Nonlinearly Weighted First-order Regression

for Denoising MC Renderings

Benedikt Bitterli
Fabrice Rousselle
Bochang Moon
José A. Iglesias-Guitián
Kenny Mitchell
David Adler
Wojciech Jarosz
Jan Novák

Disney Research
Walt Disney Animation Studios
Edinburgh Napier University
Dartmouth

Goal: study the state of the art

Input

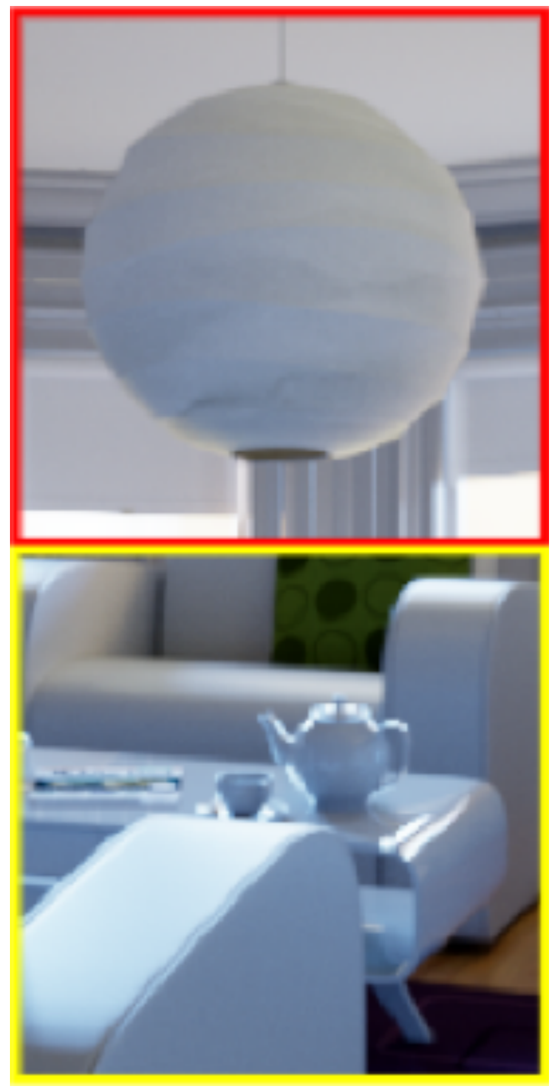

Reference

Goal: study the state of the art

RHF
Delbracio et al.
[2014]

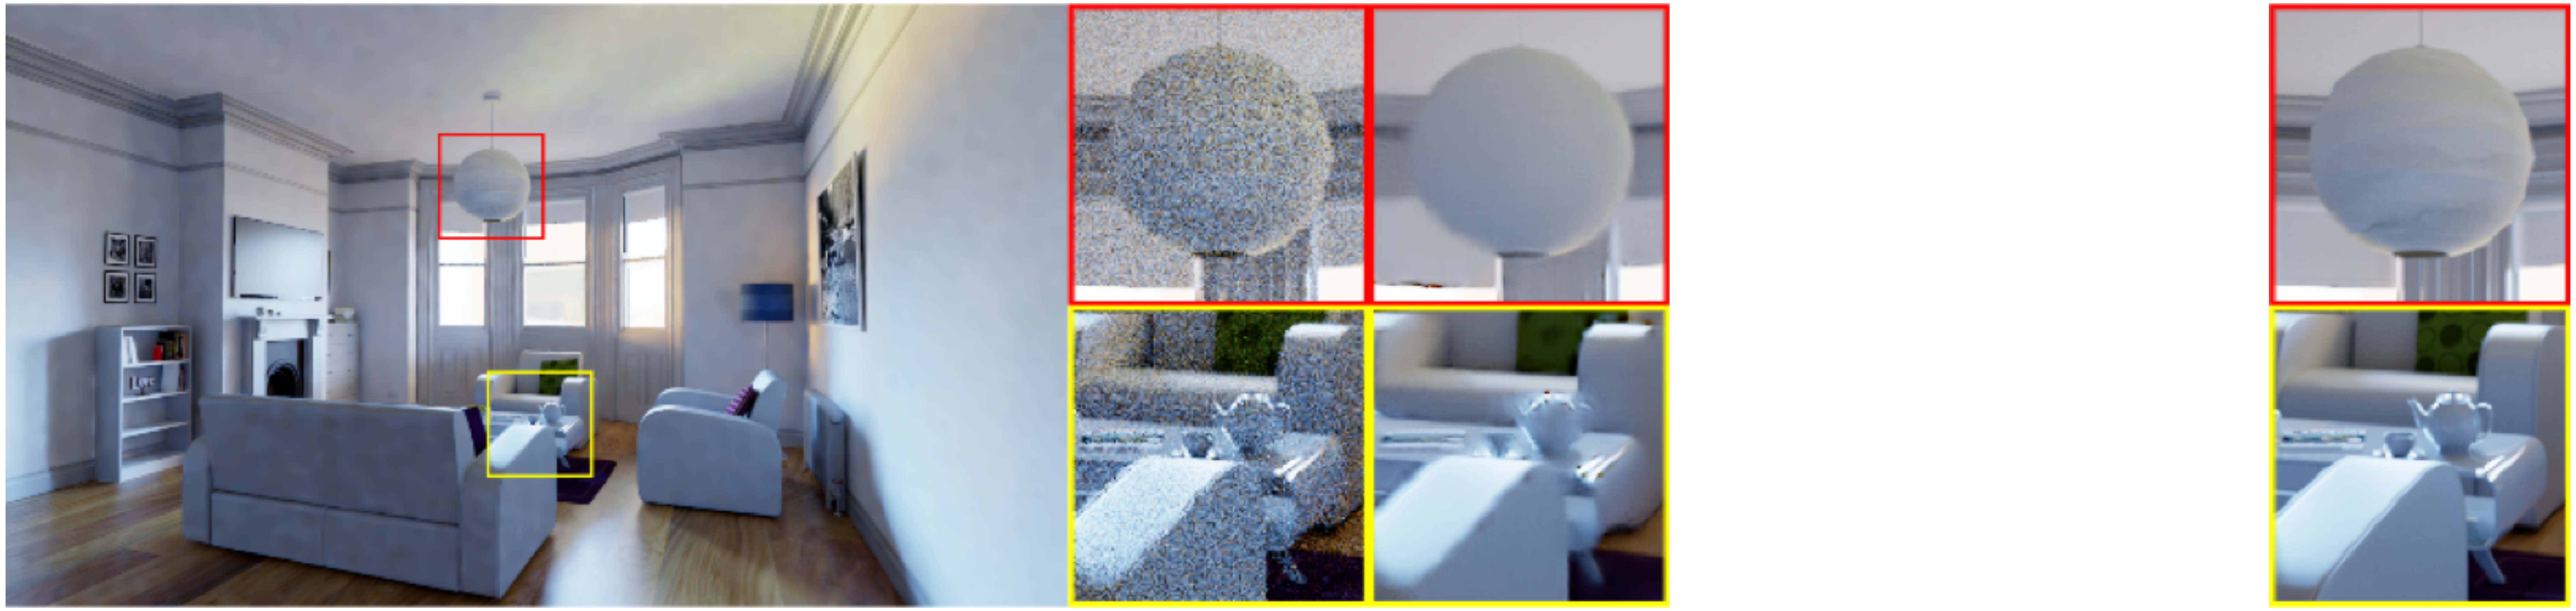

Input RHF Reference

Goal: study the state of the art

RHF
Delbracio et al.
[2014]

RDFC
Rousselle et al. [2013]

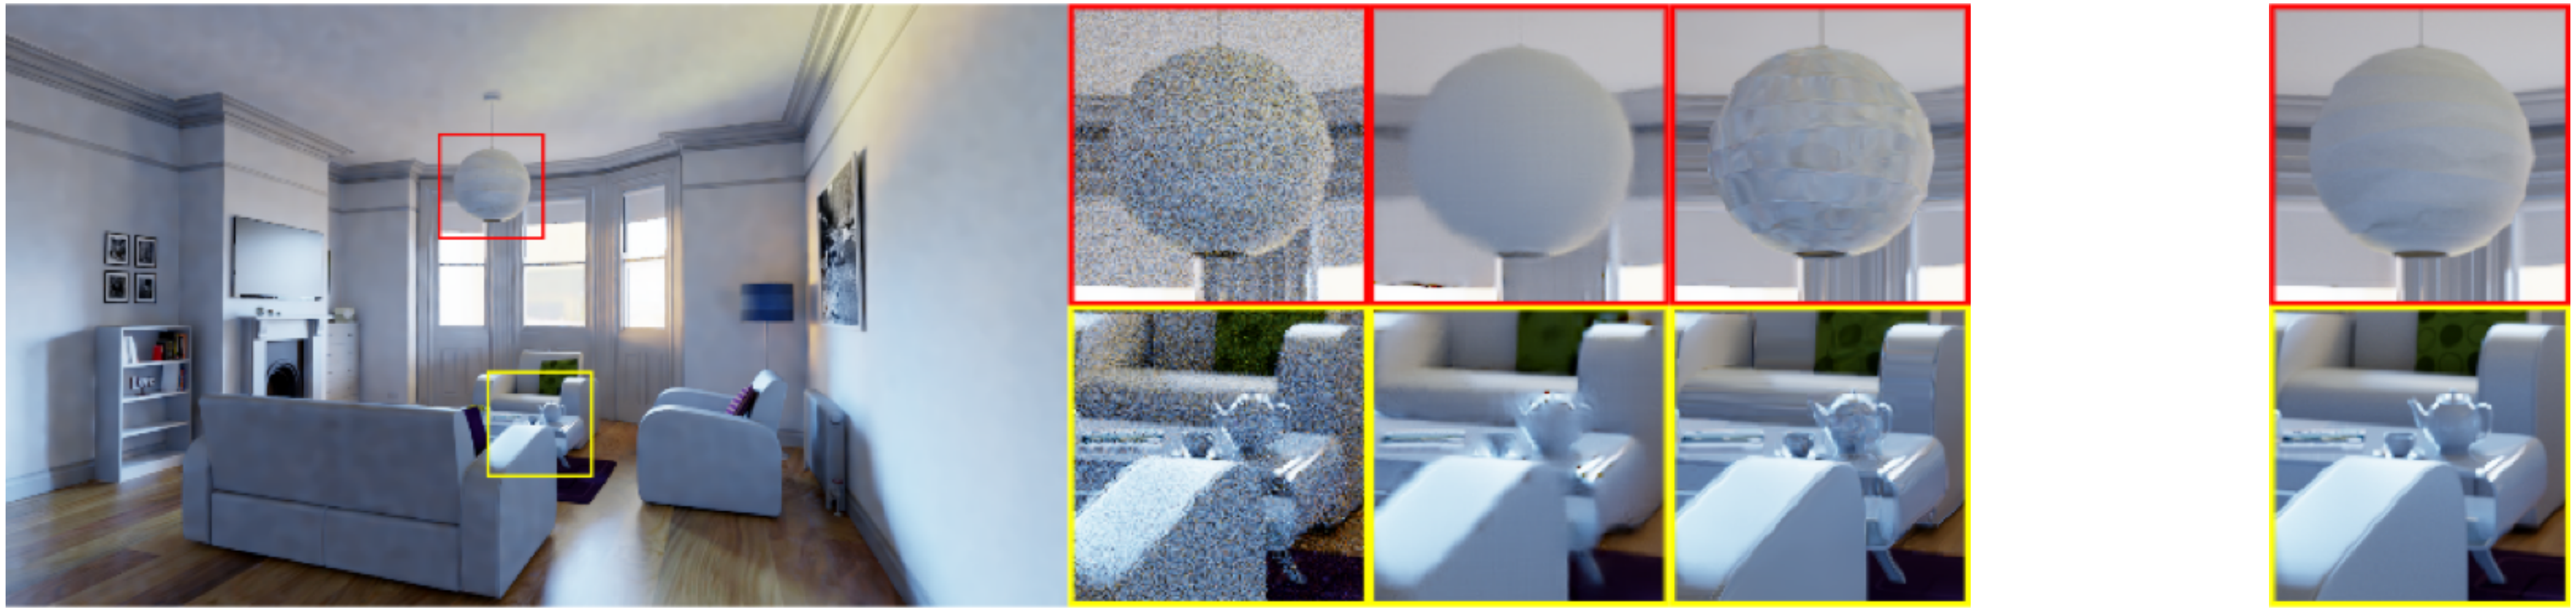

Input RHF RDFC Reference

Disney Research

Goal: study the state of the art

RHF
Delbracio et al. [2014]

RDFC
Rousselle et al. [2013]

WLR
Moon et al. [2014]

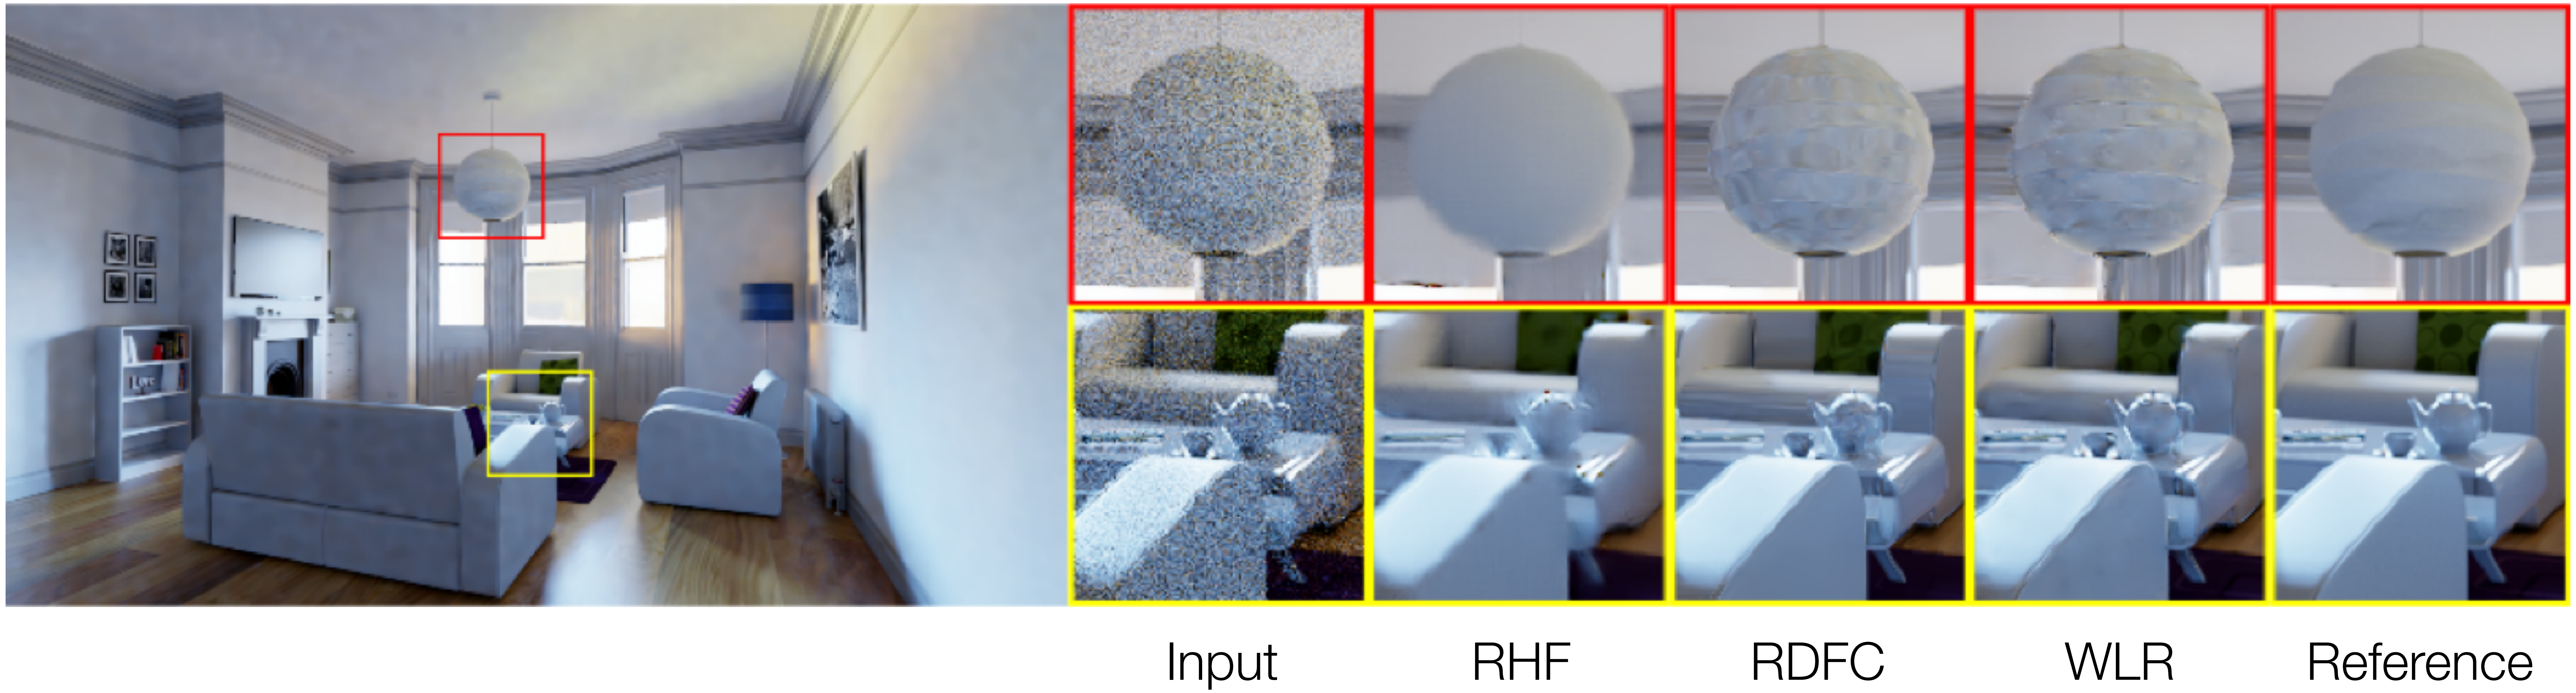

Evaluation Criteria

Essential

- Effectiveness – significant noise reduction
- Predictability – no manual tuning
- Stability – flicker-free results

Evaluation Criteria

Essential

Effectiveness – significant noise reduction

Predictability – no manual tuning

Stability – flicker-free results

Required for Production

Ease of adoption – no intrusive or expensive changes

Speed – significantly faster than rendering

Memory – significantly less than rendering

No Silver Bullet

RHF

Too blurry, needs manual tuning

Input RHF Reference

Disney Research

No Silver Bullet

RDFC

Banding artifacts

Disney Research

No Silver Bullet

WLR

- Residual noise

Input WLR Reference

Disney Research

Methodology

Theoretical analysis – how are previous filters related

Comparative analysis – how do previous filters perform

New design – how can we combine the strengths of previous filters

Filtering Framework

256’000 spp

Filtering Framework

256 spp

Filtering Framework

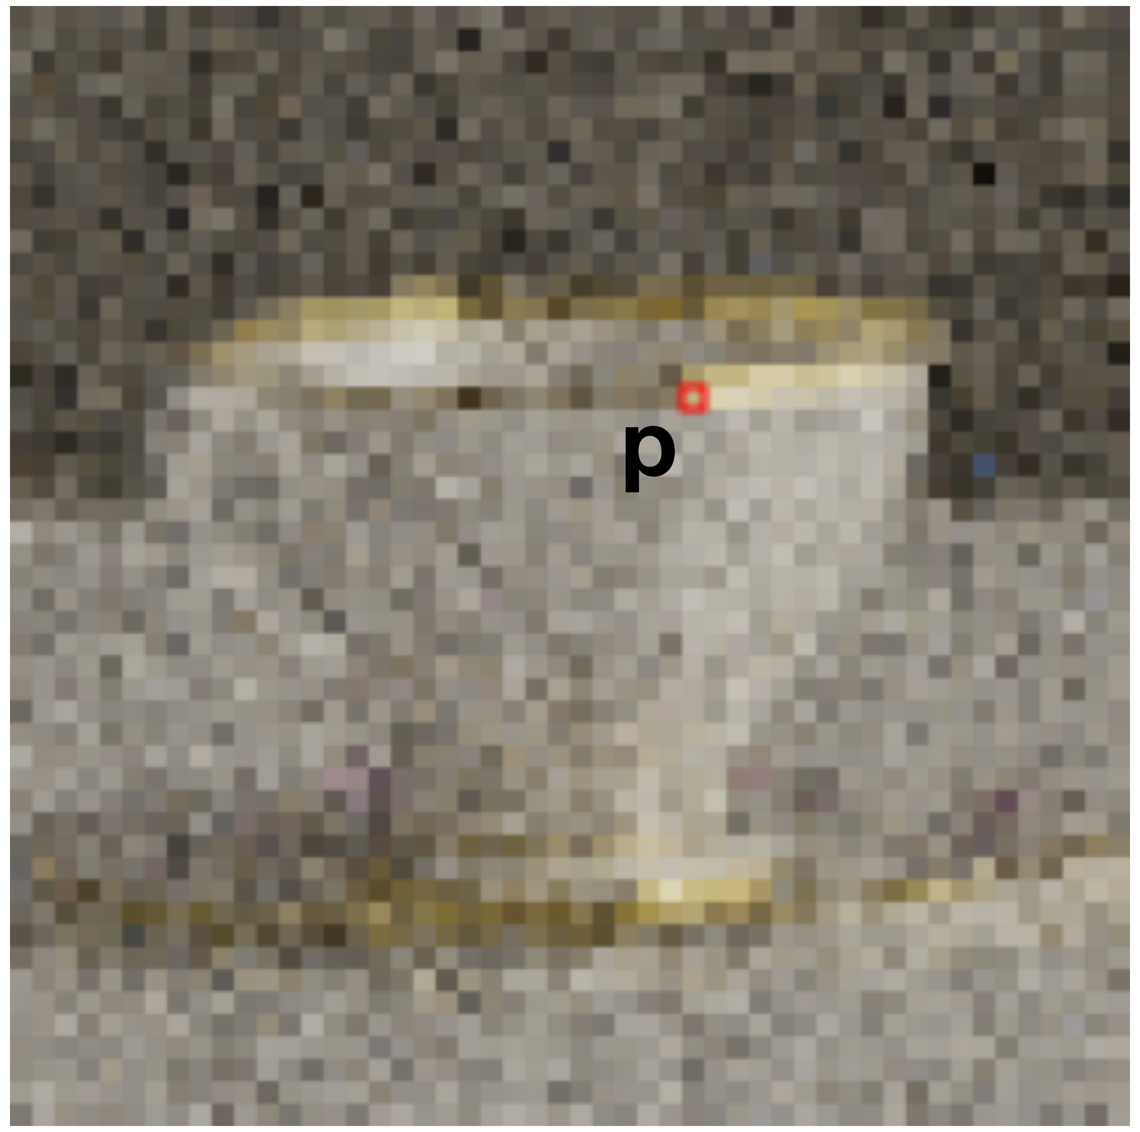

256 spp

Filtering Framework

All filters perform a regression to minimize

$$\mathrm{Error}(\hat{c}_p) =$$

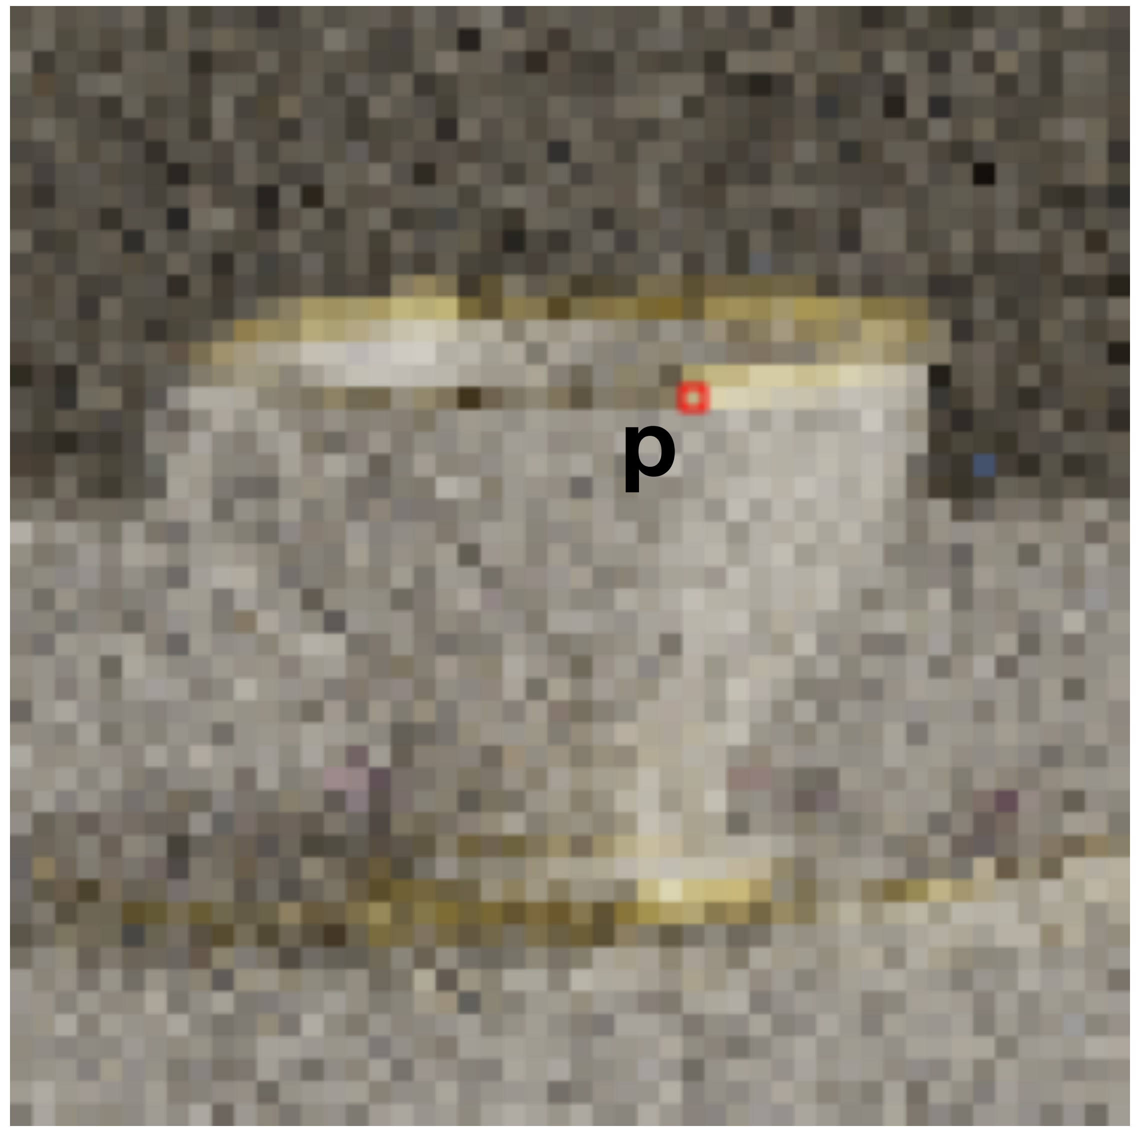

256 spp

Disney Research

Filtering Framework

All filters perform a regression to minimize

$$\text{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p}$$

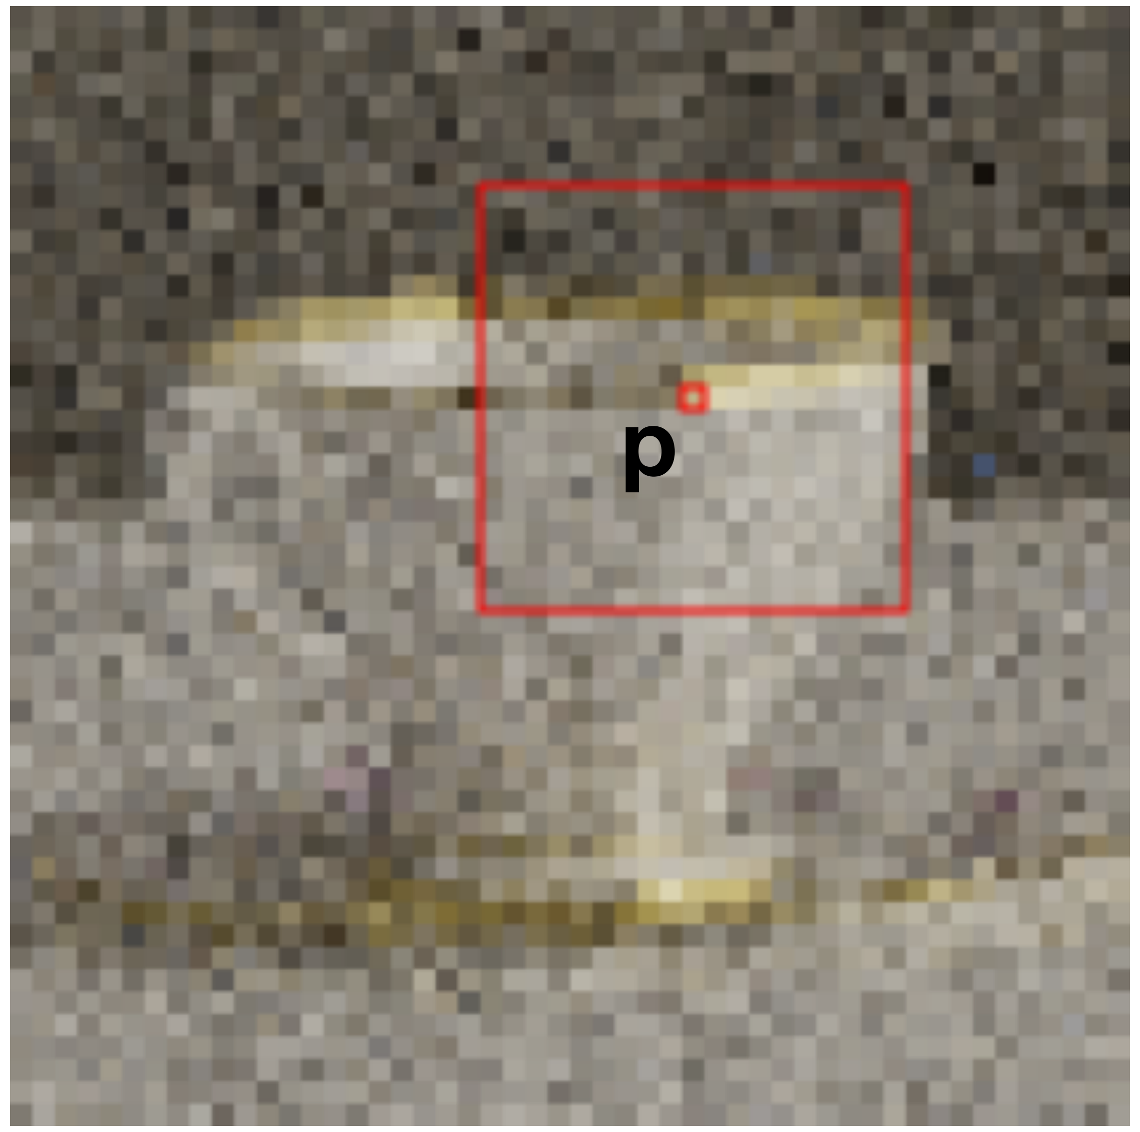

256 spp

Filtering Framework

All filters perform a regression to minimize

$$\text{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p, q))^2$$

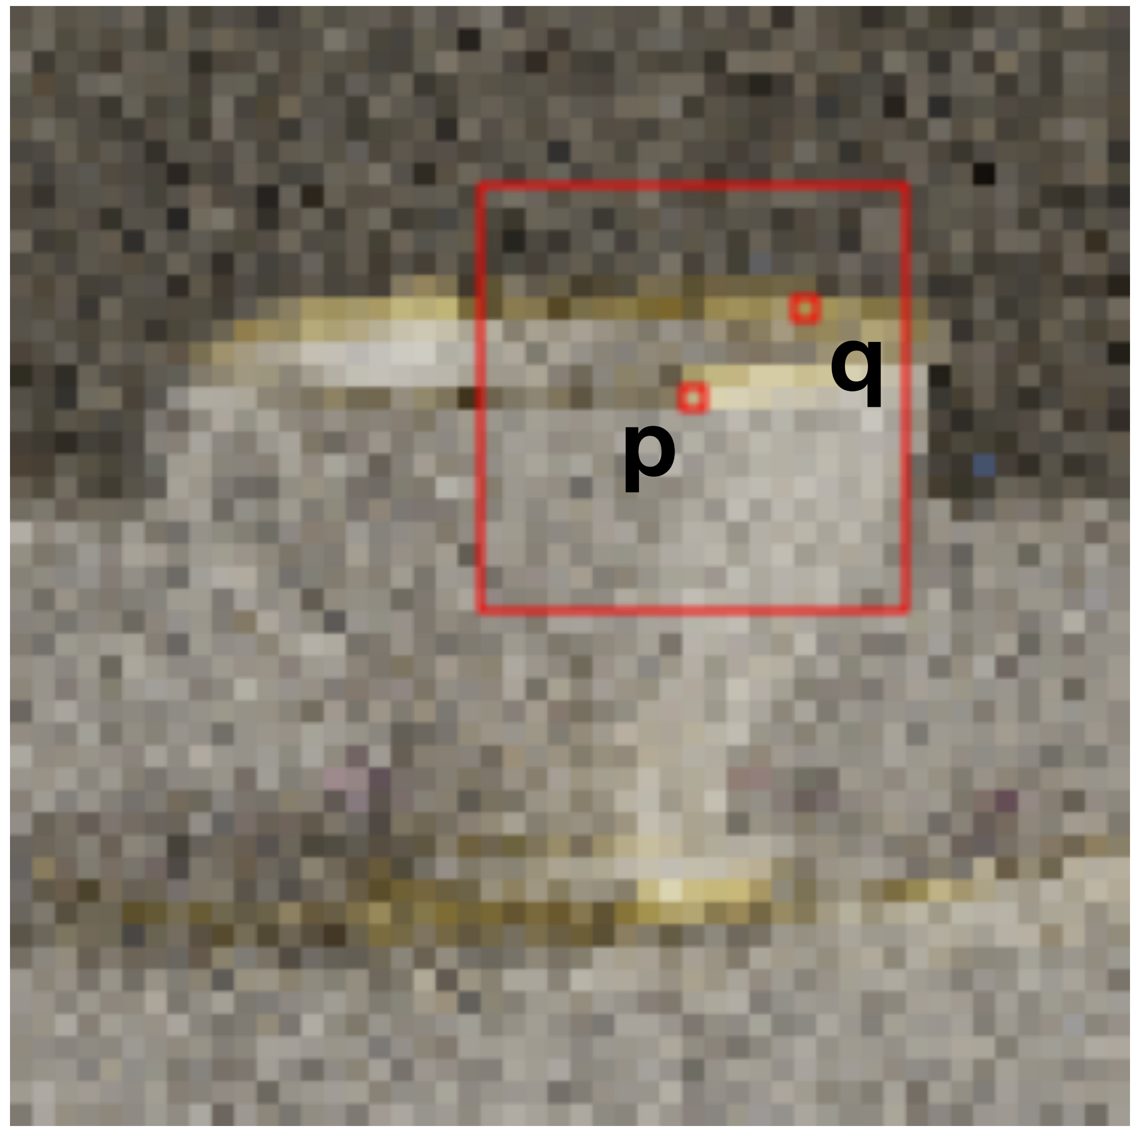

256 spp

Filtering Framework

All filters perform a regression to minimize

$$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p,q))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

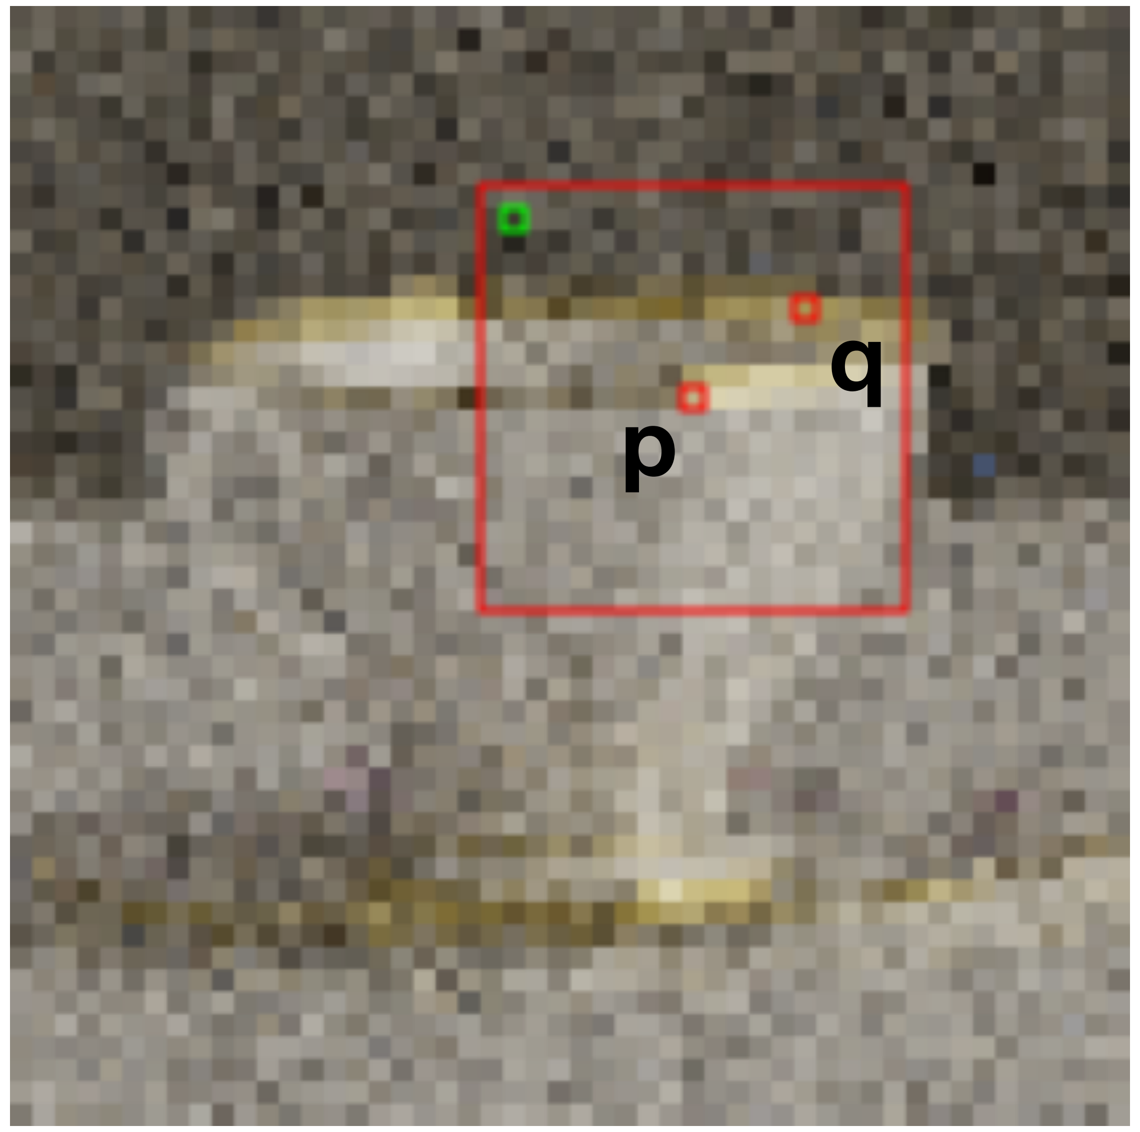

256 spp

Filtering Framework

All filters perform a regression to minimize

$$\text{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p,q))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

What ***model*** and ***weights***?

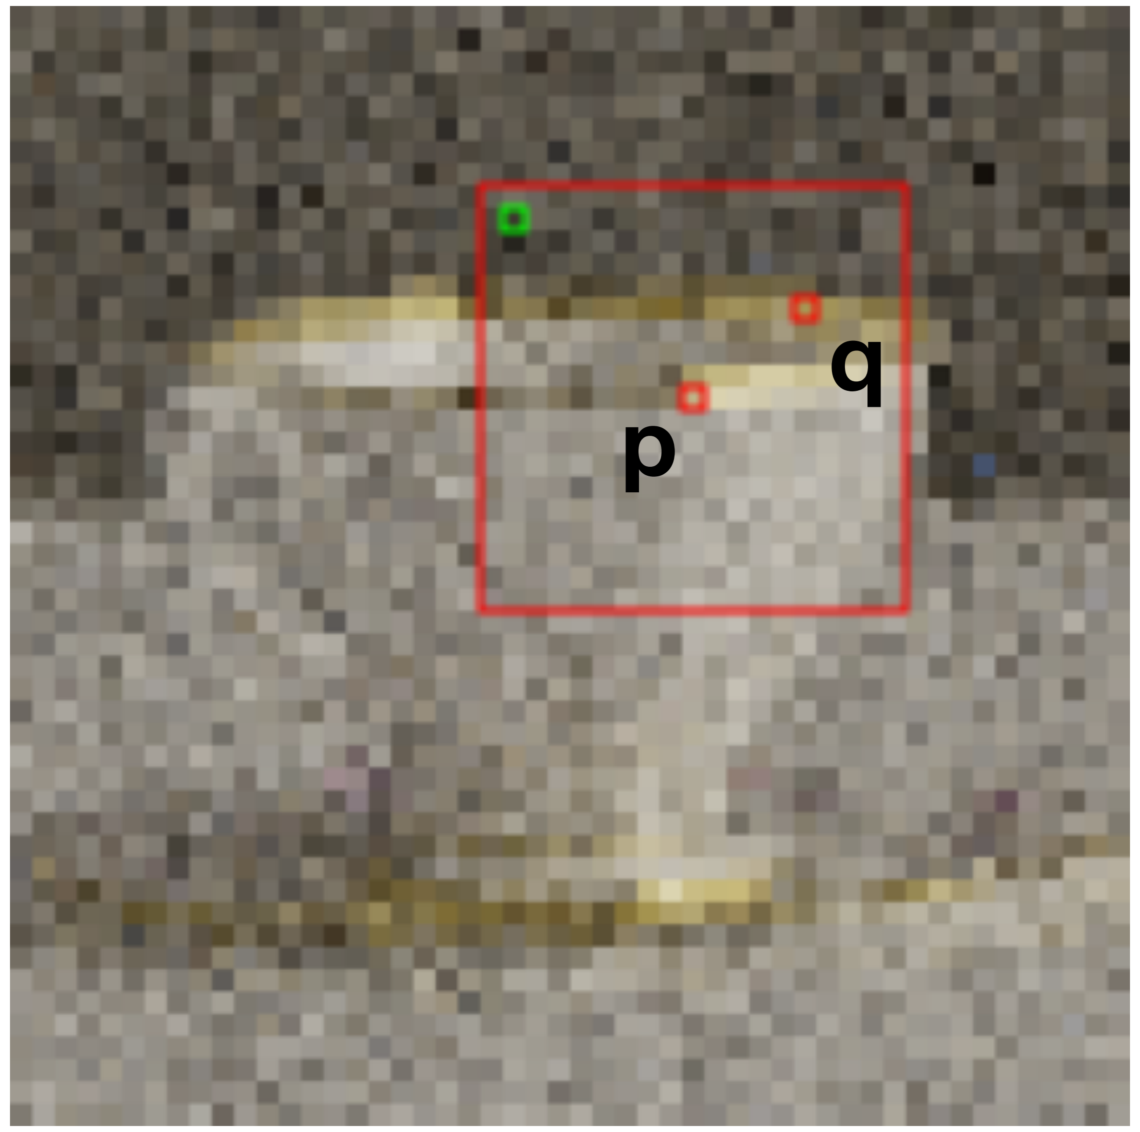

256 spp

Filtering Framework

General: $\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p,q))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$

Input – 256 spp

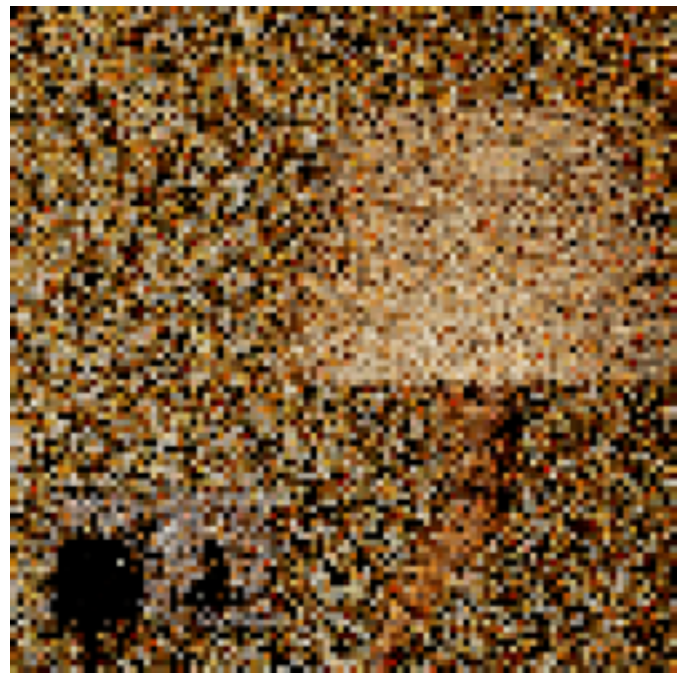

Reference

Filtering Framework — Box Filter

General: $\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p,q))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$

Input – 256 spp

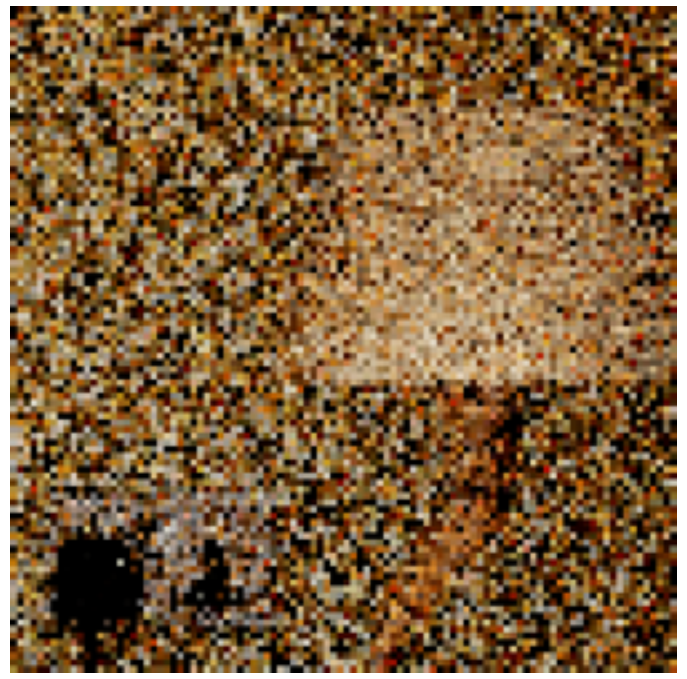

Box Filter

Reference

Filtering Framework — Box Filter

General: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p,q))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

Box: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \hat{c}_p)^2$$

Input – 256 spp

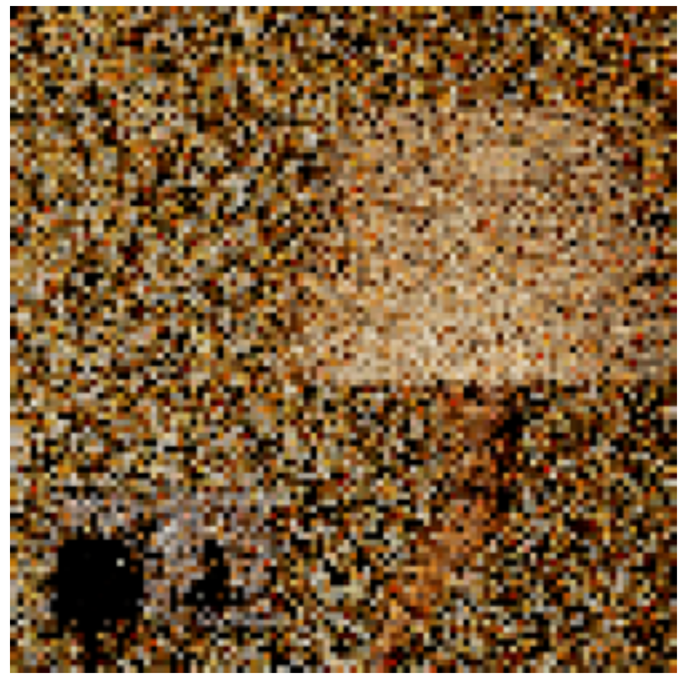

Box Filter

Reference

Filtering Framework — Box Filter

General: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p,q))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

Box: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \hat{c}_p)^2$$

Input – 256 spp

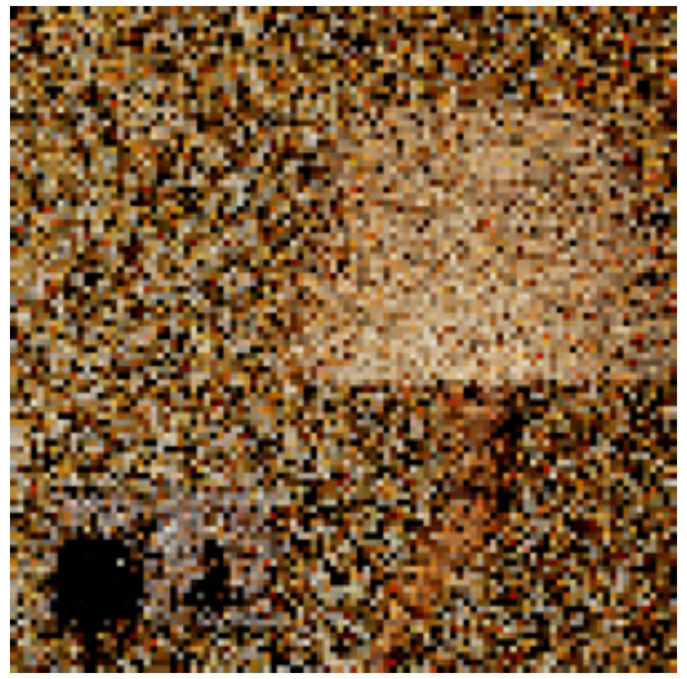

Box Filter

Weights: Uniform

Reference

Filtering Framework — Box Filter

General: $$\text{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p,q))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

Box: $$\text{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \hat{c}_p)^2$$

Input – 256 spp

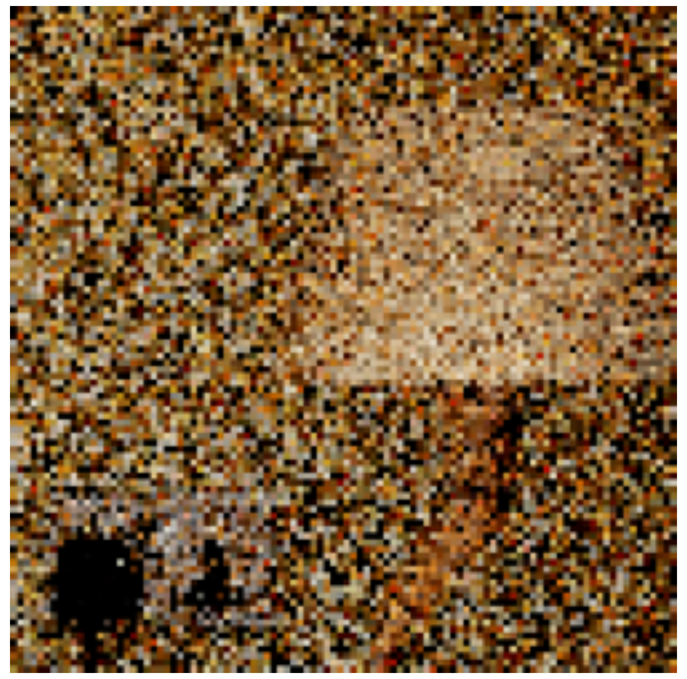

Box Filter

Reference

Weights: Uniform

Order: zero

Disney Research

Filtering Framework — NLM [Rousselle et al. 2012]

Based on the Non-Local Means filter of Buades et al. [2005]

General: $$\text{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p,q))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

NLM: $$\text{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \hat{c}_p)^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

$$\mathbf{x} = \{\text{Color}\}$$

Input – 256 spp

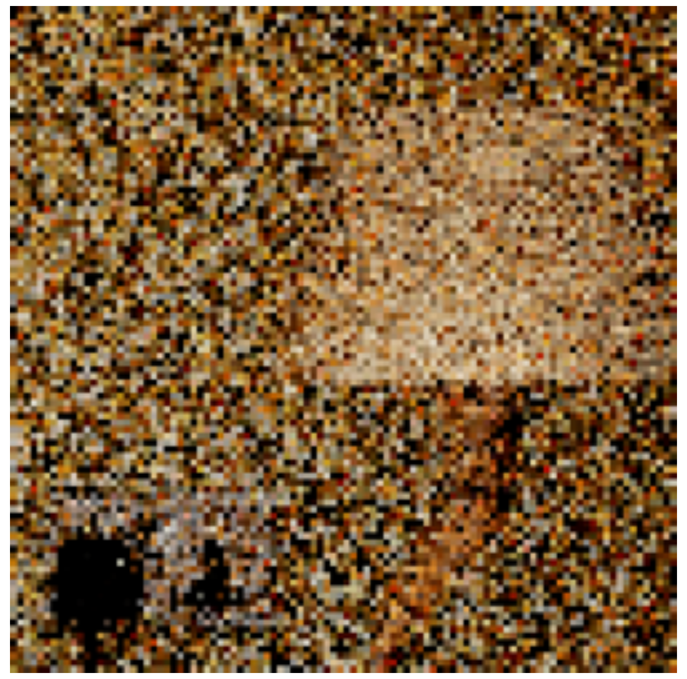

Box Filter

NLM Filter

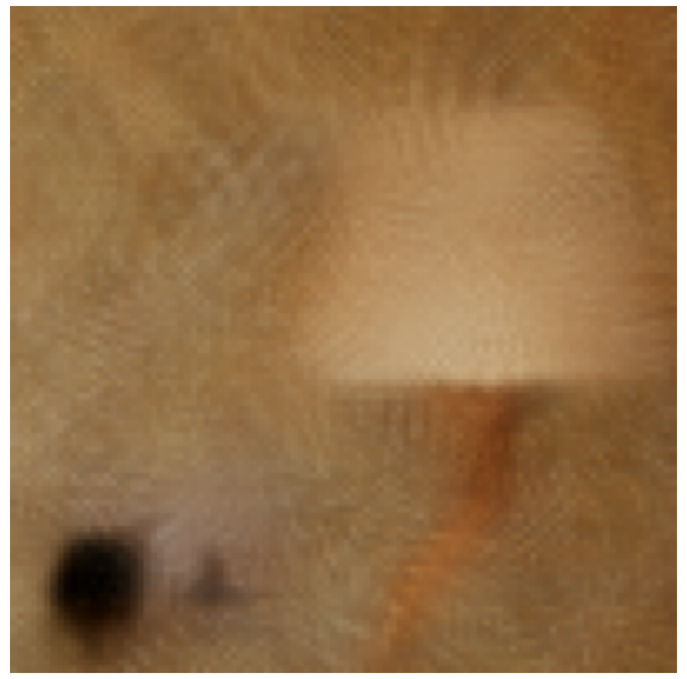

Reference

	Box Filter	NLM Filter
Weights:	Uniform	NL-Means
Order:	zero	zero

Disney Research

Filtering Framework — RDFC [Rousselle et al. 2013]

General: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p,q))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

RDFC: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \hat{c}_p)^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

$$\mathbf{x} = \{\mathrm{Color}, \mathrm{Albedo}, \mathrm{Normal}, \ldots\}$$

Input – 256 spp | Box Filter | NLM Filter

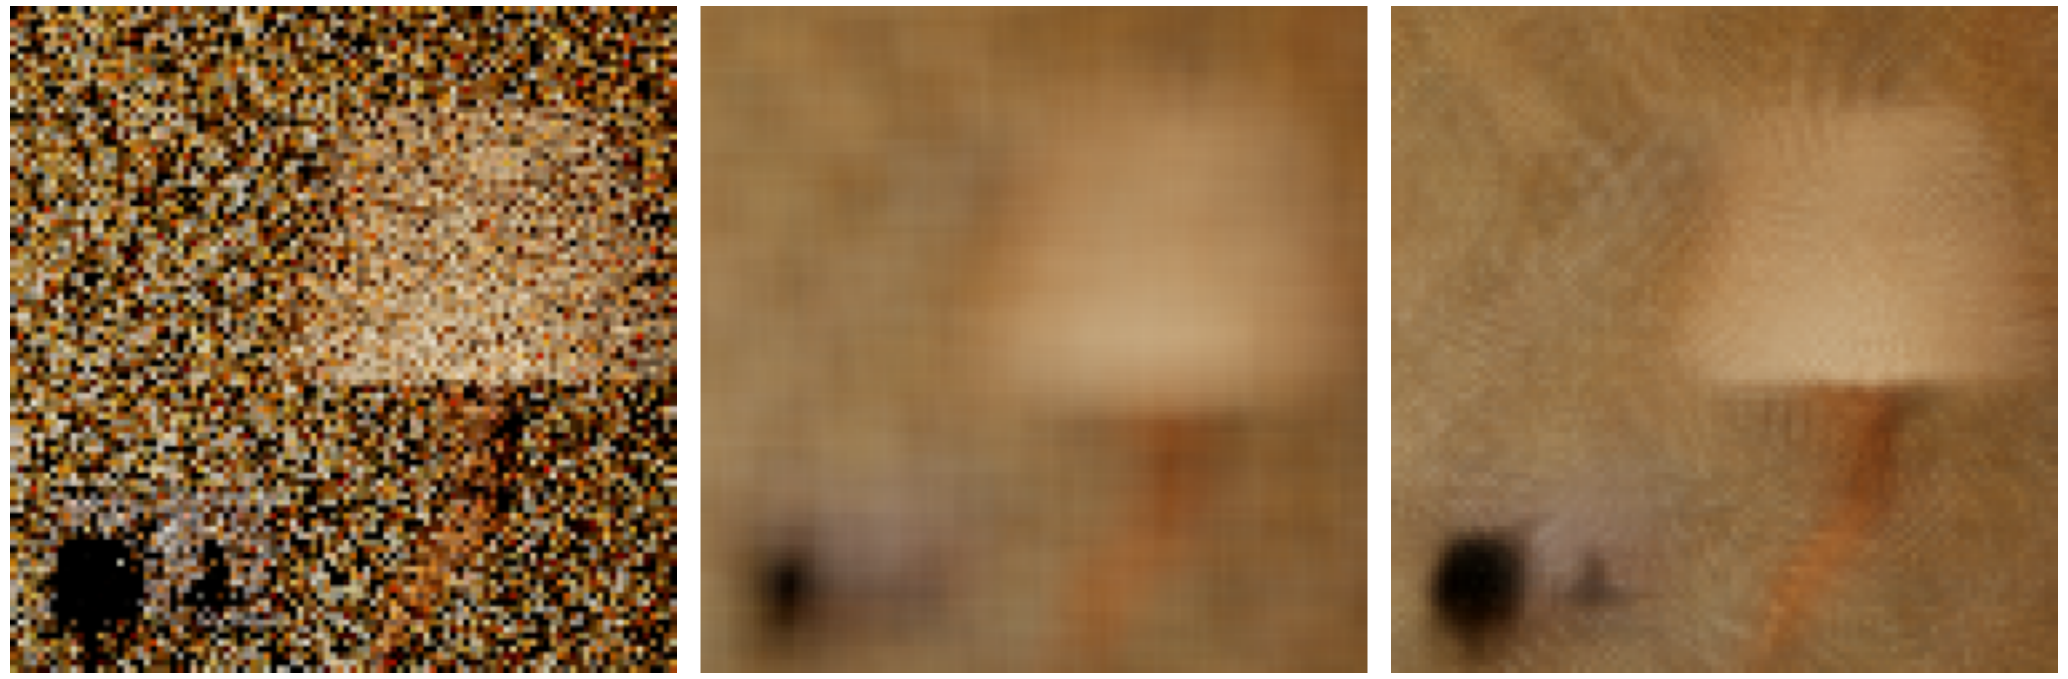

Weights: Uniform | NL-Means

Order: zero | zero

Albedo | Normal | Reference

Disney Research

Filtering Framework — RDFC [Rousselle et al. 2013]

General: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p,q))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

RDFC: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \hat{c}_p)^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

$$\mathbf{x} = \{\text{Color, Albedo, Normal, ...}\}$$

Input – 256 spp | Box Filter | NLM Filter | RDFC Filter | Albedo | Normal | Reference

	Box Filter	NLM Filter	RDFC Filter
Weights:	Uniform	NL-Means	Joint NL-Means
Order:	zero	zero	zero

Filtering Framework — RDFC [Rousselle et al. 2013]

General: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} \left(c_q - \tilde{m}_{\mathbf{x}}(p,q)\right)^2 \tilde{w}_{\mathbf{x}}(p,q)$$

RDFC: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} \left(c_q - \hat{c}_p\right)^2 \tilde{w}_{\mathbf{x}}(p,q)$$

$$\mathbf{x} = \{\text{Color}, \text{Albedo}, \text{Normal}, \ldots\}$$

Input – 256 spp

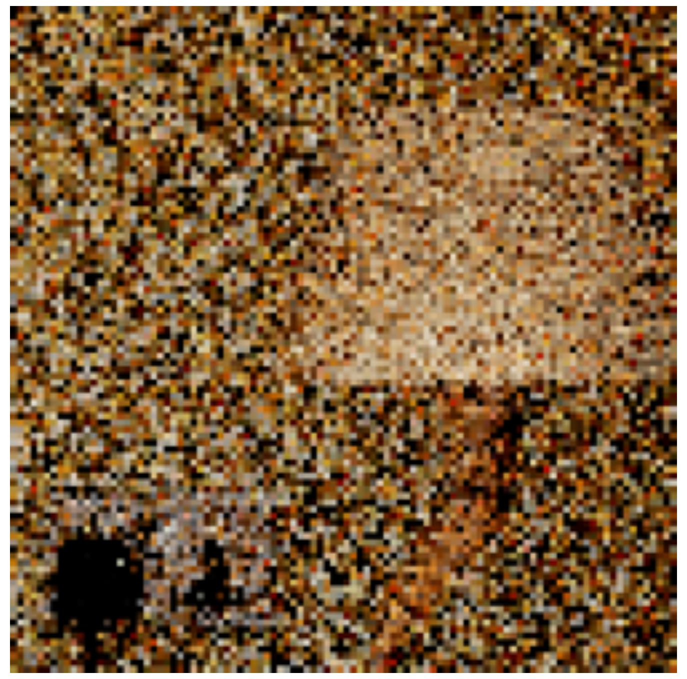

Box Filter

NLM Filter

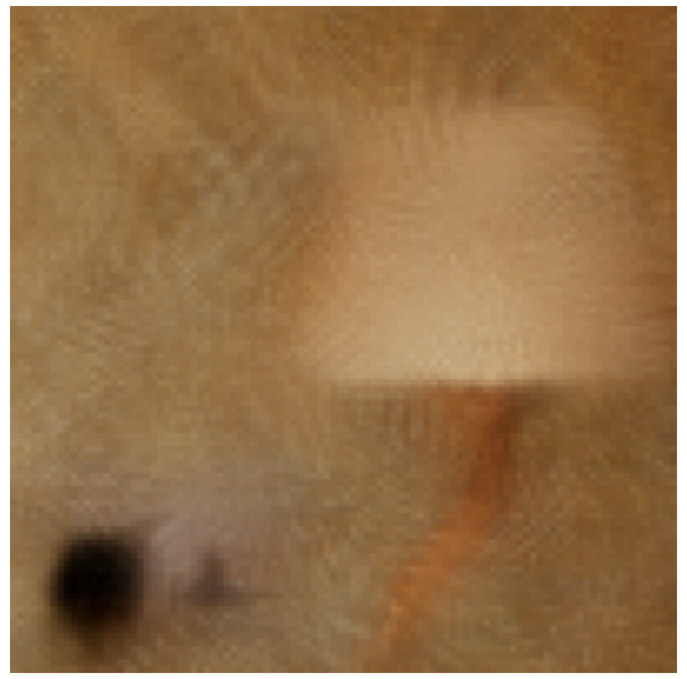

RDFC Filter

Reference

Weights:	Uniform	NL-Means	Joint NL-Means
Order:	zero	zero	zero

Filtering Framework — WLR [Moon et al. 2014]

General: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p,q))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

WLR: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \hat{c}_p - \nabla \hat{c}_p \cdot (\mathbf{x}_q - \mathbf{x}_p))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

$$\mathbf{x} = \{\cancel{\text{Color}}, \text{Albedo}, \text{Normal}, \ldots\}$$

Input – 256 spp | Box Filter | NLM Filter | RDFC Filter | WLR Filter | Reference

Weights:	Uniform	NL-Means	Joint NL-Means	Cross-bilateral
Order:	zero	zero	zero	first

Disney Research

Filtering Framework — WLR [Moon et al. 2014]

General: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p,q))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

WLR: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \hat{c}_p - \nabla \hat{c}_p \cdot (\mathbf{x}_q - \mathbf{x}_p))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

$$\mathbf{x} = \{\cancel{\mathrm{Color},}\ \mathrm{Albedo}, \mathrm{Normal}, \ldots\}$$

Input – 256 spp | Box Filter | NLM Filter | RDFC Filter | WLR Filter | Reference

	Box Filter	NLM Filter	RDFC Filter	WLR Filter
Weights:	Uniform	NL-Means	Joint NL-Means	Cross-bilateral
Order:	zero	zero	zero	first

Filtering Framework — Proposed Filter

General: $$\text{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} \left(c_q - \tilde{m}_{\mathbf{x}}(p,q)\right)^2 \tilde{w}_{\mathbf{x}}(p,q)$$

WLR: $$\text{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} \left(c_q - \hat{c}_p - \nabla \hat{c}_p \cdot (\mathbf{x}_q - \mathbf{x}_p)\right)^2 \tilde{w}_{\mathbf{y}}(p,q)$$

$\mathbf{x} = \{\cancel{\text{Color},}\ \text{Albedo}, \text{Normal}, \ldots\}$

$\mathbf{y} = \{\text{Color}\}$

Input – 256 spp | Box Filter | NLM Filter | RDFC Filter | WLR Filter | Our Filter | Reference

Weights:	Uniform	NL-Means	Joint NL-Means	Cross-bilateral	NL-Means	
Order:	zero	zero	zero	first	first	

Filtering Framework — Proposed Filter

General: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \tilde{m}_{\mathbf{x}}(p,q))^2 \, \tilde{w}_{\mathbf{x}}(p,q)$$

WLR: $$\mathrm{Error}(\hat{c}_p) = \sum_{q \in \mathcal{N}_p} (c_q - \hat{c}_p - \nabla \hat{c}_p \cdot (\mathbf{x}_q - \mathbf{x}_p))^2 \, \tilde{w}_{\mathbf{y}}(p,q)$$

$\mathbf{x} = \{\cancel{\mathrm{Color},}\ \mathrm{Albedo}, \mathrm{Normal}, \ldots\}$ $\quad\quad$ $\mathbf{y} = \{\mathrm{Color}\}$

Input – 256 spp	Box Filter	NLM Filter	RDFC Filter	WLR Filter	Our Filter	Reference

Weights:	Uniform	NL-Means	Joint NL-Means	Cross-bilateral	NL-Means
Order:	zero	zero	zero	first	first

Algorithm Overview

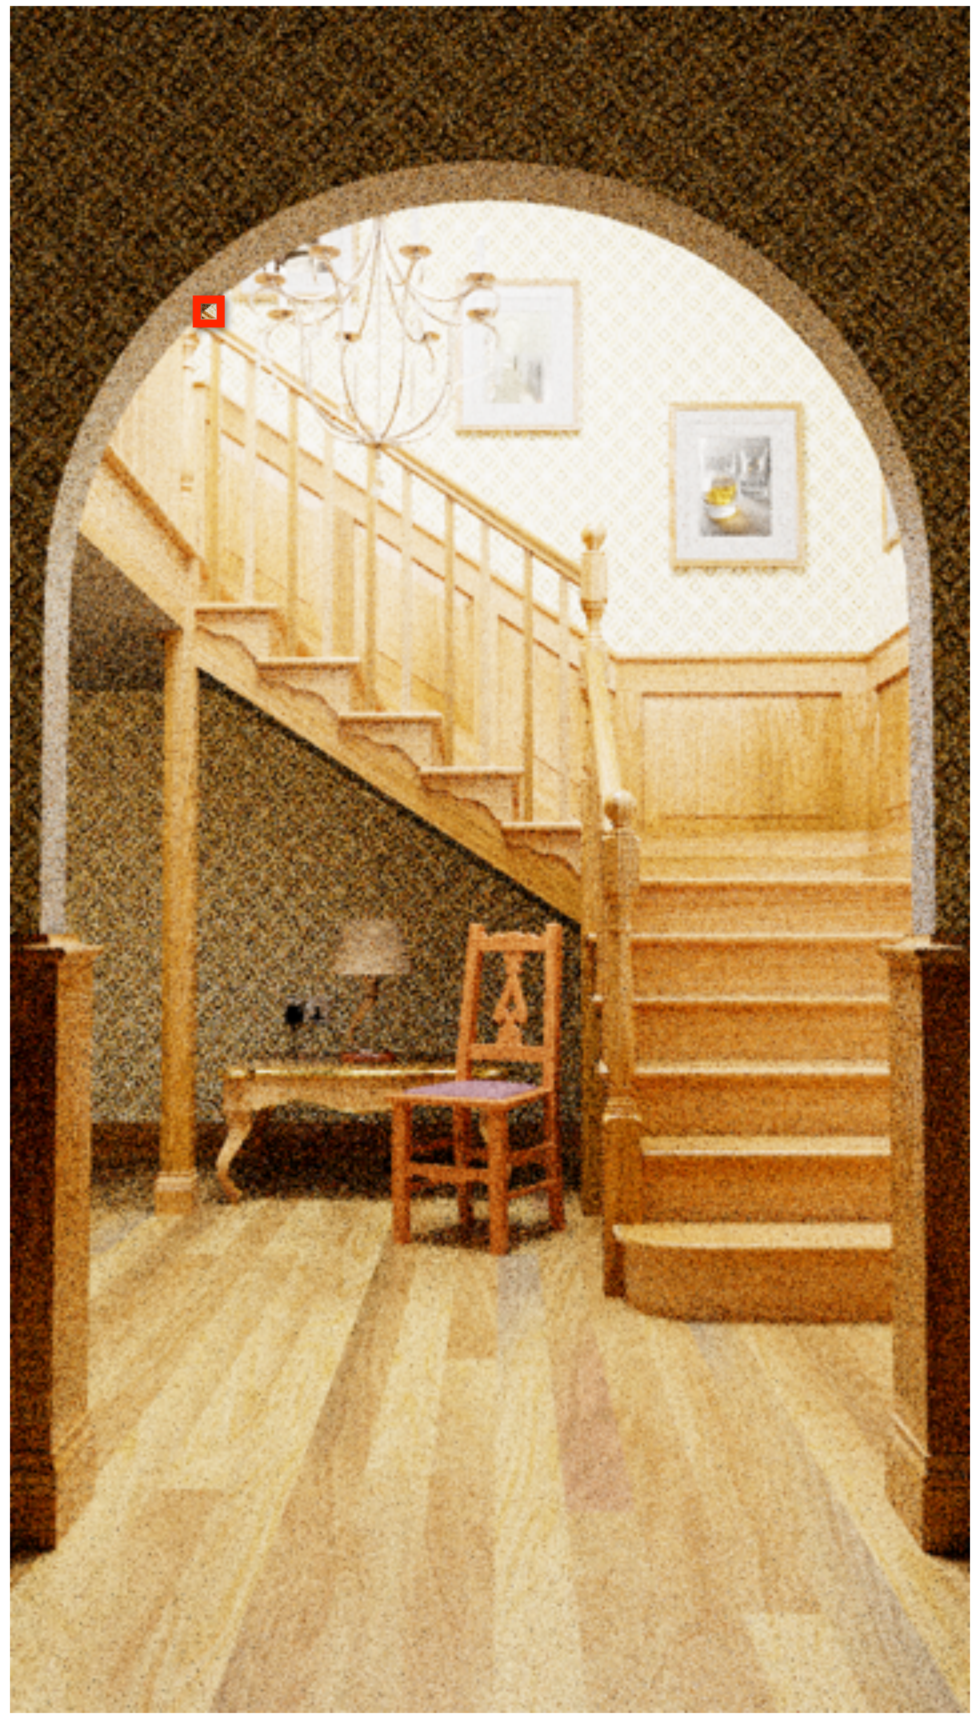

Algorithm Overview

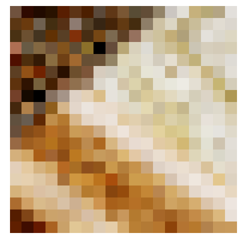

Algorithm Overview

variance

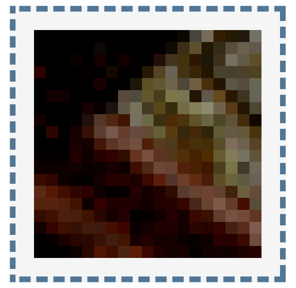

noisy input

feature buffers

Algorithm Overview

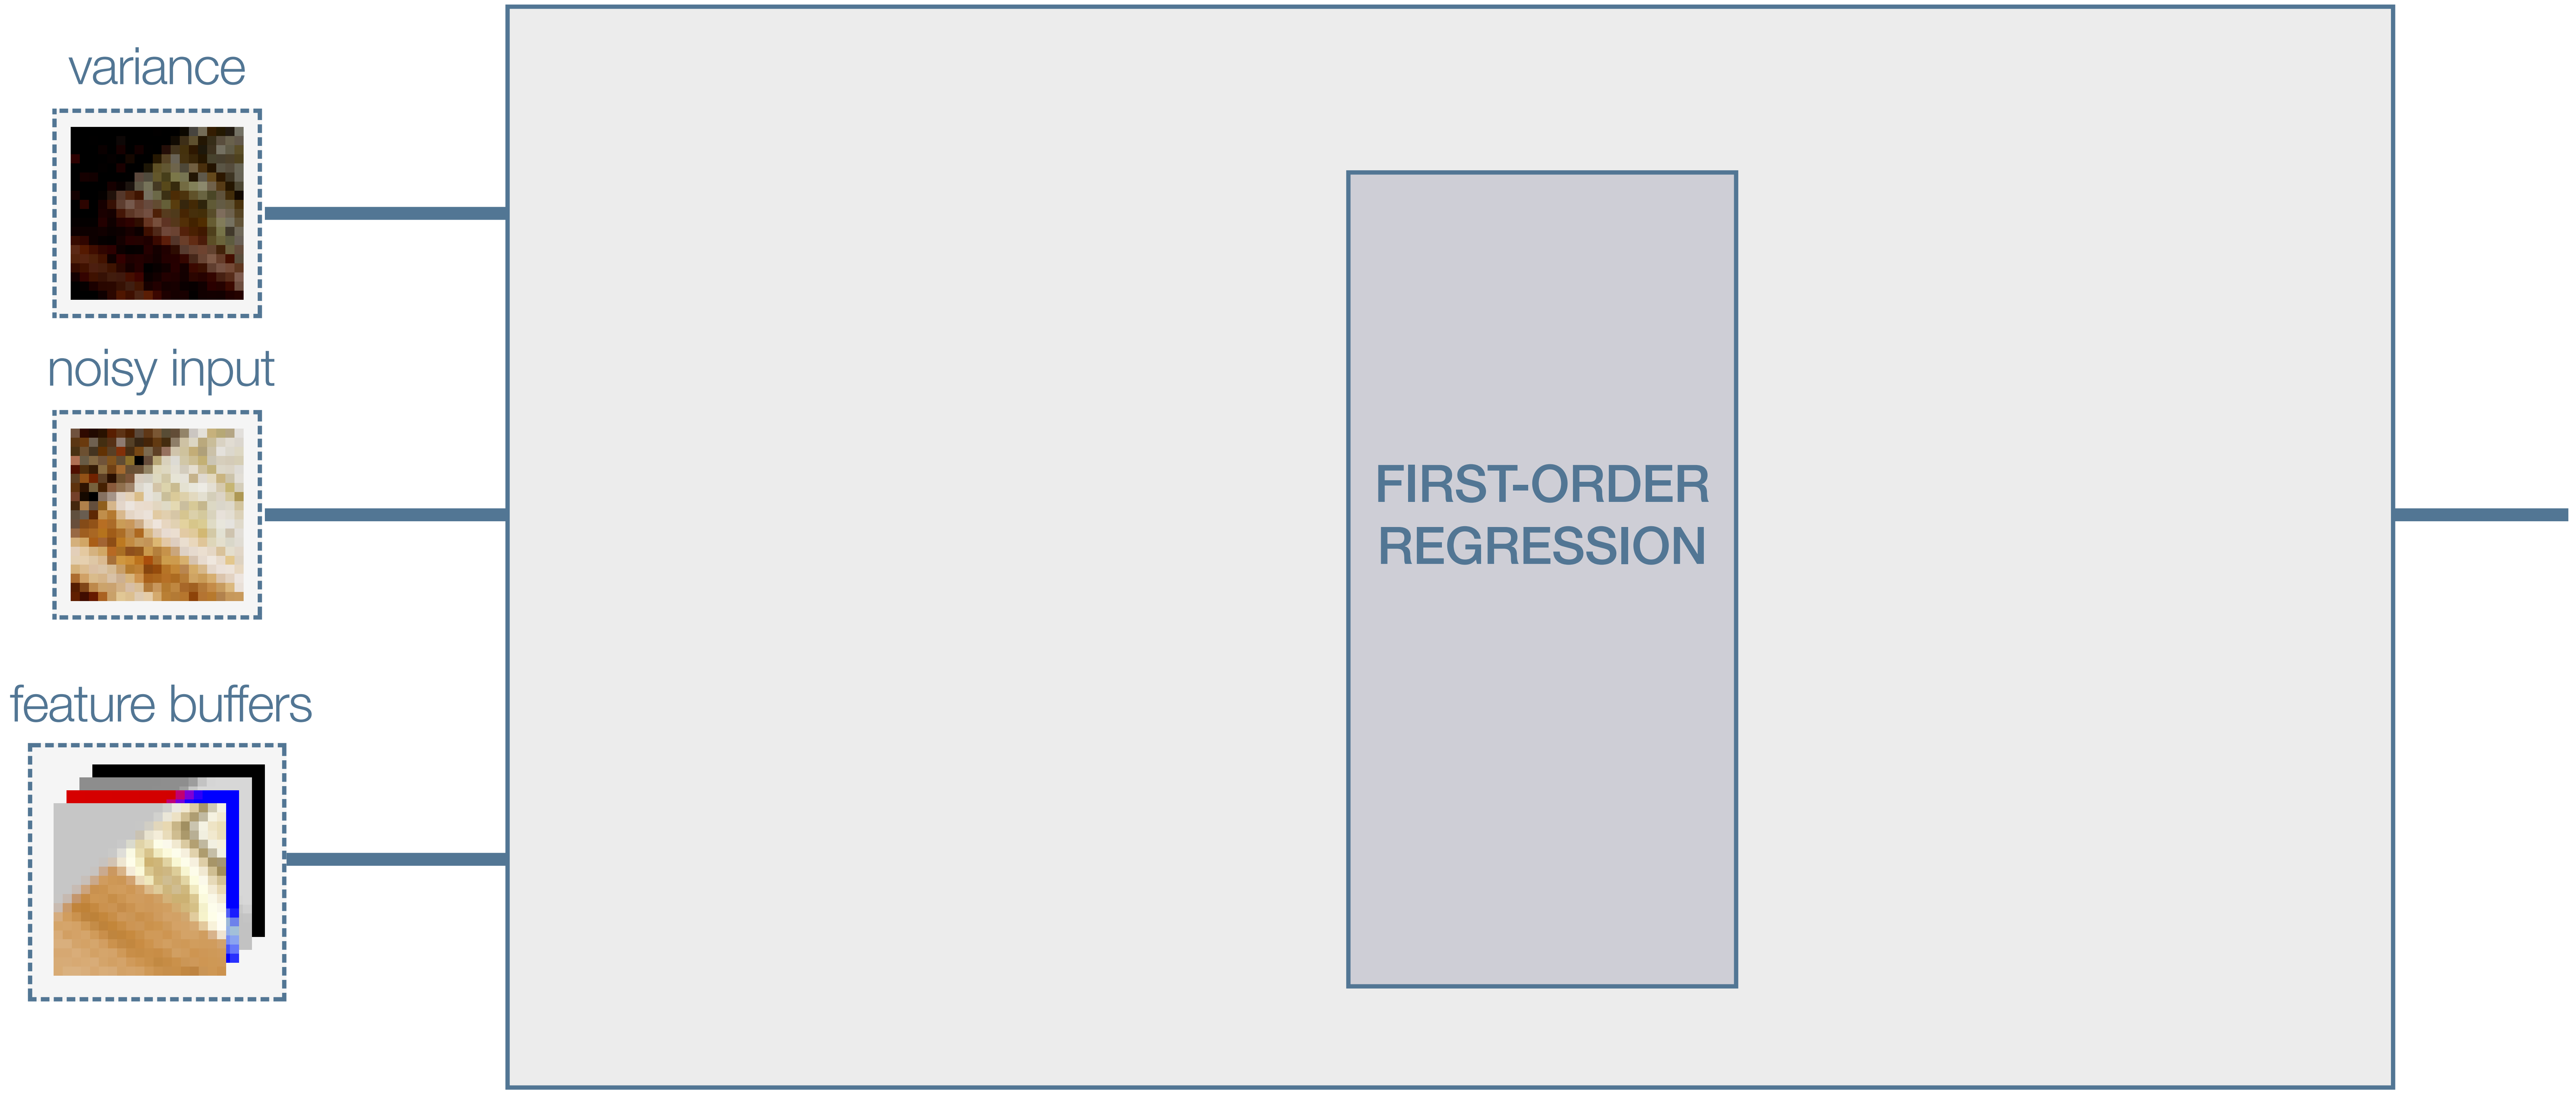

Algorithm Overview

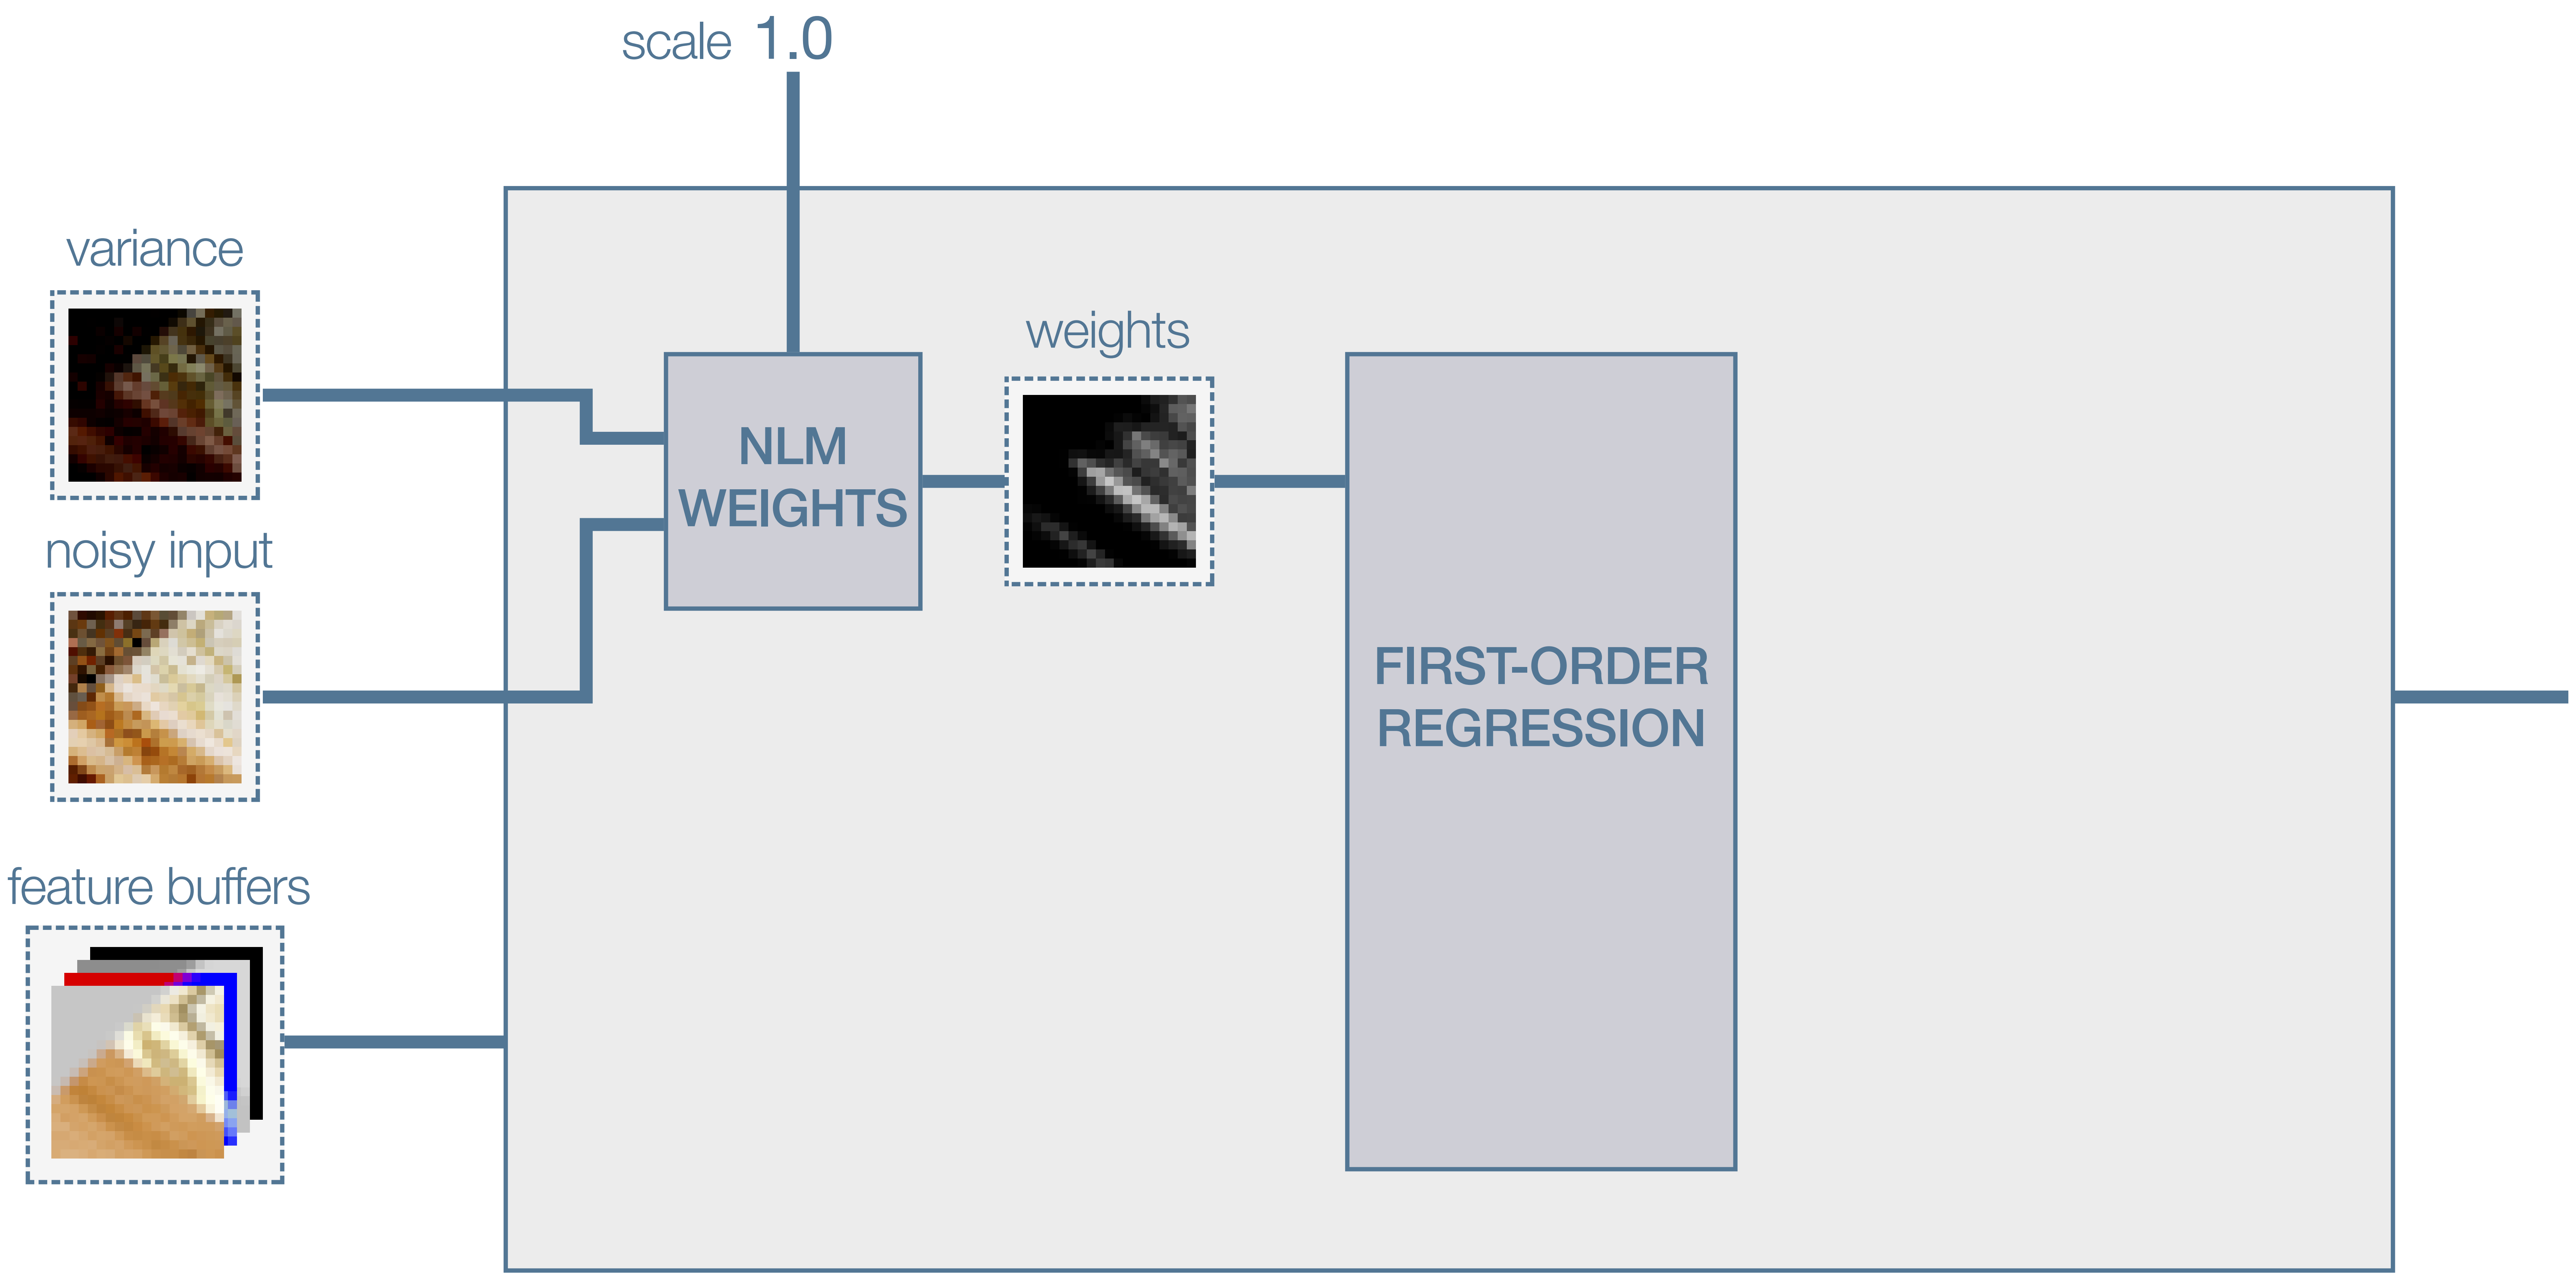

Algorithm Overview

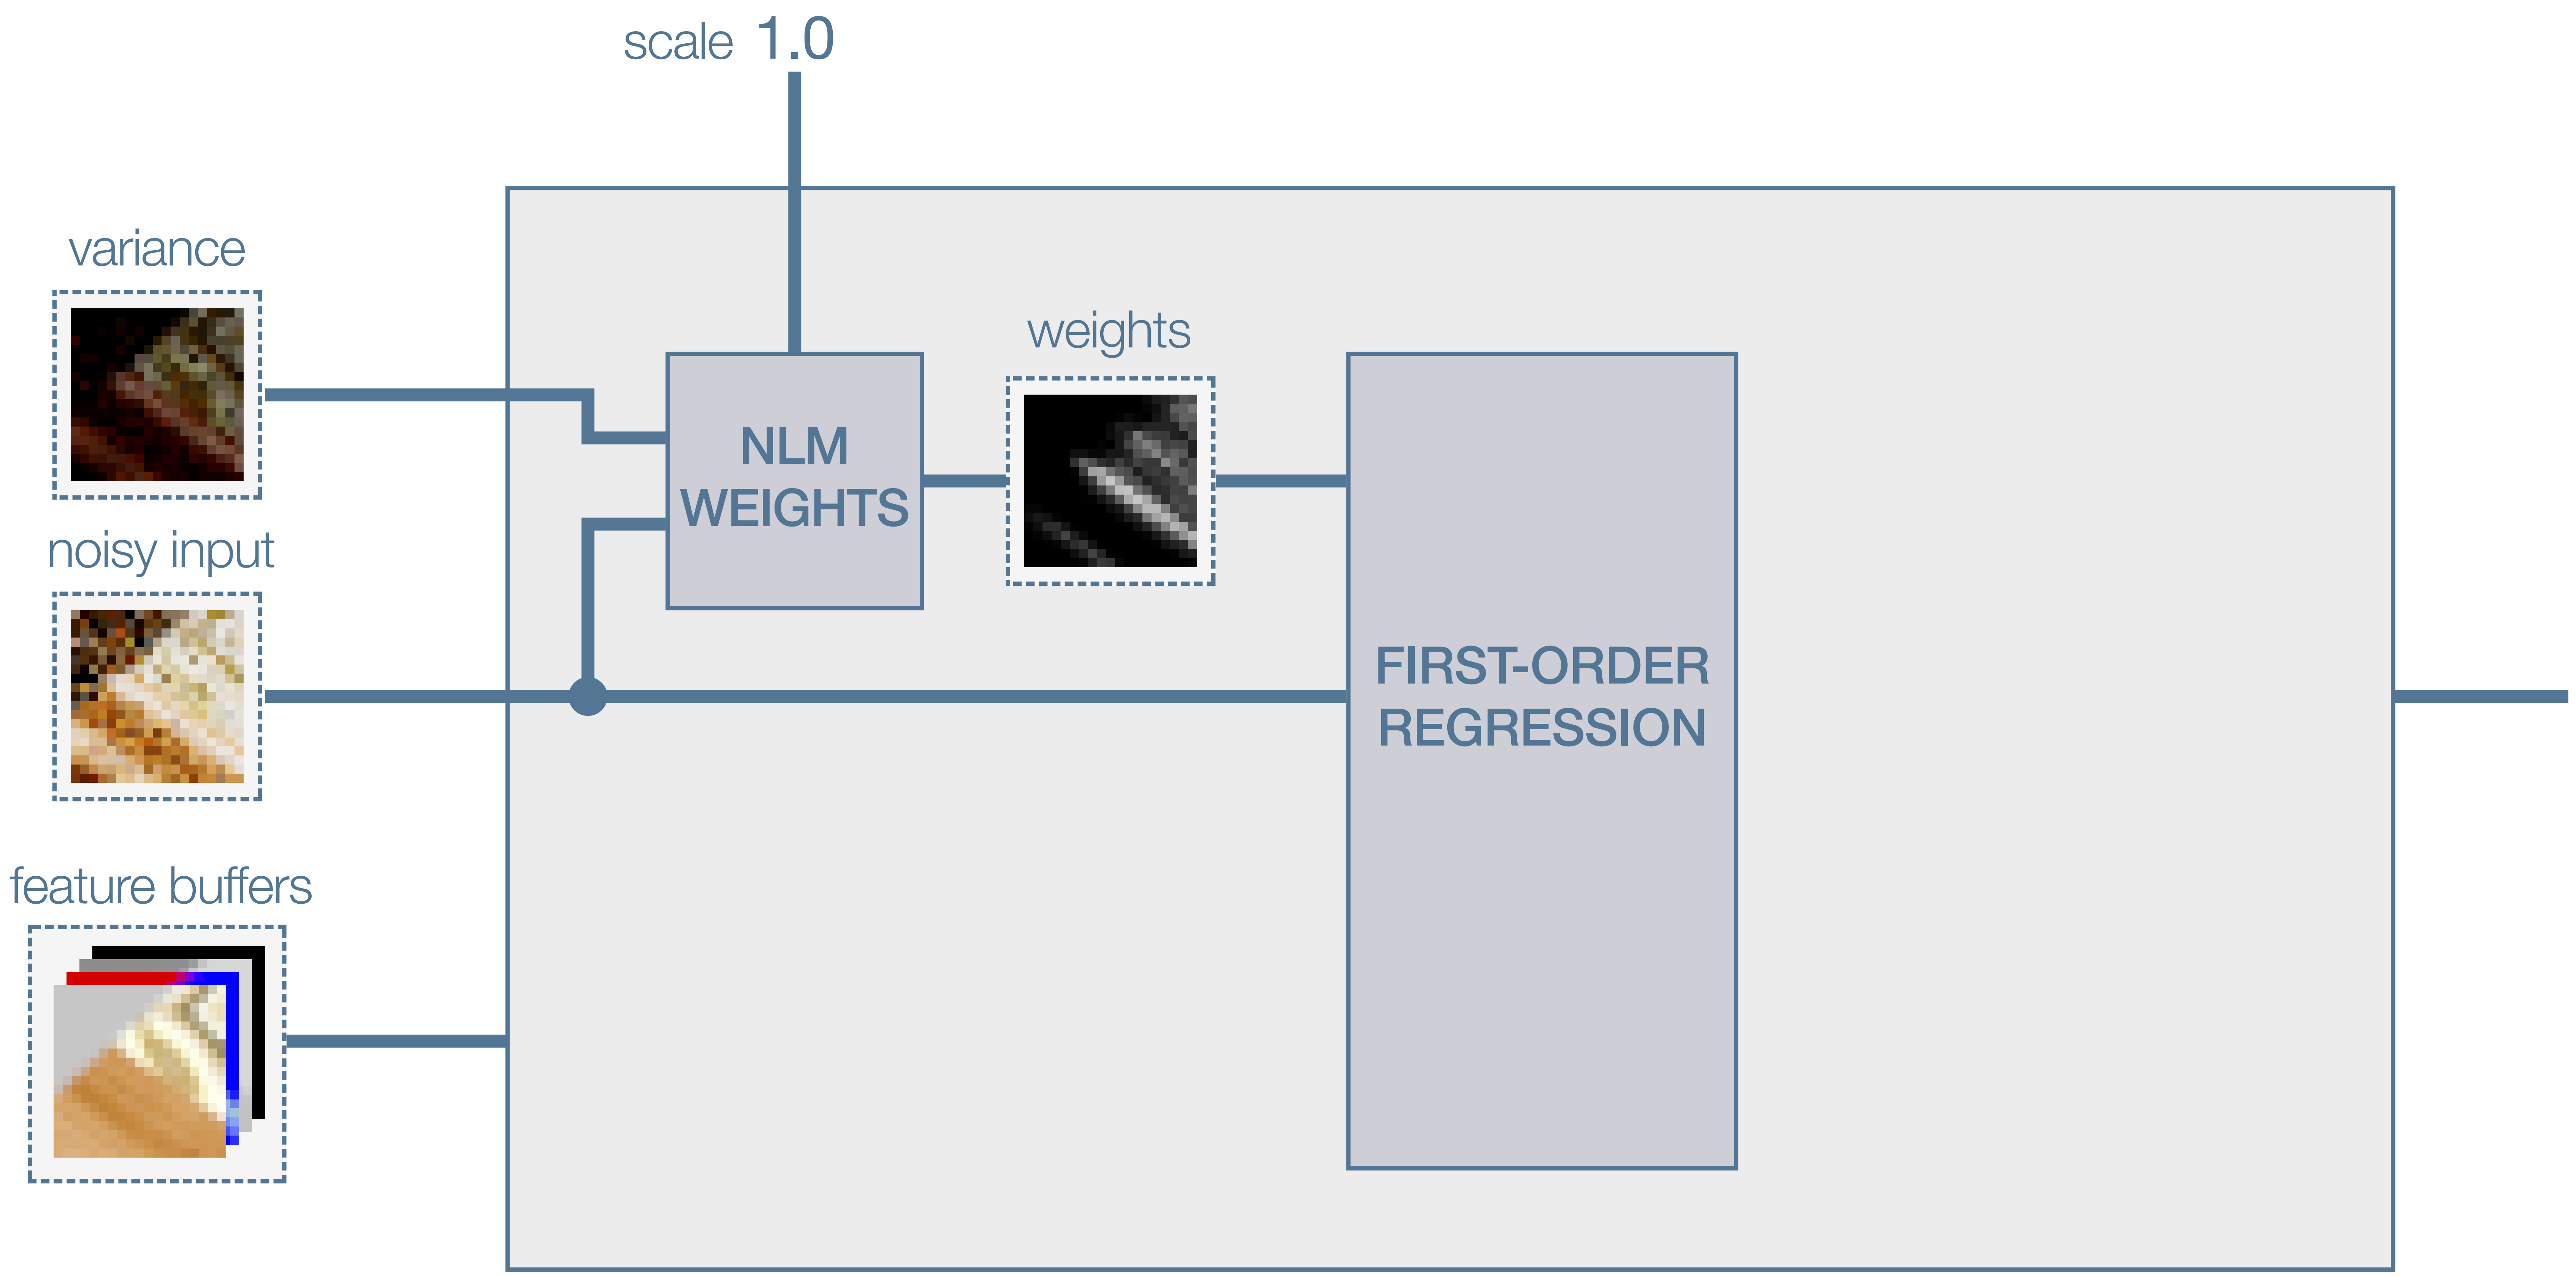

Algorithm Overview

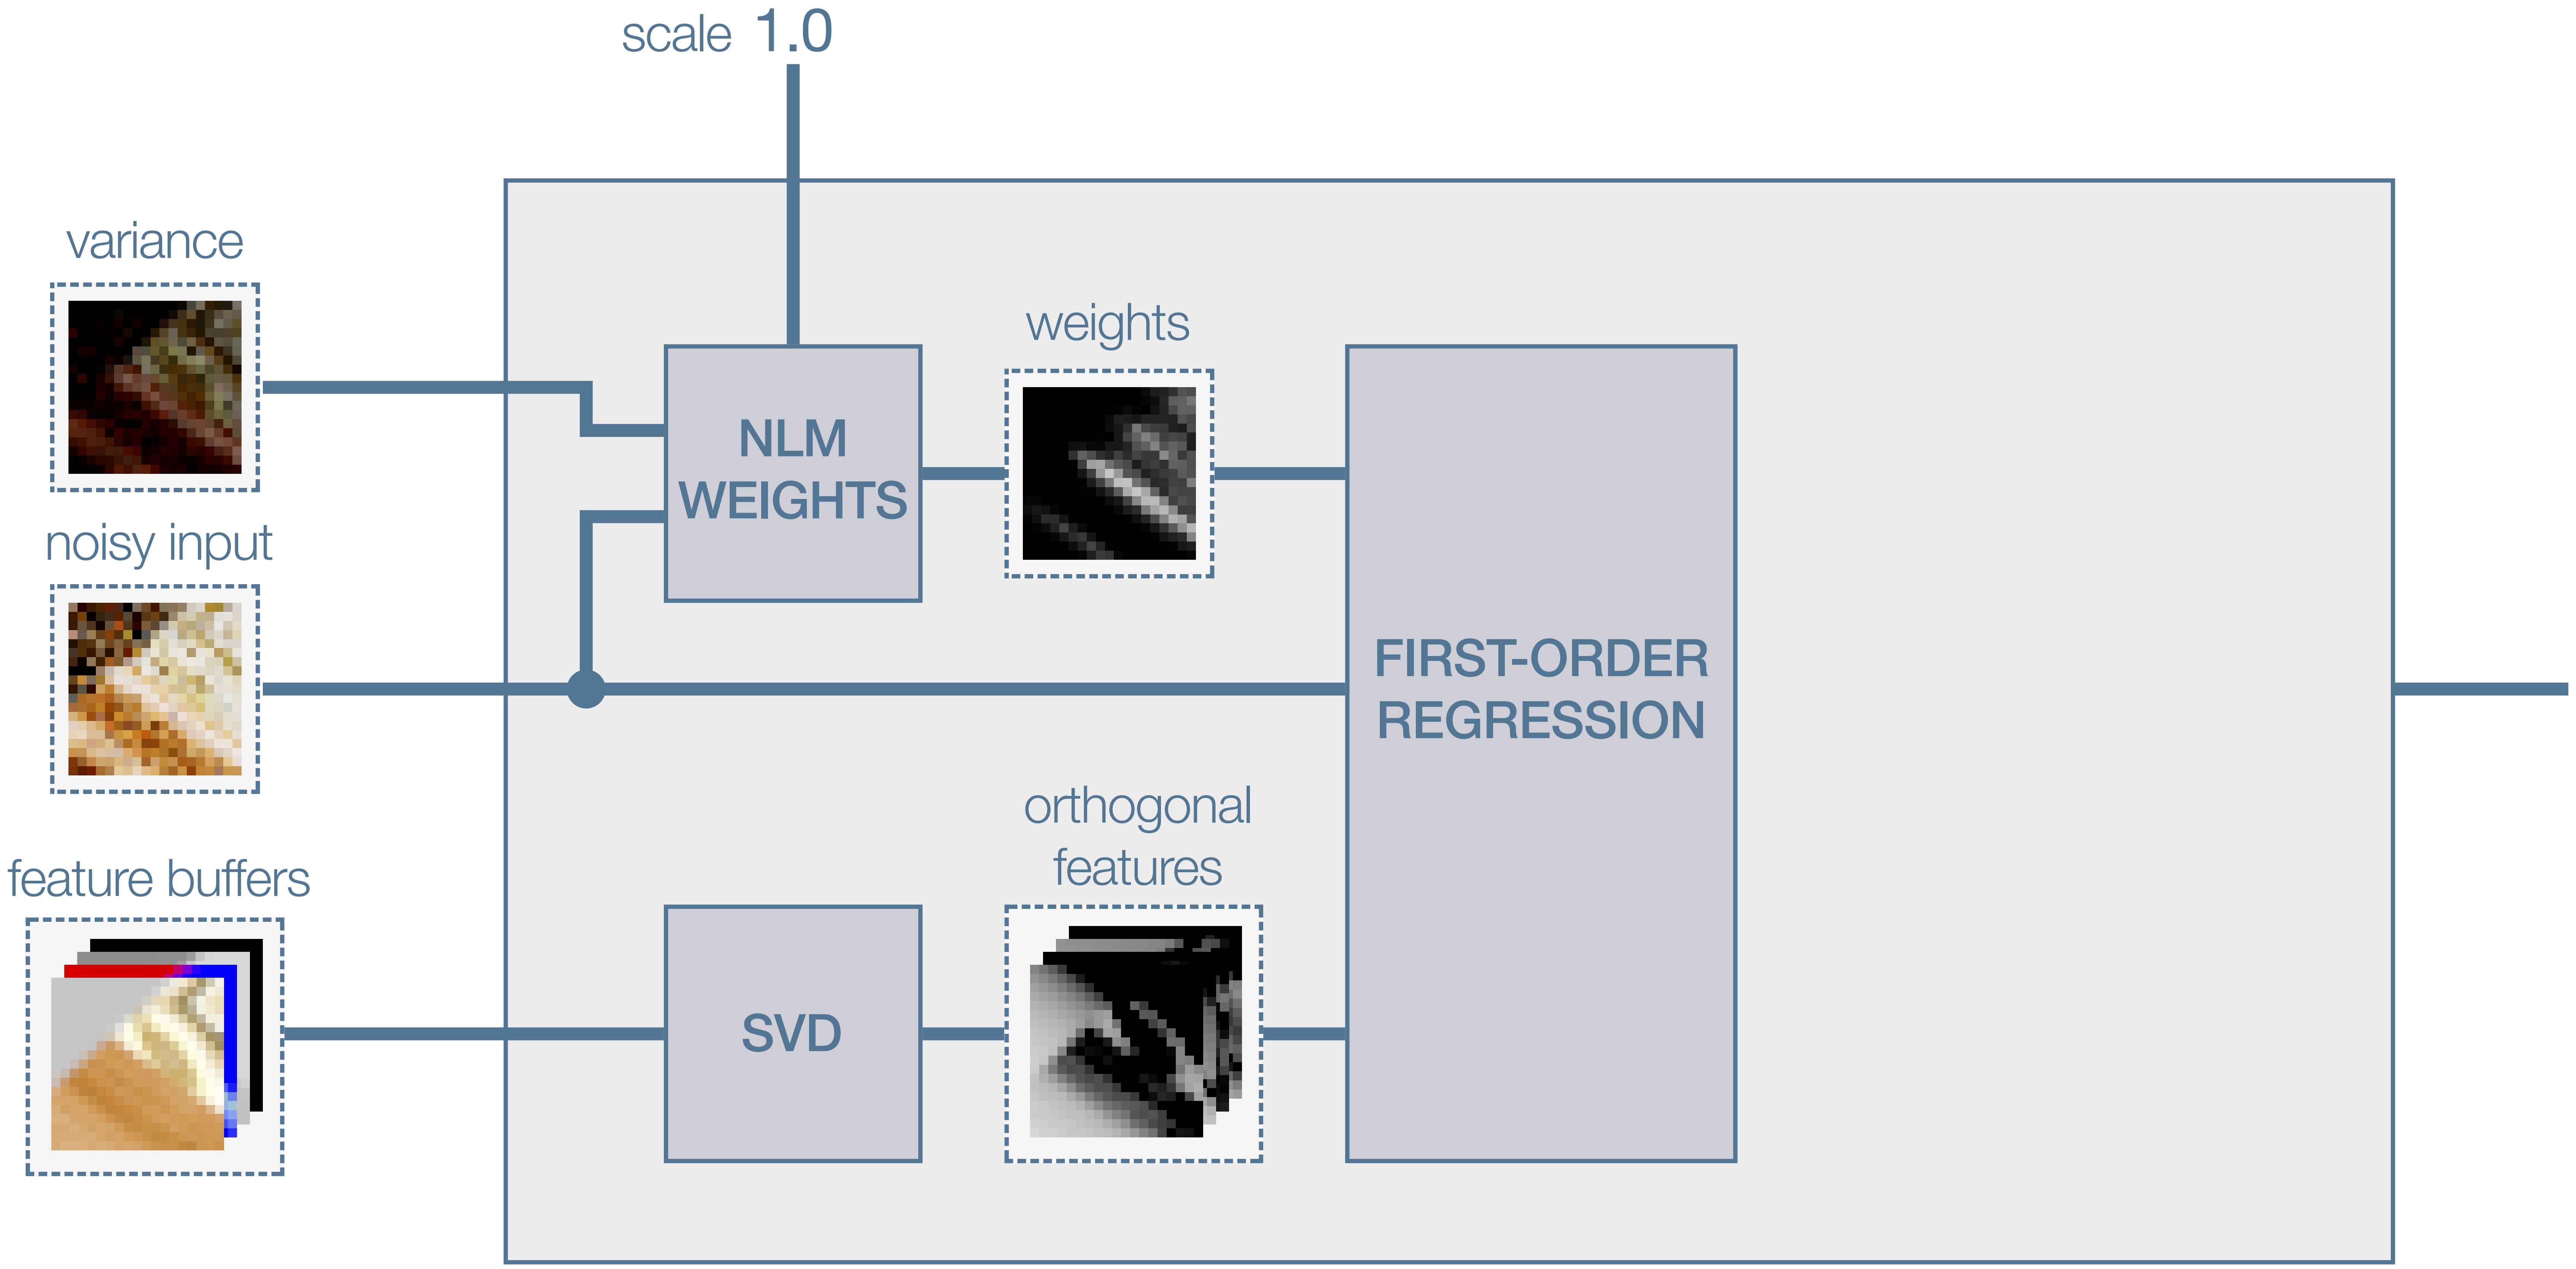

Algorithm Overview

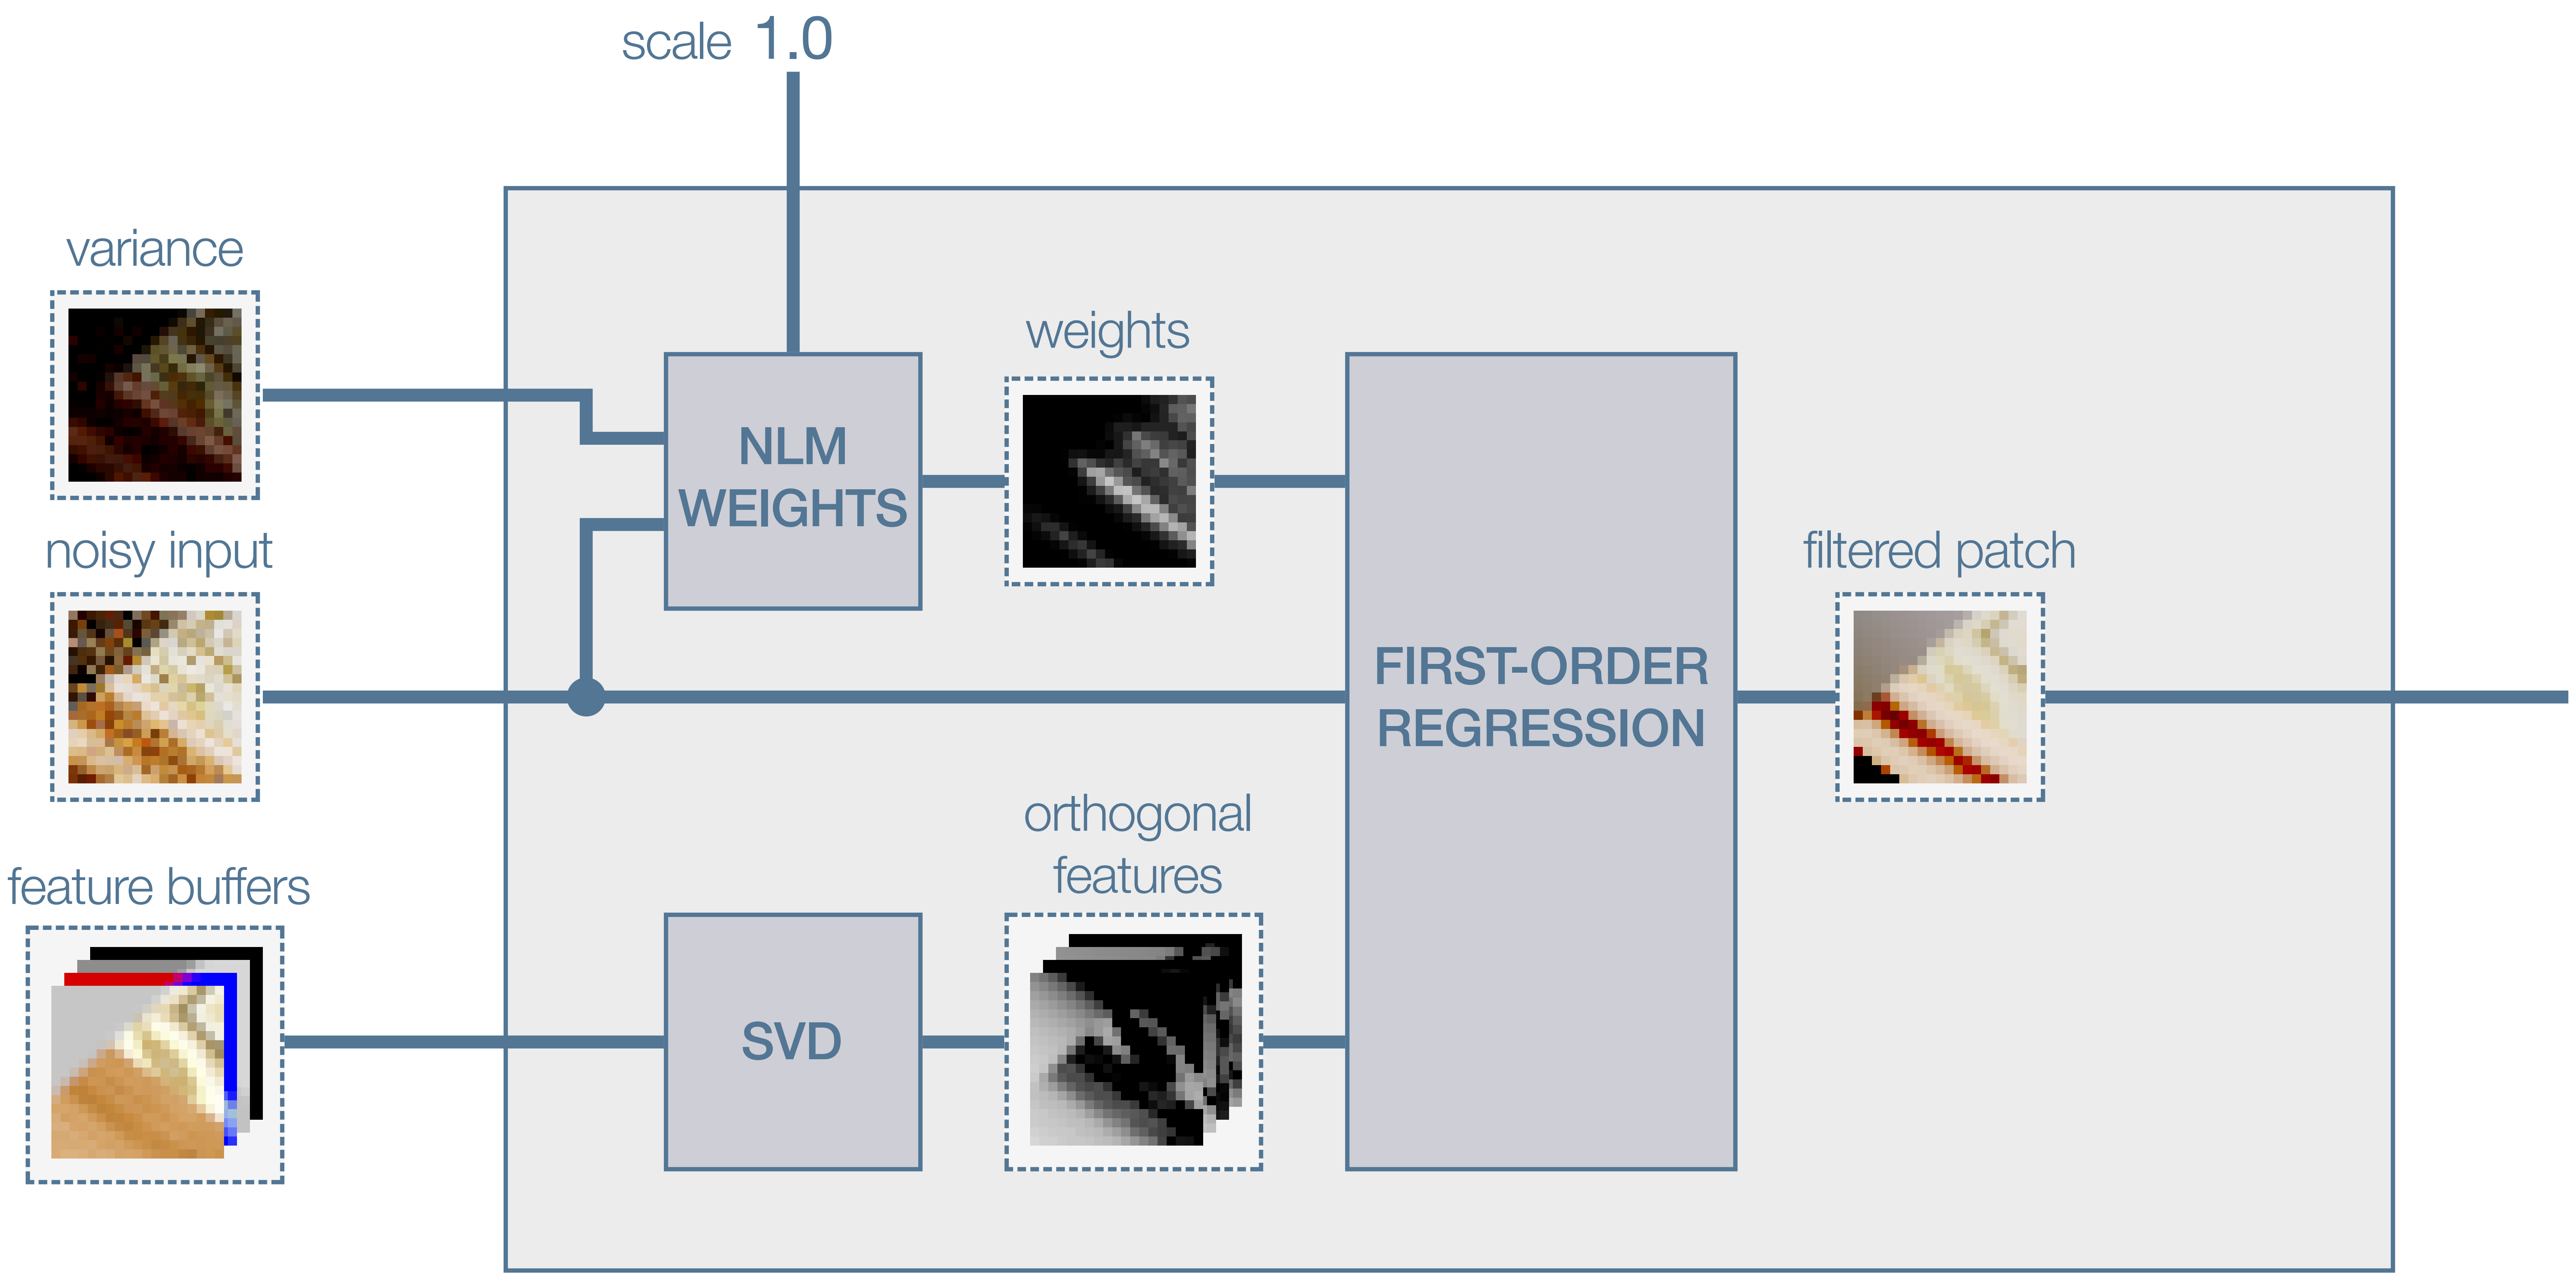

Algorithm Overview

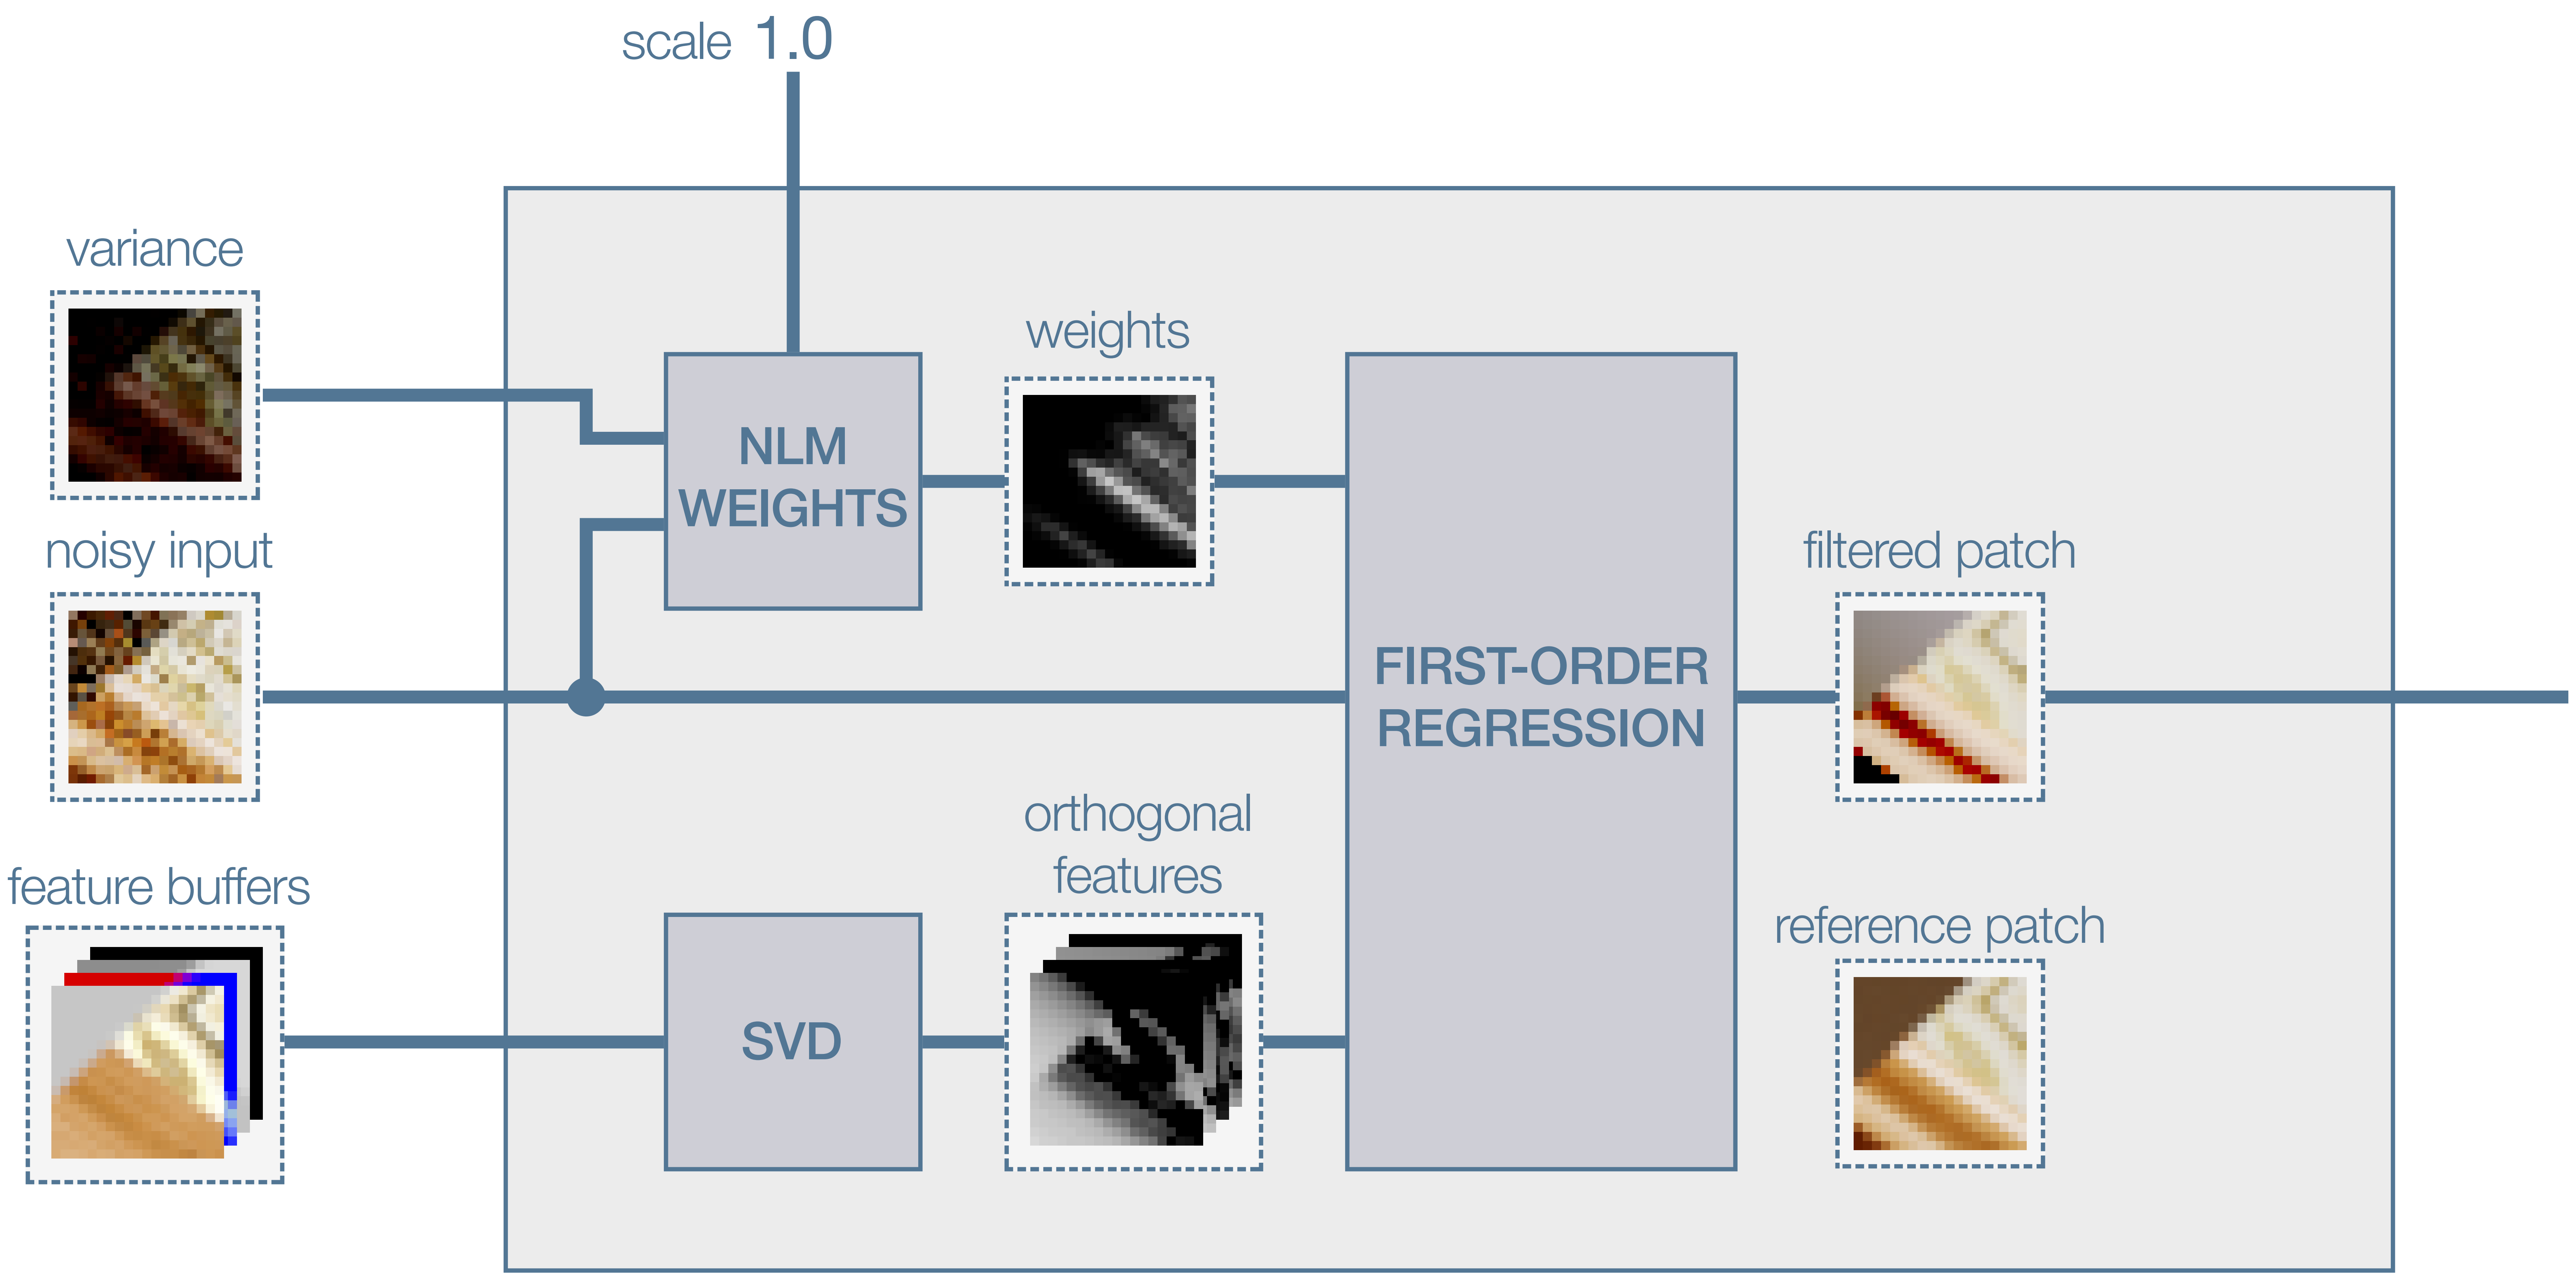

Algorithm Overview

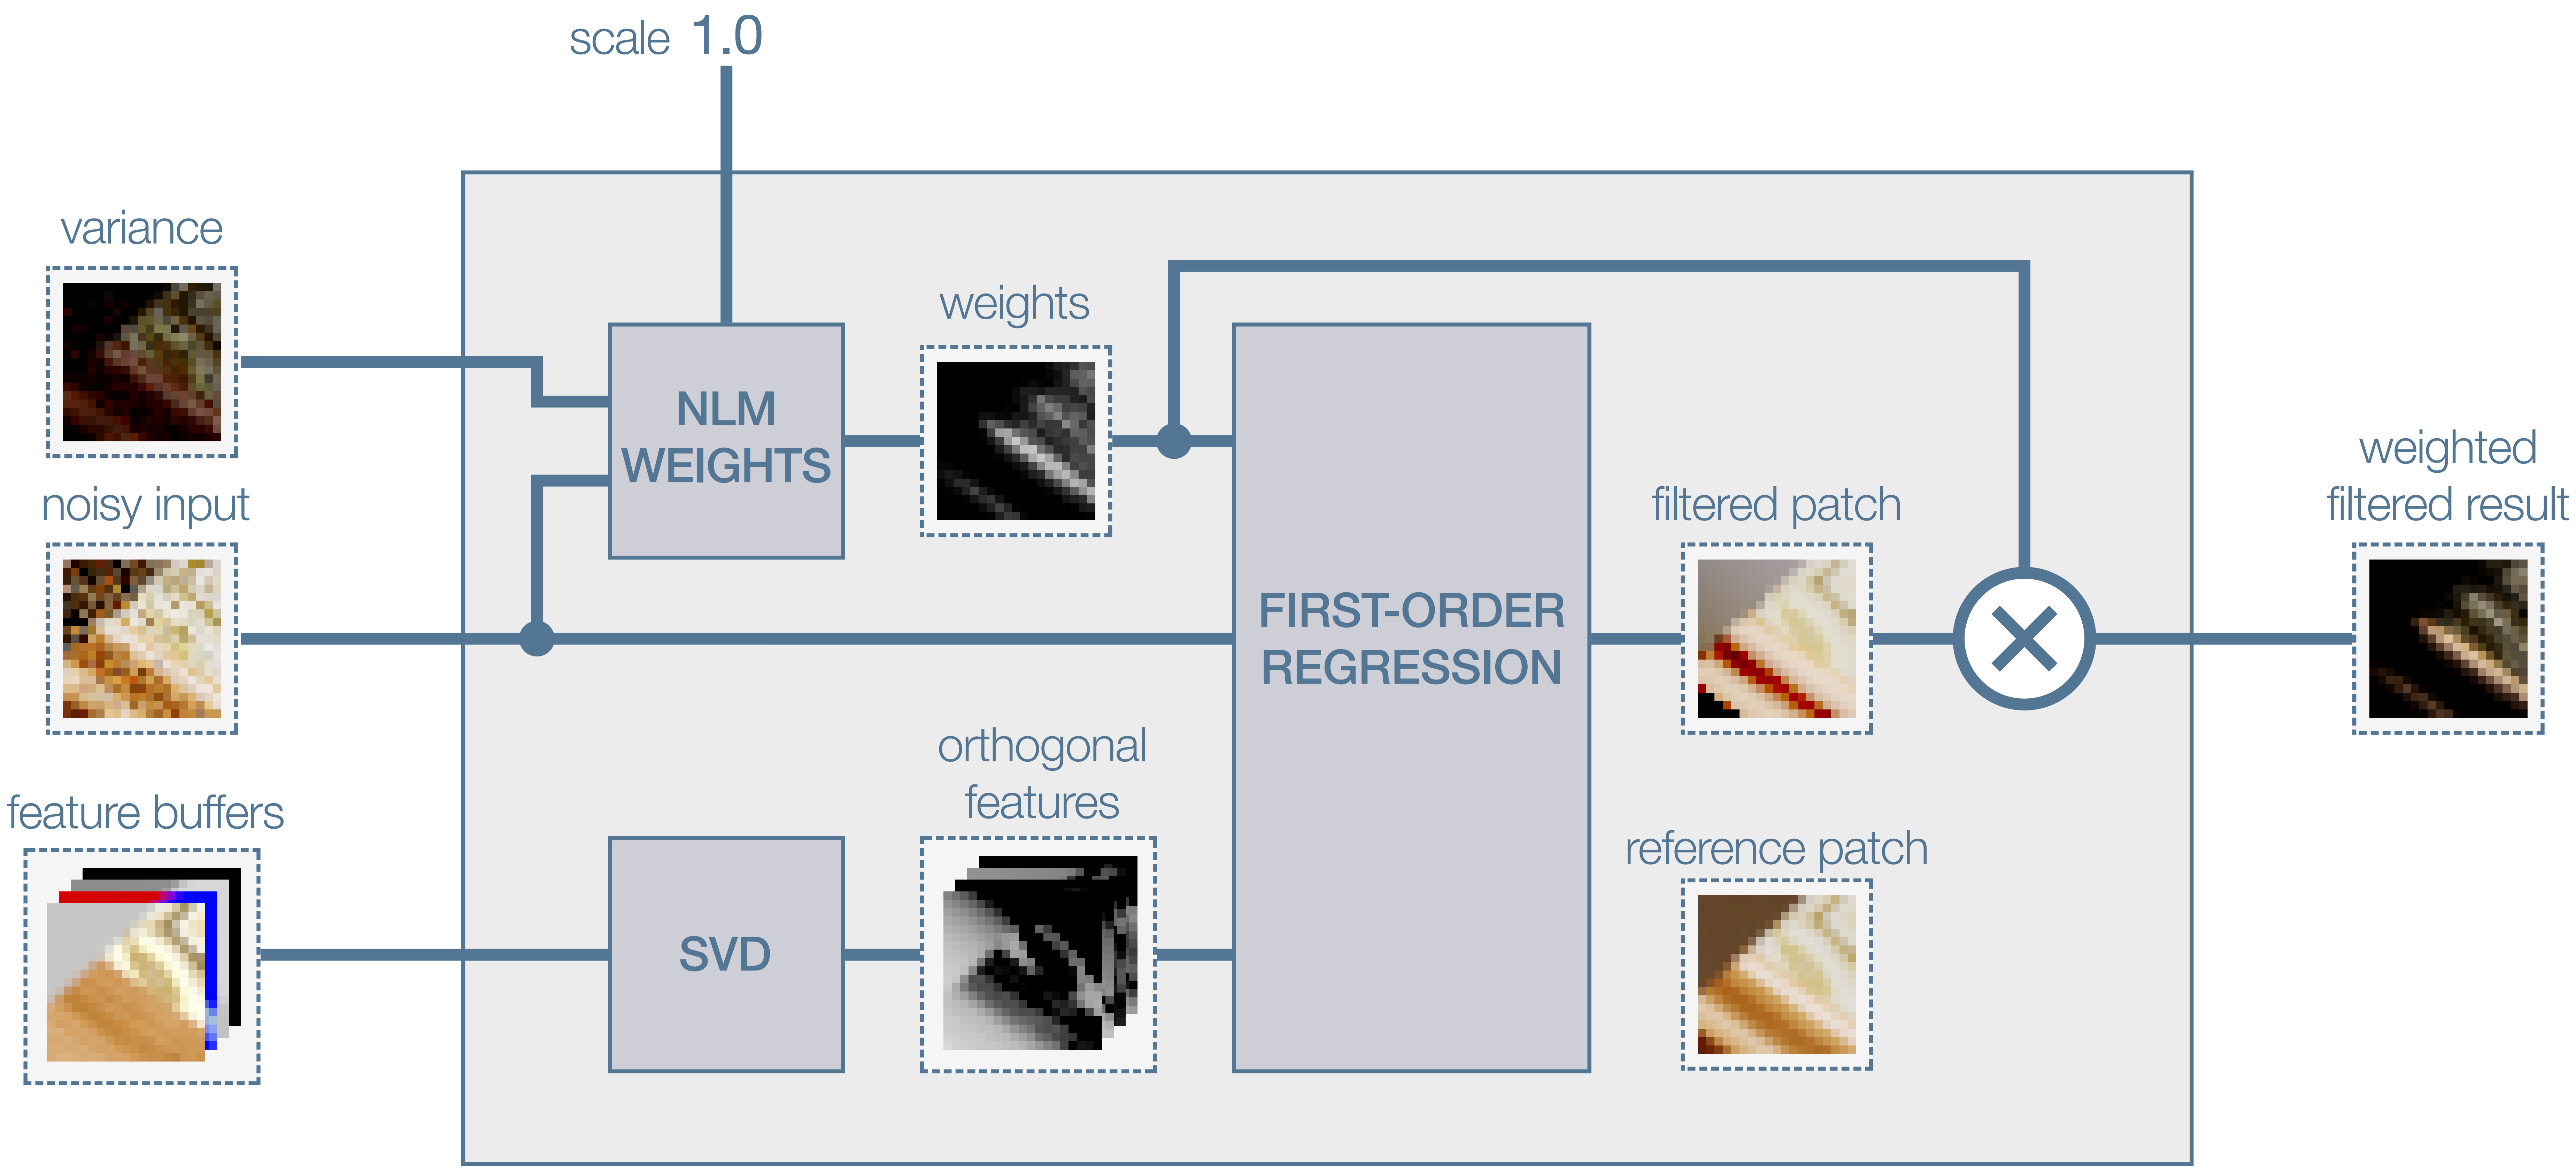

Disney Research

Algorithm Overview

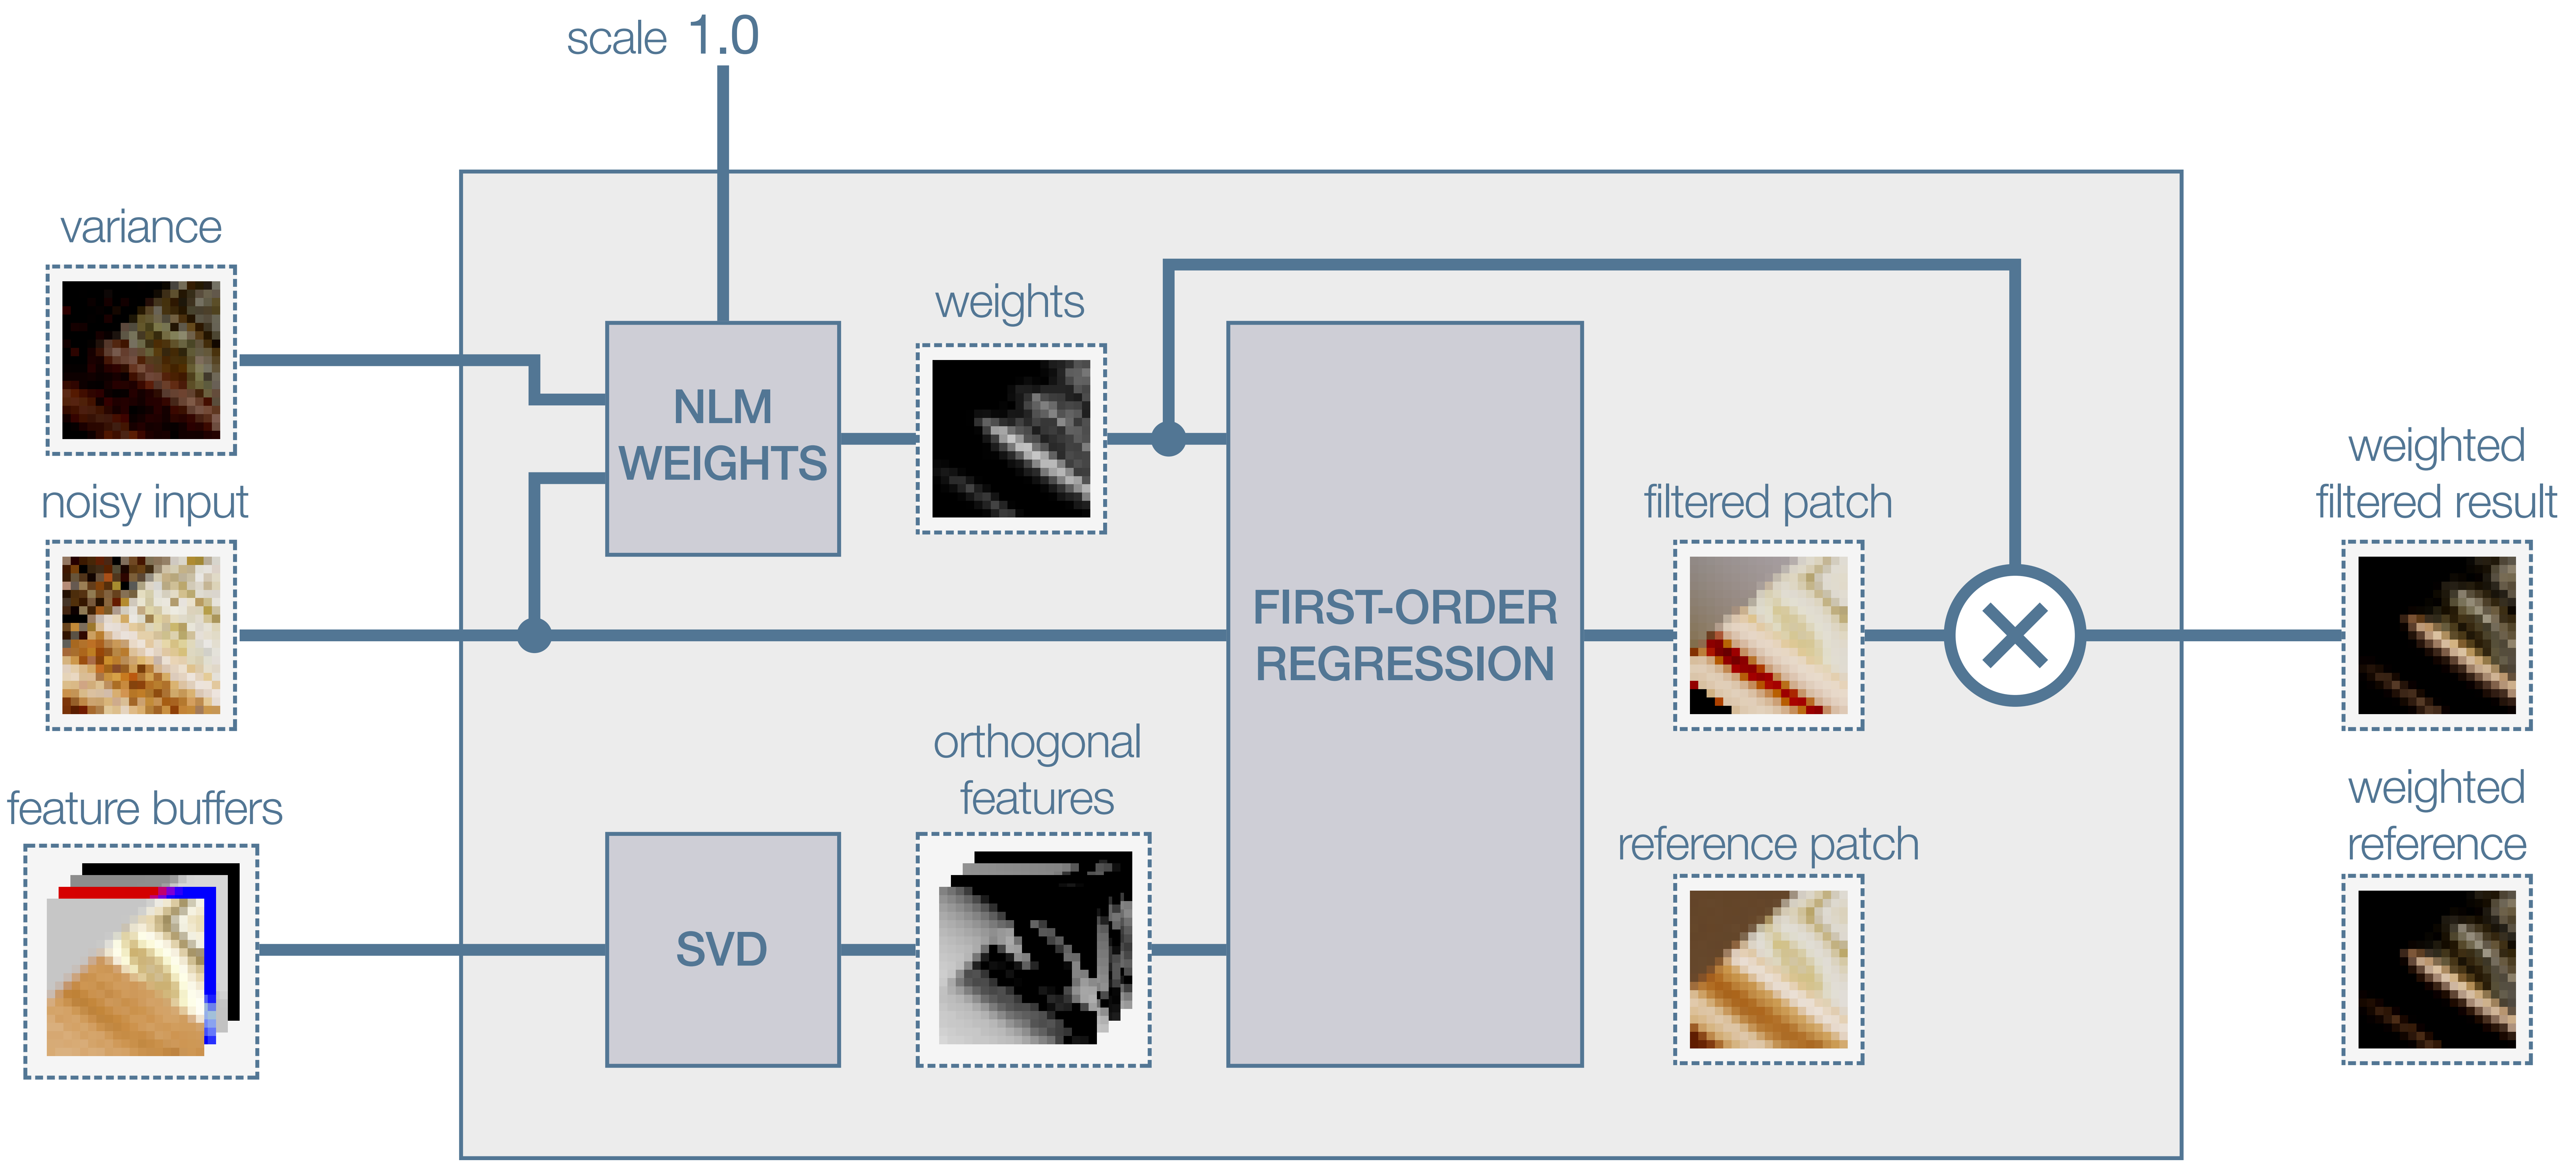

Algorithm Overview

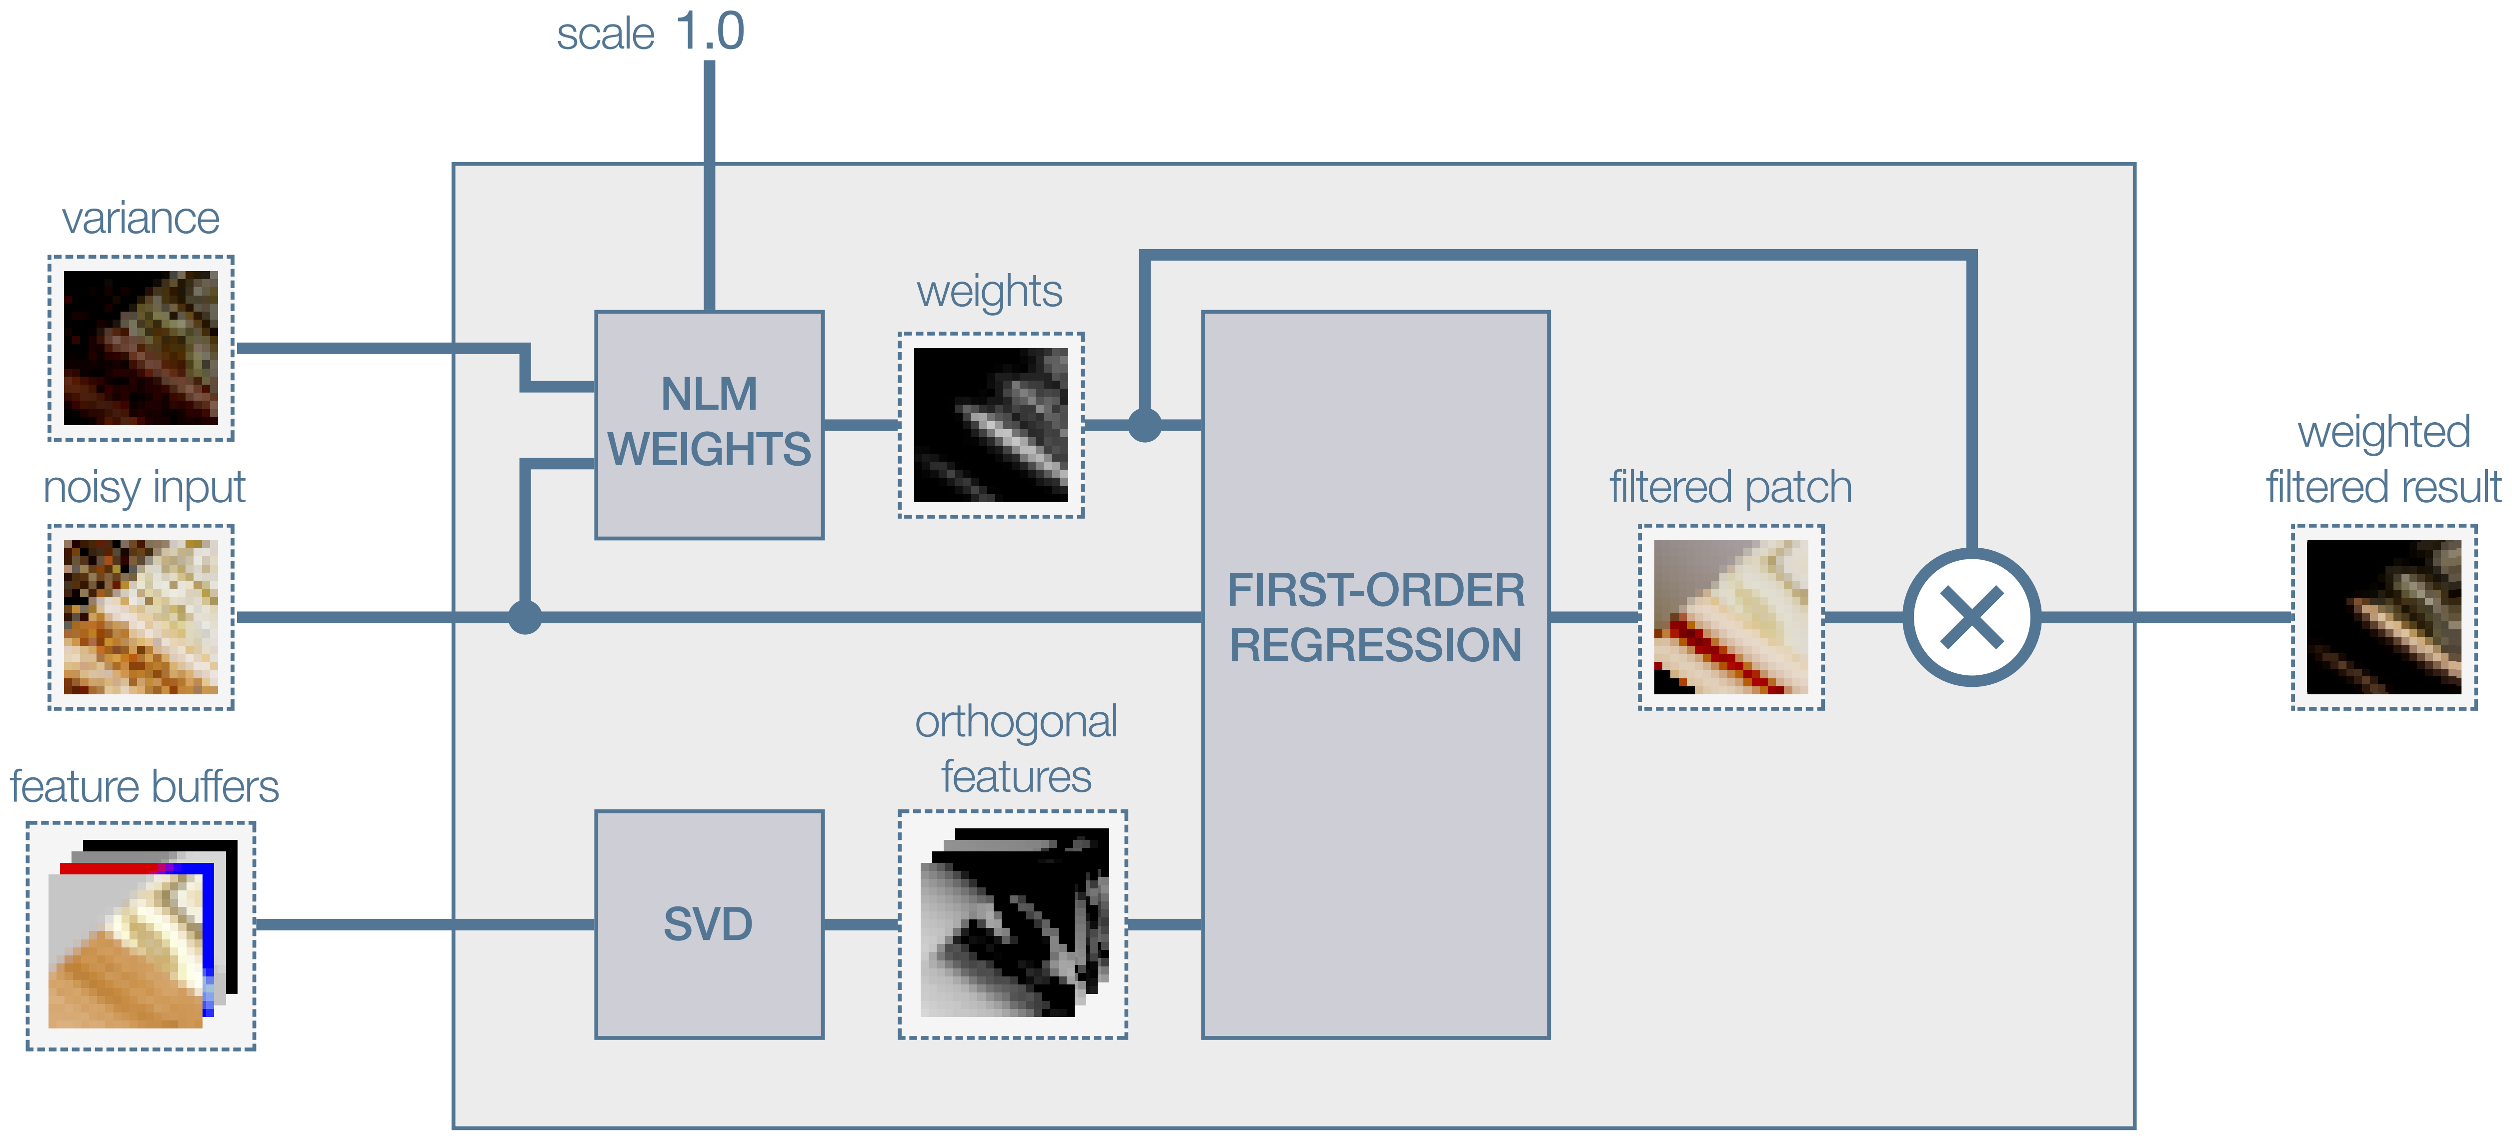

Algorithm Overview

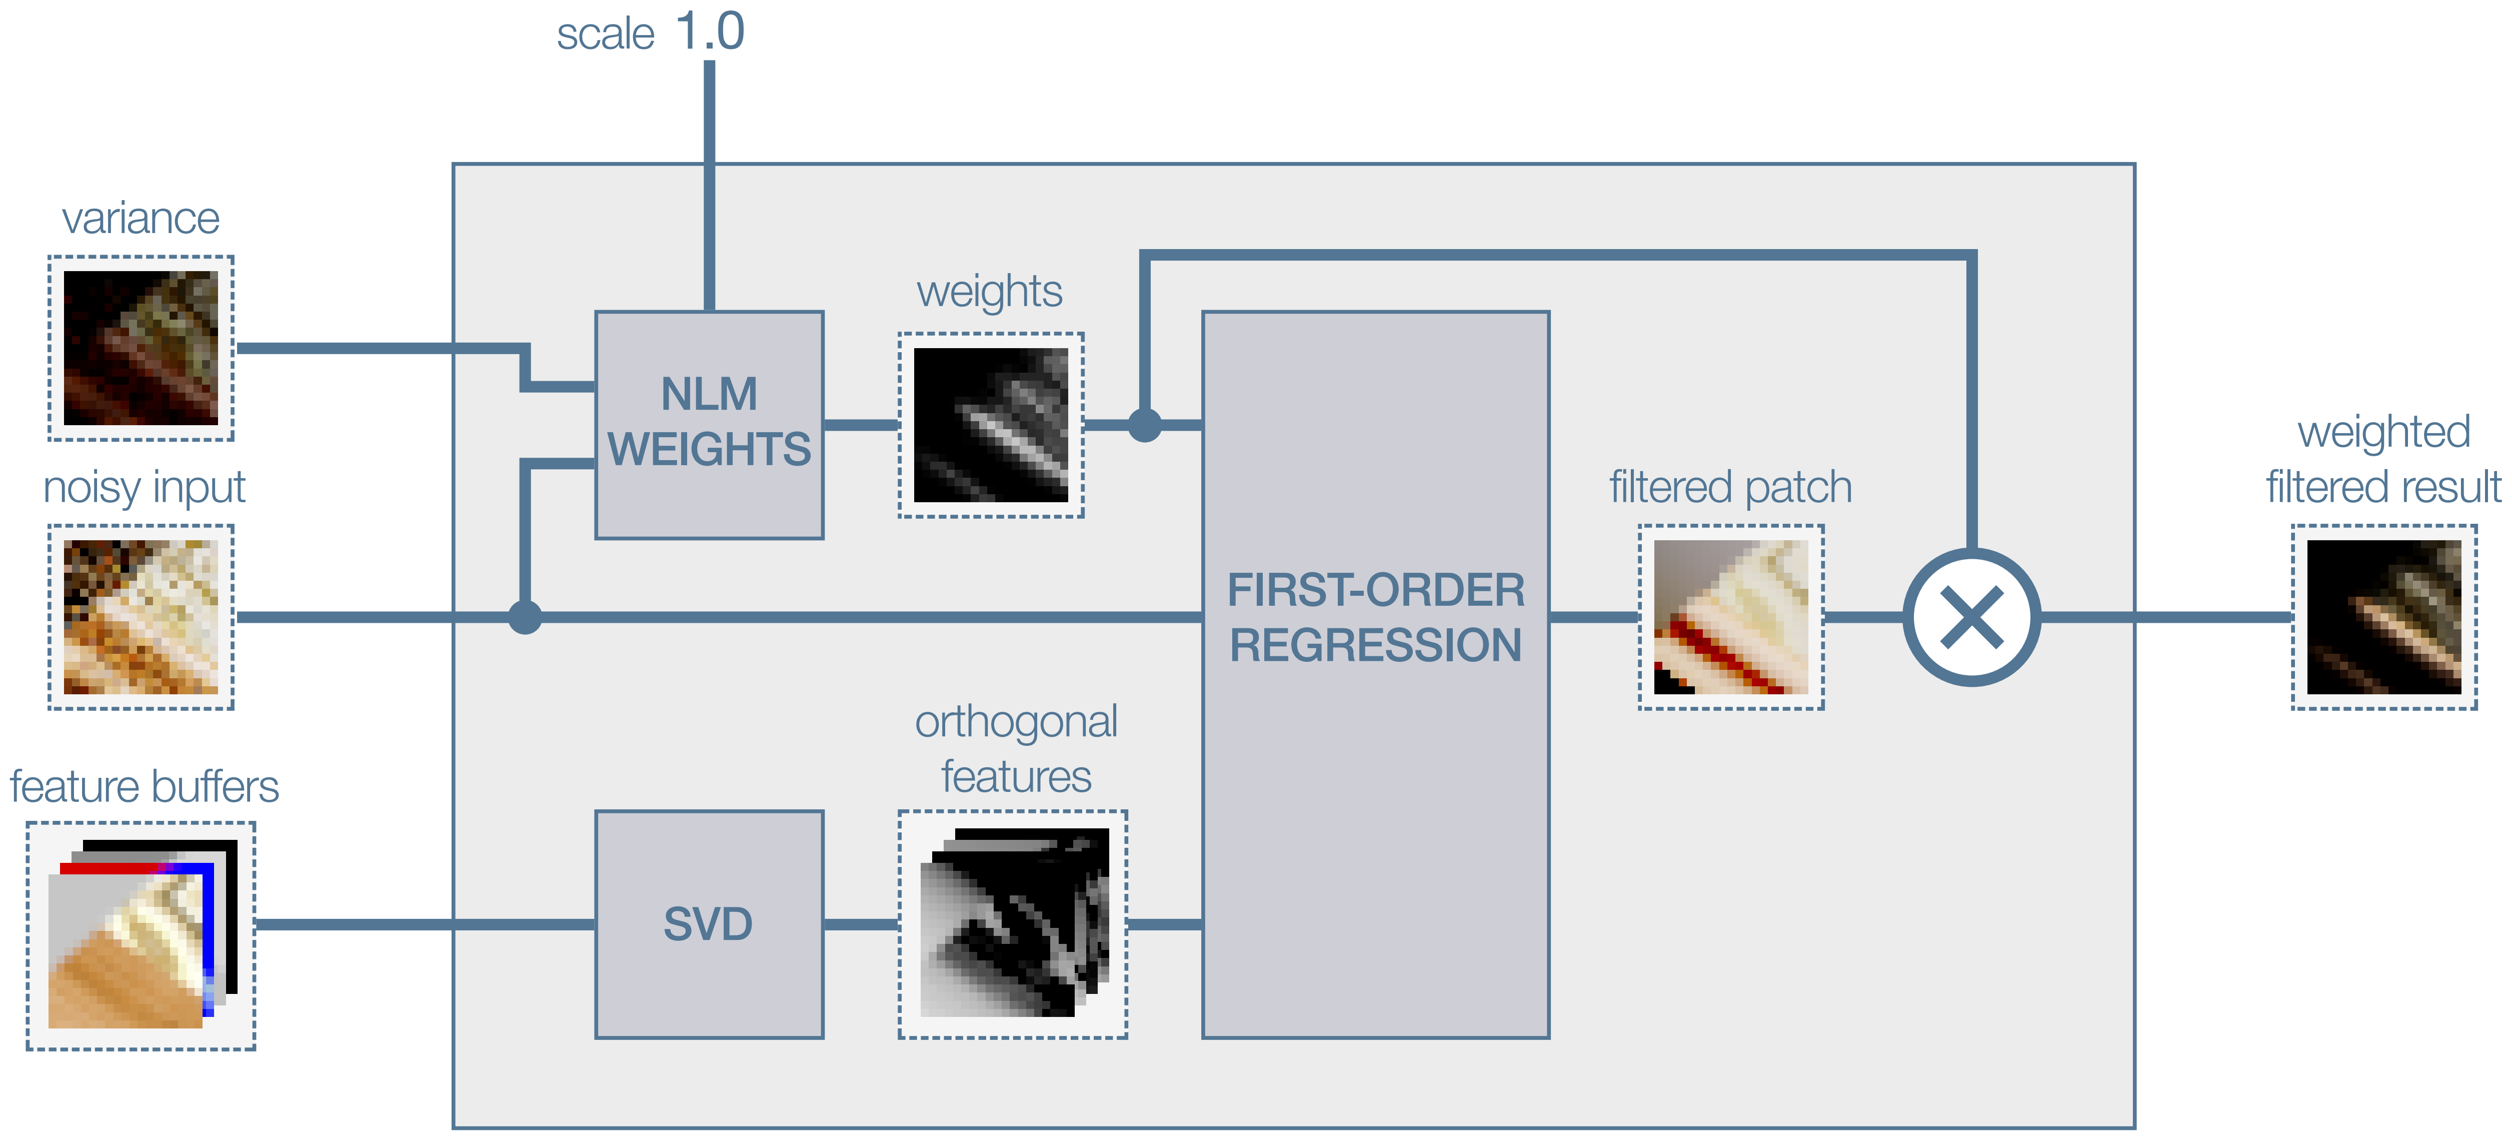

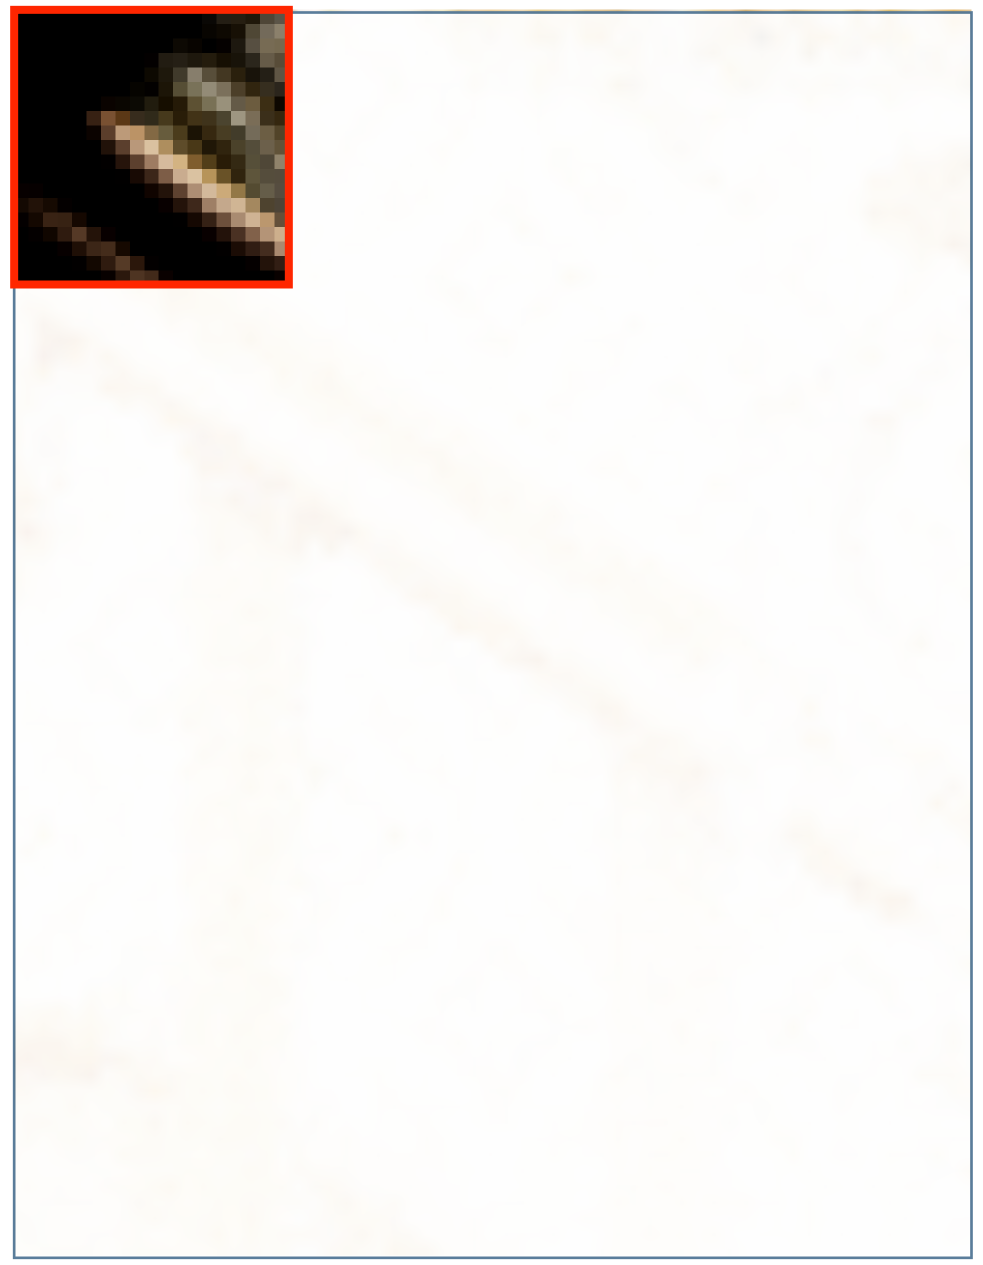

Algorithm Overview

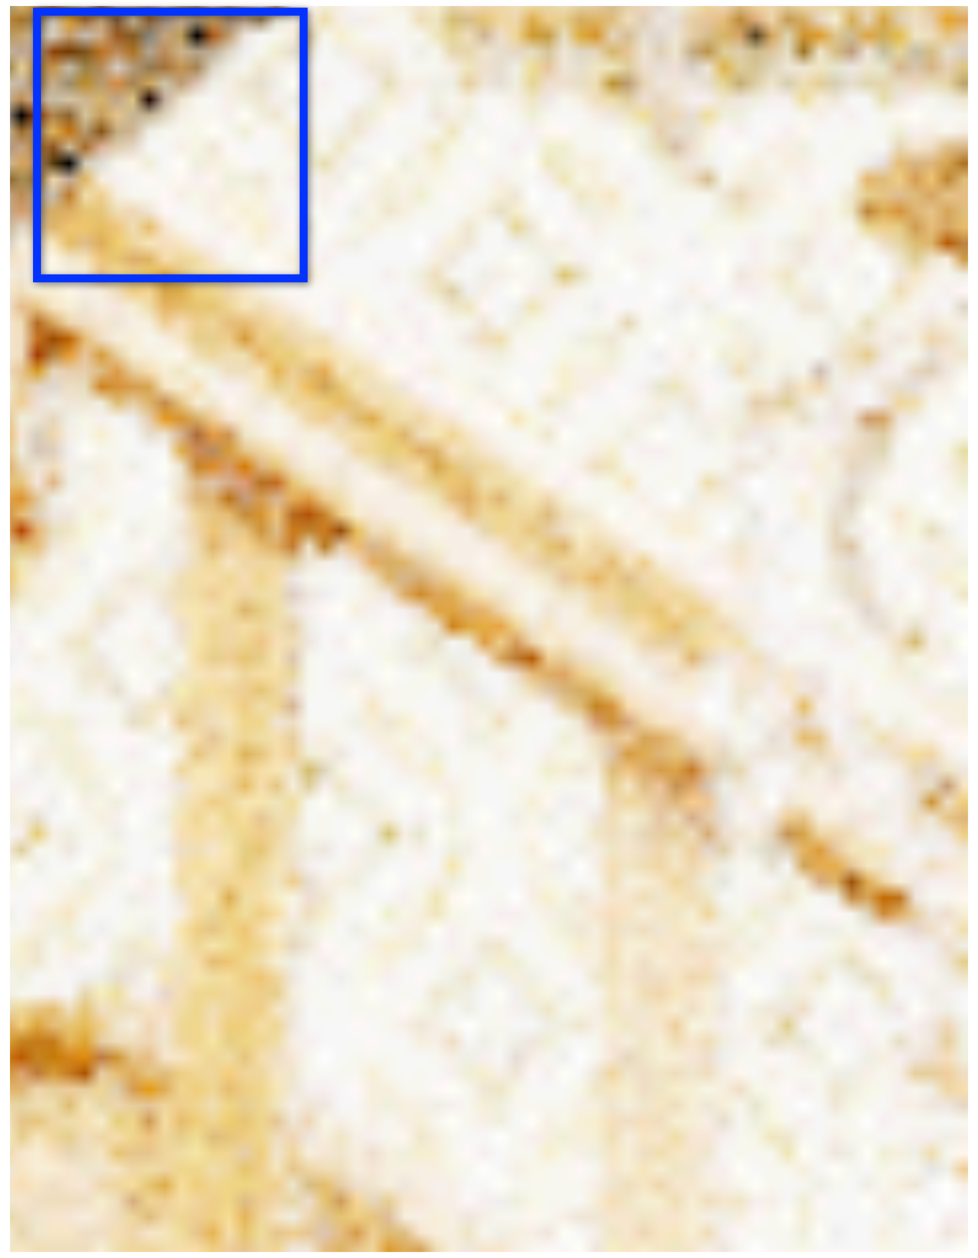

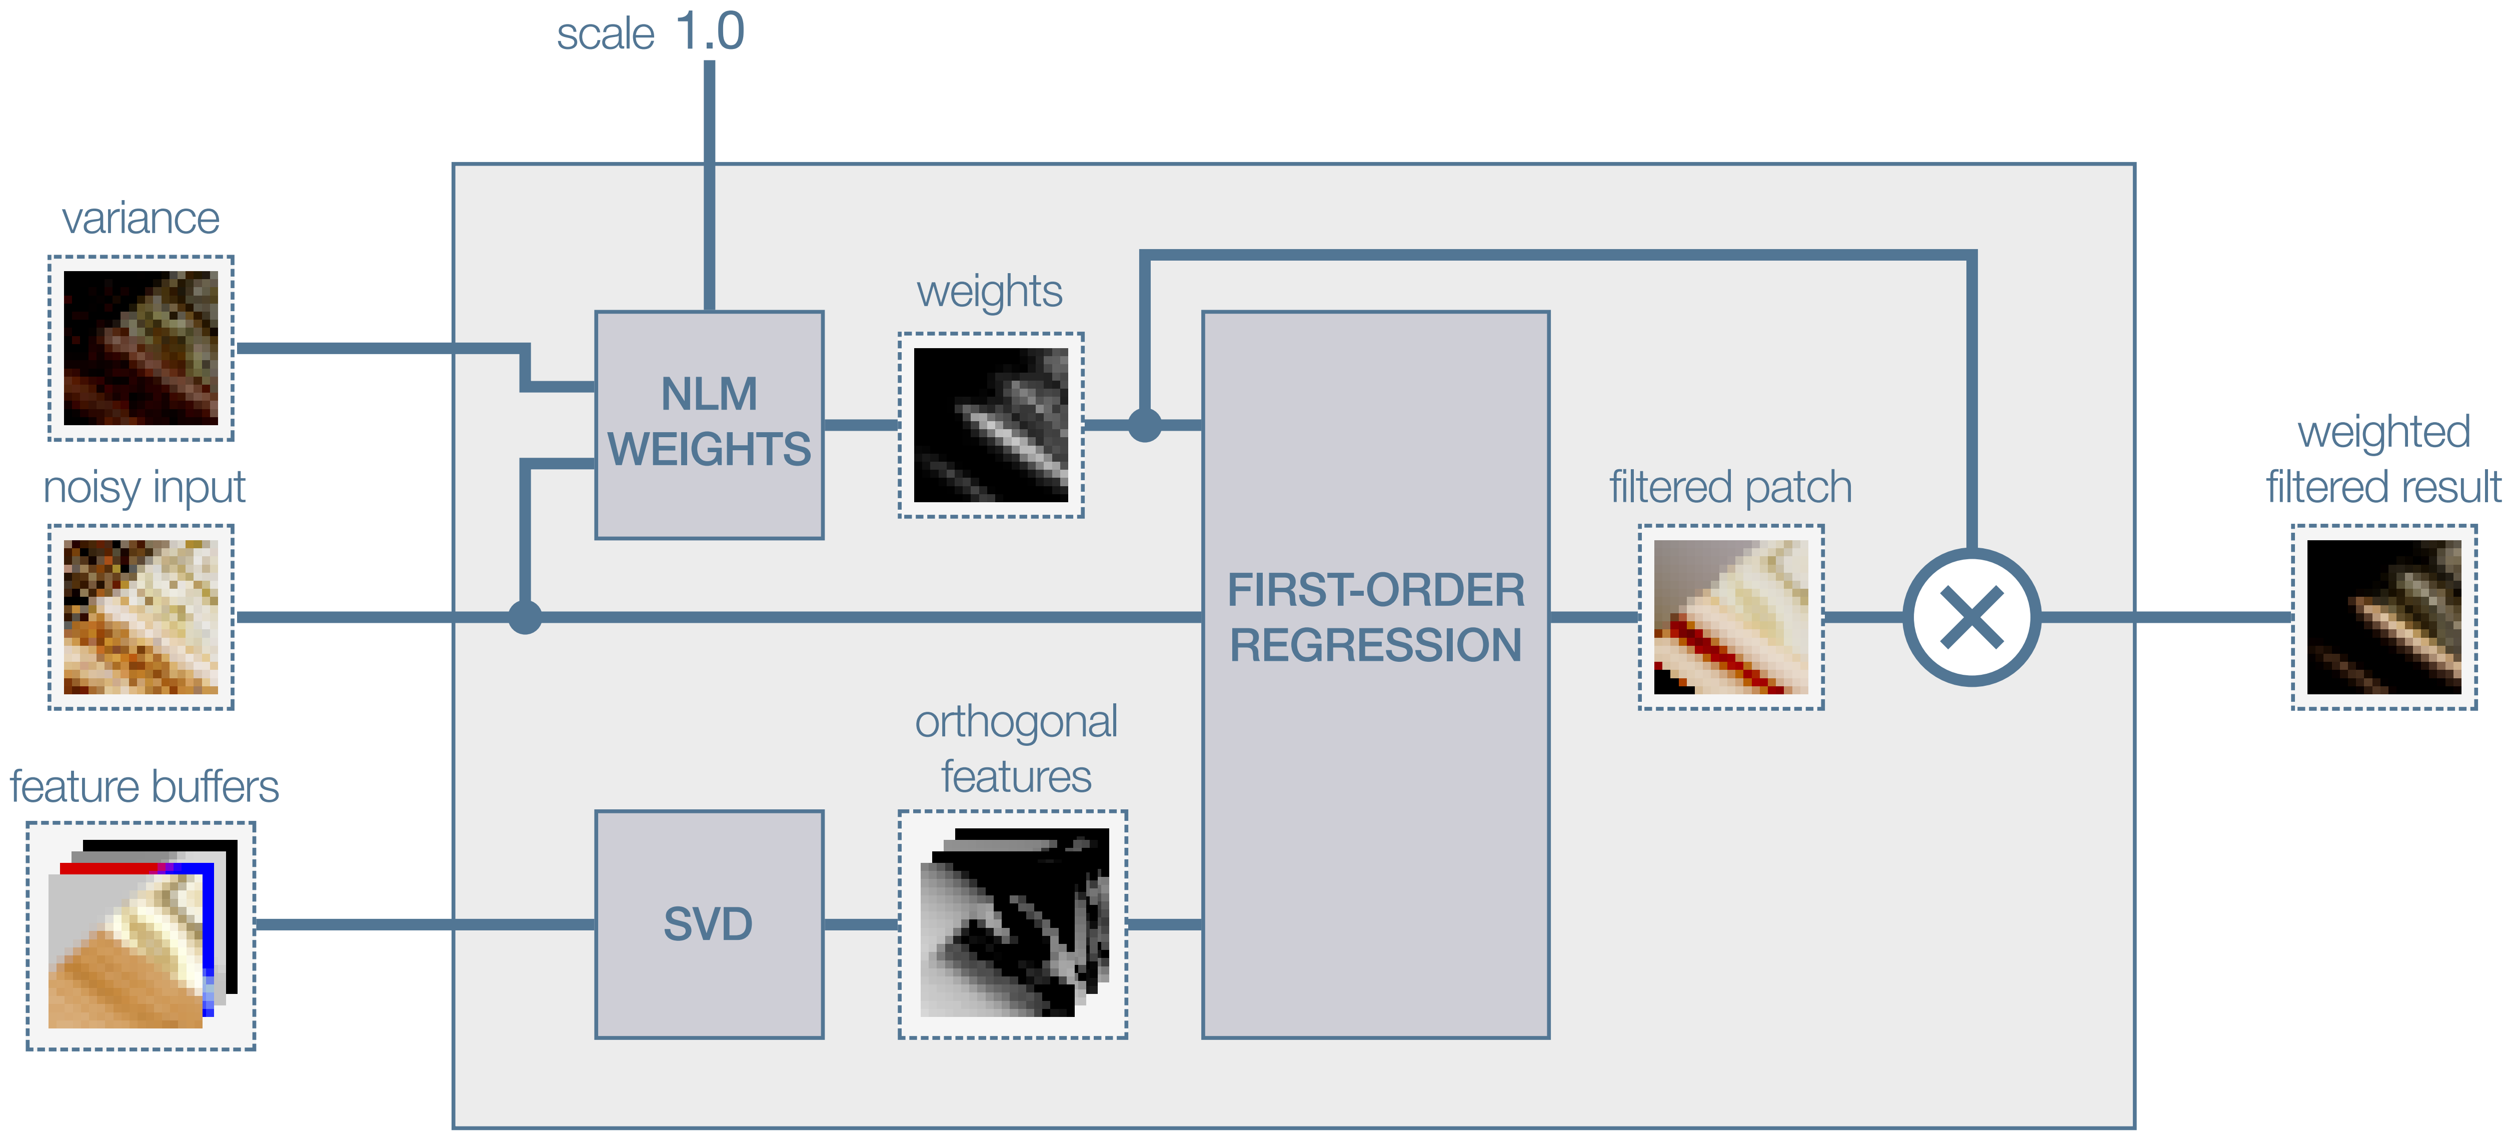

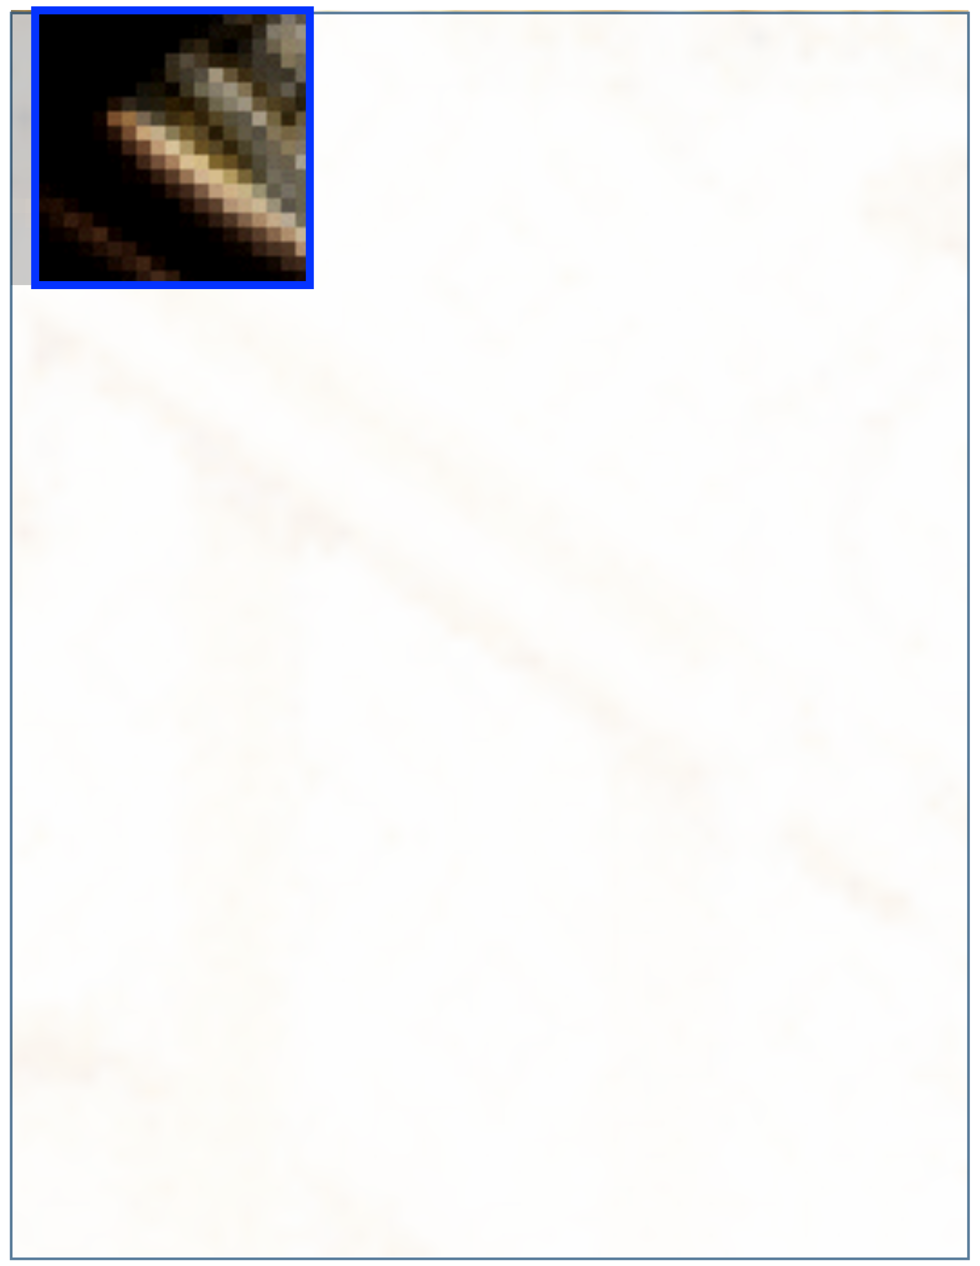

Algorithm Overview

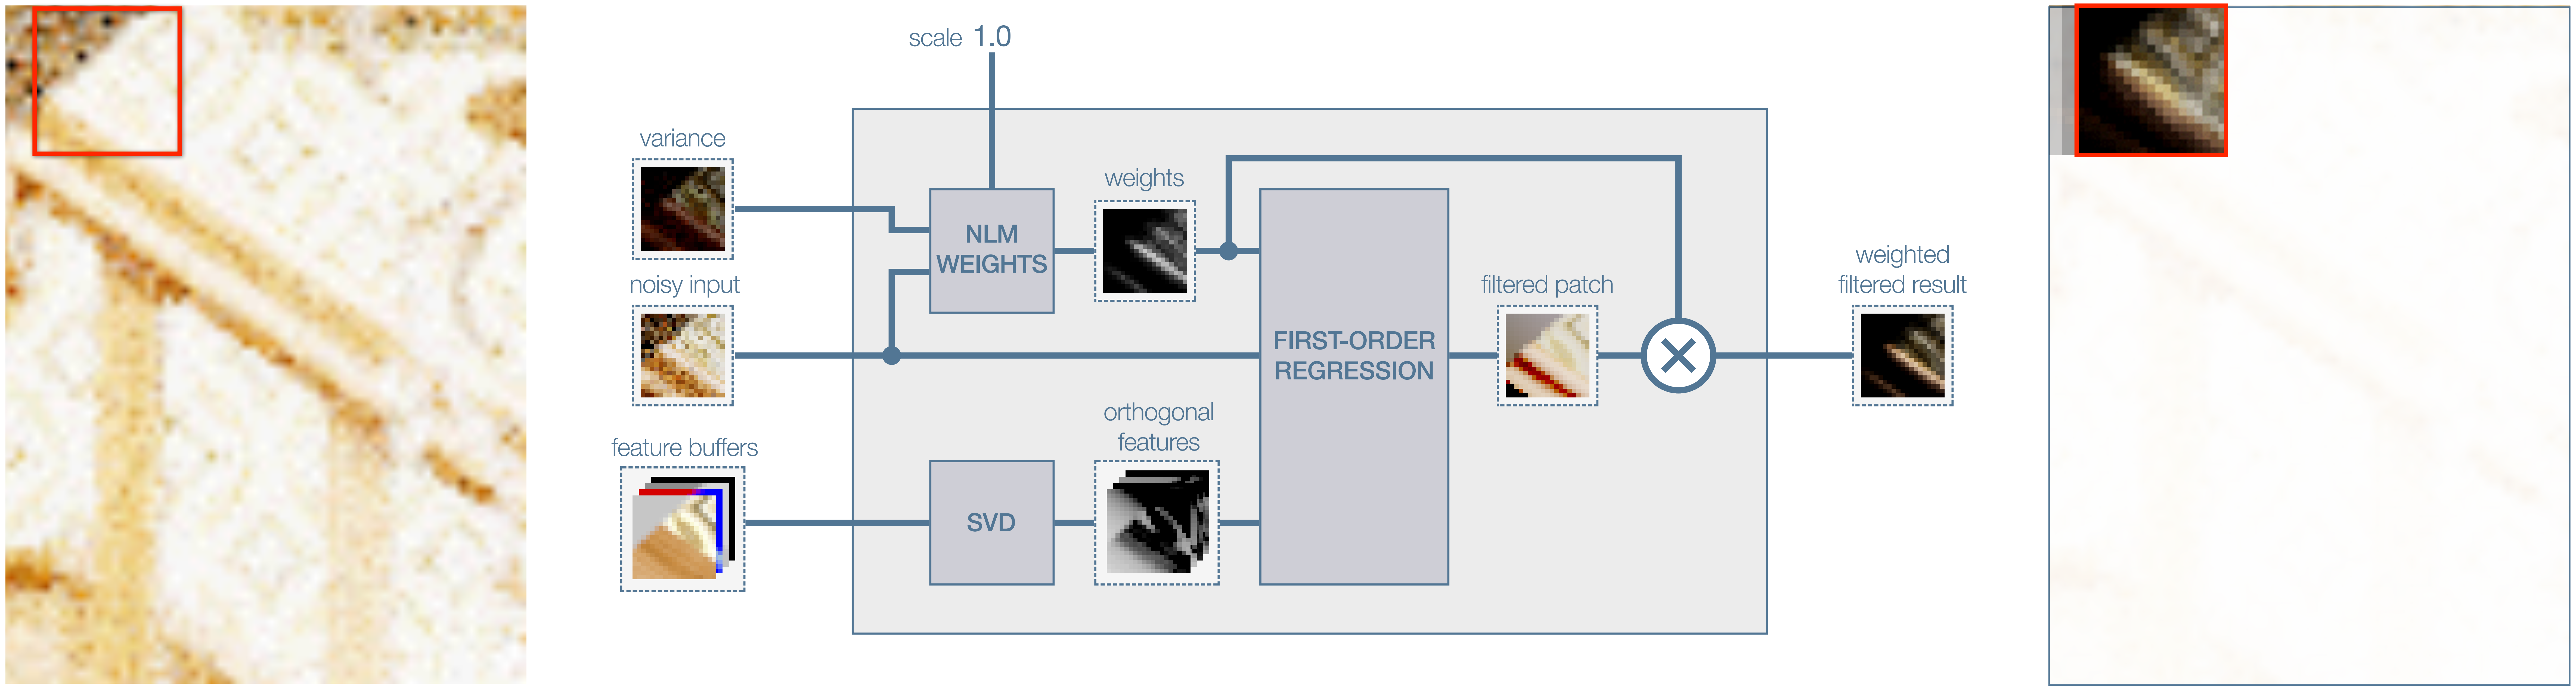

Algorithm Overview

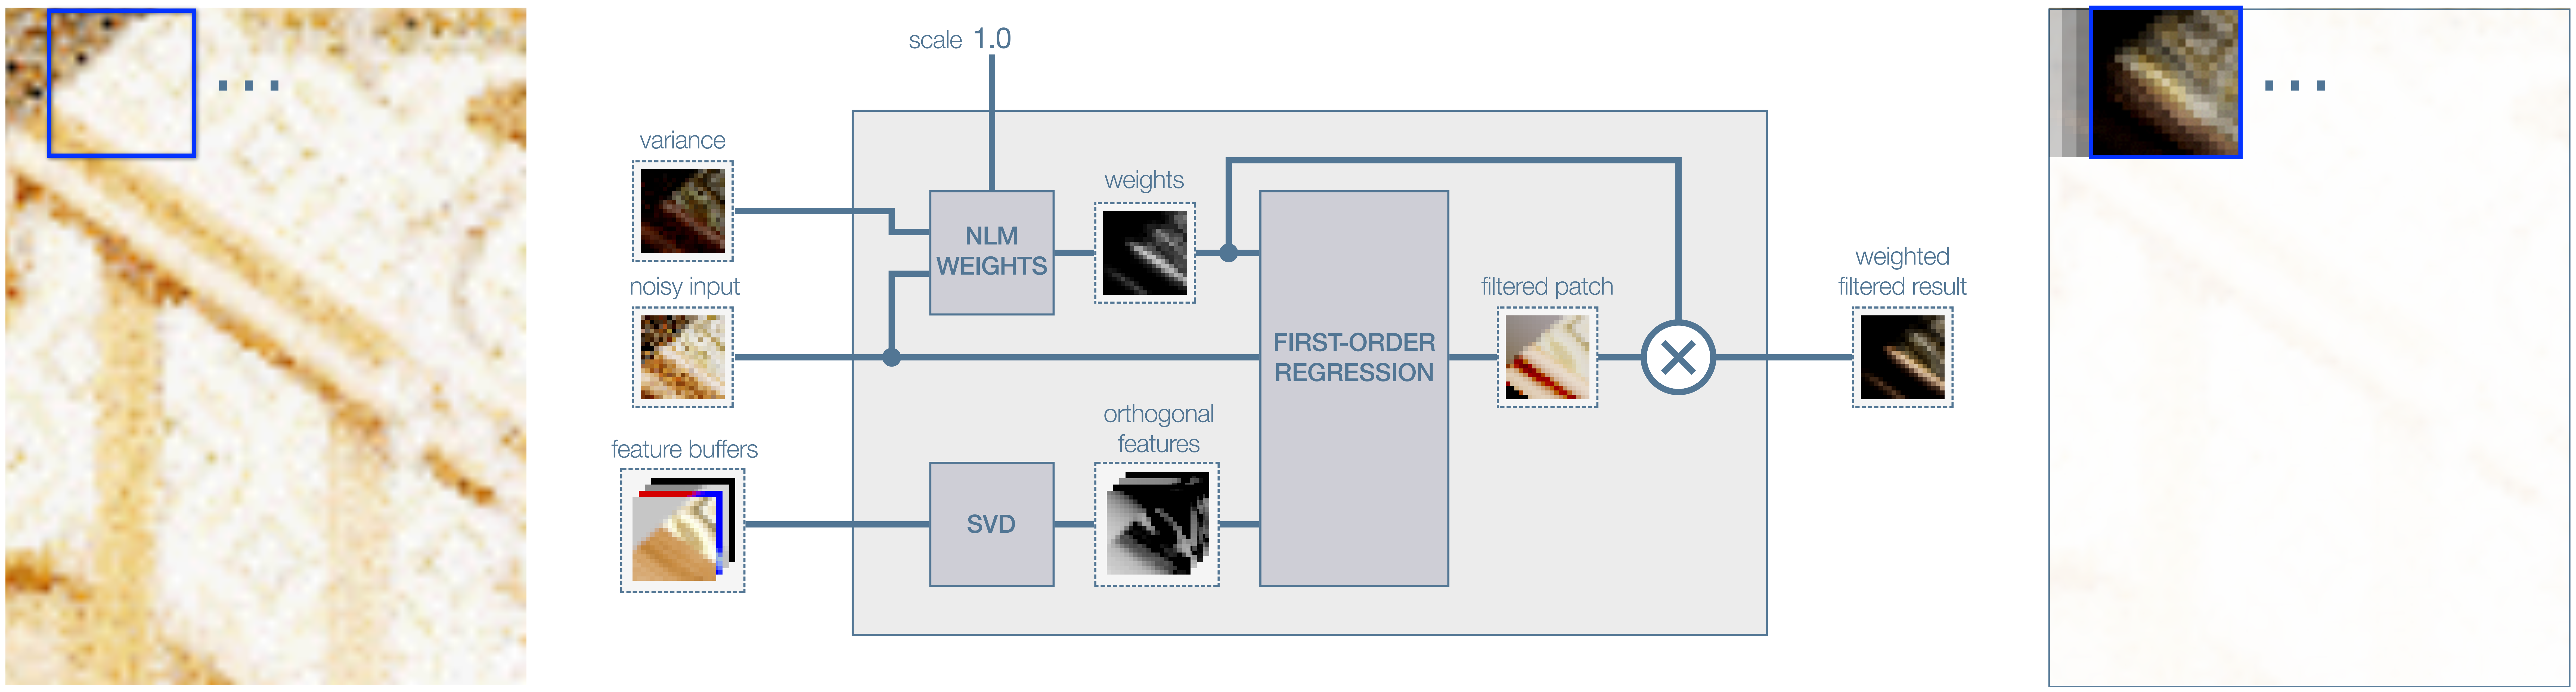

Algorithm Overview

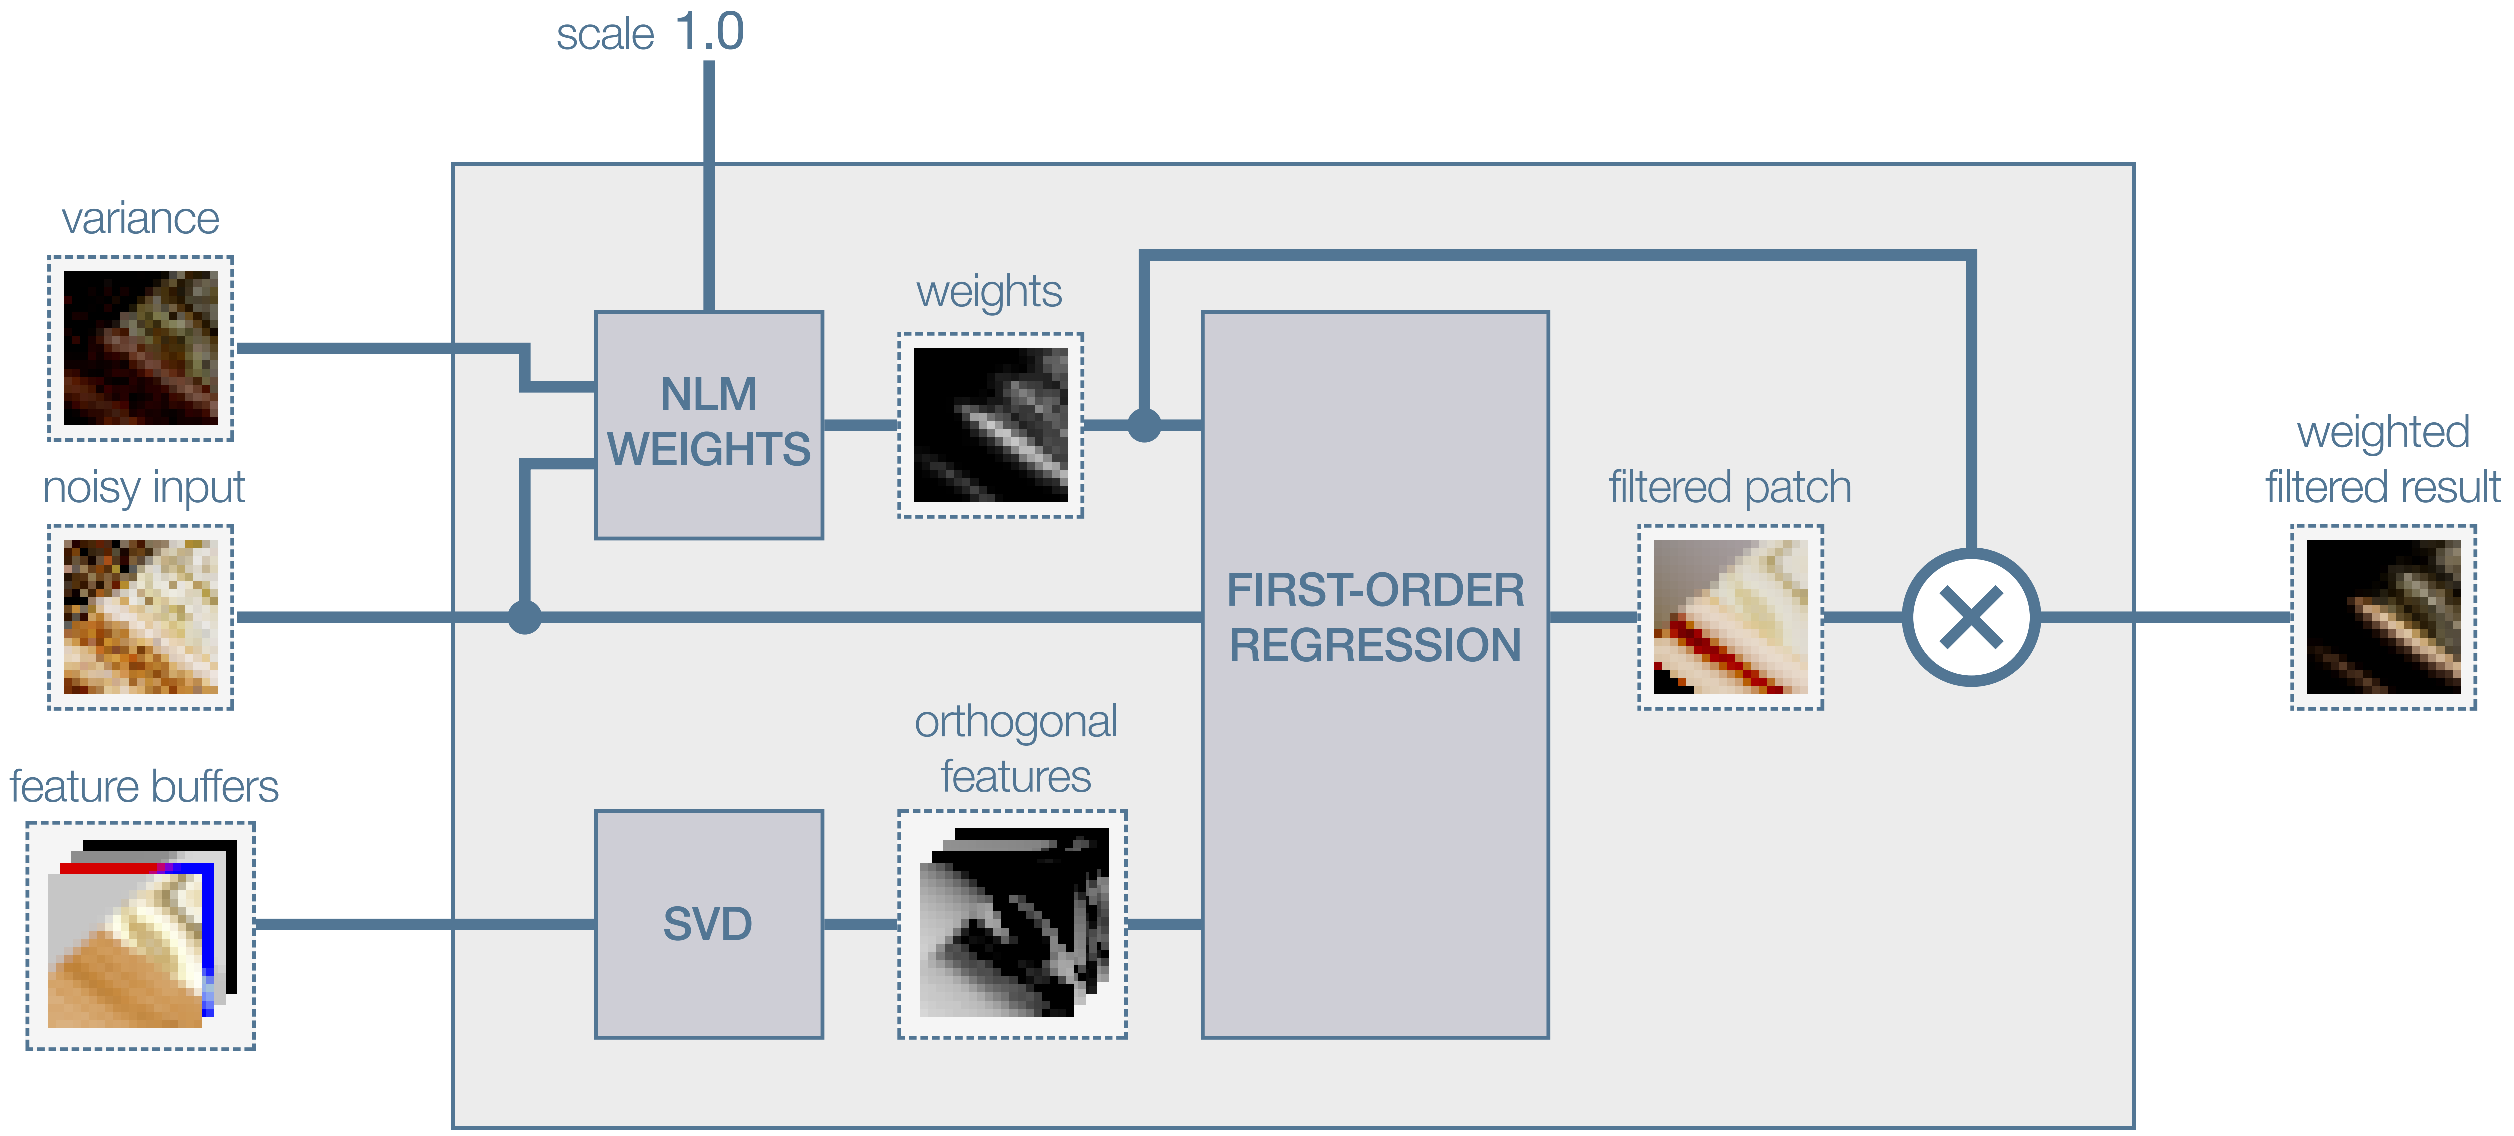

Algorithm Overview

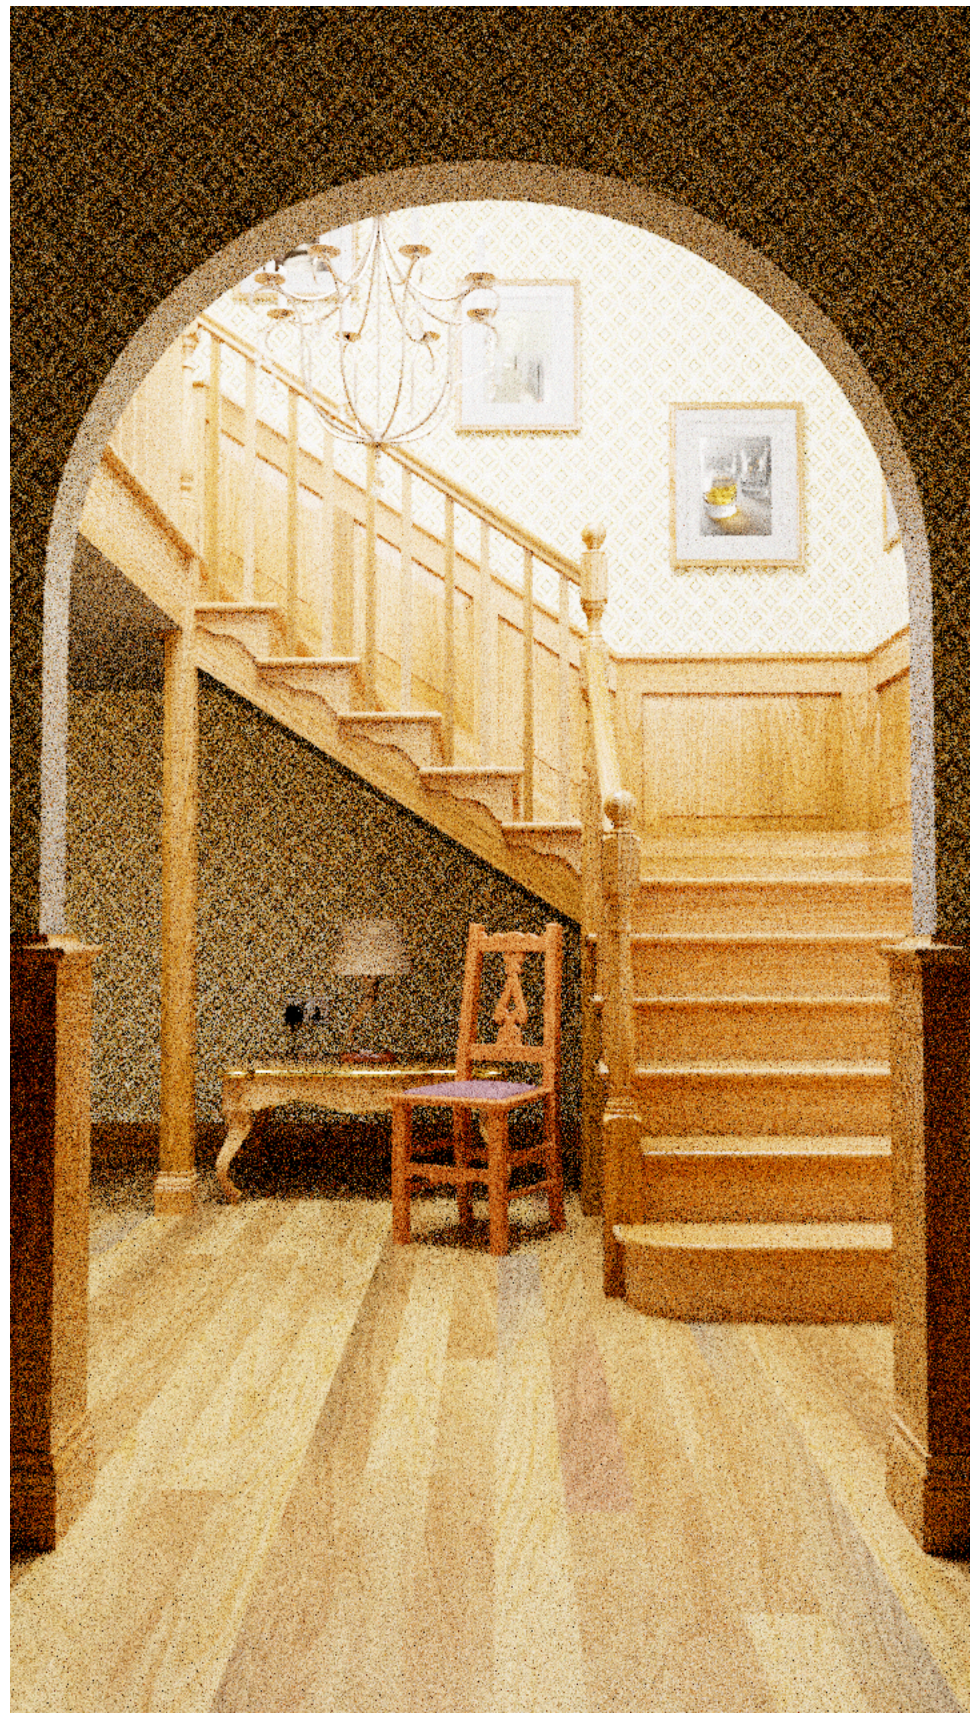

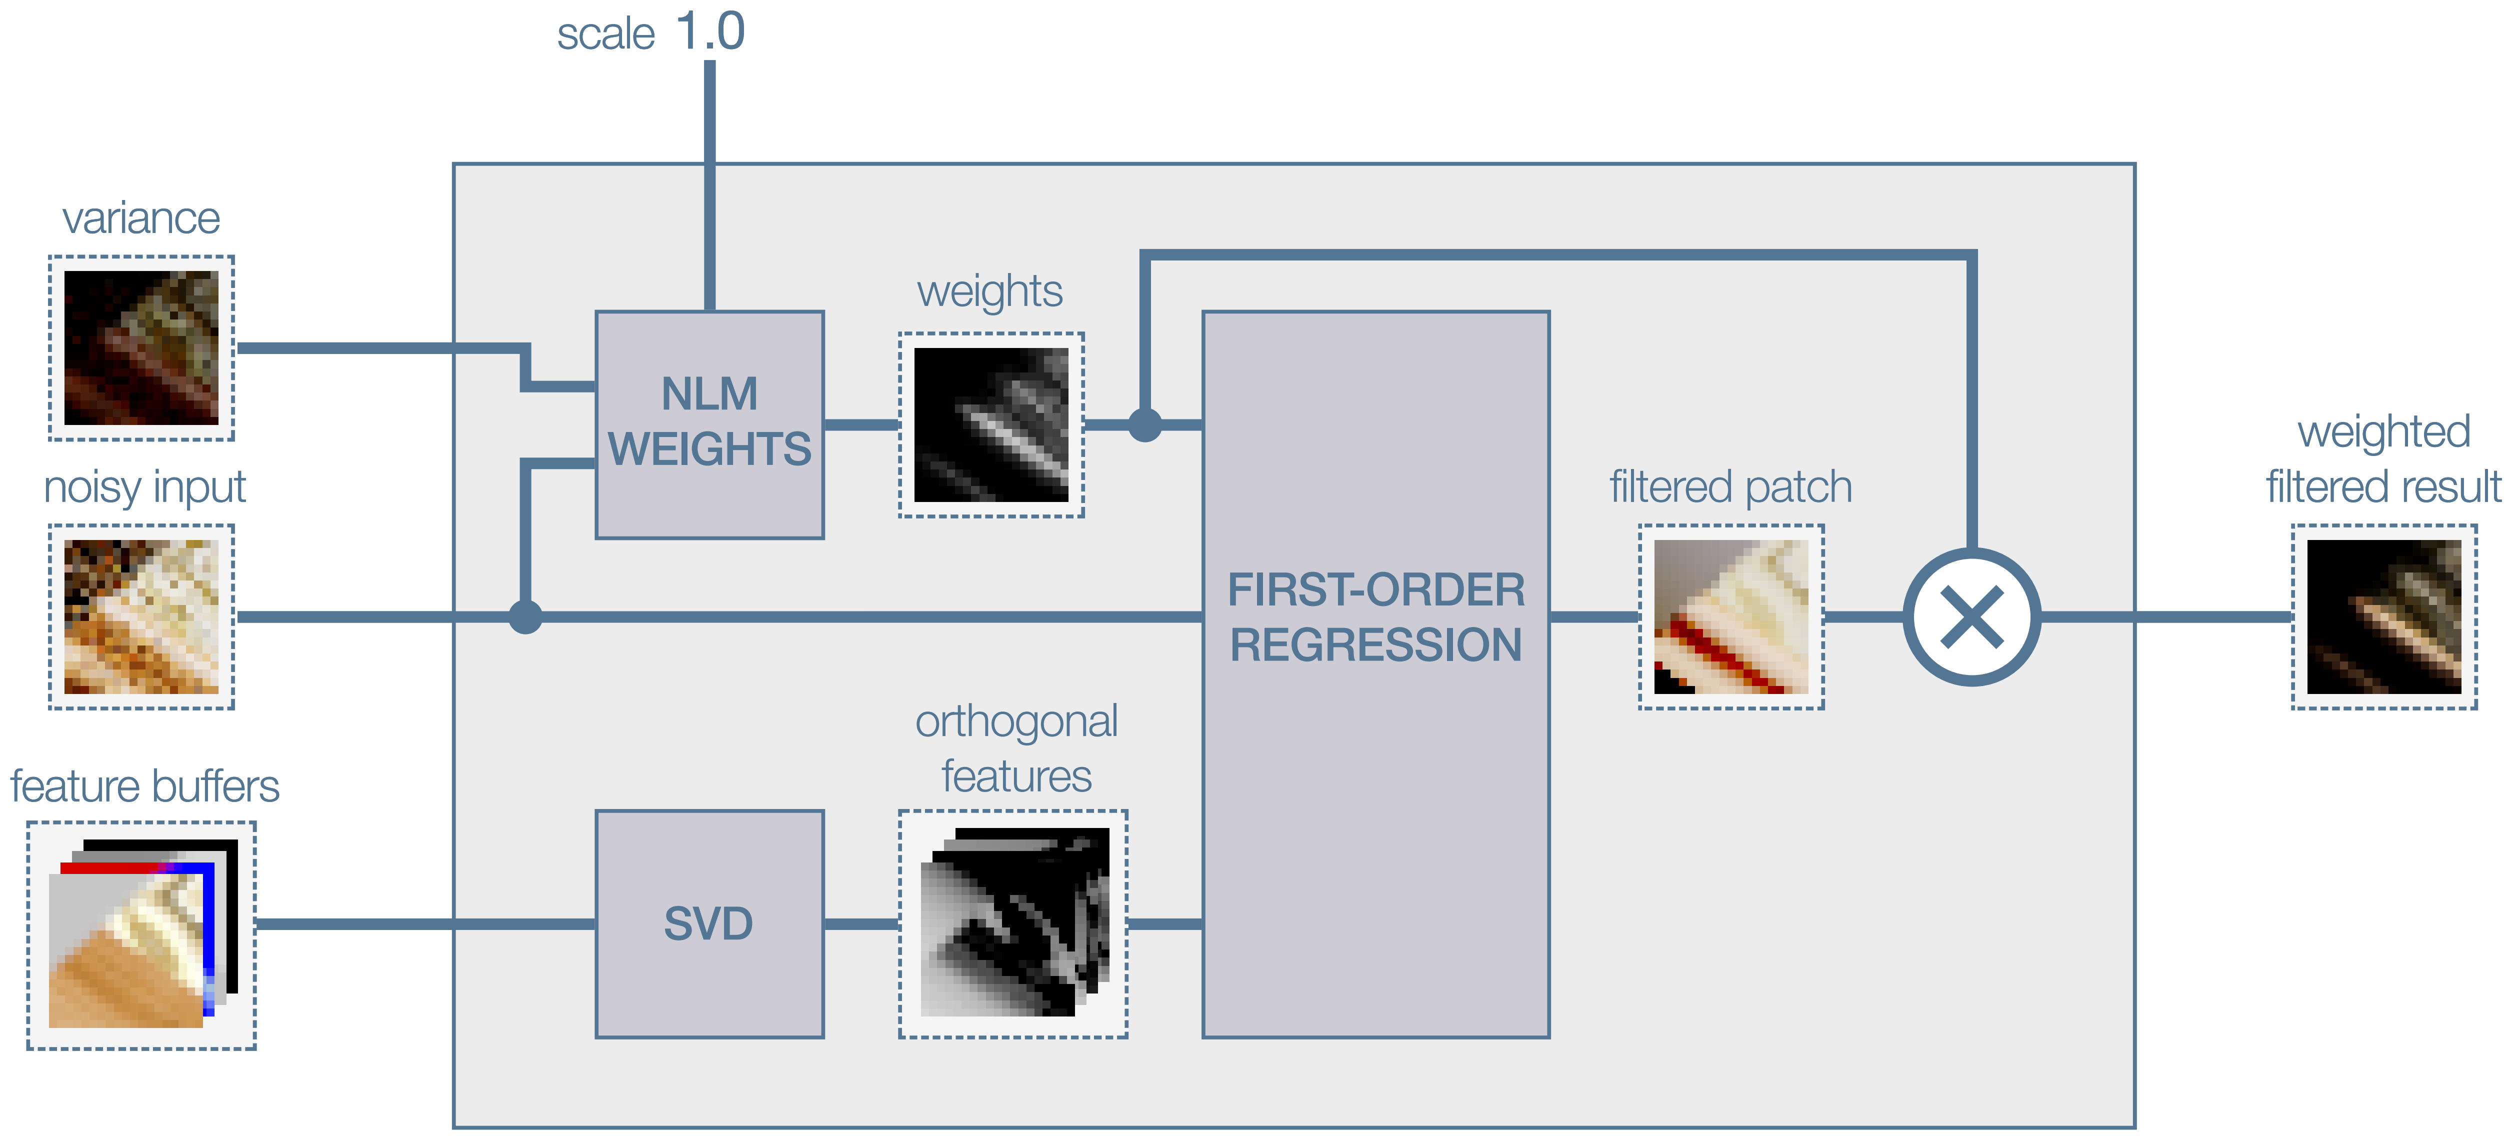

Algorithm Overview

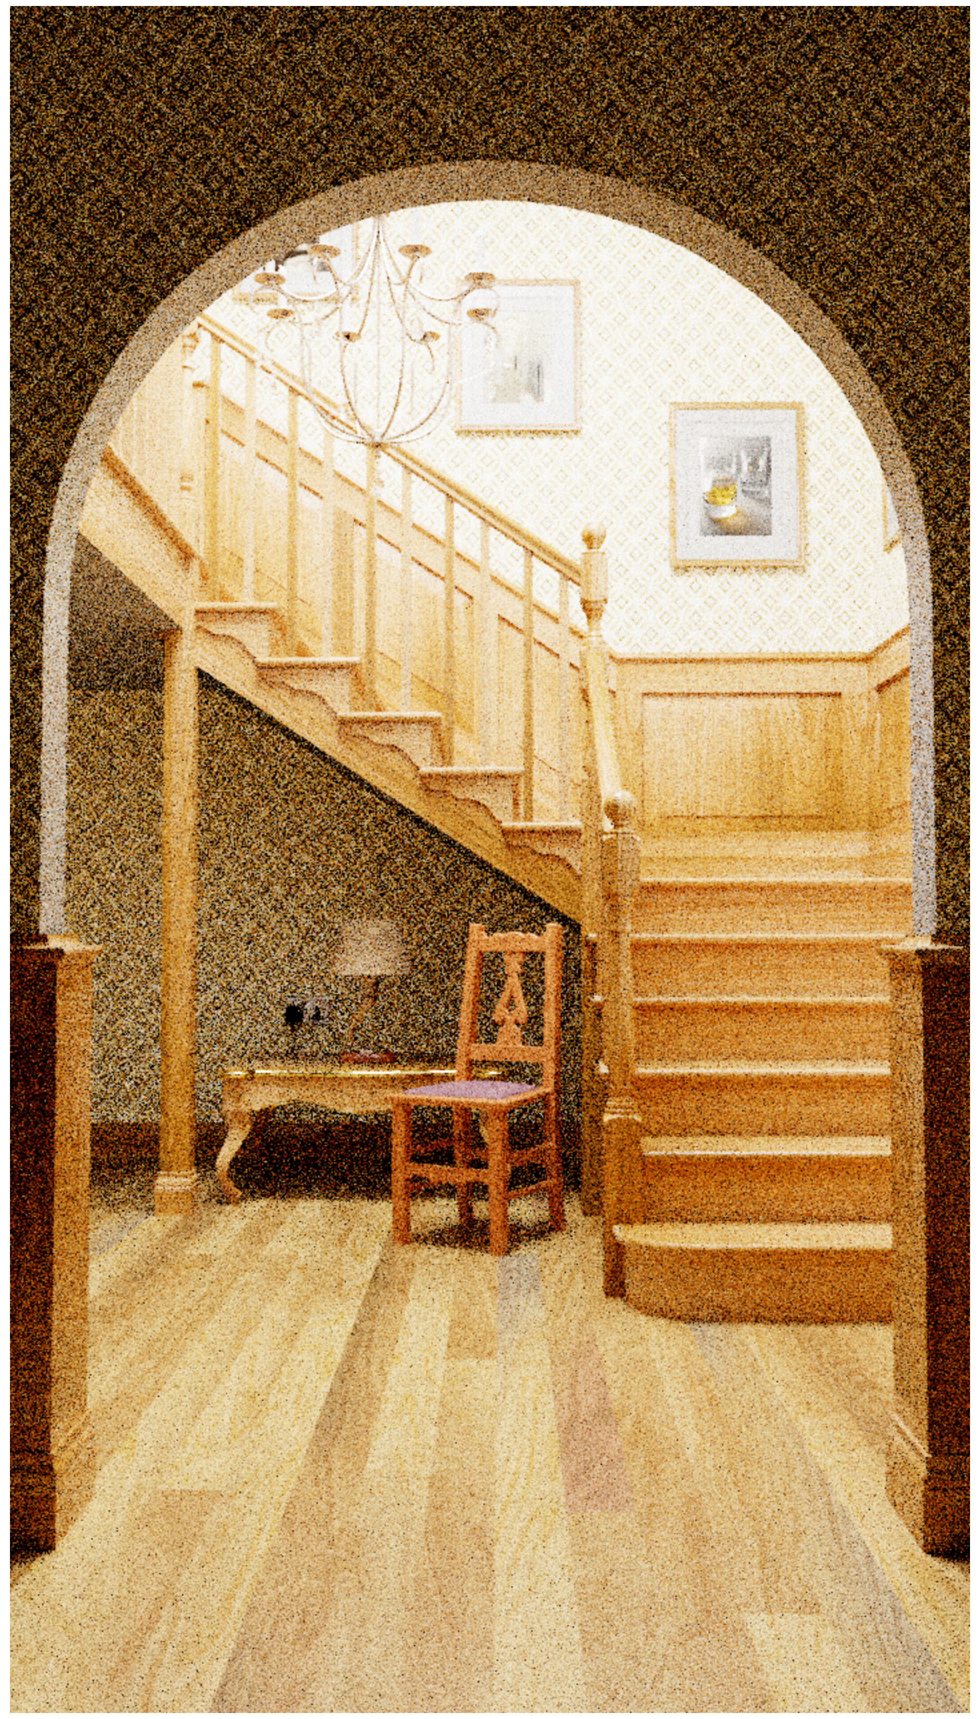

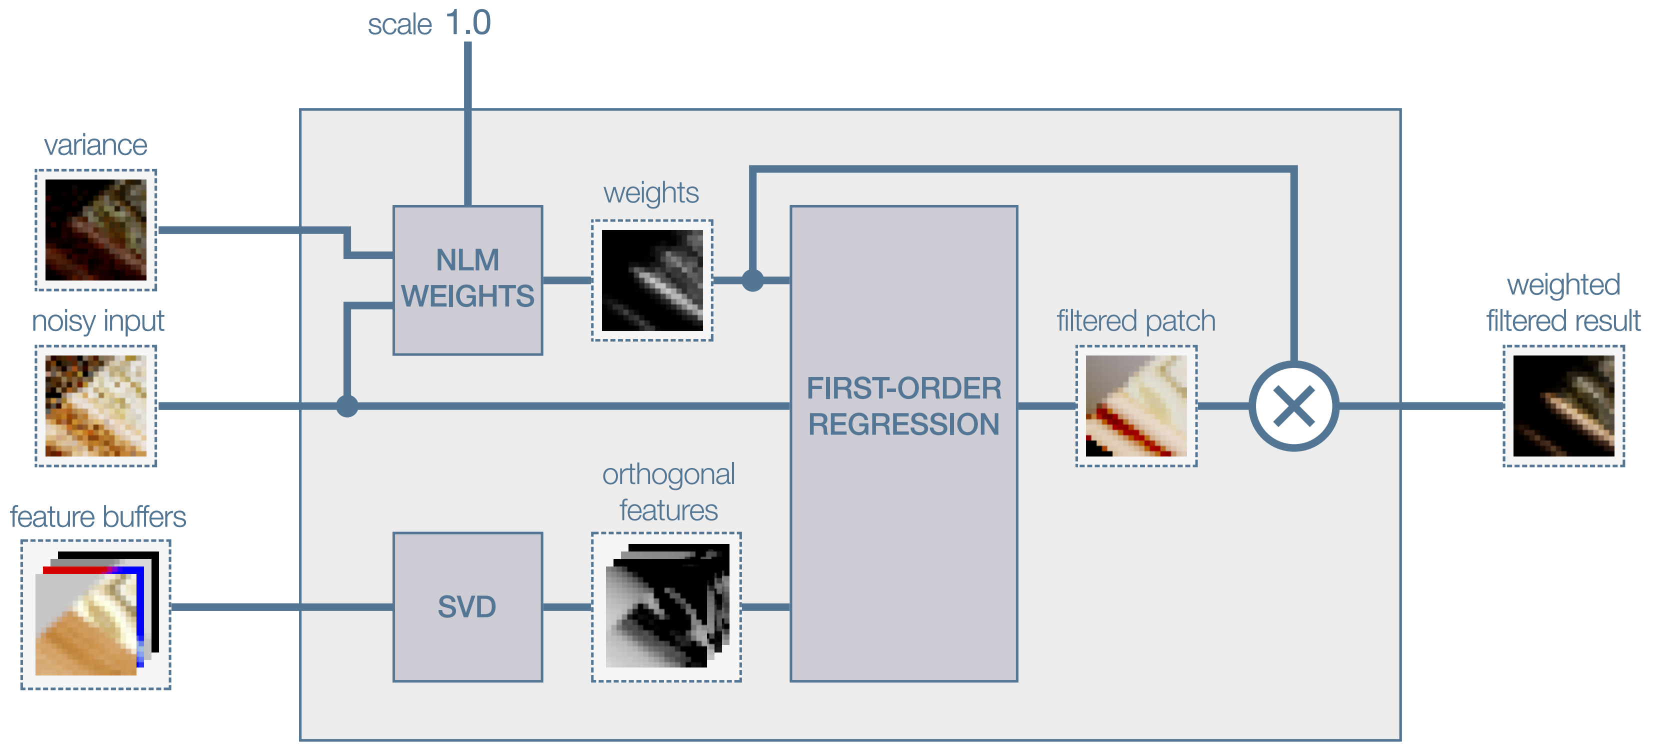

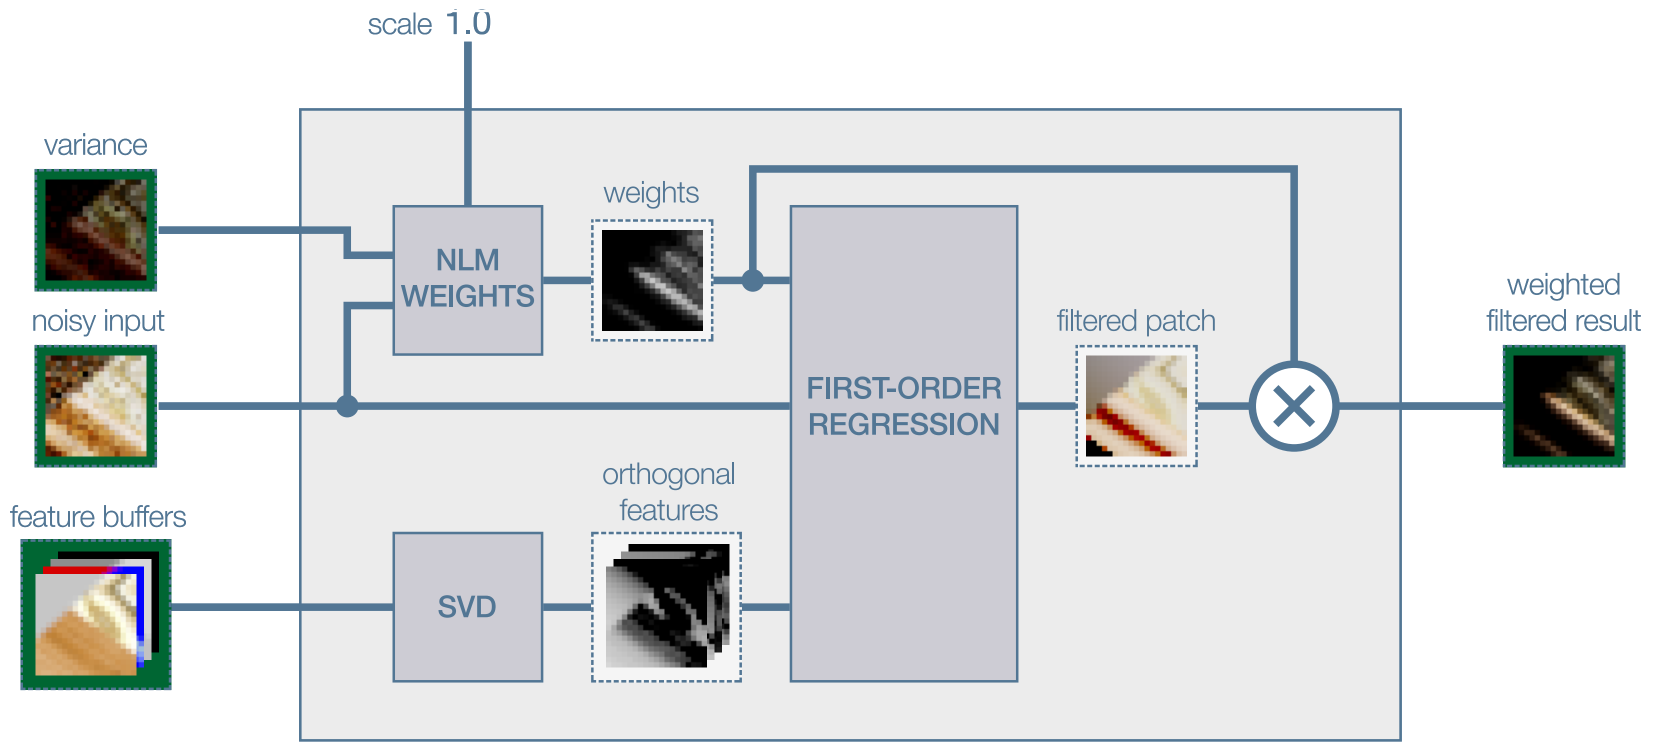

Disney Research

Algorithm Overview

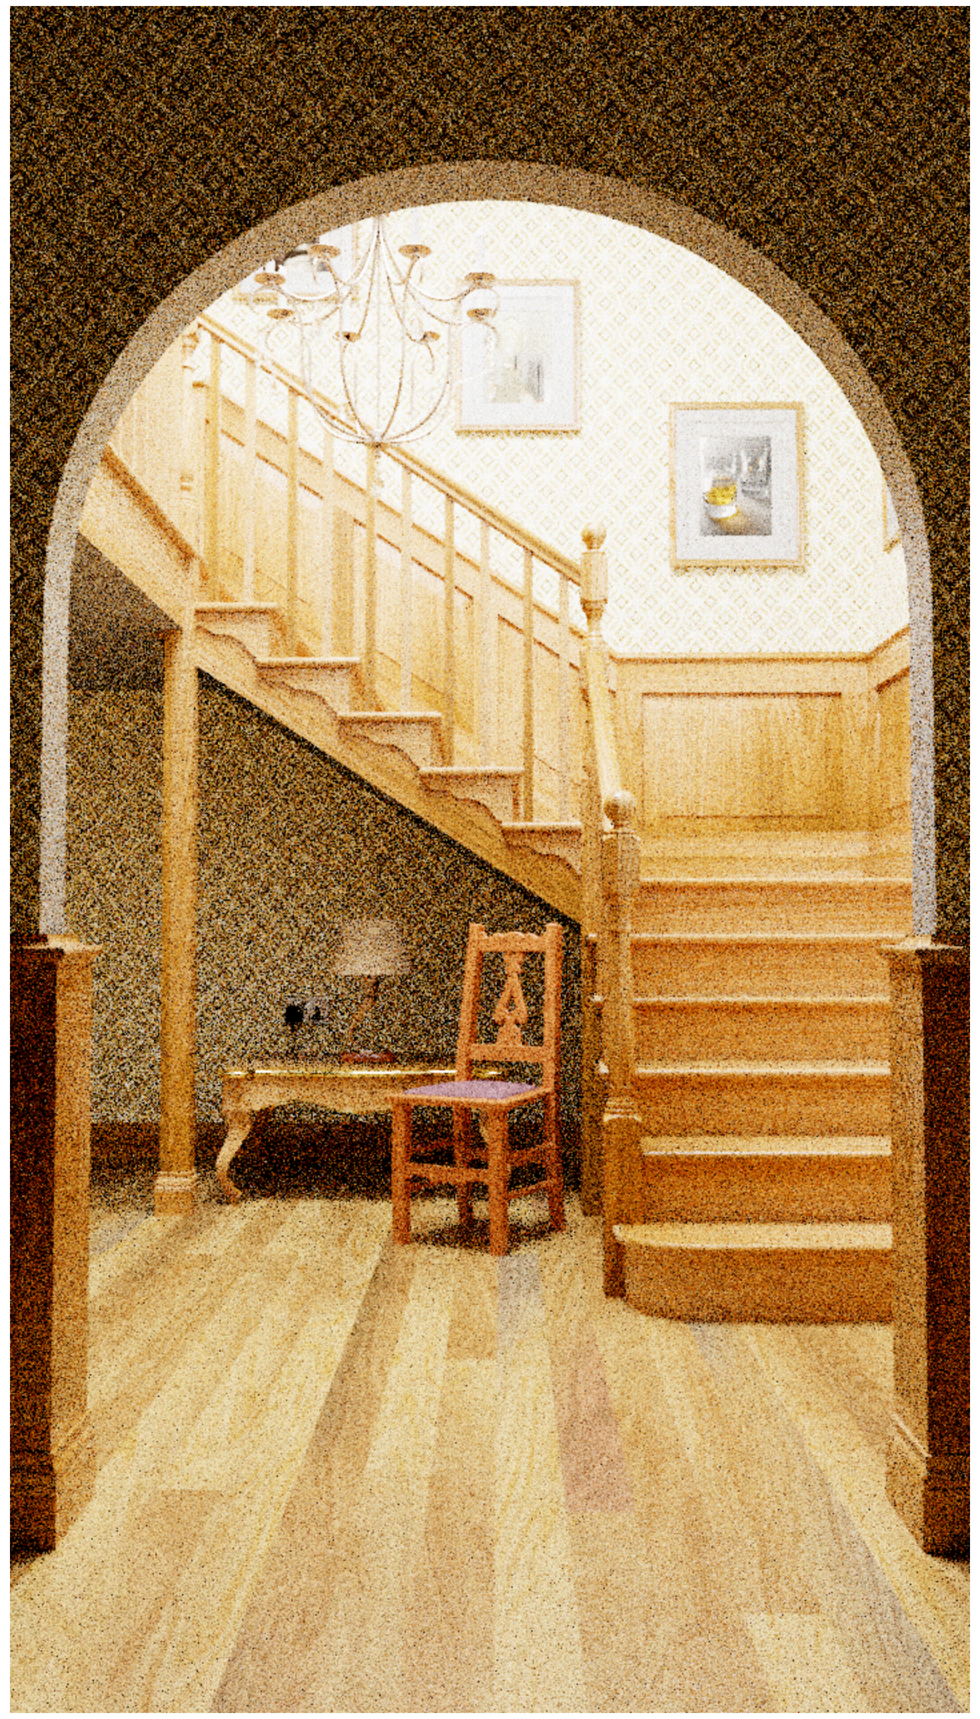

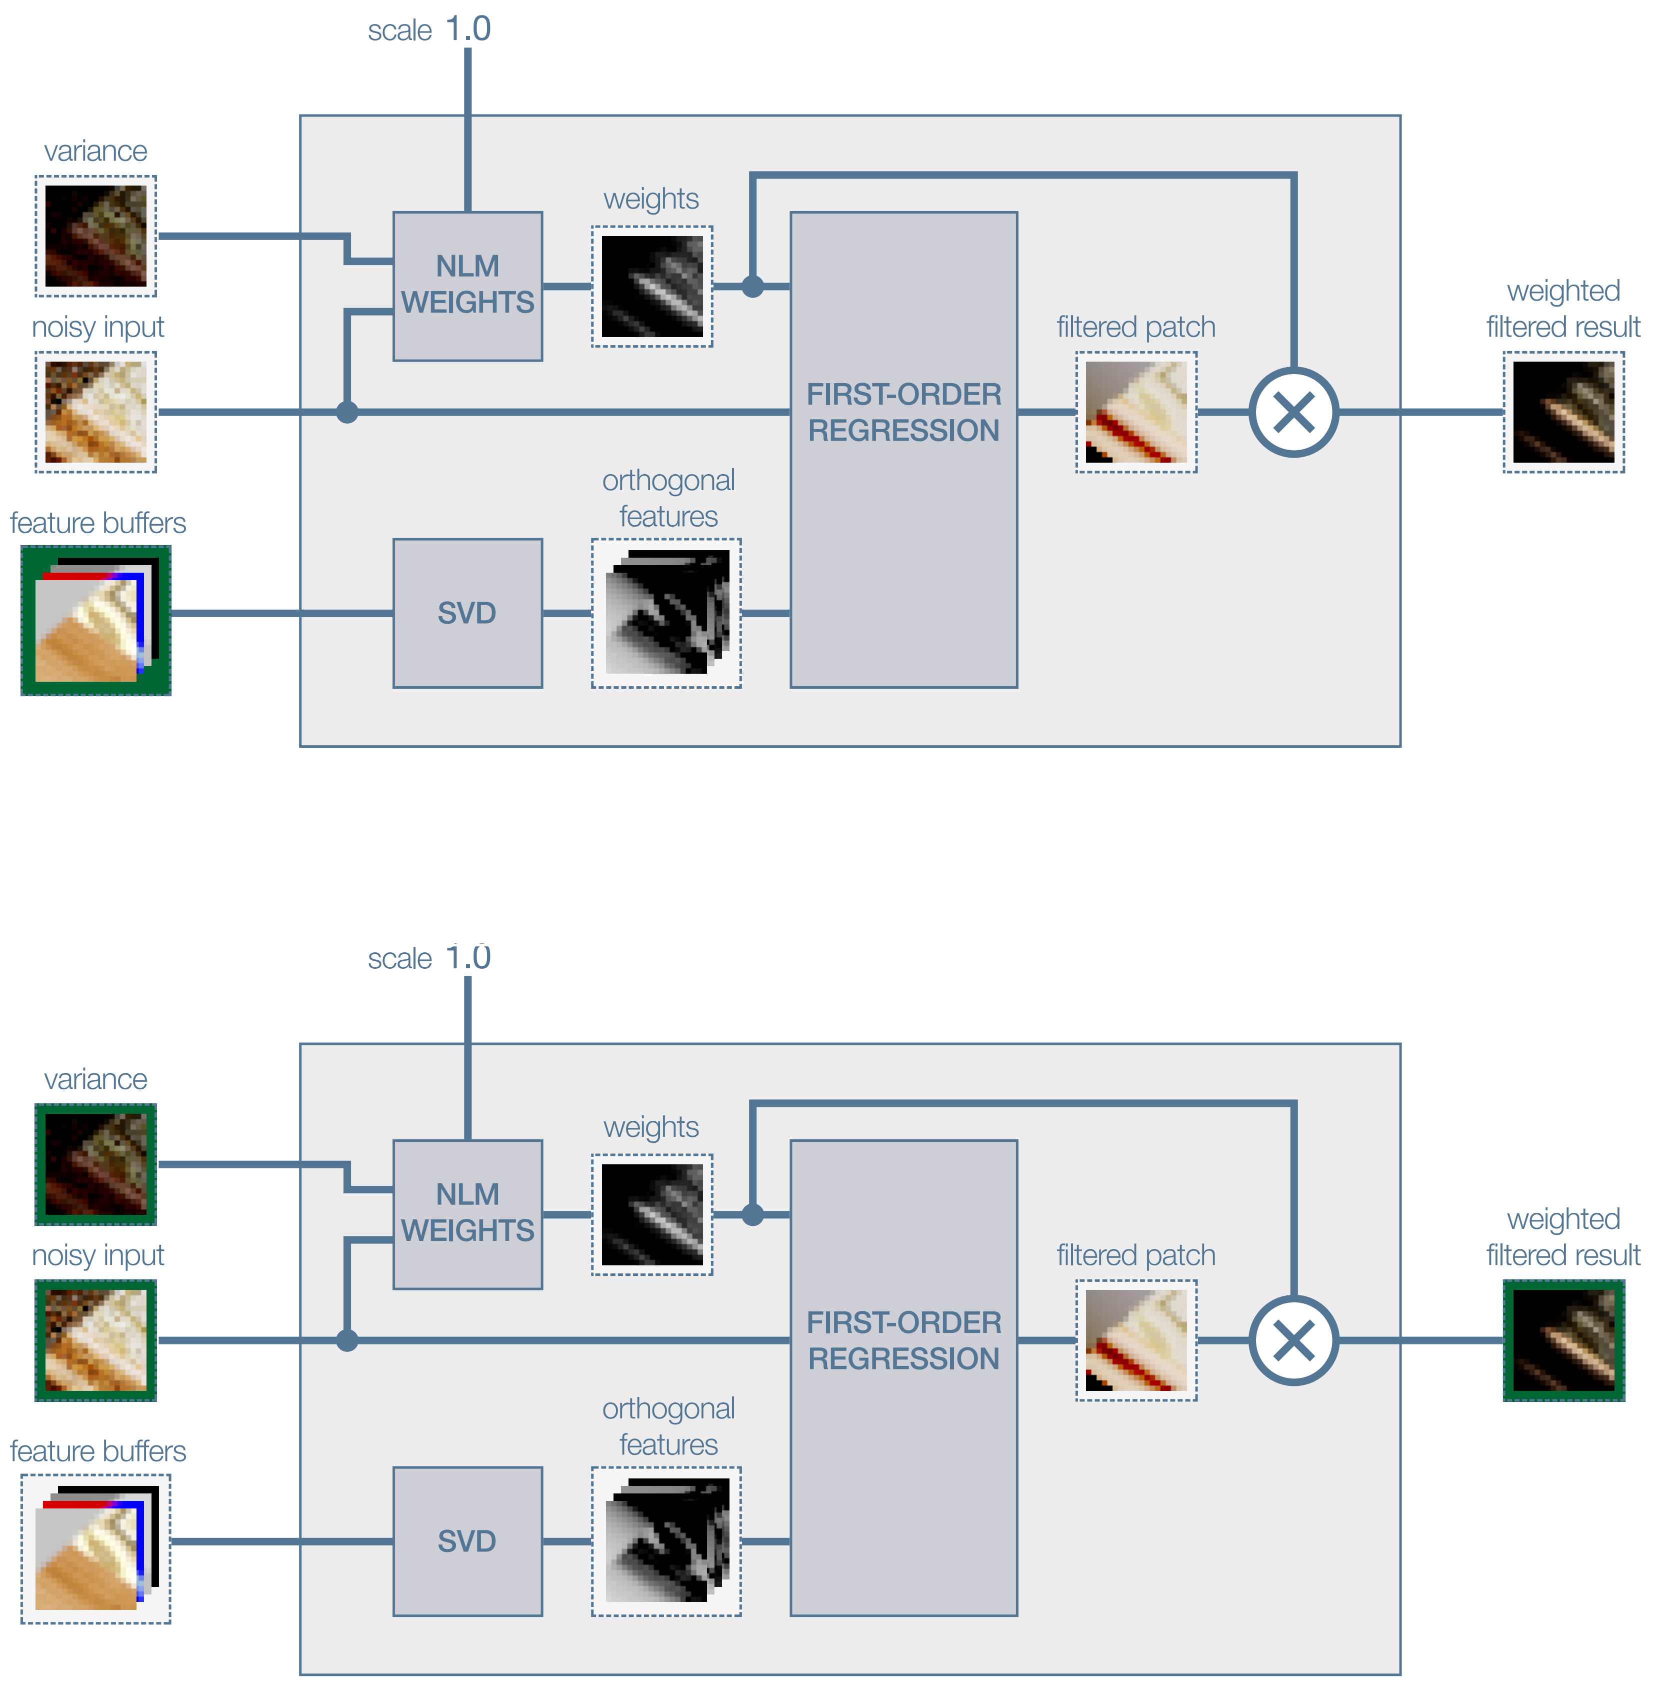

Algorithm Overview

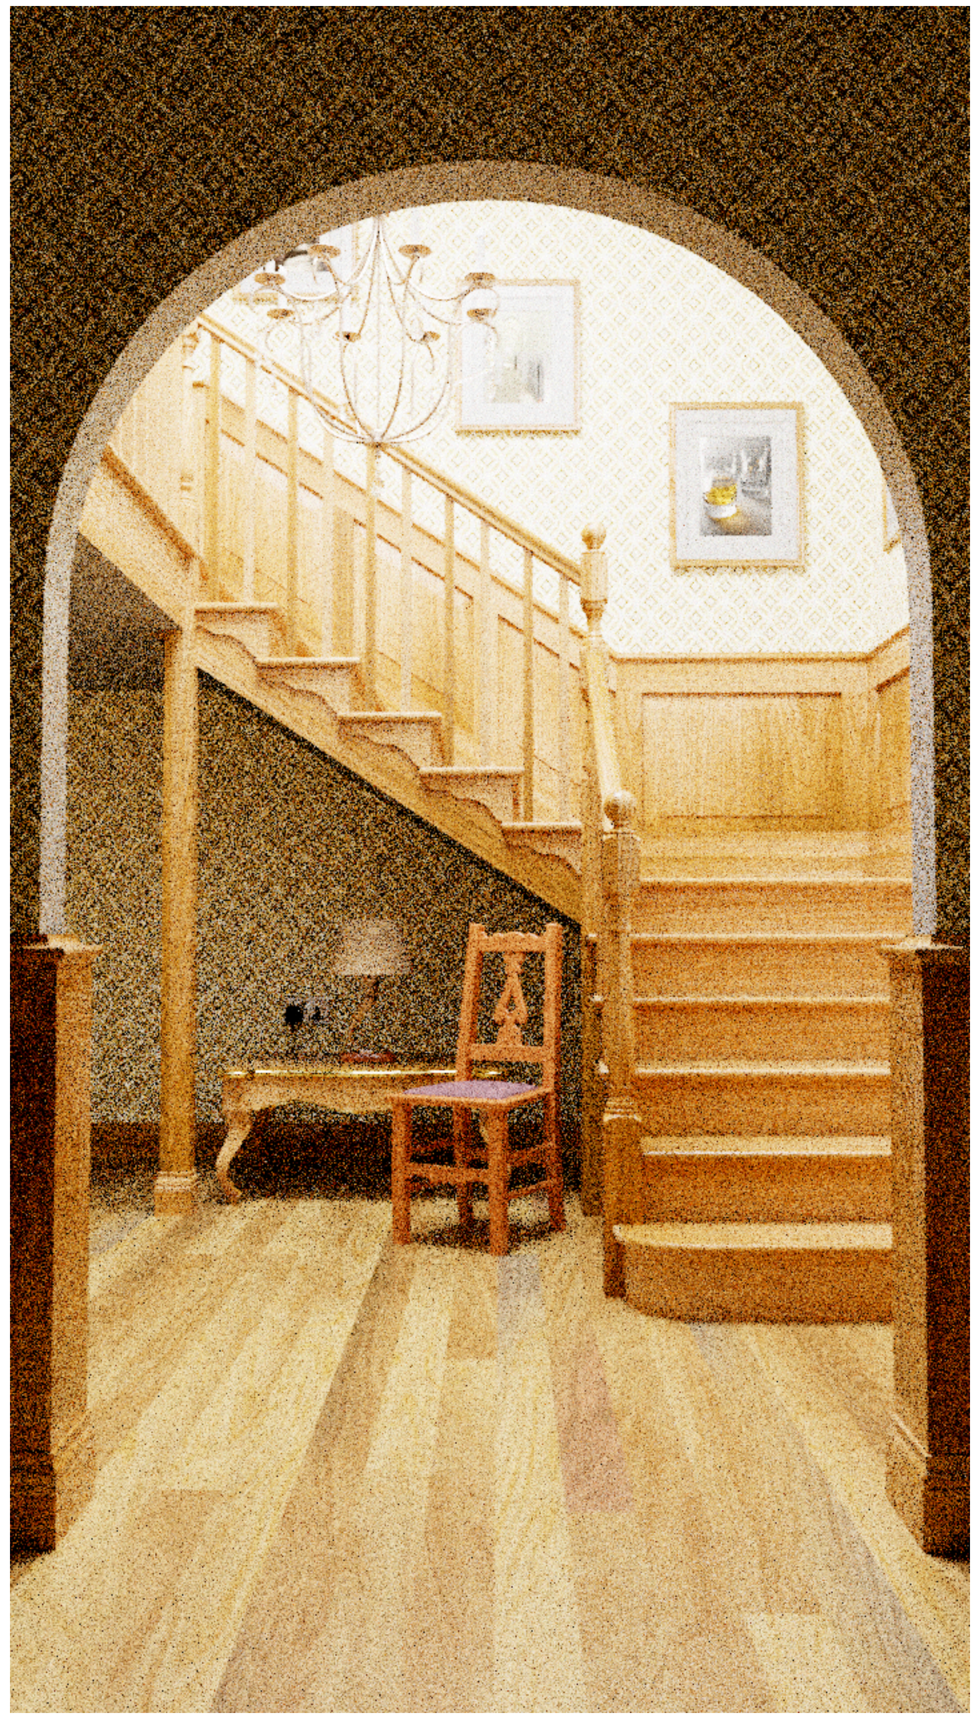

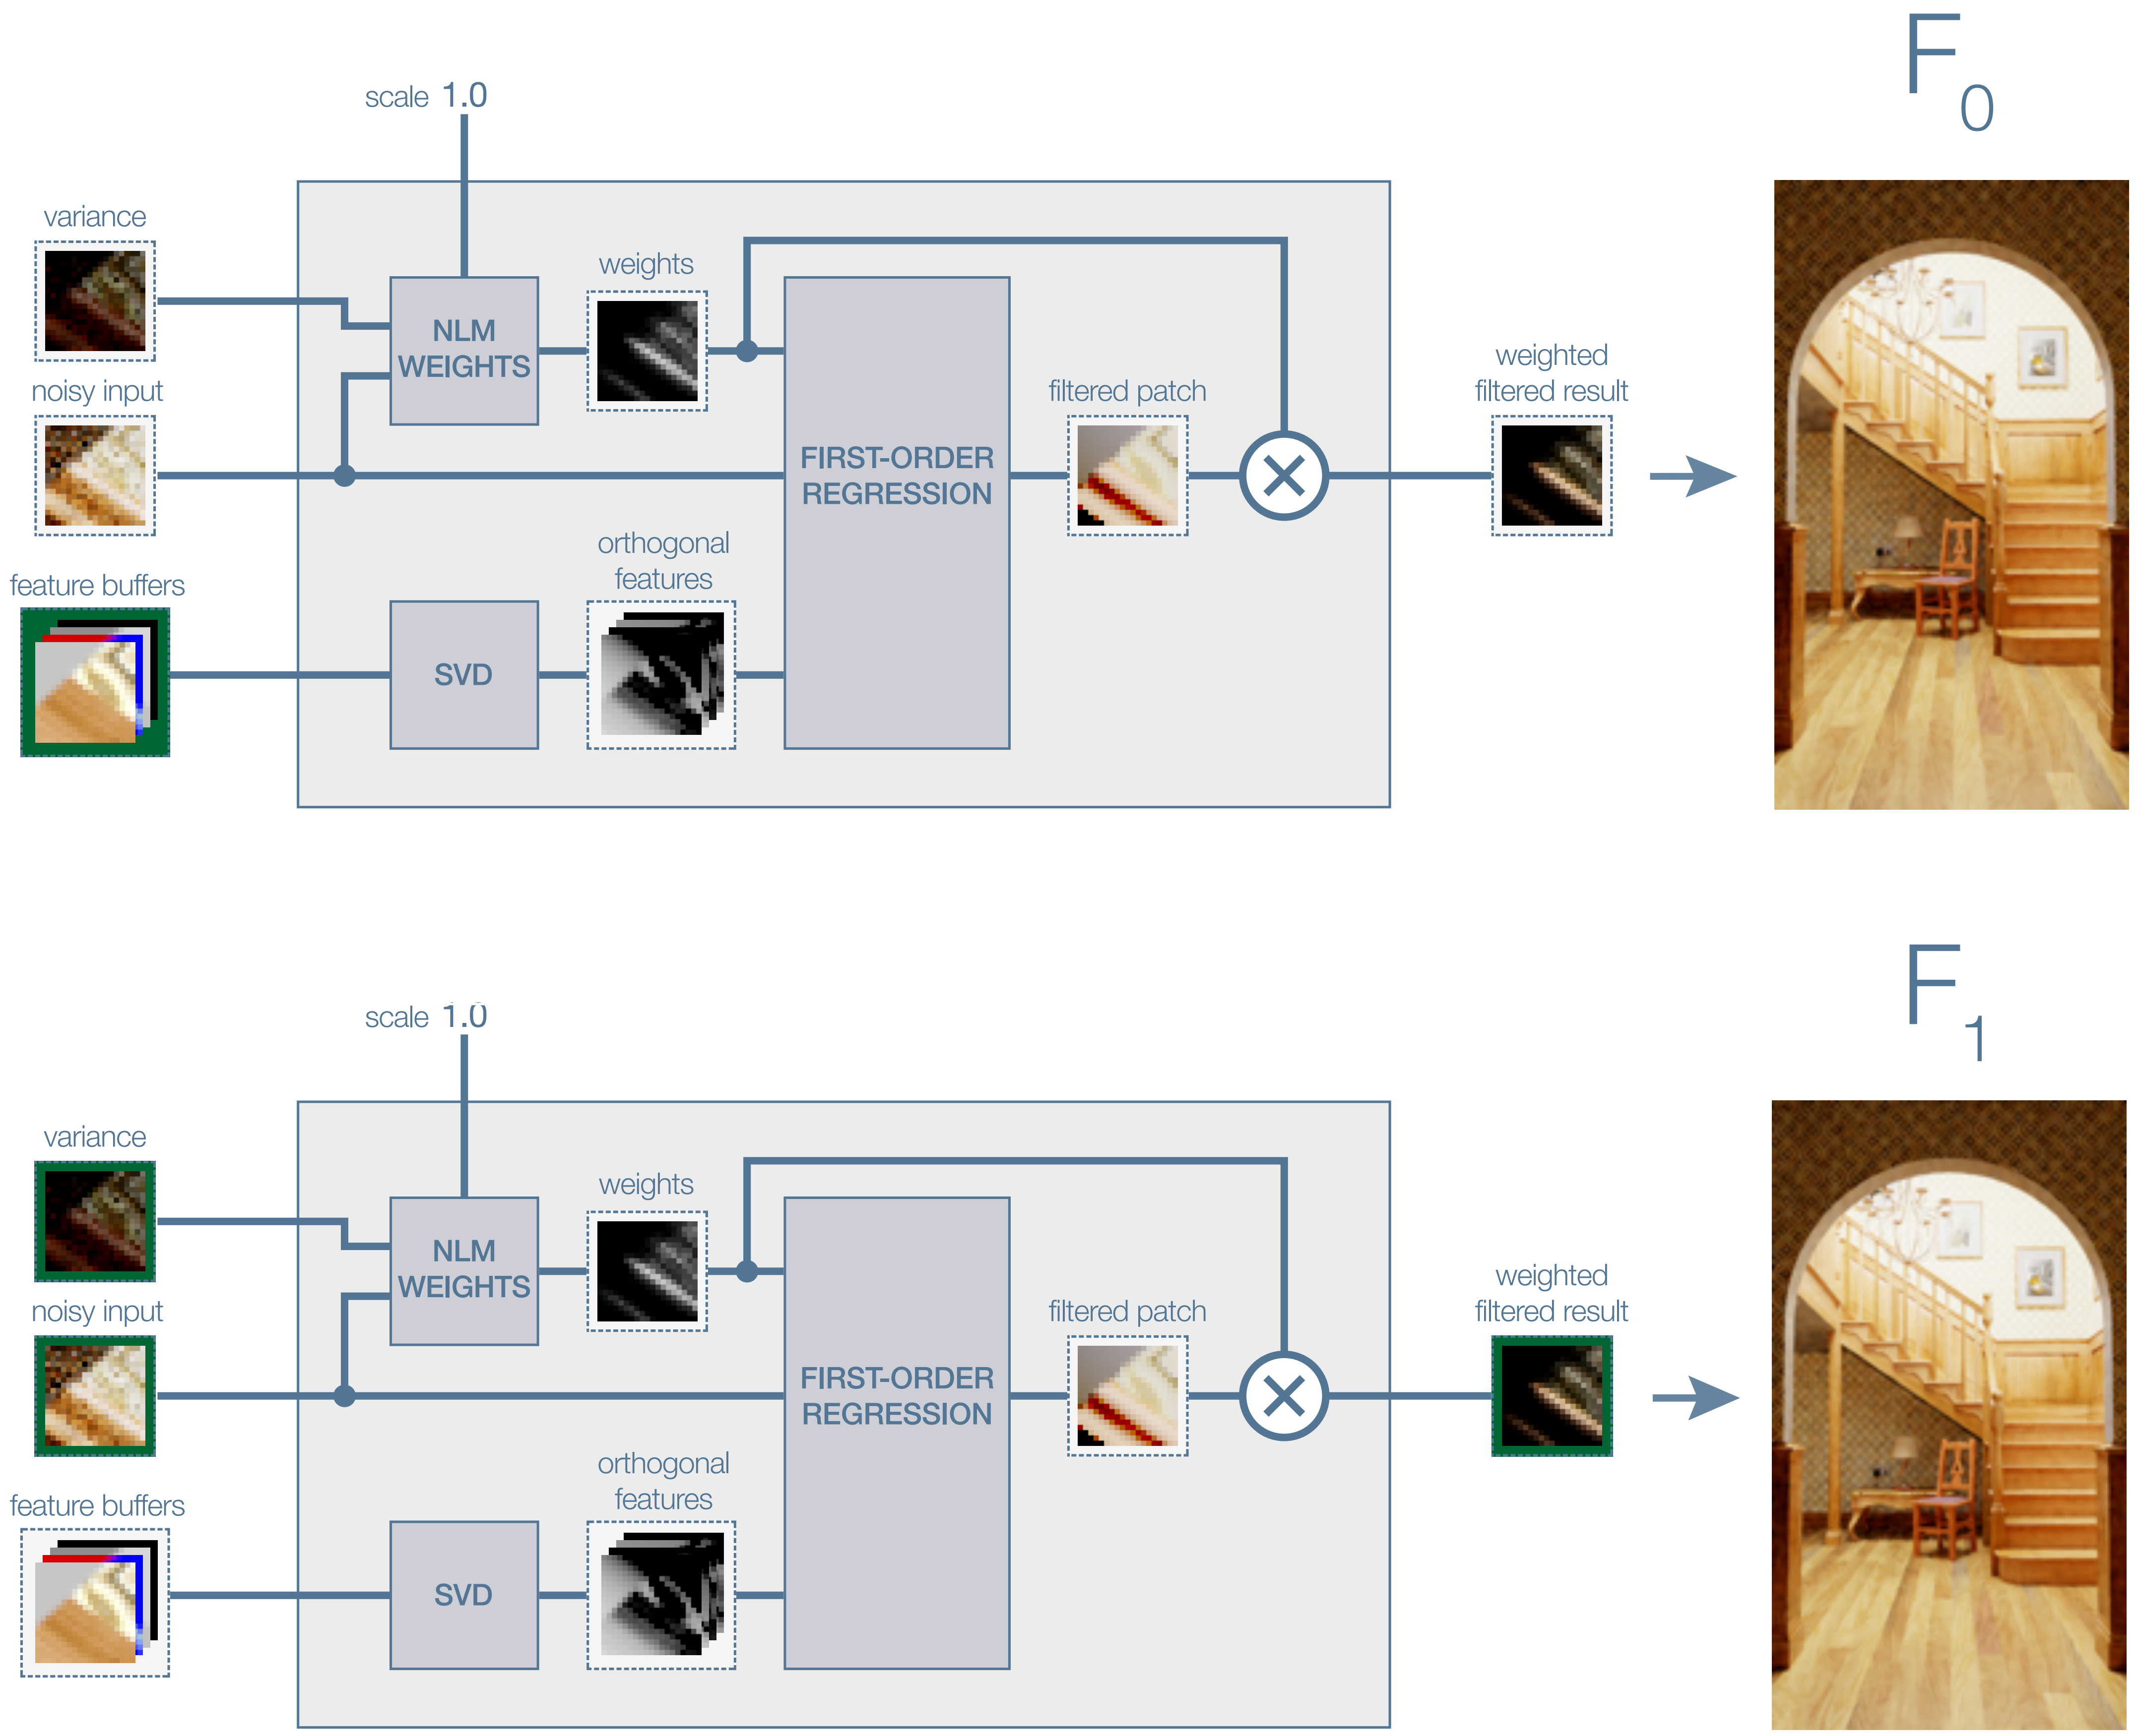

Algorithm Overview

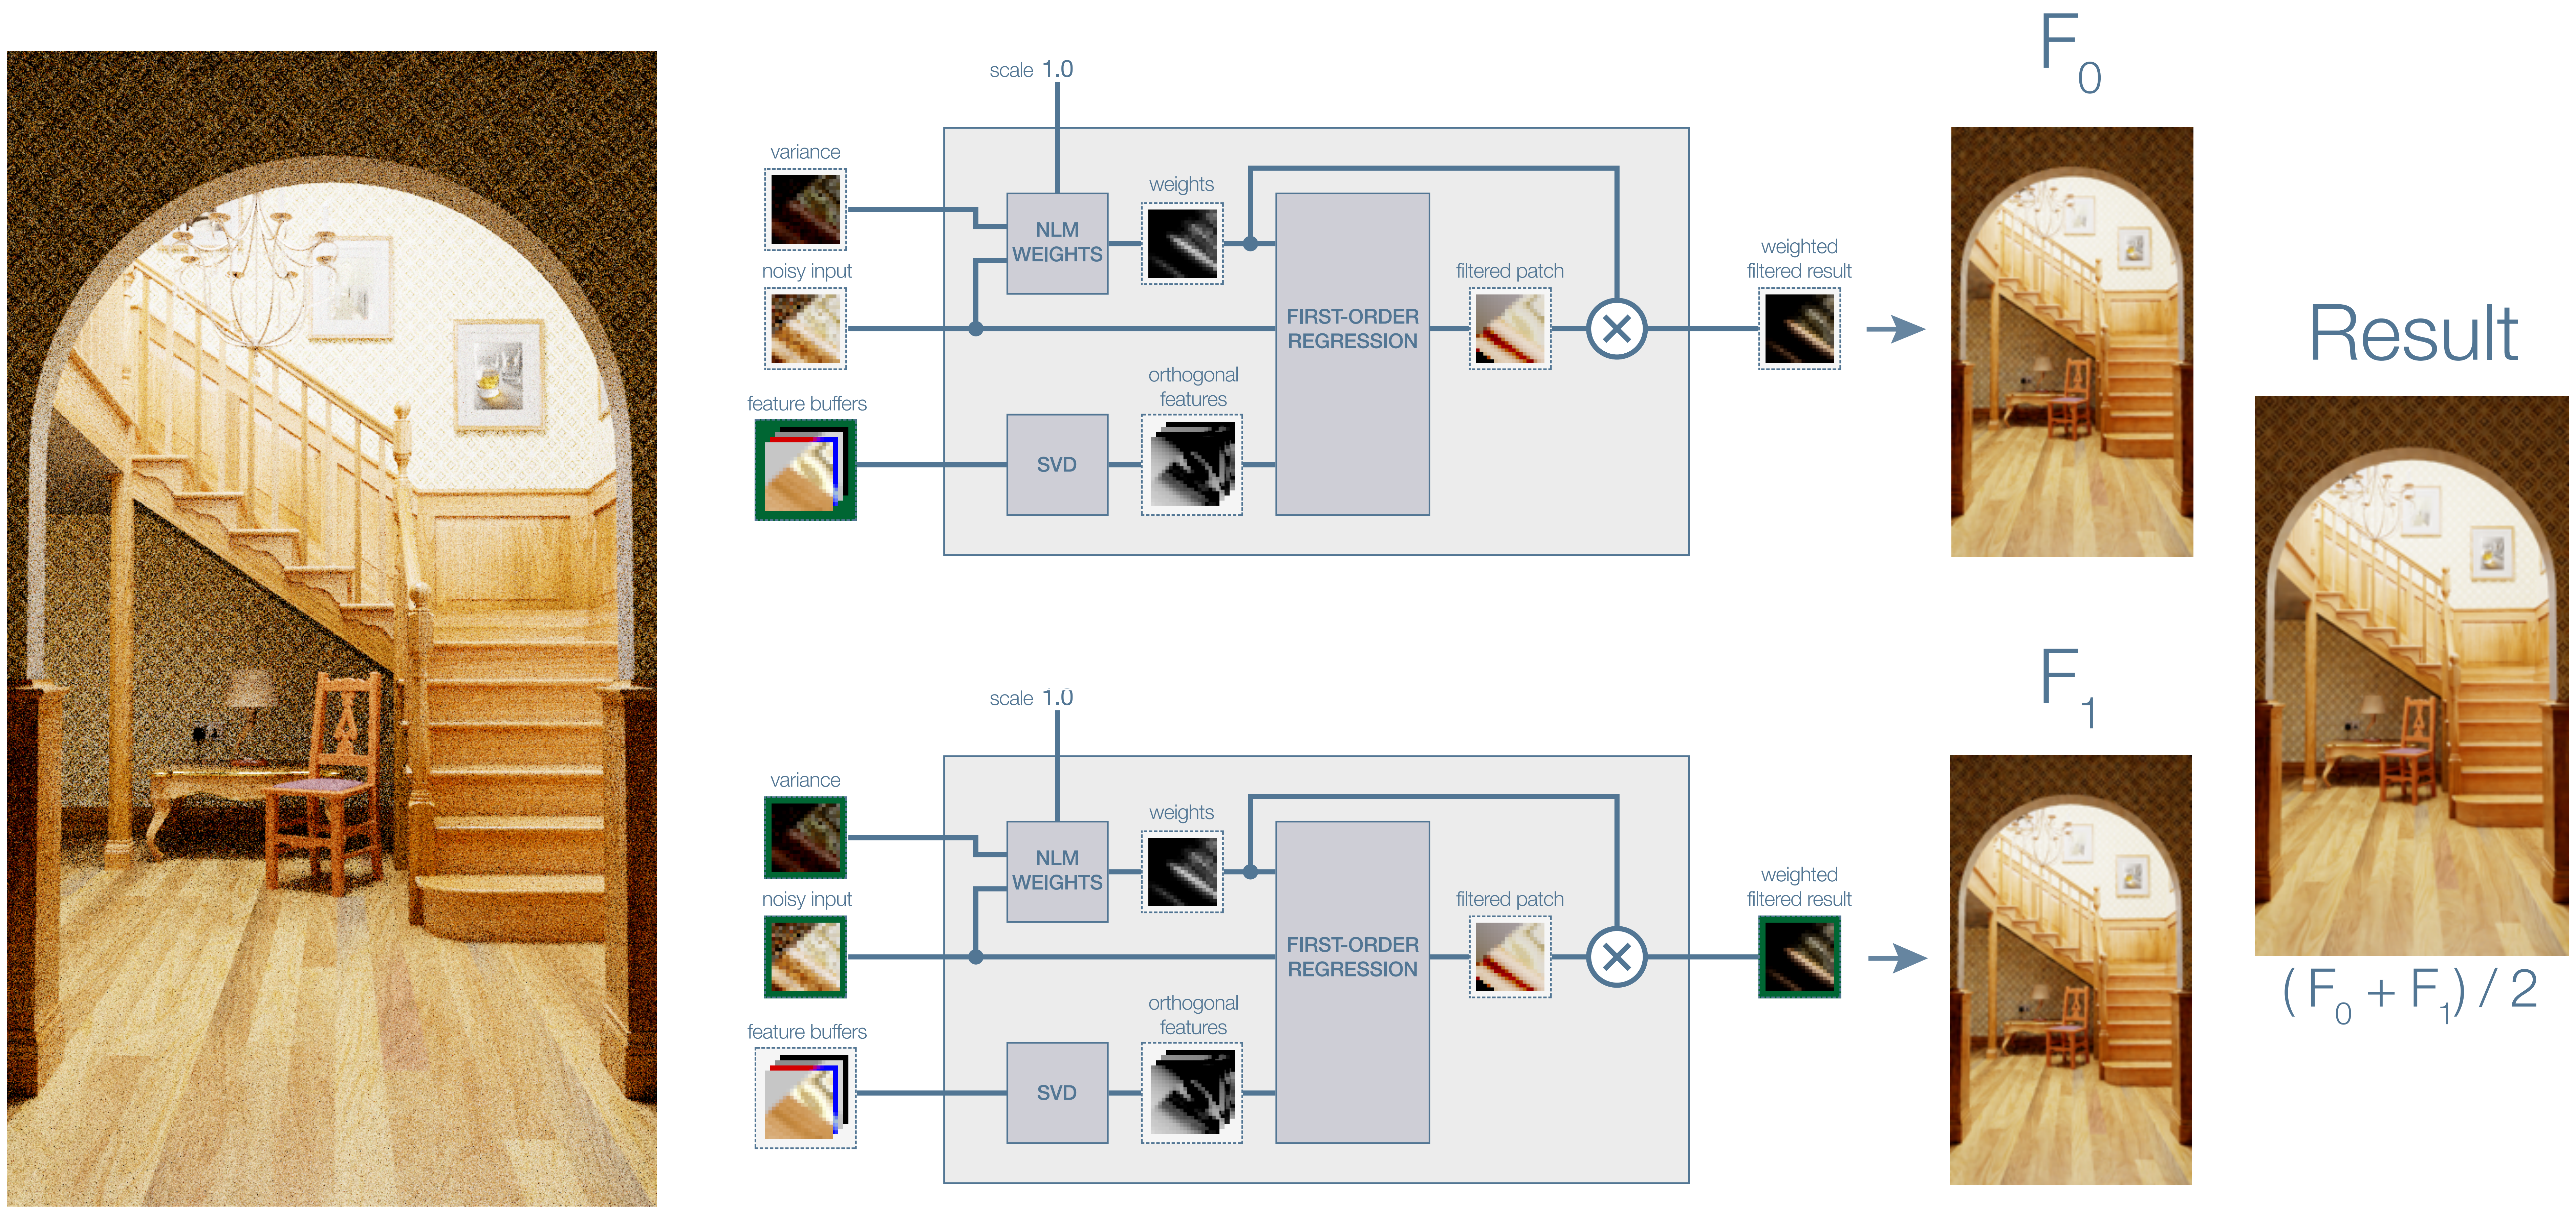

Algorithm Overview

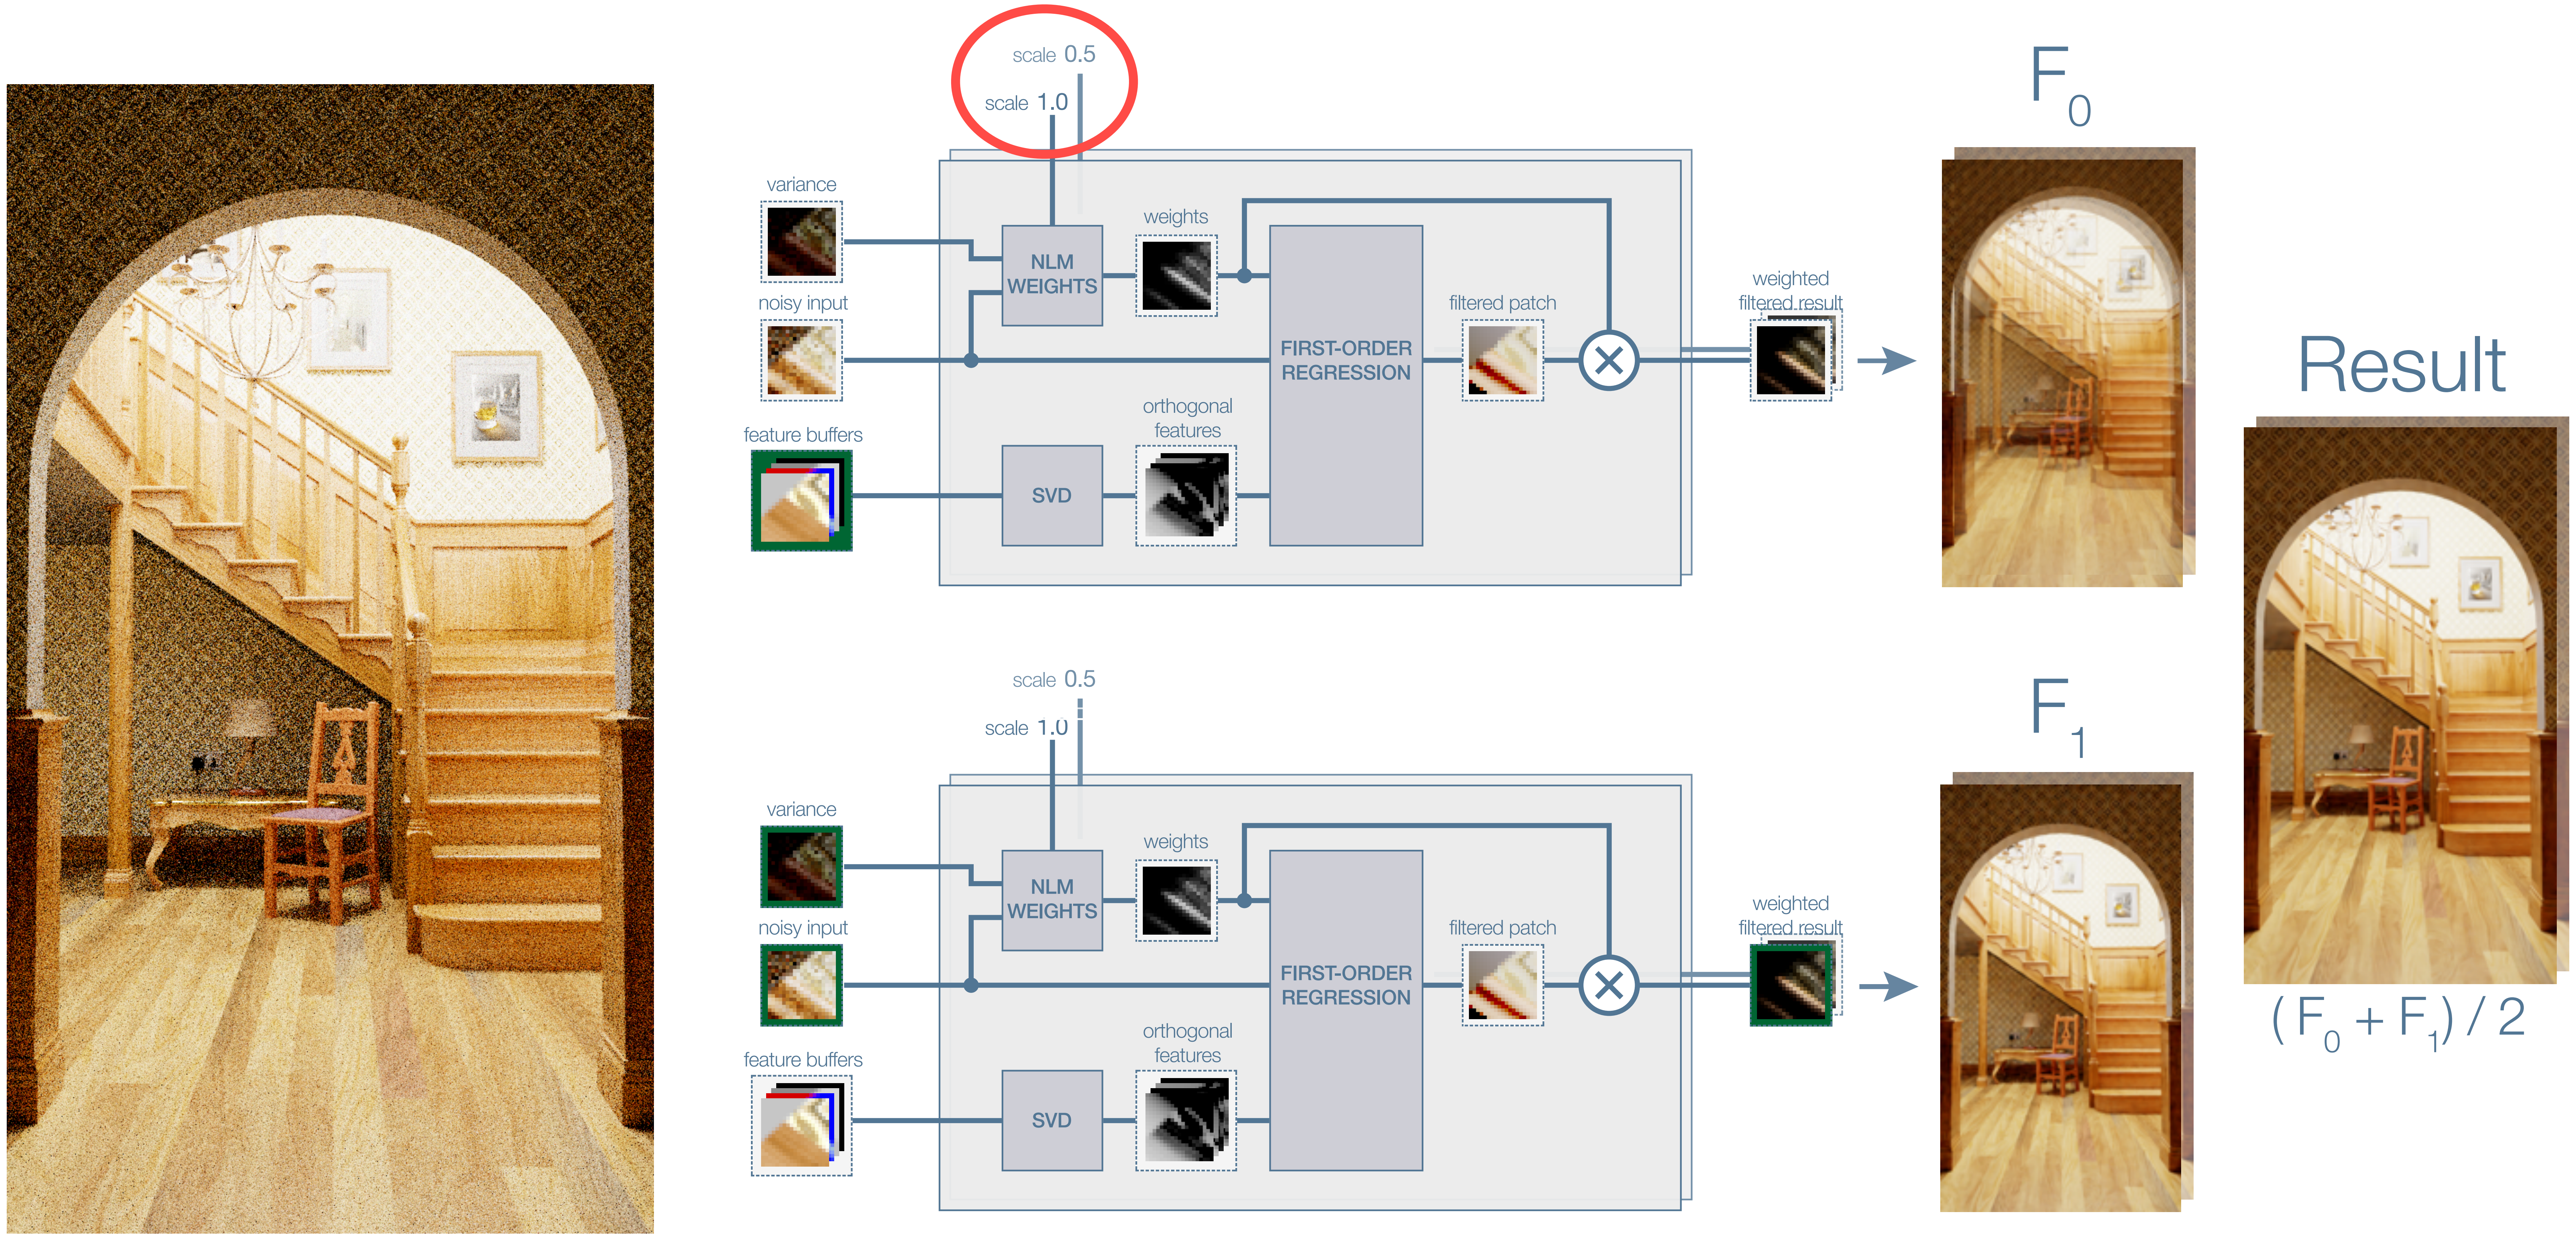

Algorithm Overview

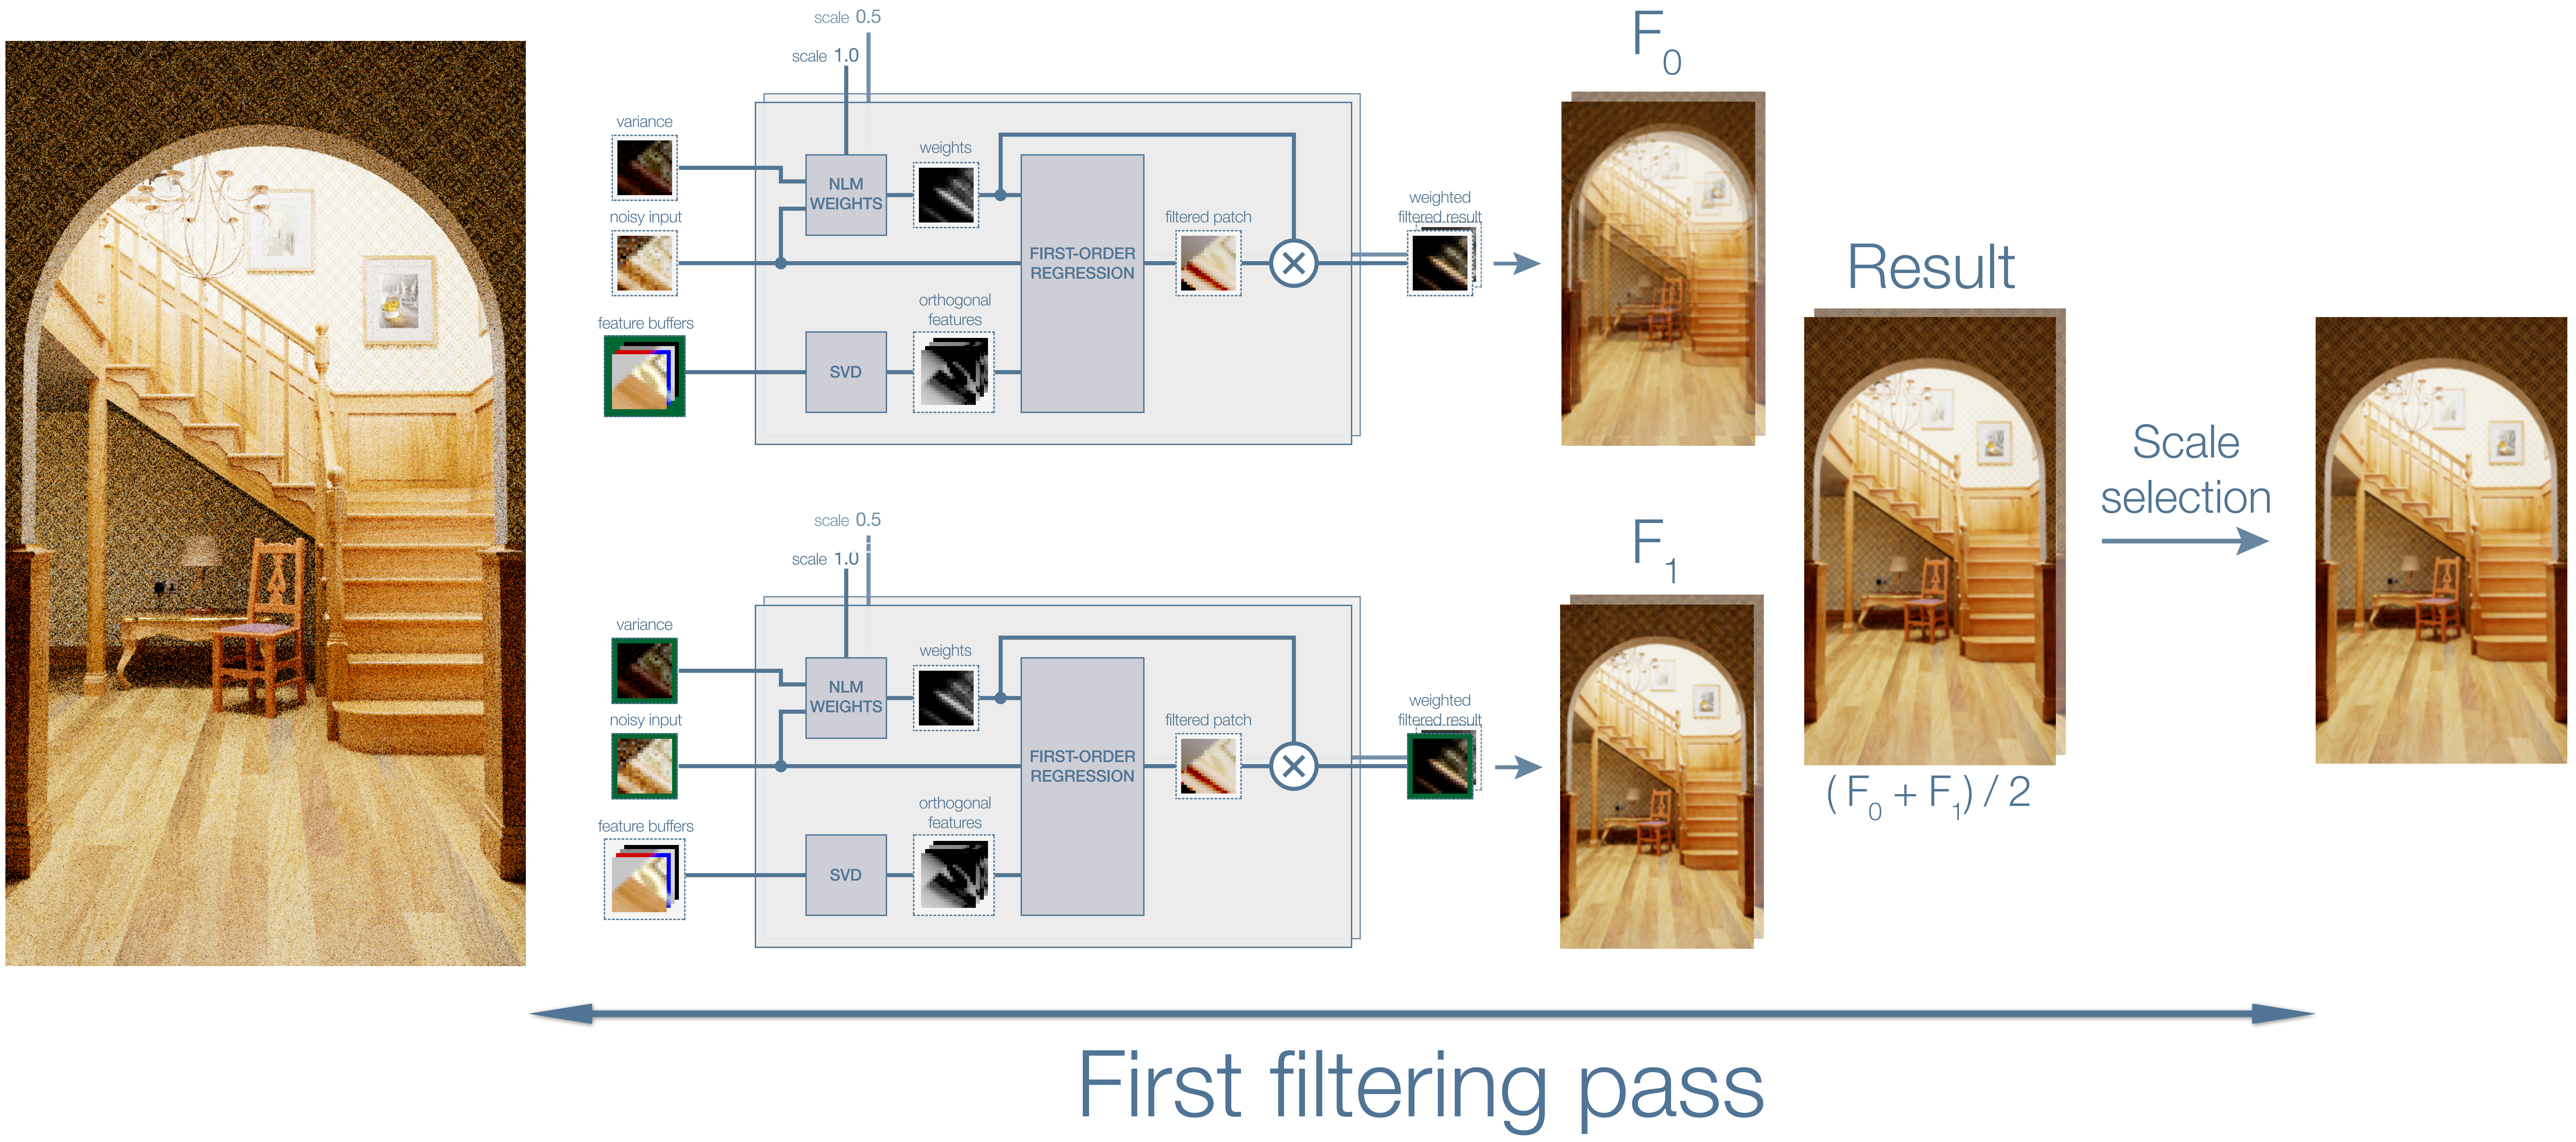

Algorithm Overview

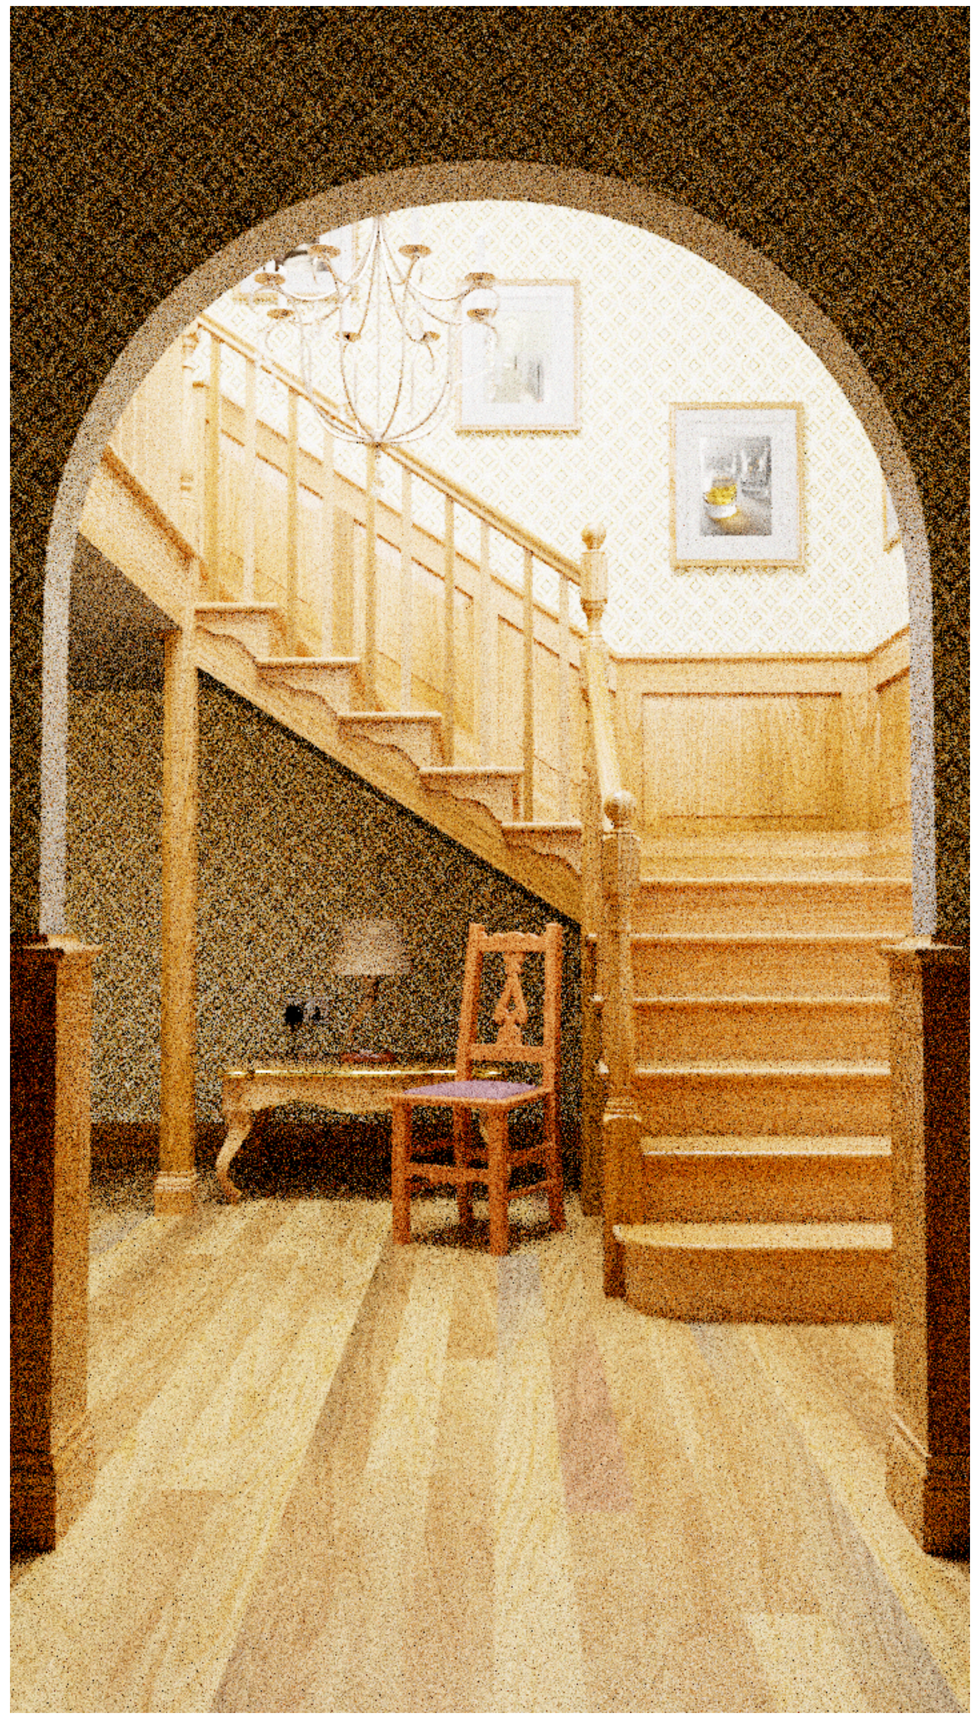

First filtering pass

Algorithm Overview

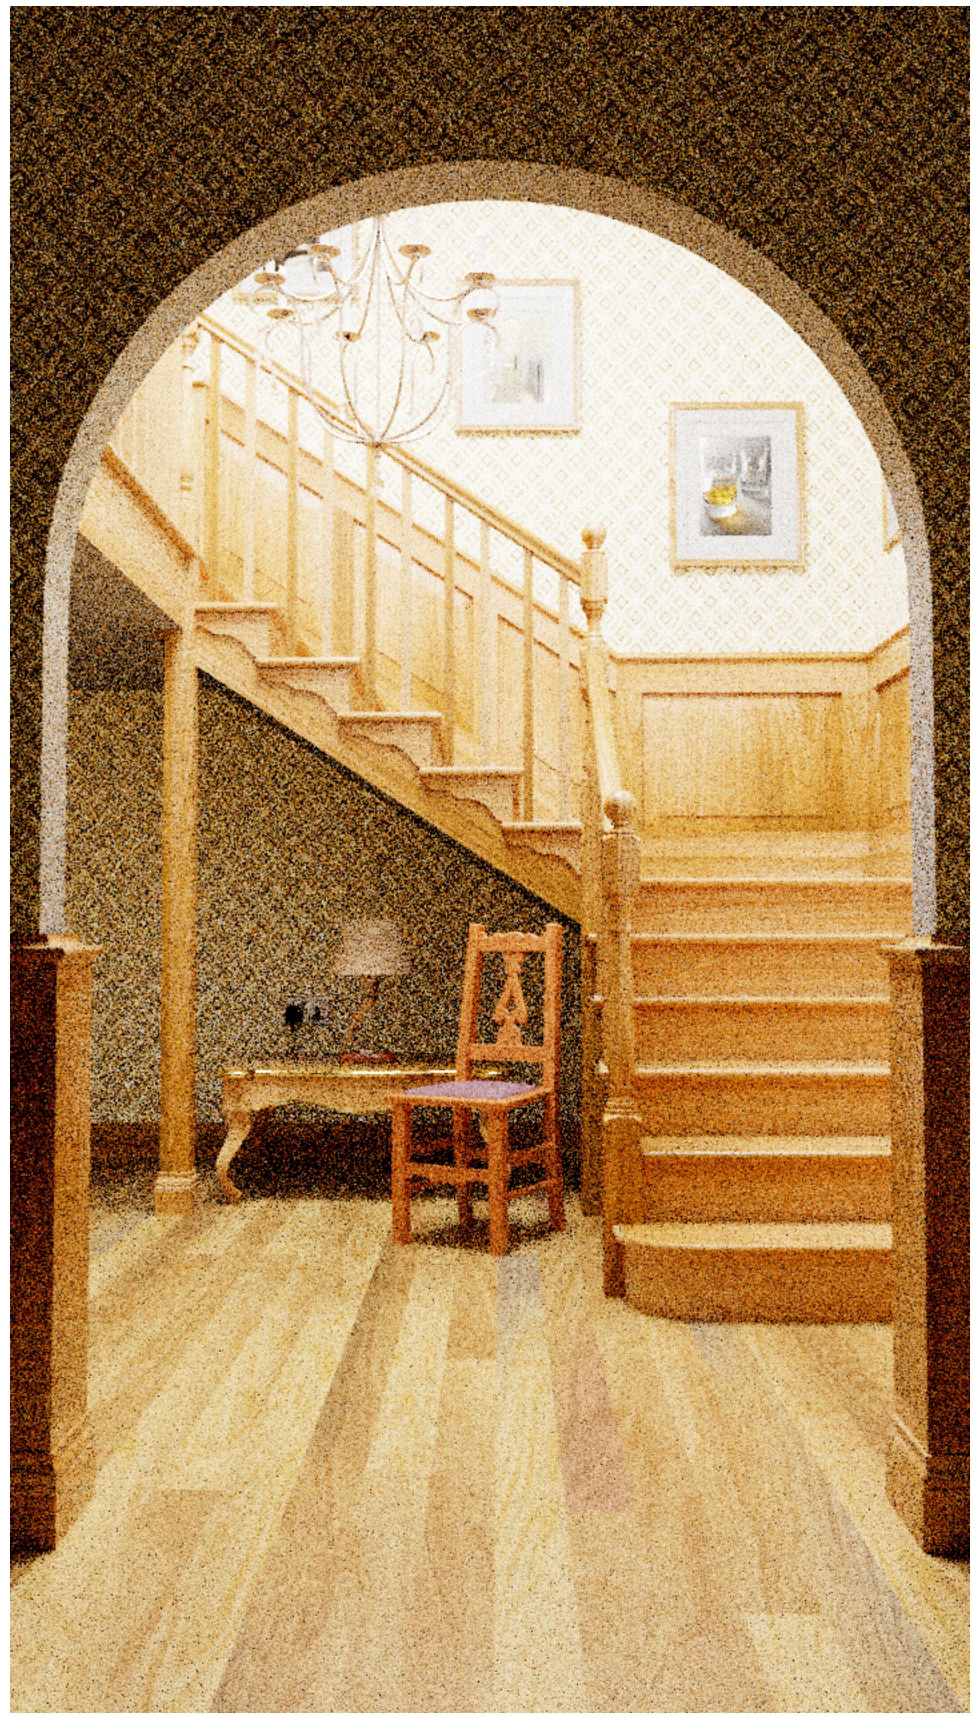

First filtering pass

Second filtering pass

≠ first pass

Results

http://drz.disneyresearch.com/~jnovak/publications/NFOR/supplementary/index.html

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: NLM | 3: NLM-MS | 4: RHF | 5: RDFC | 6: LBF | 7: WLR | 8: WLR-PF | 9: NFOR (ours) | 10: Reference

16 spp 64 spp 256 spp 1024 spp

Color buffers Structural Dissimilarity rMSE Heatmap

1: Input 2: NLM 3: NLM-MS 4: RHF 5: RDFC 6: LBF 7: WLR 8: WLR-PF 9: NFOR (ours) 10: Reference

16 spp 64 spp 256 spp 1024 spp

Color buffers Structural Dissimilarity rMSE Heatmap

1: Input 2: NLM 3: NLM-MS 4: RHF 5: RDFC 6: LBF 7: WLR 8: WLR-PF 9: NFOR (ours) 10: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: NLM | 3: NLM-MS | 4: RHF | 5: RDFC | 6: LBF | 7: WLR | 8: WLR-PF | 9: NFOR (ours) | 10: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: NLM | 3: NLM-MS | 4: RHF | 5: RDFC | 6: LBF | 7: WLR | 8: WLR-PF | 9: NFOR (ours) | 10: Reference

16 spp 64 spp 256 spp 1024 spp

Color buffers Structural Dissimilarity rMSE Heatmap

1: Input 2: NLM 3: NLM-MS 4: RHF 5: RDFC 6: LBF 7: WLR 8: WLR-PF 9: NFOR (ours)

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: NLM | 3: NLM-MS | 4: RHF | 5: RDFC | 6: LBF | 7: WLR | 8: WLR-PF | 9: NFOR (ours)

16 spp 64 spp 256 spp 1024 spp

Color buffers Structural Dissimilarity rMSE Heatmap

1: Input 2: NLM 3: NLM-MS 4: RHF 5: RDFC 6: LBF 7: WLR 8: WLR-PF 9: NFOR (ours) 10: Reference

16 spp 64 spp 256 spp 1024 spp

Color buffers Structural Dissimilarity rMSE Heatmap

1: Input 2: NLM 3: NLM-MS 4: RHF 5: RDFC 6: LBF 7: WLR 8: WLR-PF 9: NFOR (ours) 10: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: NLM | 3: NLM-MS | 4: RHF | 5: RDFC | 6: LBF | 7: WLR | 8: WLR-PF | 9: NFOR (ours) | 10: Reference

16 spp 64 spp 256 spp 1024 spp

Color buffers Structural Dissimilarity rMSE Heatmap

1: Input 2: NLM 3: NLM-MS 4: RHF 5: RDFC 6: LBF 7: WLR 8: WLR-PF 9: NFOR (ours) 10: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: NLM | 3: NLM-MS | 4: RHF | 5: RDFC | 6: LBF | 7: WLR | 8: WLR-PF | 9: NFOR (ours) | 10: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: NLM | 3: NLM-MS | 4: RHF | 5: RDFC | 6: LBF | 7: WLR | 8: WLR-PF | 9: NFOR (ours) | 10: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: NLM | 3: NLM-MS | 4: RHF | 5: RDFC | 6: LBF | 7: WLR | 8: WLR-PF | 9: NFOR (ours) | 10: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: NLM | 3: NLM-MS | 4: RHF | 5: RDFC | 6: LBF | 7: WLR | 8: WLR-PF | 9: NFOR (ours) | 10: Reference

16 spp 64 spp 256 spp 1024 spp

Color buffers Structural Dissimilarity rMSE Heatmap

1: Input 2: NLM 3: NLM-MS 4: RHF 5: RDFC 6: LBF 7: WLR 8: WLR-PF 9: Ours 10: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: NLM | 3: NLM-MS | 4: RHF | 5: RDFC | 6: LBF | 7: WLR | 8: WLR-PF | 9: Ours | 10: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: NLM | 3: NLM-MS | 4: RHF | 5: RDFC | 6: LBF | 7: WLR | 8: WLR-PF | 9: Ours | 10: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: NLM | 3: NLM-MS | 4: RHF | 5: RDFC | 6: LBF | 7: WLR | 8: WLR-PF | 9: Ours | 10: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: NLM | 3: NLM-MS | 4: RHF | 5: RDFC | 6: LBF | 7: WLR | 8: WLR-PF | 9: Ours | 10: Reference

16 spp 64 spp 256 spp 1024 spp

Color buffers Structural Dissimilarity rMSE Heatmap

1: Input 2: RenderMan Denoiser 3: RenderMan Denoiser 21.0b1 4: NFOR (ours) 5: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: RenderMan Denoiser | 3: RenderMan Denoiser 21.0b1 | 4: NFOR (ours) | 5: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: RenderMan Denoiser | 3: RenderMan Denoiser 21.0b1 | 4: NFOR (ours) | 5: Reference

16 spp 64 spp 256 spp 1024 spp

Color buffers Structural Dissimilarity rMSE Heatmap

1: Input 2: RenderMan Denoiser 3: RenderMan Denoiser 21.0b1 4: NFOR (ours) 5: Reference

16 spp | 64 spp | 256 spp | 1024 spp

Color buffers | Structural Dissimilarity | rMSE Heatmap

1: Input | 2: RenderMan Denoiser | 3: RenderMan Denoiser 21.0b1 | 4: NFOR (ours) | 5: Reference

Error Metrics

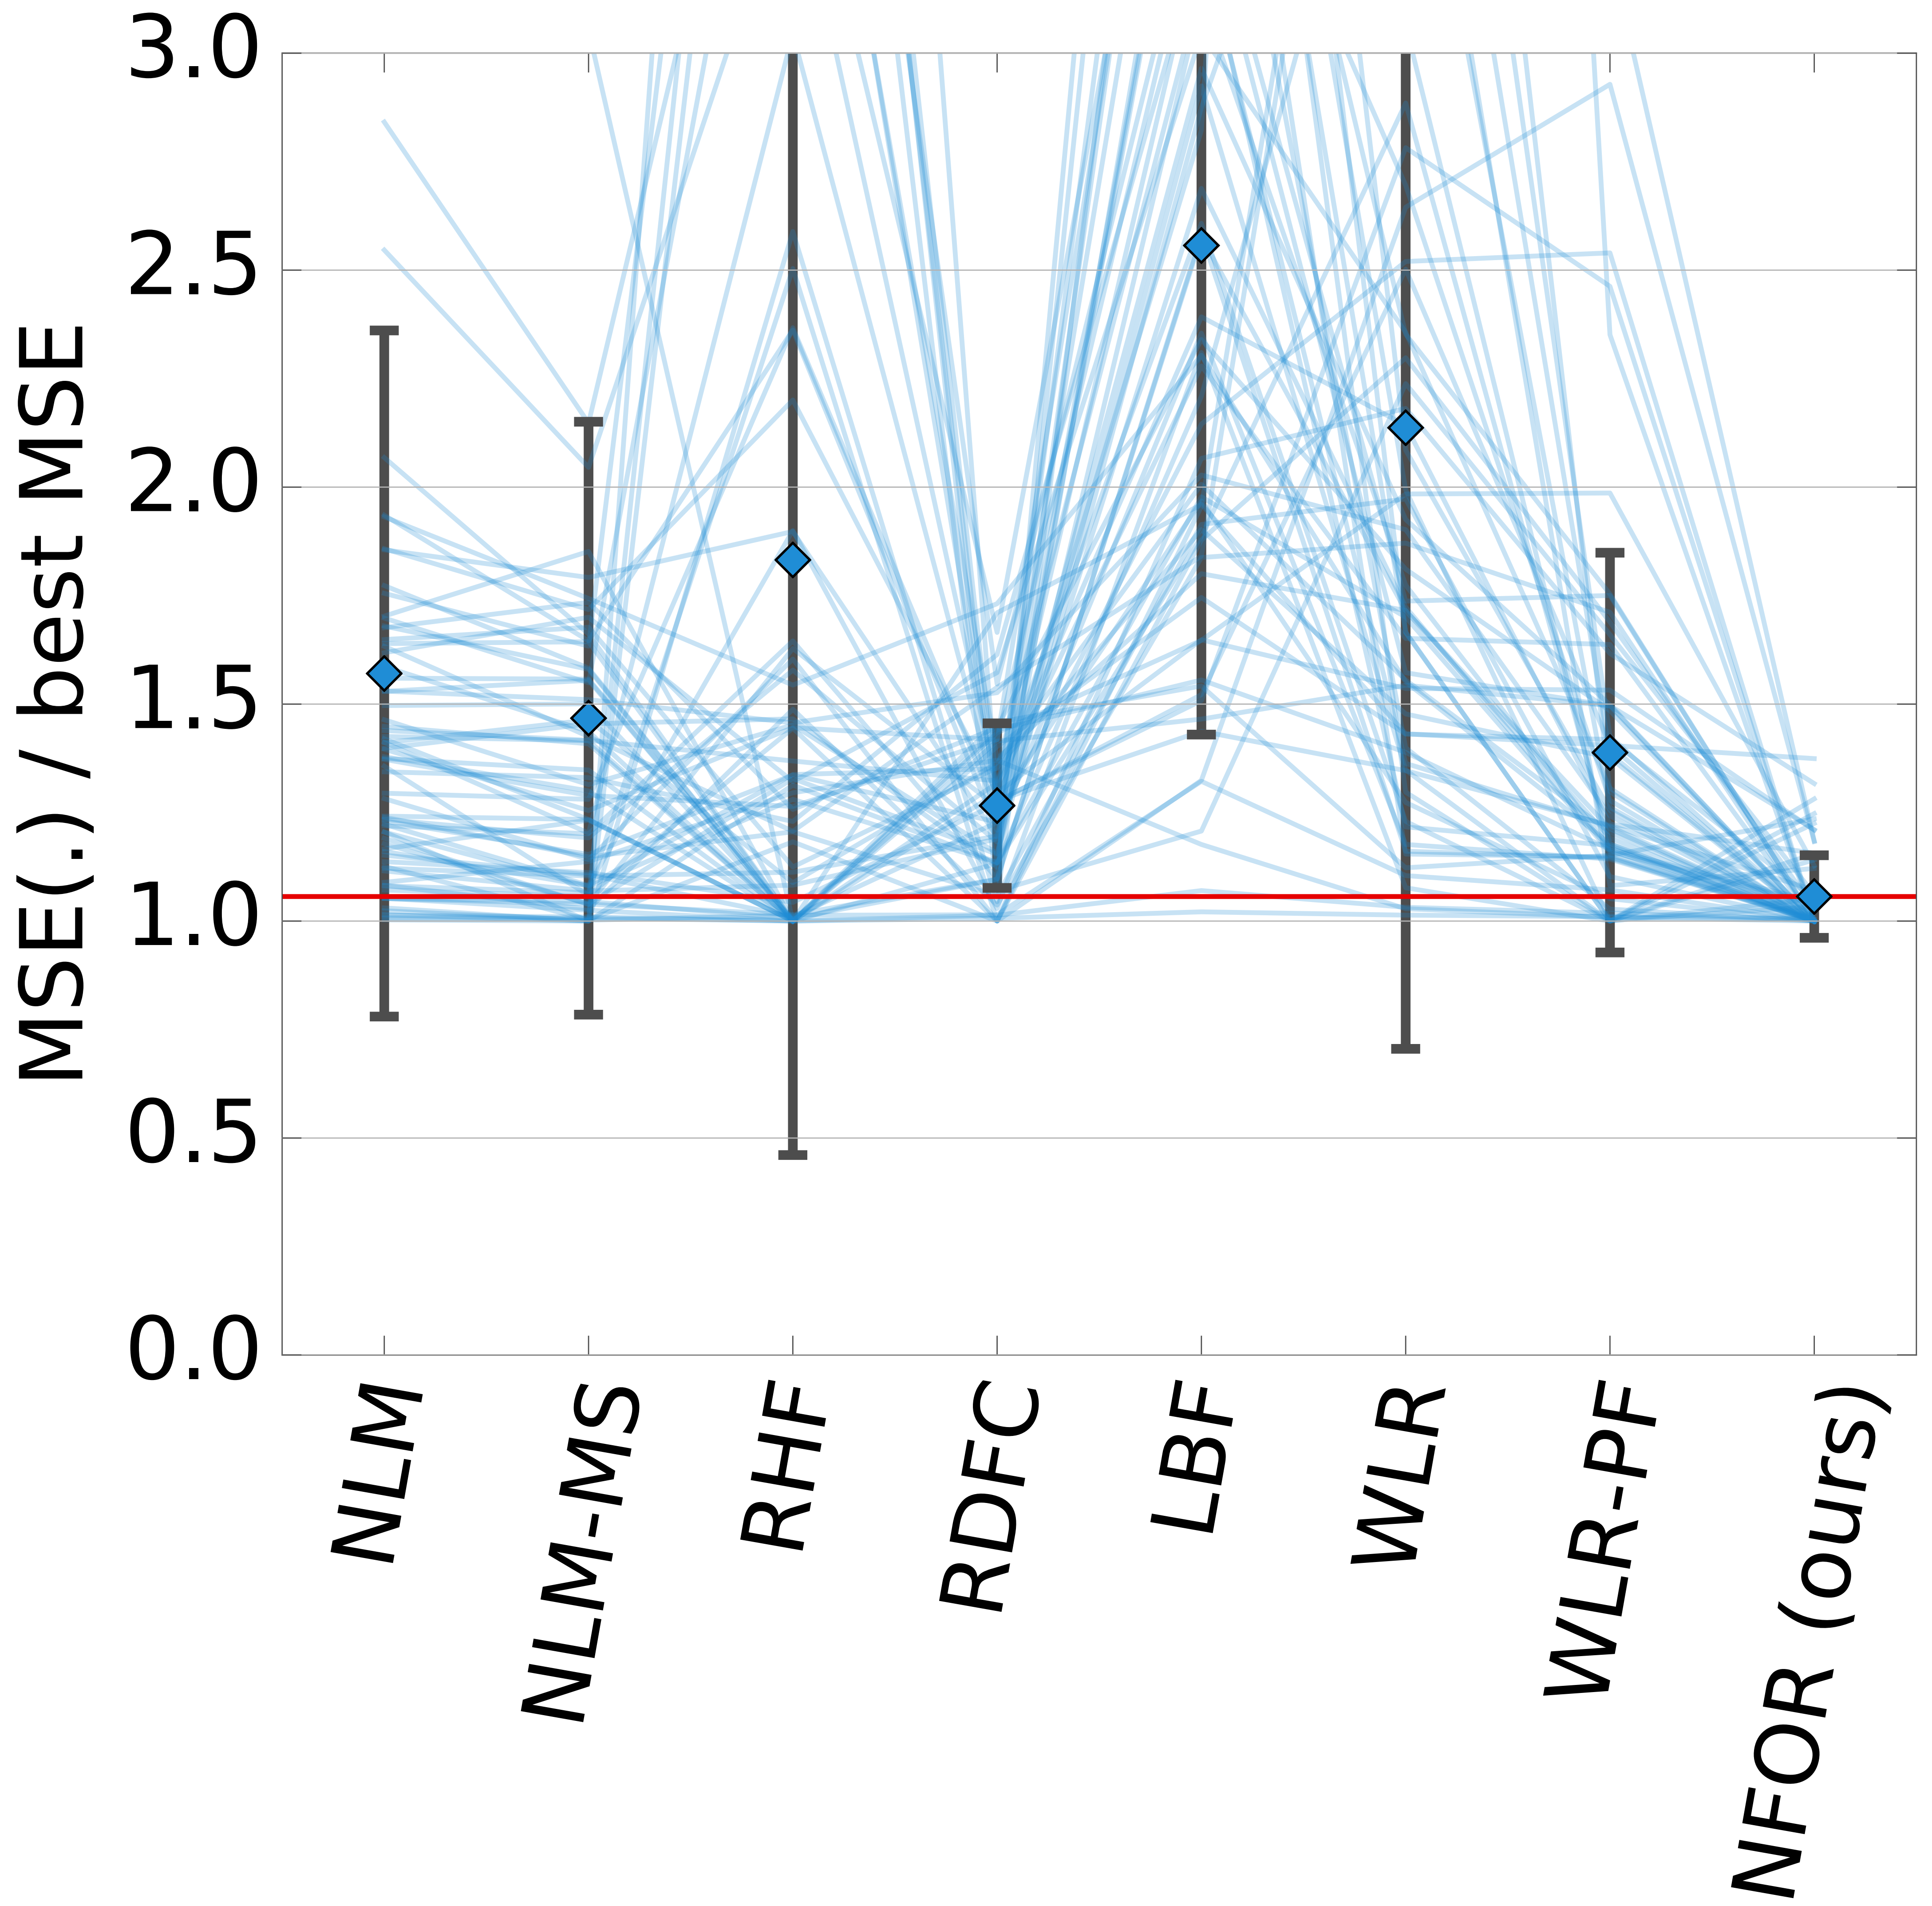

Error Metrics

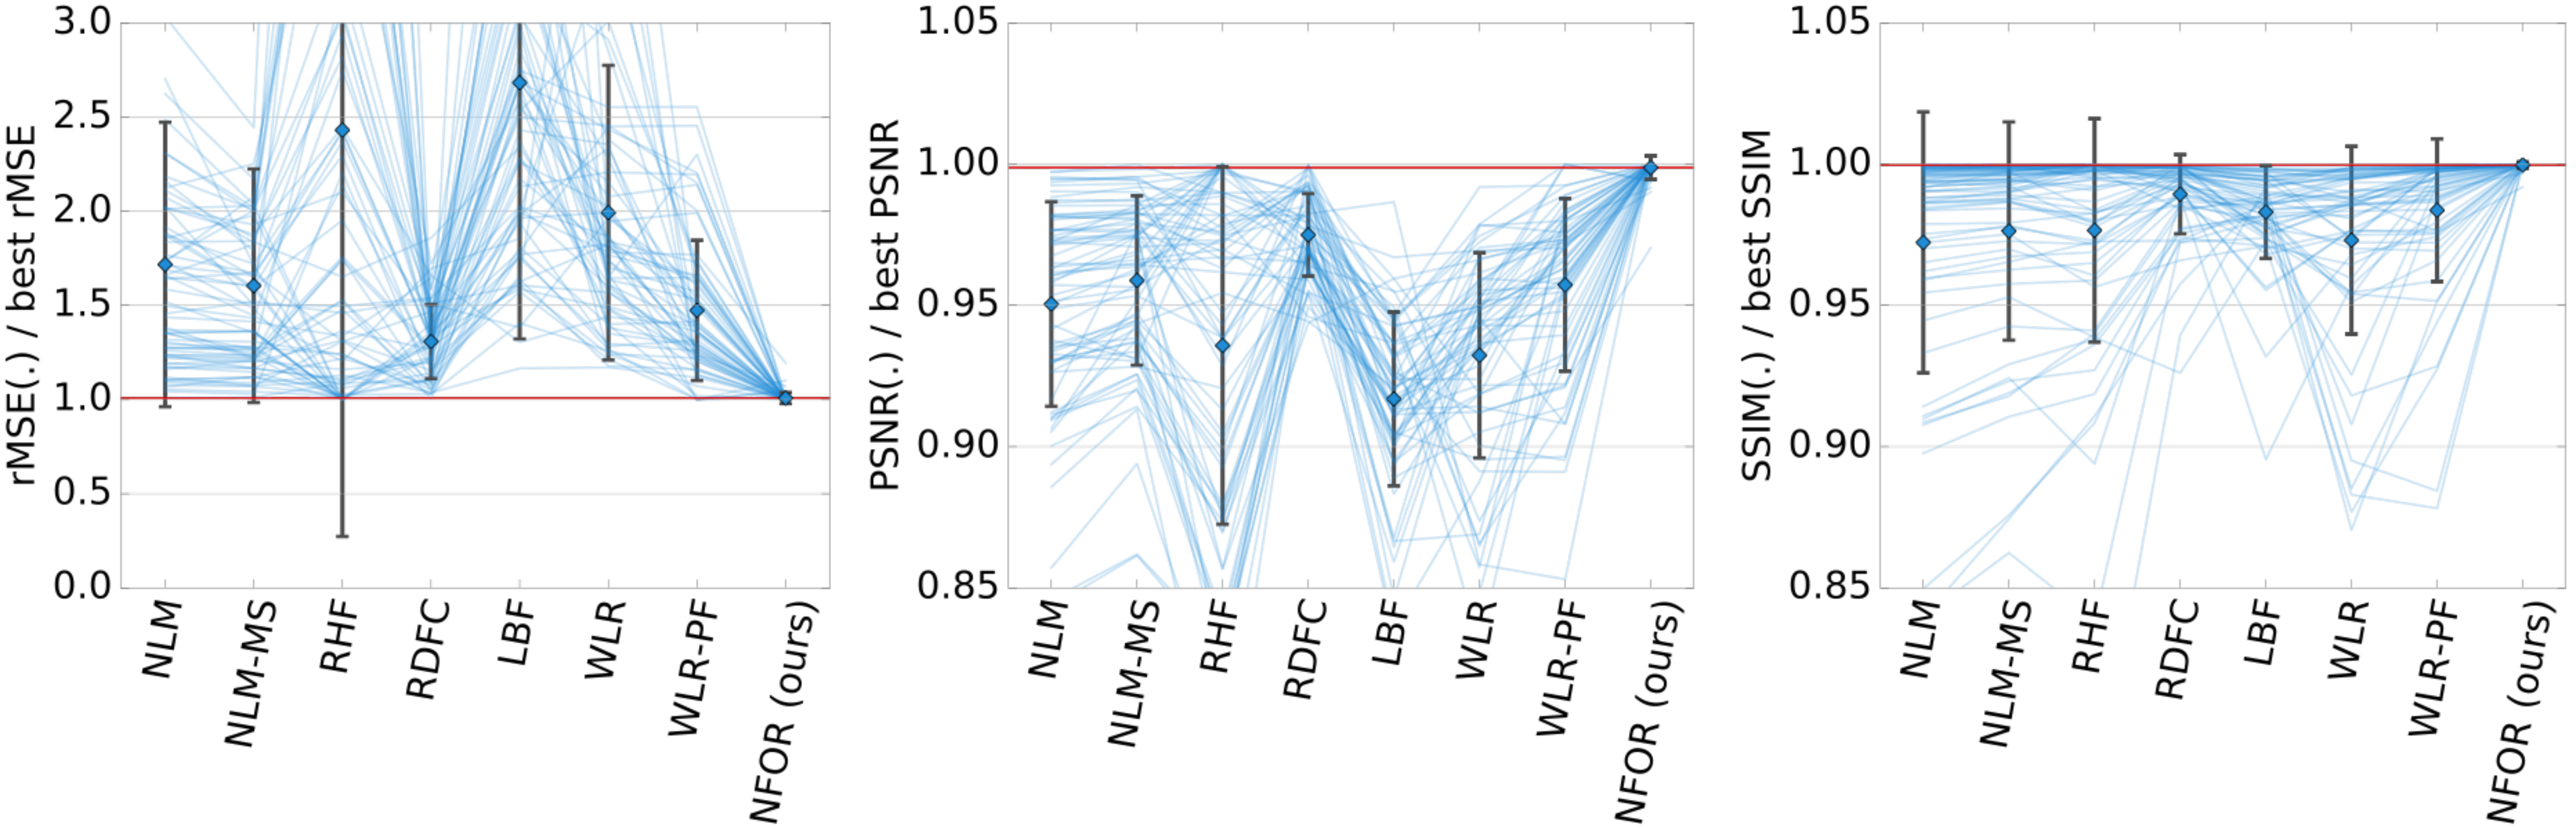

Limitations

No adaptive sampling (yet)

High overhead

Algorithm	Avg. Runtime	Avg. Memory Usage
RDFC	41s	349 MB
RHF	18s	460 MB
WLR	72s	309 MB
WLR-PF	91s	311 MB
NFOR (ours)	223s	3248 MB

Conclusion

Comparative analysis of recent denoising techniques

Novel filter with state-of-the-art results

Future work

- Comparative analysis of adaptive sampling techniques
- Sparse reconstruction to reduce overhead


Input
© Disney