Denoising Algorithms: Path to Neural Networks I



TRAINING

Philipp Slusallek Karol Myszkowski



Gurprit Singh

Previous Lecture Overview











Image-space Adaptive Sampling

Realistic Image Synthesis SS2020





Realistic Image Synthesis SS2020



Image-space Adaptive Sampling

Hachisuka et al. [2008]



Depth of field





Depth of field





Depth of field











1 scanline











Slide from Jakko Lehtinen



Lens u



Slide from Jakko Lehtinen

 \bigcirc

 \bigcirc



The trajectories of samples originating from a single **apparent surface** never intersect.



The trajectories of samples originating from a single **apparent surface** never intersect.

Slide from Jakko Lehtinen



The trajectories of samples originating from a single **apparent surface** never intersect.



The trajectories of samples originating from a single **apparent surface** never intersect.



Introduction Denoising using Data

Path to Machine Learning

Introduction **Denoising using Data**

Path to Machine Learning

MLP based Denoising

Introduction **Denoising using Data**

Path to Machine Learning

MLP based Denoising

CNN based Denosing (Next lecture)



Filtering Monte Carlo Noise From **Random Parameters**



Realistic Image Synthesis SS2020

Sen and Darabi [2012]







input Monte Carlo (8 samples/pixel)

after RPF (8 samples/pixel)

High-dimensional Monte Carlo Integration













(a) Input MC (8 spp)







(b) Dependency on (u, v) (c) Our approach (RPF)



Parameters in Monte Carlo estimator

Random parameters: $\mathbf{r} = \{r_1, r_2, ..., r_n\}$

Color: $\mathbf{c}_i \leftarrow f(\mathbf{p}_{i,1},\mathbf{p}_{i,2};\mathbf{r}_{i,1},\mathbf{r}_{i,2},\ldots,\mathbf{r}_{i,n})$

screen position



random parameters

Realistic Image Synthesis SS2020





Random parameter for each pixel :

$$\mathbf{x}_i \Leftarrow f(\mathbf{p}_{i,1},\mathbf{p})$$

$$\mathbf{x}_{i} = \{\underbrace{\mathbf{p}_{i,1}, \mathbf{p}_{i,2}}_{\text{screen position}}; \underbrace{\mathbf{r}_{i,1}, \dots, \mathbf{r}_{i,2}}_{\text{random parameters}}; \underbrace{\mathbf{p}_{i,1}, \dots, \mathbf{r}_{i,2}}_{\text{random parameters}}; \underbrace{\mathbf{p}_{i,1}, \dots, \mathbf{r}_{i,2}}_{\text{parameters}}; \underbrace{\mathbf{p}_{i,2}, \dots, \mathbf{p}_{i,2}}_{\text{parameters}}; \underbrace{\mathbf{p}_{i,2}, \dots, \mathbf{p}_{i,2}, \dots, \mathbf{p}_{i,2}, \dots, \mathbf{p}_{i,2}}_{\text{parameters}}; \underbrace{\mathbf{p}_{i,2}, \dots, \mathbf{p}_{i,2}, \dots, \mathbf{p}_{i,2},$$



 $\mathbf{p}_{i,2}; \mathbf{r}_{i,1}, \mathbf{r}_{i,2}, \dots, \mathbf{r}_{i,n}$





Realistic Image Synthesis SS2020



Random parameter for each pixel :

$$\mathbf{x}_i \Leftarrow f(\mathbf{p}_{i,1},\mathbf{p})$$



15



 $\mathbf{p}_{i,2}; \mathbf{r}_{i,1}, \mathbf{r}_{i,2}, \dots, \mathbf{r}_{i,n}$





Random parameter for each pixel :

$$\mathbf{x}_i \Leftarrow f(\mathbf{p}_{i,1},\mathbf{p})$$



15



 $\mathbf{p}_{i,2}; \mathbf{r}_{i,1}, \mathbf{r}_{i,2}, \dots, \mathbf{r}_{i,n}$





Random parameter for each pixel :

$$\mathbf{x}_i \Leftarrow f(\mathbf{p}_{i,1},\mathbf{p})$$



15



 $\mathbf{p}_{i,2}; \mathbf{r}_{i,1}, \mathbf{r}_{i,2}, \dots, \mathbf{r}_{i,n}$





Random parameter for each pixel :

$$\mathbf{x}_i \Leftarrow f(\mathbf{p}_{i,1},\mathbf{p})$$





 $\mathbf{p}_{i,2}; \mathbf{r}_{i,1}, \mathbf{r}_{i,2}, \dots, \mathbf{r}_{i,n}$



Realistic Image Synthesis SS2020



Gaussian Filtering





 $\sigma = 8$ $\sigma = 4$

$$GC[I]_{\mathbf{p}} = \sum_{\mathbf{q}\in\mathcal{S}} G_{\sigma}(\|\mathbf{p}-\mathbf{q}\|) I_{\mathbf{q}}, \quad G_{\sigma}(x) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2}{2\sigma^2}\right)$$



 $\sigma = 16$

 $\sigma = 32$

٦в

Realistic Image Synthesis SS2020



$$BF[I]_{\mathbf{p}} = \frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathbf{s}}}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_{\mathbf{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|) I_{\mathbf{q}}$$

 $W_{\mathbf{p}} = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathbf{s}}}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_{\mathbf{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)$





Bilateral filter weights at the central pixel



Realistic Image Synthesis SS2020



$$BF[I]_{\mathbf{p}} = \frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathbf{s}}}(\|\mathbf{p} - \mathbf{q}\|) \frac{G_{\sigma_{\mathbf{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)}{G_{\sigma_{\mathbf{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)} I_{\mathbf{q}}$$

 $W_{\mathbf{p}} = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathbf{s}}}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_{\mathbf{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)$





Bilateral filter weights at the central pixel



Realistic Image Synthesis SS2020



$$BF[I]_{\mathbf{p}} = \frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathbf{s}}}(\|\mathbf{p} - \mathbf{q}\|) \frac{G_{\sigma_{\mathbf{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)}{|I_{\mathbf{q}}|} I_{\mathbf{q}}$$

 $W_{\mathbf{p}} = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathbf{s}}}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_{\mathbf{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)$





Realistic Image Synthesis SS2020



$$BF[I]_{\mathbf{p}} = \frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathbf{s}}}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_{\mathbf{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|) I_{\mathbf{q}}$$

 $W_{\mathbf{p}} = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathbf{s}}}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_{\mathbf{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)$







Realistic Image Synthesis SS2020





$$W_{\mathbf{p}} = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathbf{s}}}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_{\mathbf{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)$$





Realistic Image Synthesis SS2020




$$W_{\mathbf{p}} = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathbf{s}}}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_{\mathbf{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)$$





Realistic Image Synthesis SS2020



$$BF[I]_{\mathbf{p}} = \frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathbf{s}}}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_{\mathbf{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|) I_{\mathbf{q}}$$

$$W_{\mathbf{p}} = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_{\mathbf{s}}}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_{\mathbf{r}}}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)$$

$$I_{\mathbf{p}}$$



al Filtering



Realistic Image Synthesis SS2020



Bilateral vs Gaussian Filtering

$\sigma_s \backslash \sigma_r$ 0.05

0.2



16

4

8





Realistic Image Synthesis SS2020

0.8

GC

18



Bilateral Filtering of Features $w_{ij} = \exp\left[-\frac{1}{2\sigma_{\mathbf{p}}^2} \sum_{1 < k < 2} (\bar{\mathbf{p}}_{i,k} - \bar{\mathbf{p}}_{j,k})^2\right] \times$ $\exp\left[-\frac{1}{2\sigma_{\mathbf{c}}^2}\sum_{1\leq k\leq 2}\alpha_k(\mathbf{\bar{c}}_{i,k}-\mathbf{\bar{c}}_{j,k})^2\right]\times$ $\exp\left[-\frac{1}{2\sigma_{\mathbf{f}}^2}\sum_{1 < k < m} \beta_k (\bar{\mathbf{f}}_{i,k} - \bar{\mathbf{f}}_{j,k})^2\right],$



Realistic Image Synthesis SS2020



Dependency on Random Parameters



Input Monte Carlo (8 spp)

Dependency of color on random parameters $(D_{\mathbf{c}}^{\mathbf{r}})$

Dependency of color on screen position $(D_{\mathbf{c}}^{\mathbf{p}})$



Fractional dependency on random parameters $(W_{\mathbf{c}}^{\mathbf{r}})$

Reference MC (512 spp)



Realistic Image Synthesis SS2020







Bilateral Weights





Realistic Image Synthesis SS2020





Pixels, Random Params, Features



(a) Screen position (b) Random parameters



(c) World space coords.

22





Pixels, Random Params, Features



(d) Surface normals



Realistic Image Synthesis SS2020

(e) Texture value

(f) Sample color







Pixels, Random Params, Features



(a) Screen position

(b) Random parameters

(c) World space coords.



(d) Surface normals

(e) Texture value

(f) Sample color

The algorithm computes the statistical dependency of (c-f) on the random parameters in (b)







Random Parameter Filtering



(a) Reference (b) MC Input



Realistic Image Synthesis SS2020

(d) no clustering (e) no DoF params

25

(c) RPF



Random Parameter Filtering



(a) $W^{\mathbf{r},1}_{\mathbf{c},k}$ and $W^{\mathbf{r},2}_{\mathbf{c},k}$



Realistic Image Synthesis SS2020



(b)

(c) Our output (RPF)



Mutual information between two random variables:

$$\mu(X;Y) = \sum_{x \in X} \sum_{y \in Y} p(x,y) \log \frac{p(x,y)}{p(x)p(y)}$$

where, these probabilities are computed over the neighborhood of samples around a given pixel



Statistical Dependency

Realistic Image Synthesis SS2020



Functional dependency of the k-th scene parameter:

$$D_{\mathbf{f},k}^{\mathbf{r}} = \sum_{1 \le l \le n} D_{\mathbf{f},k}^{\mathbf{r},l} = \sum_{1 \le l \le n} \mu(\overline{\mathbf{f}}_{\mathcal{N},k}; \overline{\mathbf{r}}_{\mathcal{N},l})$$

$$\begin{split} D_{\mathbf{f},k}^{\mathbf{p}} &= \sum_{1 \leq l \leq 2} D_{\mathbf{f},k}^{\mathbf{p},l} = \sum_{1 \leq l \leq 2} \mu(\bar{\mathbf{f}}_{\mathcal{N},k}; \bar{\mathbf{p}}_{\mathcal{N},l}), \\ D_{\mathbf{c},k}^{\mathbf{r}} &= \sum_{1 \leq l \leq n} D_{\mathbf{c},k}^{\mathbf{r},l} = \sum_{1 \leq l \leq n} \mu(\bar{\mathbf{c}}_{\mathcal{N},k}; \bar{\mathbf{r}}_{\mathcal{N},l}), \\ D_{\mathbf{c},k}^{\mathbf{p}} &= \sum_{1 \leq l \leq 2} D_{\mathbf{c},k}^{\mathbf{p},l} = \sum_{1 \leq l \leq 2} \mu(\bar{\mathbf{c}}_{\mathcal{N},k}; \bar{\mathbf{p}}_{\mathcal{N},l}). \end{split}$$



Statistical Dependency



Realistic Image Synthesis SS2020



$W_{\mathbf{c}}^{\mathbf{f},k} = \frac{D_{\mathbf{c}}^{\mathbf{f},k}}{D_{\mathbf{c}}^{\mathbf{r}} + D_{\mathbf{c}}^{\mathbf{p}} + D_{\mathbf{c}}^{\mathbf{f}}}$

Statistical Dependency $D_{\mathbf{f},k}^{\mathbf{r}} = \sum D_{\mathbf{f},k}^{\mathbf{r},l} = \sum \mu(\bar{\mathbf{f}}_{\mathcal{N},k}; \bar{\mathbf{r}}_{\mathcal{N},l})$ $1 \leq l \leq n$ $1 \leq l \leq n$

 $1{\leq}k{\leq}3$



 $D_{\mathbf{c}}^{\mathbf{r}} = \sum D_{\mathbf{c},k}^{\mathbf{r}}, \quad D_{\mathbf{c}}^{\mathbf{p}} = \sum D_{\mathbf{c},k}^{\mathbf{p}}, \quad D_{\mathbf{c}}^{\mathbf{f}} = \sum D_{\mathbf{c},k}^{\mathbf{f}},$ $1{\leq}k{\leq}3$ $1 \le k \le 3$

Realistic Image Synthesis SS2020



Weighted Average Bilateral Filtering





 $\mathbf{c}_{i,k}' = rac{\sum_{j \in \mathcal{N}} w_{ij} \mathbf{c}_{j,k}}{\sum_{j \in \mathcal{N}} w_{ij}}$









(a) MC Input (8 spp) (b) Our approach (RPF) (c) $\alpha_k = 0, \beta_k = 0$





Results





Results



) (c) $\alpha_k = 0, \beta_k = 0$ (d) $\alpha_k = 1, \beta_k = 0$



(e) $\alpha_k = 0, \beta_k = 1$ (f) $\alpha_k = 1, \beta_k = 1$

Realistic Image Synthesis SS2020



Results













Multi-Layer Perceptrons







History of Neural Networks

- model for neural networks
- accelerated the training of multi-layer networks.



In 1943, McCulloch and Pitts created a computational

• In 1975, Werbos's back propagation algorithm generally

• In 1980s, Recurrent Neural Networks were developed











Classifiers









Realistic Image Synthesis SS2020







Complex classifier



Complex Classifiers



What features can produce this decision rule?



Realistic Image Synthesis SS2020

Complex classifier







1











1





Realistic Image Synthesis SS2020









Realistic Image Synthesis SS2020













Multi-layer Perceptron

 x_1

1







Multi-layer Perceptron f \sum $\sum \rightarrow f$ f \sum

 x_1

1







Realistic Image Synthesis SS2020





Multi-layer Perceptron f \sum $\sum \rightarrow f$ f \sum

 x_1

1







Realistic Image Synthesis SS2020

44







Multi-layer Perceptron ff

44

Realistic Image Synthesis SS2020



Multi-layer Perceptron w_{11} w_{10} w_{21} f w_{20} w_{31} w_{30}





Realistic Image Synthesis SS2020

44



Multi-layer Perceptron w_{11} w_{10} w_{21} w_{20} w_{31} w_{30} $x_1 w_{11}$ w_{10} $x_1 w_{21} + w_{20}$





45

Realistic Image Synthesis SS2020
































Hidden layers





Output layers

Realistic Image Synthesis SS2020





Hidden layers





Output layers

48

Realistic Image Synthesis SS2020





Matrix of first layer weights

w_{11}	w_{10}
w_{21}	w_{20}
w_{31}	w_{30}



Output layers

"Features" are outputs of perceptrons



Matrix of second layer weights

 w_1 w_2 w_3









Matrix of first layer weights

w_{11}	w_{10}
w_{21}	w_{20}
w_{31}	w_{30}



Output layers

"Features" are outputs of perceptrons



Matrix of second layer weights













Matrix of first layer weights

w_{11}	w_{10}
w_{21}	w_{20}
w_{31}	w_{30}



Output layers

"Features" are outputs of perceptrons



Matrix of second layer weights











Input features







Realistic Image Synthesis SS2020

Perceptron: Step function with linear decision boundary











These outputs are now input features to the next layer









These outputs are now input features to the next layer

"Features" are now decision boundaries (partitions)











- These outputs are now input features to the next layer
- "Features" are now decision boundaries (partitions)
- All linear combination of those partitions give complex partitions











Layer 2

These complex outputs become the features for the new layer

52













Deep Neural Networks



53

Realistic Image Synthesis SS2020



Computational Graph representation of Neural Networks



Realistic Image Synthesis SS2020



Fully connected layers





Realistic Image Synthesis SS2020









Realistic Image Synthesis SS2020



Fully connected layers







Neural Networks



Realistic Image Synthesis SS2020







N represents number of pixels in an image



Neural Networks













Realistic Image Synthesis SS2020



59







UNIVERSITÄT DES SAARLANDES \odot

Realistic Image Synthesis SS2020

Neural Networks









Realistic Image Synthesis SS2020

Two-layer model

Fully connected layers











Realistic Image Synthesis SS2020

Two-layer model

Fully connected layers



Reference











Realistic Image Synthesis SS2020

Two-layer model

Fully connected layers

*











Realistic Image Synthesis SS2020

Two-layer model

Fully connected layers

*













Two-layer model

What can be a loss function ?



Realistic Image Synthesis SS2020

Reference













Two-layer model

What can be a loss function ?



Realistic Image Synthesis SS2020

Reference

















67

Realistic Image Synthesis SS2020







Realistic Image Synthesis SS2020

Gradient Descent Algorithm for back propagation



68









Gradient Descent Algorithm for back propagation



68

Realistic Image Synthesis SS2020









Gradient Descent Algorithm for back propagation



Realistic Image Synthesis SS2020









Gradient Descent Algorithm for back propagation



Realistic Image Synthesis SS2020



Back Propagation Slides courtesy: <u>Stanford Online Course</u>





69









Realistic Image Synthesis SS2020







Realistic Image Synthesis SS2020









Realistic Image Synthesis SS2020







Realistic Image Synthesis SS2020






Realistic Image Synthesis SS2020







Realistic Image Synthesis SS2020



Back Propagation Slides courtesy: <u>Stanford Online Course</u>

Backpropagation: a simple example f(x, y, z) = (x + y)ze.g. x = -2, y = 5, z = -4







Back Propagation Slides courtesy: <u>Stanford Online Course</u>

Backpropagation: a simple example

$$f(x, y, z) = (x + y)z$$

e.g. x = -2, y = 5, z = -4

q = x + y	$rac{\partial q}{\partial x}=1, rac{\partial q}{\partial y}=1$
f = qz	$rac{\partial f}{\partial q}=z, rac{\partial f}{\partial z}=q$
Want: $\frac{\partial f}{\partial x}$	$, \frac{\partial f}{\partial y}, \frac{\partial f}{\partial z}$





77



Machine Learning for Filtering Monte Carlo Noise

Kalantari et al. [SIGGRAPH 2015]



Realistic Image Synthesis SS2020







Realistic Image Synthesis SS2020

 $\hat{\mathbf{c}}_{i} = \frac{\sum_{j \in \mathcal{N}(i)} d_{i,j} \bar{\mathbf{c}}_{j}}{\sum_{j \in \mathcal{N}(i)} d_{i,j}} \quad , \quad \hat{\mathbf{c}} = \{\hat{c}_{r}, \hat{c}_{g}, \hat{c}_{b}\}$









Realistic Image Synthesis SS2020

 $\hat{\mathbf{c}}_{i} = \frac{\sum_{j \in \mathcal{N}(i)} d_{i,j} \bar{\mathbf{c}}_{j}}{\sum_{j \in \mathcal{N}(i)} d_{i,j}} \quad , \quad \hat{\mathbf{c}} = \{\hat{c}_{r}, \hat{c}_{g}, \hat{c}_{b}\}$









Realistic Image Synthesis SS2020

 $\hat{\mathbf{c}}_{i} = \frac{\sum_{j \in \mathcal{N}(i)} d_{i,j} \bar{\mathbf{c}}_{j}}{\sum_{j \in \mathcal{N}(i)} d_{i,j}} \quad , \quad \hat{\mathbf{c}} = \{\hat{c}_{r}, \hat{c}_{g}, \hat{c}_{b}\}$ Pixel neighborhood









Realistic Image Synthesis SS2020

 $\hat{\mathbf{c}}_{i} = \frac{\sum_{j \in \mathcal{N}(i)} d_{i,j} \bar{\mathbf{c}}_{j}}{\sum_{j \in \mathcal{N}(i)} d_{i,j}}, \quad \hat{\mathbf{c}} = \{\hat{c}_{r}, \hat{c}_{g}, \hat{c}_{b}\}$ **Pixel neighborhood** 79







Realistic Image Synthesis SS2020

Filter weights $=\frac{\sum_{j\in\mathcal{N}(i)}d_{i,j}\bar{\mathbf{c}}_j}{\sum_{j\in\mathcal{N}(i)}d_{i,j}}$ $\hat{\mathbf{c}} = \{\hat{c}_r, \hat{c}_g, \hat{c}_b\}$ **Pixel neighborhood** 79



Filter weights $\hat{\mathbf{c}}_{i} = \frac{\sum_{j \in \mathcal{N}(i)} d_{i,j} \bar{\mathbf{c}}_{j}}{\sum_{j \in \mathcal{N}(i)} d_{i,j}}$ Pixel neighborhood



Realistic Image Synthesis SS2020

Filter weights





Filter weights $\hat{\mathbf{c}}_{i} = \frac{\sum_{j \in \mathcal{N}(i)} d_{i,j} \bar{\mathbf{c}}_{j}}{\sum_{j \in \mathcal{N}(i)} d_{i,j}}$ Pixel neighborhood



Realistic Image Synthesis SS2020

Filter weights

For cross Bilateral filters:





Filter weights $\hat{\mathbf{c}}_{i} = \frac{\sum_{j \in \mathcal{N}(i)} d_{i,j} \bar{\mathbf{c}}_{j}}{\sum_{j \in \mathcal{N}(i)} d_{i,j}}$ Pixel neighborhood $d_{i,j}$ =



Filter weights

For cross Bilateral filters:

$$= \exp\left[-\frac{\|\bar{\mathbf{p}}_{i} - \bar{\mathbf{p}}_{j}\|^{2}}{2\alpha_{i}^{2}}\right] \times \exp\left[-\frac{D(\bar{\mathbf{c}}_{i}, \bar{\mathbf{c}}_{j})}{2\beta_{i}^{2}}\right] \\ \times \prod_{k=1}^{K} \exp\left[-\frac{D_{k}(\bar{\mathbf{f}}_{i,k}, \bar{\mathbf{f}}_{j,k})}{2\gamma_{k,i}^{2}}\right],$$



Realistic Image Synthesis SS2020



Filter weights $\hat{\mathbf{c}}_{i} = \frac{\sum_{j \in \mathcal{N}(i)} d_{i,j} \bar{\mathbf{c}}_{j}}{\sum_{j \in \mathcal{N}(i)} d_{i,j}}$ Pixel neighborhood $d_{i,j}$ =



Filter weights

For cross Bilateral filters:

$$= \exp\left[-\frac{\|\bar{\mathbf{p}}_{i} - \bar{\mathbf{p}}_{j}\|^{2}}{2\alpha_{i}^{2}}\right] \times \exp\left[-\frac{D(\bar{\mathbf{c}}_{i}, \bar{\mathbf{c}}_{j})}{2\beta_{i}^{2}}\right] \\ \times \prod_{k=1}^{K} \exp\left[-\frac{D_{k}(\bar{\mathbf{f}}_{i,k}, \bar{\mathbf{f}}_{j,k})}{2\gamma_{k,i}^{2}}\right],$$





Filter weights $\hat{\mathbf{c}}_{i} = \frac{\sum_{j \in \mathcal{N}(i)} d_{i,j} \bar{\mathbf{c}}_{j}}{\sum_{j \in \mathcal{N}(i)} d_{i,j}}$ Pixel neighborhood



Filter weights

For cross Bilateral filters:



Sen and Darabi [2012]





For cross Bilateral filters:

$$d_{i,j} = \exp\left[-\frac{\|ar{\mathbf{p}}_i - ar{\mathbf{p}}_j\|^2}{2lpha_i^2}
ight] imes \exp\left[-rac{M_i \|ar{\mathbf{p}}_i - ar{\mathbf{p}}_j\|^2}{2lpha_i^2}
ight] imes \exp\left[-rac{D_k(ar{\mathbf{f}}_{i,k}, ar{\mathbf{f}}_{j,k})}{2\gamma_{k,i}^2}
ight]$$



Filter weights

 $\left[-\frac{D(\bar{\mathbf{c}}_i, \bar{\mathbf{c}}_j)}{2\beta_i^2}\right]$

,

Sen and Darabi [2012]



82



For cross Bilateral filters:

$$d_{i,j} = \exp\left[-\frac{\|\bar{\mathbf{p}}_i - \bar{\mathbf{p}}_j\|^2}{2\alpha_i^2}\right] \times \exp\left[-\frac{K}{2\gamma_{k,i}^2}\right] \times \exp\left[-\frac{D_k(\bar{\mathbf{f}}_{i,k}, \bar{\mathbf{f}}_{j,k})}{2\gamma_{k,i}^2}\right]$$



(a) Screen position

(b) Random parameters

(c) World space coords.

(d) Surface normals



Filter weights

 $-\frac{D(\bar{\mathbf{c}}_i, \bar{\mathbf{c}}_j)}{2\beta^2}$

,

(e) Texture value 83

(f) Sample color





For cross Bilateral filters:





(a) Screen position

(b) Random parameters

(c) World space coords.

(d) Surface normals



Realistic Image Synthesis SS2020

Filter weights

(e) Texture value 83

(f) Sample color









(a) Screen position

(b) Random parameters

(c) World space coords.

(d) Surface normals



Realistic Image Synthesis SS2020

(e) Texture value 83

(f) Sample color



















Neural Network Approach

- Feed-forward Neural network
- Best part: We can learn weights in a training phase
- Back propagation: Important for training weights
- For Back propagation, the Loss function should be differentiable and
- all the intermediate functionals should be differentiable.









One Hidden-layer model

Relative Mean Square Error:

$$E_{i} = \frac{n}{2} \sum_{q \in \{r,g,b\}} \frac{(\hat{c}_{i,q} - c_{i,q})^{2}}{c_{i,q}^{2} + \varepsilon}$$



Realistic Image Synthesis SS2020





One Hidden-layer model

Relative Mean Square Error:





$$\frac{\partial E_i}{\partial w_{t,s}^l} = \sum_{m=1}^M \left[\sum_{q \in \{r,g,b\}} \left[\frac{\partial E_{i,q}}{\partial \hat{c}_{i,q}} \frac{\partial \hat{c}_{i,q}}{\partial \theta_{m,i}} \right] \frac{\partial \theta_{m,q}}{\partial w_{t,q}^l} \right]$$

$$\frac{\partial E_i}{\partial \hat{c}_{i,q}} = ???$$



Realistic Image Synthesis SS2020



One Hidden-layer model

Relative Mean Square Error:





Realistic Image Synthesis SS2020

$$\frac{\partial E_i}{\partial w_{t,s}^l} = \sum_{m=1}^M \left[\sum_{q \in \{r,g,b\}} \left[\frac{\partial E_{i,q}}{\partial \hat{c}_{i,q}} \frac{\partial \hat{c}_{i,q}}{\partial \theta_{m,i}} \right] \frac{\partial \theta_{m,q}}{\partial w_{t,q}^l} \right]$$

$$\frac{\partial E_i}{\partial \hat{c}_{i,q}} = n \frac{\hat{c}_{i,q} - c_{i,q}}{c_{i,q}^2 + \epsilon}$$





Results



Realistic Image Synthesis SS2020







Our result with a cross-bilateral filter (4 spp)









Our result with a non-local means filter (4 spp)







Introduction to CNNs

Introduction to CNNs

Kernel Predicting Denoising

Introduction to CNNs

Kernel Predicting Denoising

Sample-based MC Denoising



No zero padding



Convolution





Convolution



No zero padding









No zero padding



Stride-1 Convolution











No zero padding



Realistic Image Synthesis SS2020

Stride-1 Convolution

94




No zero padding



Stride-1 Convolution

-2	4	3	-2

95









No zero padding



Stride-1 Convolution

95

Realistic Image Synthesis SS2020







Stride-1 Convolution

-2	4	3	-2
-1	-2	3	-3

0

96







No zero padding



Realistic Image Synthesis SS2020

Stride-1 Convolution



0









Stride-2 Convolution



97









Stride-2 Convolution





Zero Padding and Strides

1D image to illustrate the strides and zero padding





98





Zero Padding and Strides

1D image to illustrate the strides and zero padding





98





1D image to illustrate the strides and zero padding





Strides





Max Pooling / Down Sampling

1	1	2	4
5	6	7	8
3	2	1	0
1	2	3	4



6	8
3	4

100





Overview on Convolutional Neural Networks



Image Courtesy: Mathworks (online tutorial)



Multi-layer Perceptron vs. CNNs



Realistic Image Synthesis SS2020



Multi-layer Perceptron vs. CNNs

Multi-layer perceptron



All nodes are fully connected in all layers

In theory, should be able to achieve good quality results in small number of layers.

Number of weights to be learnt are very high



CNNs



Weights are shared across layers

Requires significant number of layers to capture all the features (e.g. Deep CNNs)

Relatively small number of weights required









Introduction to CNNs

Kernel-Predicting Denoising

Kernel-Predicting Networks for Denoising Monte-Carlo Renderings



Realistic Image Synthesis SS2020

Bako et al. [2017]



Kernel was pre-selected to be joint bilateral filter



106





Kernel was pre-selected to be joint bilateral filter

- Unable to explicitly capture all details



106





- Kernel was pre-selected to be joint bilateral filter
 - Unable to explicitly capture all details
 - lacked flexibility to handle wide range of MC noise in production scenes







- Kernel was pre-selected to be joint bilateral filter
 - Unable to explicitly capture all details
 - lacked flexibility to handle wide range of MC noise in production scenes

Fixed







- Kernel was pre-selected to be joint bilateral filter
 - Unable to explicitly capture all details
 - lacked flexibility to handle wide range of MC noise in production scenes

Fixed

- can cause unstable weights causing bright ringing and color artifacts









- Kernel was pre-selected to be joint bilateral filter
 - Unable to explicitly capture all details
 - lacked flexibility to handle wide range of MC noise in production scenes

Fixed

- can cause unstable weights causing bright ringing and color artifacts

Too many parameters to optimize









Requirements



- The function must be flexible to capture complex relationship between input data and reference colors over wide range of scenarios.
- Choice of loss function is crucial. Should capture perceptual aspects of the scene.

To avoid overfitting, large dataset required





Denoising a raw, noisy color buffer causes overblurring



Using a Vanilla CNN

108





Denoising a raw, noisy color buffer causes overblurring

- difficulty in distinguishing scene details and MC noise



Using a Vanilla CNN

108





Denoising a raw, noisy color buffer causes overblurring - difficulty in distinguishing scene details and MC noise

High dynamic range



Realistic Image Synthesis SS2020

Using a Vanilla CNN





Denoising a raw, noisy color buffer causes overblurring

- difficulty in distinguishing scene details and MC noise

High dynamic range



Using a Vanilla CNN

- can cause unstable weights causing bright ringing and color artifacts



Realistic Image Synthesis SS2020



Vanilla CNN



Ours

Input (32 spp)



Realistic Image Synthesis SS2020

Vanilla CNN

Ours

Ref. (1K spp)





θ

Denoising Model $\widehat{\boldsymbol{\theta}}_p = \operatorname*{argmin}_{\boldsymbol{\rho}} \ell(\overline{\mathbf{c}}_p, g(\mathbf{X}_p; \boldsymbol{\theta}))$ Denoised function with parameters Reference image $\widehat{\mathbf{c}}_p = g(\mathbf{X}_p; \widehat{\boldsymbol{\theta}}_p)$ $\ell(\overline{\mathbf{c}}, \widehat{\mathbf{c}})$

Denoised value



Realistic Image Synthesis SS2020

Loss function

Computational Model

$$\widehat{\boldsymbol{\theta}}_{p} = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \sum_{q \in \mathcal{N}(p)} \left(\mathbf{c}_{q} - \boldsymbol{\theta}^{\top} \boldsymbol{\phi}(\mathbf{x}_{q}) \right)^{2} \boldsymbol{\omega}(\mathbf{x}_{p}, \mathbf{x}_{q})$$
Neighborhood

$$\widehat{\mathbf{c}}_p = g(\mathbf{X}_p; \widehat{\boldsymbol{\theta}}_p)$$

Denoised value

$$\widehat{\mathbf{c}}_p = \widehat{\boldsymbol{\theta}}_p^\top \phi(\mathbf{x}_p)$$

Final denoised value



$$\phi: \mathbb{R}^{3+\bar{D}} \to \mathbb{R}^{\bar{M}}$$

 $\omega(\mathbf{x}_p, \mathbf{x}_q)$ Kernel weights



Realistic Image Synthesis SS2020



Direct Prediction Network

Direct prediction convolution network: outputs denoised image

$\widehat{\mathbf{c}}_p = g_{\text{direct}}$



$$\mathbf{z}_{t}(\mathbf{X}_{p};\boldsymbol{\theta}) = \mathbf{z}_{p}^{L}$$





Direct Prediction Network

Direct prediction convolution network: outputs denoised image

$$\widehat{\mathbf{c}}_p = g_{\text{direct}}(\mathbf{X}_p; \boldsymbol{\theta}) = \mathbf{z}_p^L$$

Issues:

The constrained nature and complexity of the problem makes optimization difficult.

The magnitude and variance of stochastic gradients computed during training can be large, which slows convergence of training loss.









Kernel Prediction Network

Kernel prediction convolution network: outputs learned kernel weights

 $w_{pq} = \frac{1}{\sum_{q' \in Q}}$

Denoised color values:

 $\widehat{\mathbf{c}}_p = g_{\text{weighter}}$



$$\begin{split} \exp([\mathbf{z}_{p}^{L}]_{q}) & 0 \leq w_{pq} \leq 1 \\ \in \mathcal{N}(p) \exp([\mathbf{z}_{p}^{L}]_{q'}) & \text{Softmax activation to enform on the set of the set$$

$$_{\mathrm{ed}}(\mathbf{X}_{p};\boldsymbol{\theta}) = \sum_{q \in \mathcal{N}(p)} \mathbf{c}_{q} w_{pq}$$

Realistic Image Synthesis SS2020



Kernel Prediction Network

$$w_{pq} = \frac{\exp([\mathbf{z}_{p}^{L}]_{q})}{\sum_{q' \in \mathcal{N}(p)} \exp([\mathbf{z}_{p}^{L}]_{q'})}$$
$$0 \le w_{pq} \le 1$$

Final color estimate always lies within the convex hull of the respective neighborhood (avoid color shifts).

Ensures well-behaved gradients of the error w.r.t the kernel weights



$$\widehat{\mathbf{c}}_p = g_{\text{weighted}}(\mathbf{X}_p; \boldsymbol{\theta}) = \sum_{q \in \mathcal{N}(p)} \mathbf{c}_q w_p$$

Realistic Image Synthesis SS2020



Proposed Architecture











Diffuse/Specular components

Each component is denoised separately

Diffuse components are well-behaved and typically has small ranges

albedo is factored out to allow large rate

Specular components are challenging due to high dynamic ranges: uses logarithmic transform

 $\mathbf{c}_{\text{specular}} = \mathbf{1}$





ange kernels
$$\tilde{\mathbf{c}}_{diffuse} = \mathbf{c}_{diffuse} \oslash (\mathbf{f}_{albedo} + \epsilon)$$

$$og(1 + c_{specular})$$

Training Dataset: 600 frames
















8-hidden layers used with 100 kernels of 5x5 in each layer for each network

For KPCN (kernel-predicting network), output kernel size used = 21

Weights for 128 app and 32 spp networks were initialized using Xavier method

Diffuse and specular components were independently trained with L1 loss metric



Training







Learning rate of DPCN vs. KPCN





Realistic Image Synthesis SS2020









Realistic Image Synthesis SS2020

NFOR (log)









Realistic Image Synthesis SS2020

Input (32 spp)



Input (32 spp)



122





Input (32 spp)





w/o Decomposition, w/o Albedo divide















w/o Decomposition, w/o Albedo divide

w/ Decomposition, w/o Albedo divide





w/o Decomposition, w/ Albedo divide

w/ Decomposition, w/ Albedo divide

Ref. (2K spp)

Realistic Image Synthesis SS2020





Ours

Input (32 spp)



Ours

Input (32 spp)



Realistic Image Synthesis SS2020

NFOR (log)

Ours

Ref. (8K spp)

Also works on Piper short movie frames





Interactive Reconstruction of Monte Carlo Sequences



Realistic Image Synthesis SS2020

Chaitanya et al. [2017]













Motivation: Interactive Reconstruction

Limited to a few rays per pixel @ 1080p @ 30Hz

Never enough to reconstruct an image

Deep learning approach for interactive graphics













Motivation: Interactive Reconstruction

Limited to a few rays per pixel @ 1080p @ 30Hz

Never enough to reconstruct an image

Deep learning approach for interactive graphics









Realistic Image Synthesis SS2020



Handle generic effects:

- Soft shadows
- Diffuse and specular reflections
- Global illumination (one-bounce)
- No Motion blur or depth of field



Problem Statement

129





















System setup: Path tracing



System setup: Path tracing



Rasterize primary hits in G-buffers

Path-tracing from the primary paths

- 1 ray for direct shadows
- 2 rays for indirect (sample + connect)

1 direct + 1 indirect path (spp)

Denoising Autoencoder (DAE)



Train auto encoders to reconstruct image from 1spp



Recurrent Autoencoder [Chaitanya et al. 2017]



Fig. 2. Architecture of our recurrent autoencoder. The input is 7 scalar values per pixel (noisy RGB, normal vector, depth, roughness). Each encoder stage has a convolution and 2×2 max pooling. A decoder stage applies a 2×2 nearest neighbor upsampling, concatenates the per-pixel feature maps from a skip connection (the spatial resolutions agree), and applies two sets of convolution and pooling. All convolutions have a 3×3 -pixel spatial support. On the right we visualize the internal structure of the recurrent RCNN connections. I is the new input and h refers to the hidden, recurrent state that persists between animation frames.





Encoder and decoder stages for dimensionality reduction





Realistic Image Synthesis SS2020





Encoder and decoder stages for dimensionality reduction





Decoder



Encoder and decoder stages for dimensionality reduction







Encoder and decoder stages for dimensionality reduction



Skip connections to reintroduce lost information



136



Auxillary Features





Untextured color





View space normals

Linearize depth







Training sequences





SponzaDiffuse

SponzaGlossy



Classroom

Training sequences





SponzaDiffuse

SponzaGlossy



Classroom





DAE 1spp approx. 70 ms + approx. 60 ms





Reference 1024 spp approx. 240 ms





Recurrent Denoising Autoencoder

Feedback loops to retain important information after every encoding stage







Recurrent Denoising Autoencoder

Feedback loops to retain important information after every encoding stage







Recurrent Neural Networks vs. Simple Feed-Forward NN



Recurrent Neural Network



Realistic Image Synthesis SS2020

Feed-Forward Neural Network







An unrolled recurrent neural network.



Realistic Image Synthesis SS2020

Source link





Fully convolutional blocks to support arbitrary image resolution

6 RNN blocks, one per pool layer in the encoder



Design:

- 1 conv layer (3x3) for current features
- 2 conv layers (3x3) for previous features








Realistic Image Synthesis SS2020







Realistic Image Synthesis SS2020







148





CNNs, fixed input, fixed output

one to one





148





CNNs, fixed input, fixed output

one to one





149





CNNs, fixed input, fixed output





149





CNNs, fixed input, fixed output



e.g., image captioning takes an image as input and outputs a sentence of words

Sequence output









Sequence output



Realistic Image Synthesis SS2020

e.g., to know the sentiments of a sentence







Sequence output



Sequence input, Sequence output. e.g. Machine translation







Sequence output



Sequence input, Sequence output. e.g. Machine translation

We video classification where i. Ф

152



Input is a sequence of 7 frames

128x128 random image crop per sequence

Play the sequence forward/backward

Each frame advance the camera or random seed



Training



153





Spatial Loss to emphasize more the dark regions





Realistic Image Synthesis SS2020

Loss Functions









Realistic Image Synthesis SS2020

Loss Functions

Temporal loss









Realistic Image Synthesis SS2020

Loss Functions

Imporal lossHigh frequency error norm left
for stable edges
$$V_{i}$$
 $\left(\left| \frac{\partial P_{i}}{\partial t} - \frac{\partial T_{i}}{\partial t} \right| \right)$ $L_{g} = \frac{1}{N} \sum_{i}^{N} |\nabla P_{i} - \nabla T_{i}|$

156

OSS



Final Loss is a weighted averaged of above losses





Realistic Image Synthesis SS2020

Loss Functions

Emporal loss
High frequency error norm less
for stable edges

$$L_g = \frac{1}{N} \sum_{i}^{N} |\nabla P_i - \nabla T_i|$$

$$+w_gL_g+w_tL_t$$

OSS

Training Loss depends on **Auxiliary Features**



Training loss



Auxiliary Features

Epochs









Temporal Stability



Realistic Image Synthesis SS2020



Recurrent autoencoder with temporal AA

Recurrent autoencoder

Autoencoder with skips



Recurrent autoencoder with temporal AA

Recurrent autoencoder

Autoencoder with skips





1





1



Recurrent autoencoder





Recurrent autoencoder





Introduction to CNNs

Kernel Predicting Denoising

Sample-based MC Denoising (next lecture)

Acknowledgments



Realistic Image Synthesis SS2020

Thanks to Chaitanya and colleagues for making their slides publicly available.



