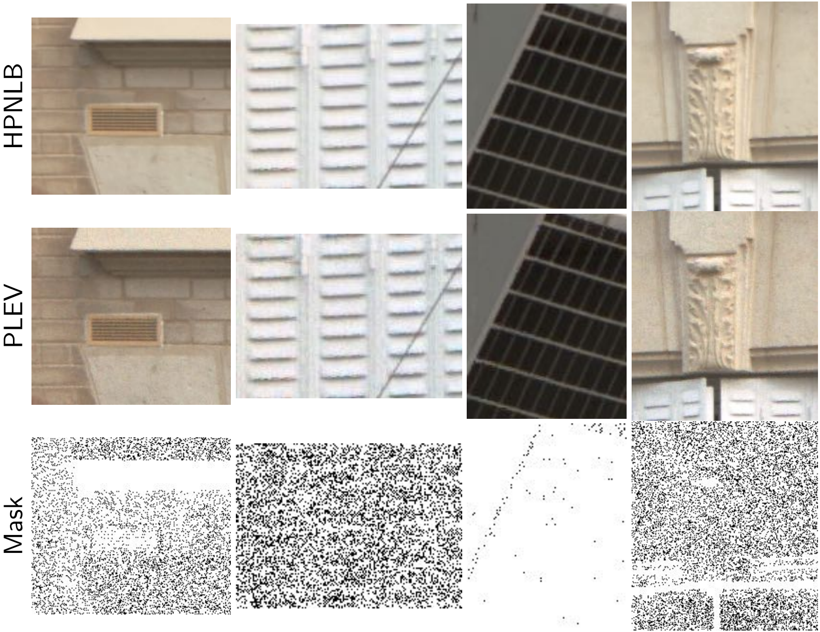

Results real data

Results real data

Results real data

Results real data