



**SIGGRAPH  
ASIA 2023  
SYDNEY**

# Efficient Graphics Representation with Differentiable Indirection

Sayantana Datta<sup>1,2</sup>

Carl Marshall<sup>2</sup>

Derek Nowrouzezahrai<sup>1,3</sup>

Zhao Dong<sup>2</sup>

Zhengqin Li<sup>2</sup>

<sup>1</sup>McGill University

<sup>2</sup>Meta

<sup>3</sup>MILA

## Connecting STORIES

The 16th ACM SIGGRAPH Conference and Exhibition on  
Computer Graphics and Interactive Techniques in Asia

CONFERENCE 12 - 15 December 2023

EXHIBITION 13 - 15 December 2023

ICC, Sydney, Australia

[ASIA.SIGGRAPH.ORG/2023](https://asia.siggraph.org/2023)

Sponsored by



Organized by




# What is Differentiable Indirection?

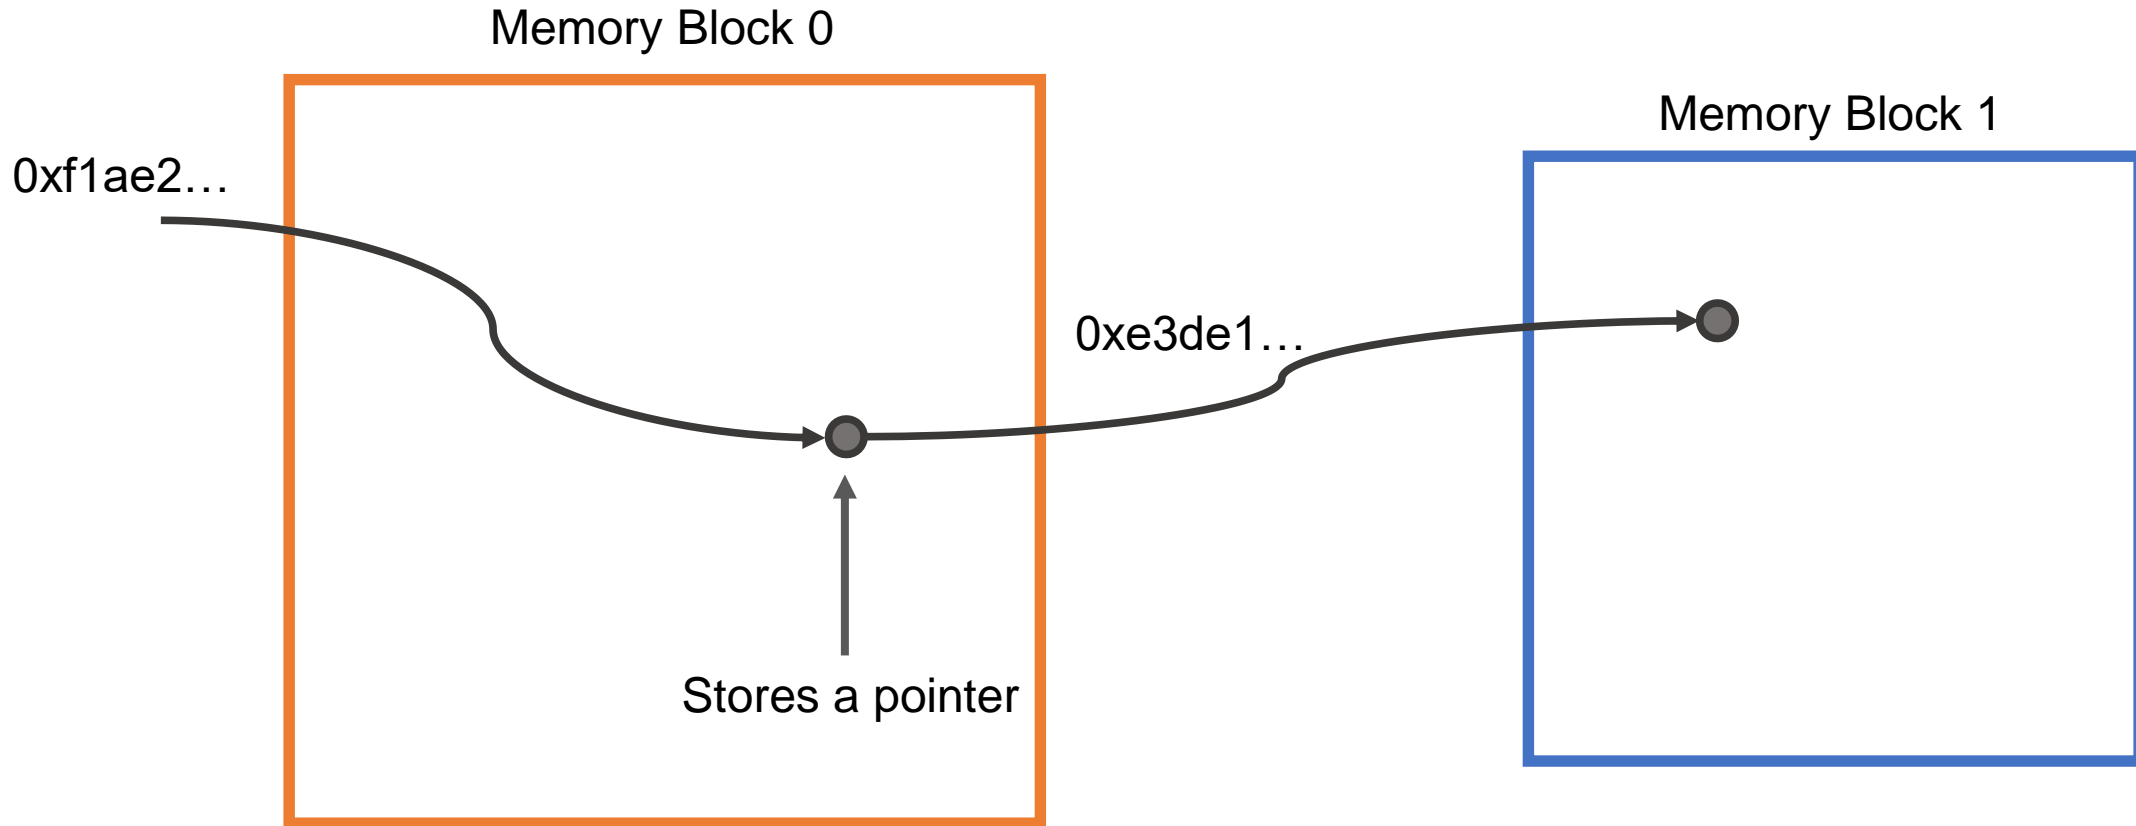
# What is Differentiable Indirection?

Memory Block 0

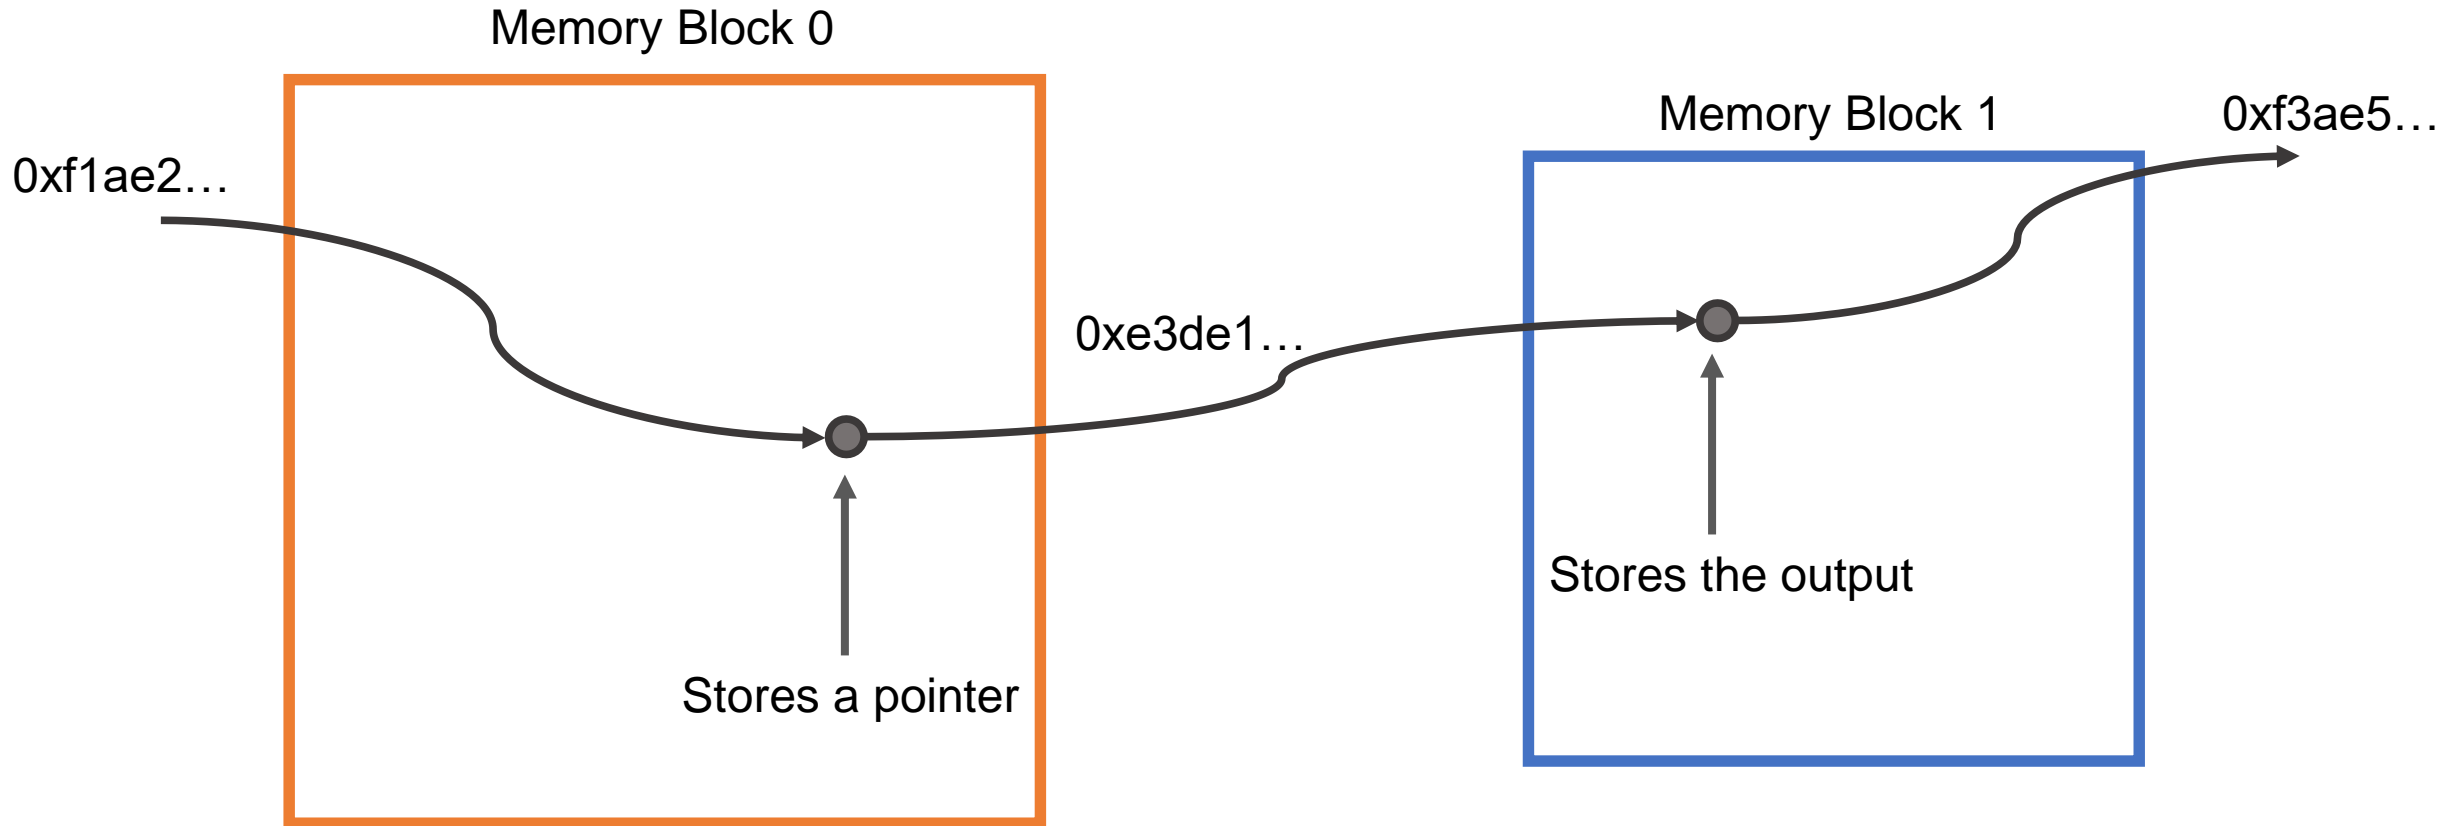
0xf1ae2...



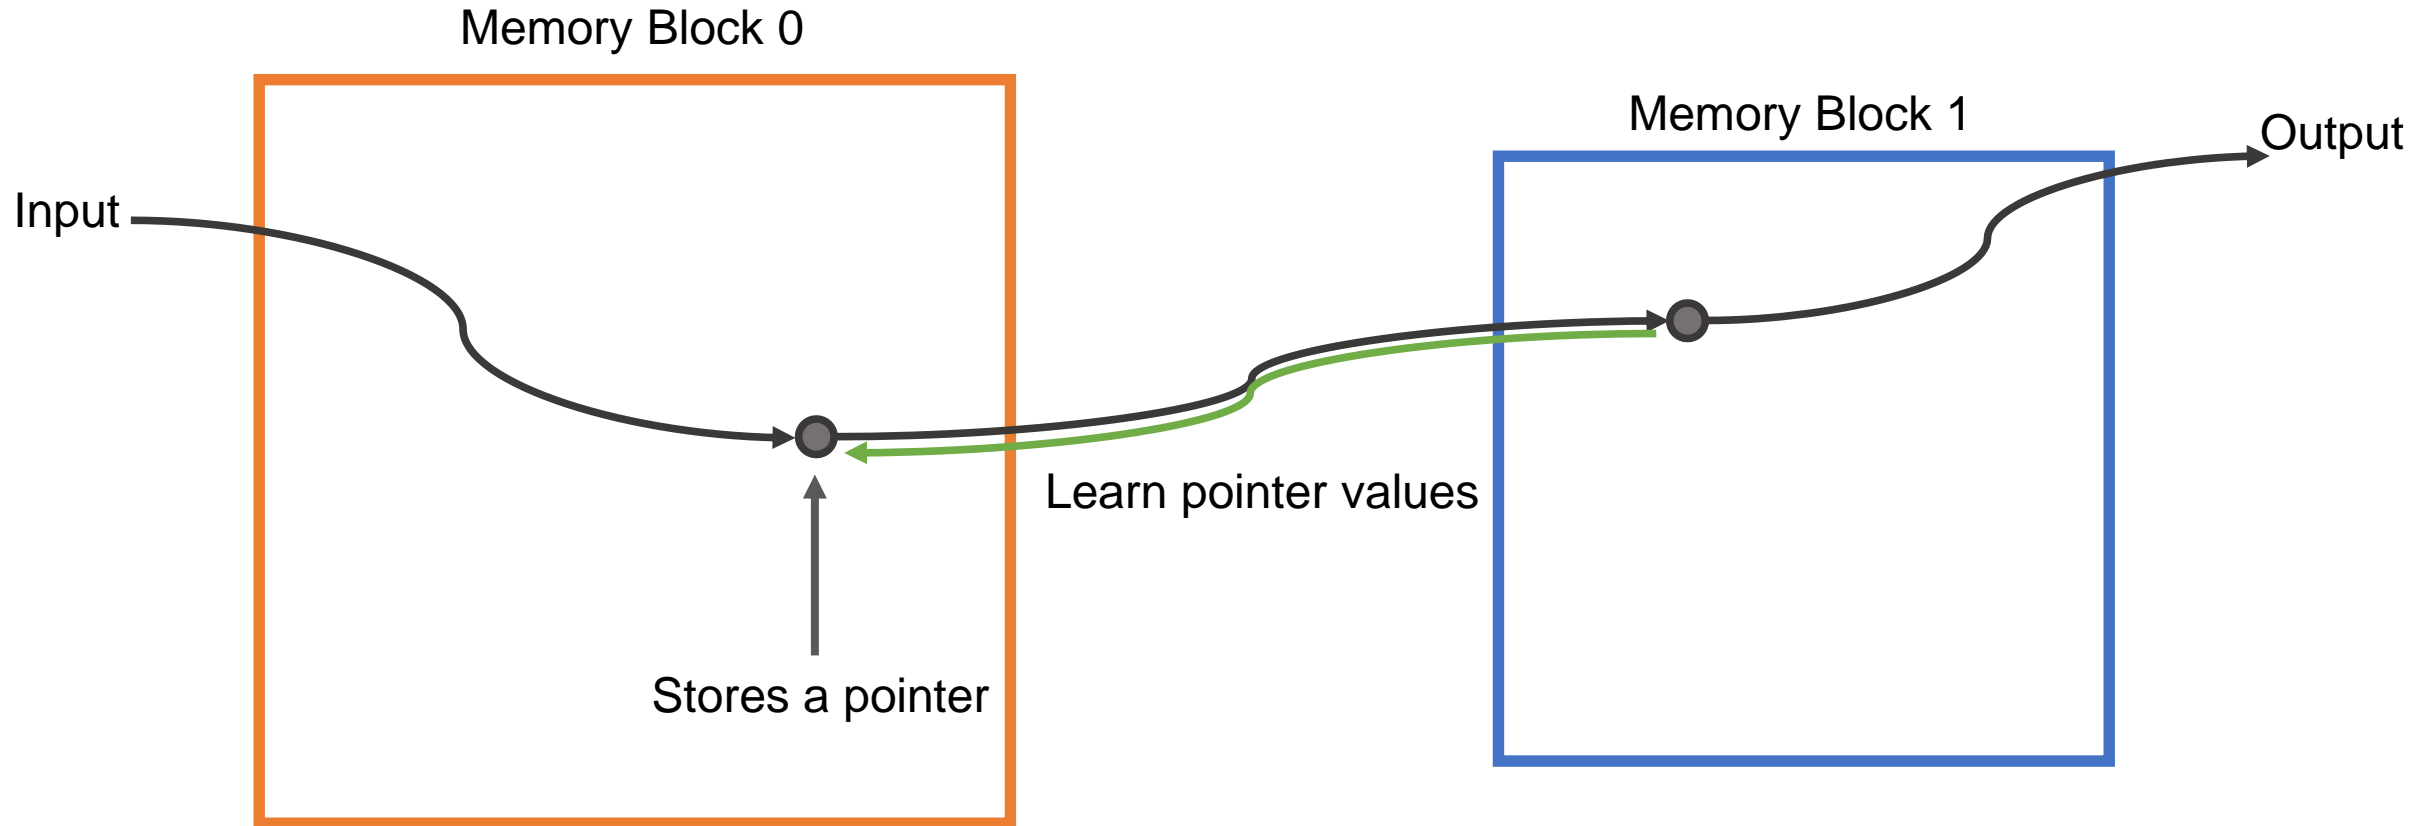
# What is Differentiable Indirection?



# What is Differentiable Indirection?



# What is Differentiable Indirection?



# Why use Differentiable Indirection?

# Why use Differentiable Indirection?


- Efficiency:



# Why use Differentiable Indirection?

- Efficiency:


MLP



# Why use Differentiable Indirection?

- Efficiency:


MLP



+	-
×	÷

  
Compute efficient


Differentiable indirection




# Why use Differentiable Indirection?

- Efficiency:


MLP



Compute efficient




Feature grids



Space efficient



Differentiable indirection




# Why use Differentiable Indirection?

- Efficiency:
  - Order of magnitude less compute than MLP.
  - Order of magnitude more compact than regular grid.
- Flexibility:

# Why use Differentiable Indirection?

- Efficiency:
  - Order of magnitude less compute than MLP.
  - Order of magnitude more compact than regular grid.
- Flexibility:


SDF



# Why use Differentiable Indirection?

- Efficiency:
  - Order of magnitude less compute than MLP.
  - Order of magnitude more compact than regular grid.
- Flexibility:

SDF




Texture  
Compression



# Why use Differentiable Indirection?

- Efficiency:
  - Order of magnitude less compute than MLP.
  - Order of magnitude more compact than regular grid.
- Flexibility:


SDF



Texture  
Compression




Compression  
and Filtering



# Why use Differentiable Indirection?

- Efficiency:
  - Order of magnitude less compute than MLP.
  - Order of magnitude more compact than regular grid.
- Flexibility:


SDF



Texture  
Compression



Compression  
and Filtering



Disney  
Shading






# Why use Differentiable Indirection?

- Efficiency:
  - Order of magnitude less compute than MLP.
  - Order of magnitude more compact than regular grid.
- Flexibility:


SDF



Texture  
Compression




Compression  
and Filtering



Disney  
Shading




Radiance Field  
Compression



# Why use Differentiable Indirection?

- Efficiency:
  - Order of magnitude less compute than MLP.
  - Order of magnitude more compact than regular grid.
- Flexibility:

Graphics and  
beyond...




# Why use Differentiable Indirection?

- Efficiency:
  - Order of magnitude less compute than MLP.
  - Order of magnitude more compact than regular grid.
- Flexibility:
  - Variety of applications across graphics pipeline and beyond.
- Adaptivity:

# Why use Differentiable Indirection?

- Efficiency:
  - Order of magnitude less compute than MLP.
  - Order of magnitude more compact than regular grid.
- Flexibility:
  - Variety of applications across graphics pipeline and beyond.
- **Adaptivity:**

Adaptive resolution with differentiable indirection.




# Why use Differentiable Indirection?

- Efficiency:
  - Order of magnitude less compute than MLP.
  - Order of magnitude more compact than regular grid.
- Flexibility:
  - Variety of applications across graphics pipeline and beyond.

- **Adaptivity:**

Adaptive resolution with spatial partitioning e.g. Oct/Kd-trees.

Adaptive resolution with differentiable indirection.




# Why use Differentiable Indirection?

- Efficiency:
  - Order of magnitude less compute than MLP.
  - Order of magnitude more compact than regular grid.
- Flexibility:
  - Variety of applications across graphics pipeline and beyond.


- **Adaptivity:**

Adaptive resolution with spatial partitioning e.g. Oct/Kd-trees.



Require multiple indirections.

Adaptive resolution with differentiable indirection.




# Why use Differentiable Indirection?


- Efficiency:
  - Order of magnitude less compute than MLP.
  - Order of magnitude more compact than regular grid.
- Flexibility:
  - Variety of applications across graphics pipeline and beyond.

- **Adaptivity:**

Adaptive resolution with spatial partitioning e.g. Oct/Kd-trees.




Require multiple indirections.



Bandwidth efficient



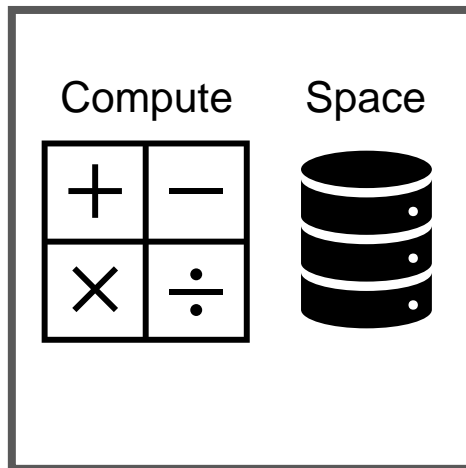
Adaptive resolution with differentiable indirection.




# Why use Differentiable Indirection?

- Efficiency:
  - Order of magnitude less compute than MLP.
  - Order of magnitude more compact than regular grid.
- Flexibility:
  - Variety of applications across graphics pipeline and beyond.
- Adaptivity:
  - Adaptive spatial resolution without tree-structures. More coherent memory access.

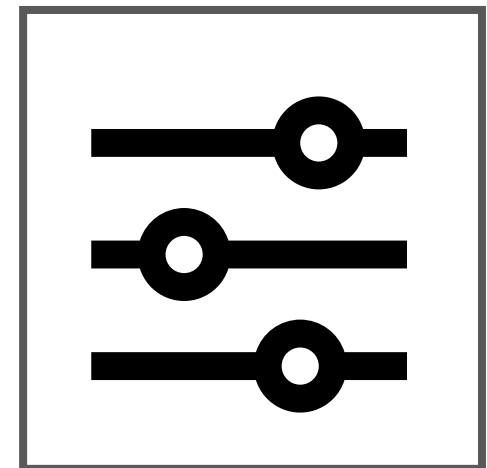
Efficient



Flexible



Adaptive resolution





# Technical Summary

# Technical Summary

Primary



# Technical Summary

Primary



Cascaded


# Technical Summary : Forward pass

Primary




Cascaded




# Technical Summary : Forward pass

Primary




Cascaded




# Technical Summary : Forward pass

Primary




Cascaded




# Technical Summary : Forward pass

Primary




Cascaded




Constraining non-linearity - bounds output values between 0 to 1.

# Technical Summary : Forward pass

Primary



Cascaded



Constraining non-linearity - bounds output values between 0 to 1.


Inference optimizations:

- Bake-in non-linearity in array cells.




# Technical Summary : Forward pass

Primary



Cascaded



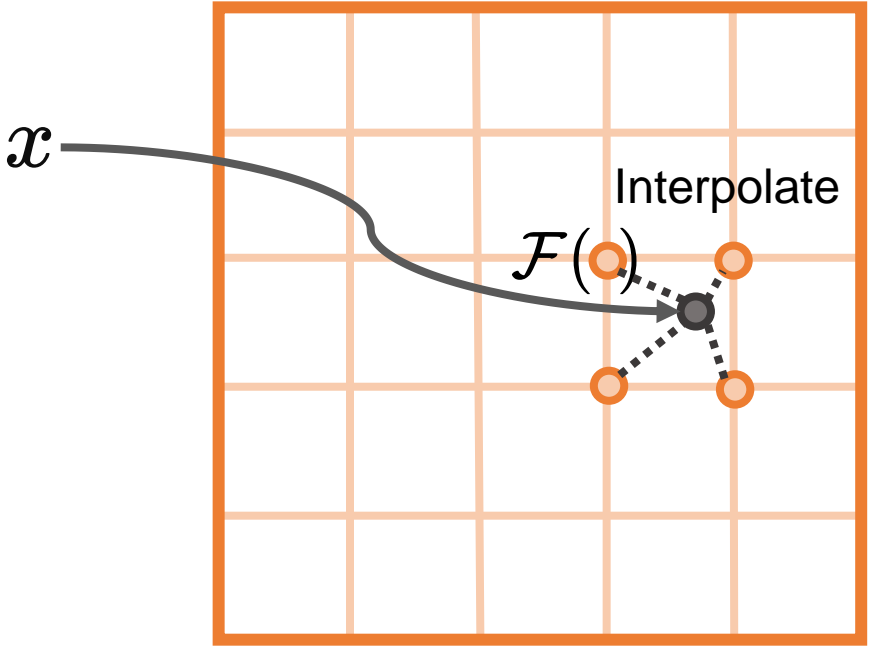
Constraining non-linearity - bounds output values between 0 to 1.

Inference optimizations:


- Bake-in non-linearity in array cells.
- Quantize array cells.

# Technical Summary : Forward pass

Primary



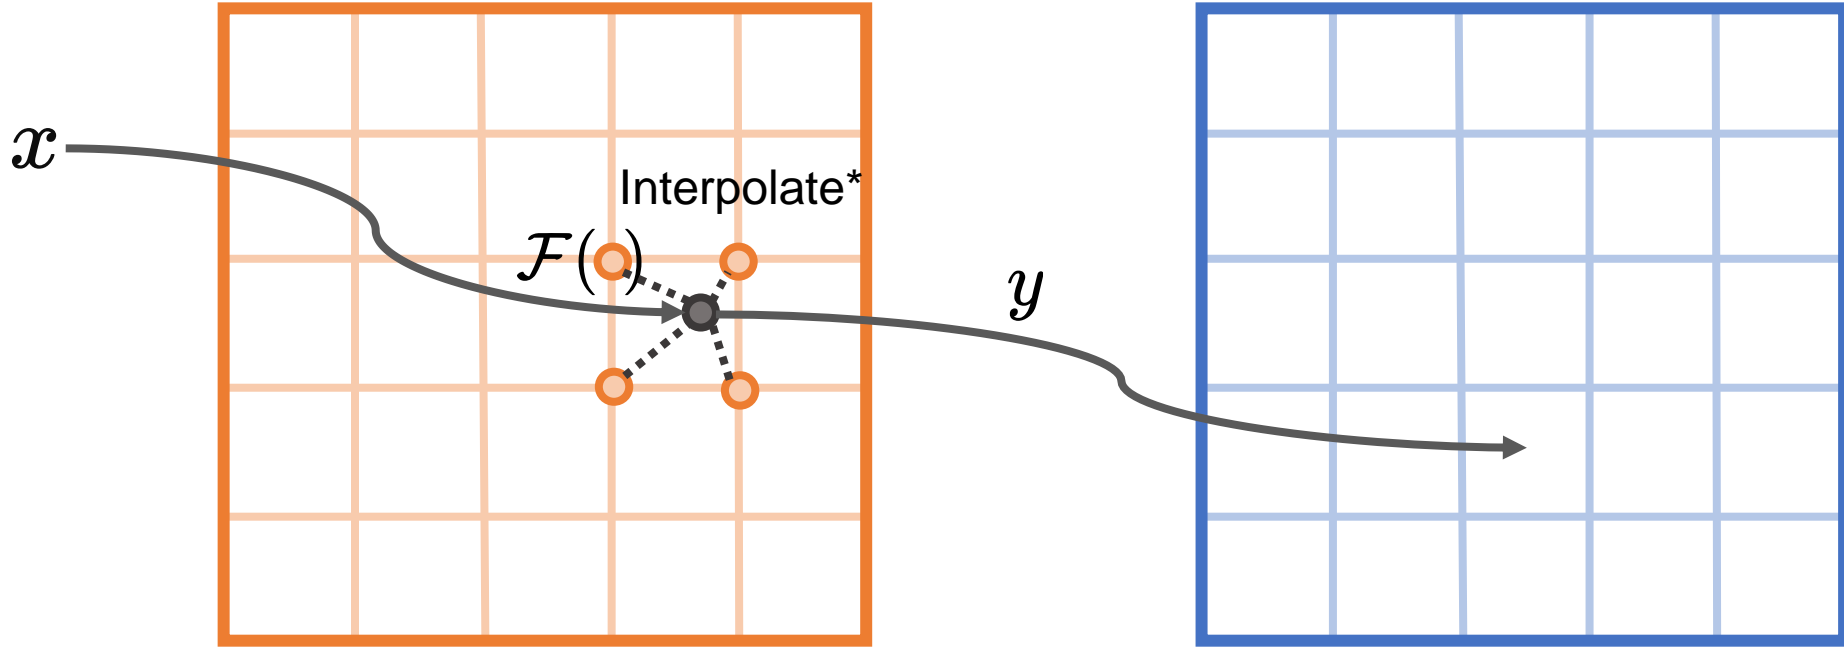
Cascaded



# Technical Summary : Forward pass

Primary

Cascaded




\* Linear (tested) or higher order interpolants.

# Technical Summary : Forward pass


Primary

Cascaded

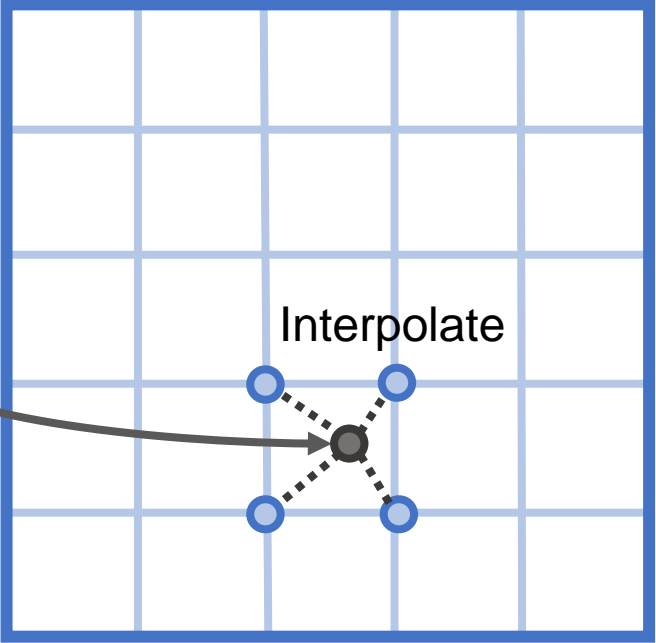


# Technical Summary

Primary




Cascaded



# Technical Summary : Forward pass

Primary


Cascaded



# Technical Summary : Backward pass


Primary

Cascaded

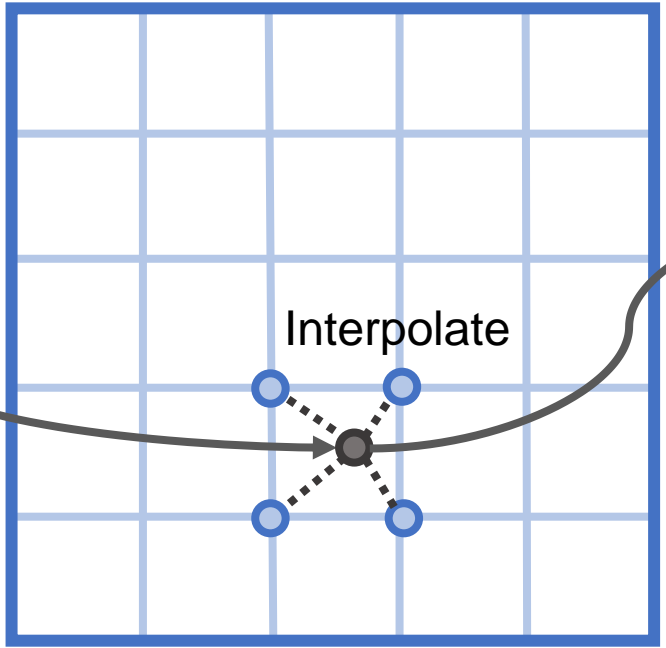


# Technical Summary : Backward pass


Primary




Cascaded



Output



Target




$\mathcal{L}$



# Technical Summary : Backward pass

Primary

Cascaded




  $\frac{d\mathcal{L}}{d\mathcal{O}}$


# Technical Summary : Backward pass

Primary

Cascaded




Update cascaded:  $\frac{d\mathcal{L}}{d\mathcal{O}}$




# Technical Summary : Backward pass

Primary




Cascaded



Update cascaded:  $\frac{d\mathcal{L}}{dO} \cdot \frac{dO}{dc_i}$


Output



Target




$\mathcal{L}$



# Technical Summary : Backward pass

Primary

Cascaded



Update cascaded:  $\frac{d\mathcal{L}}{d\mathcal{O}} \cdot \frac{d\mathcal{O}}{dc_i}$

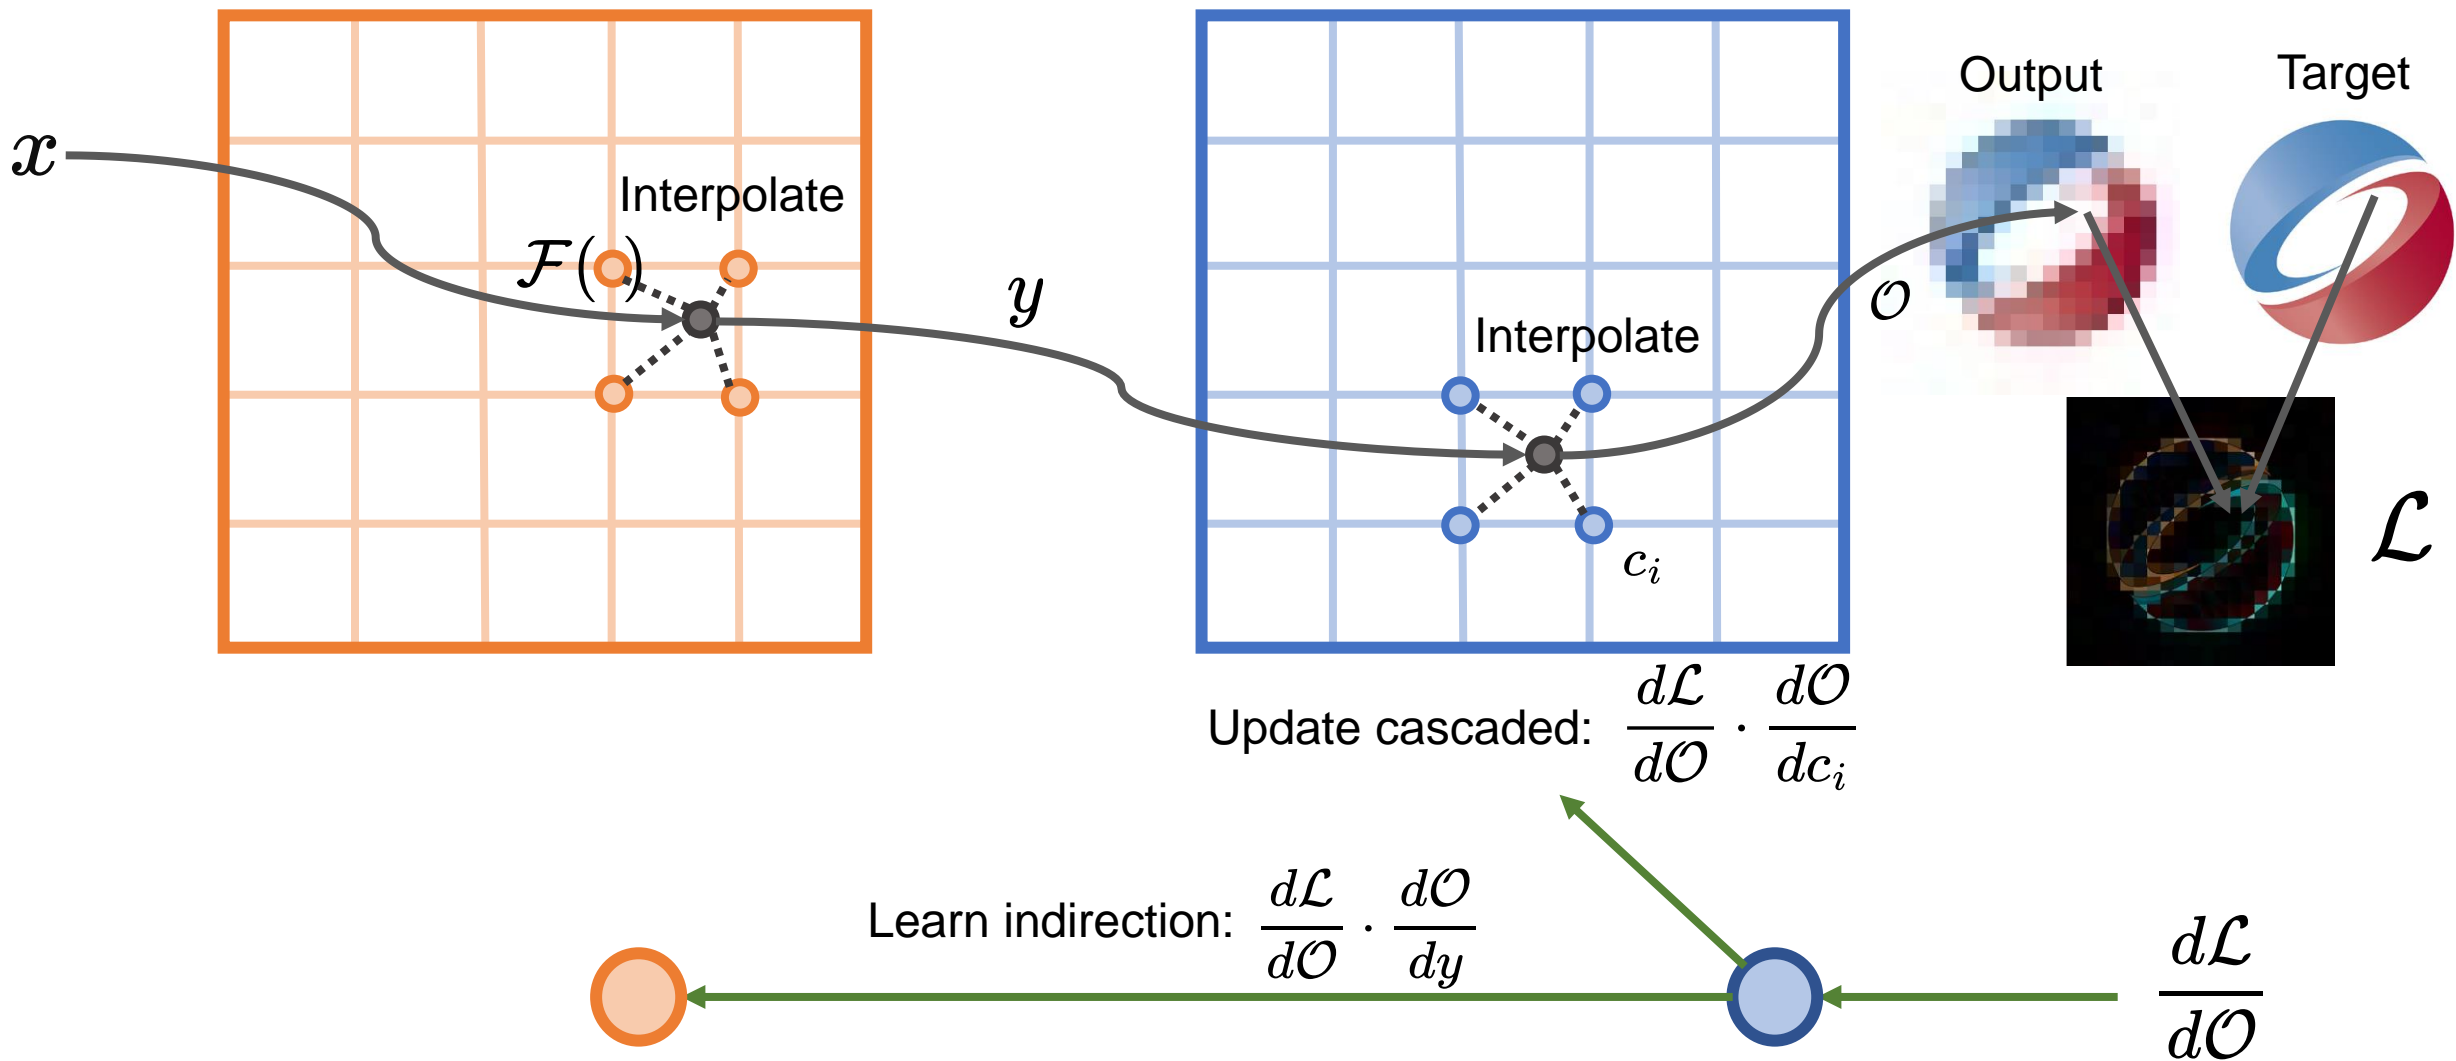
[Muller et. al 2022, Kanwar et. al 2022...]

$\frac{d\mathcal{L}}{d\mathcal{O}}$

# Technical Summary : Backward pass

Primary


Cascaded



# Technical Summary : Backward pass

Primary

Cascaded




Update primary:  $\frac{d\mathcal{L}}{d\mathcal{O}} \cdot \frac{d\mathcal{O}}{dy} \cdot \frac{dy}{dg_j}$

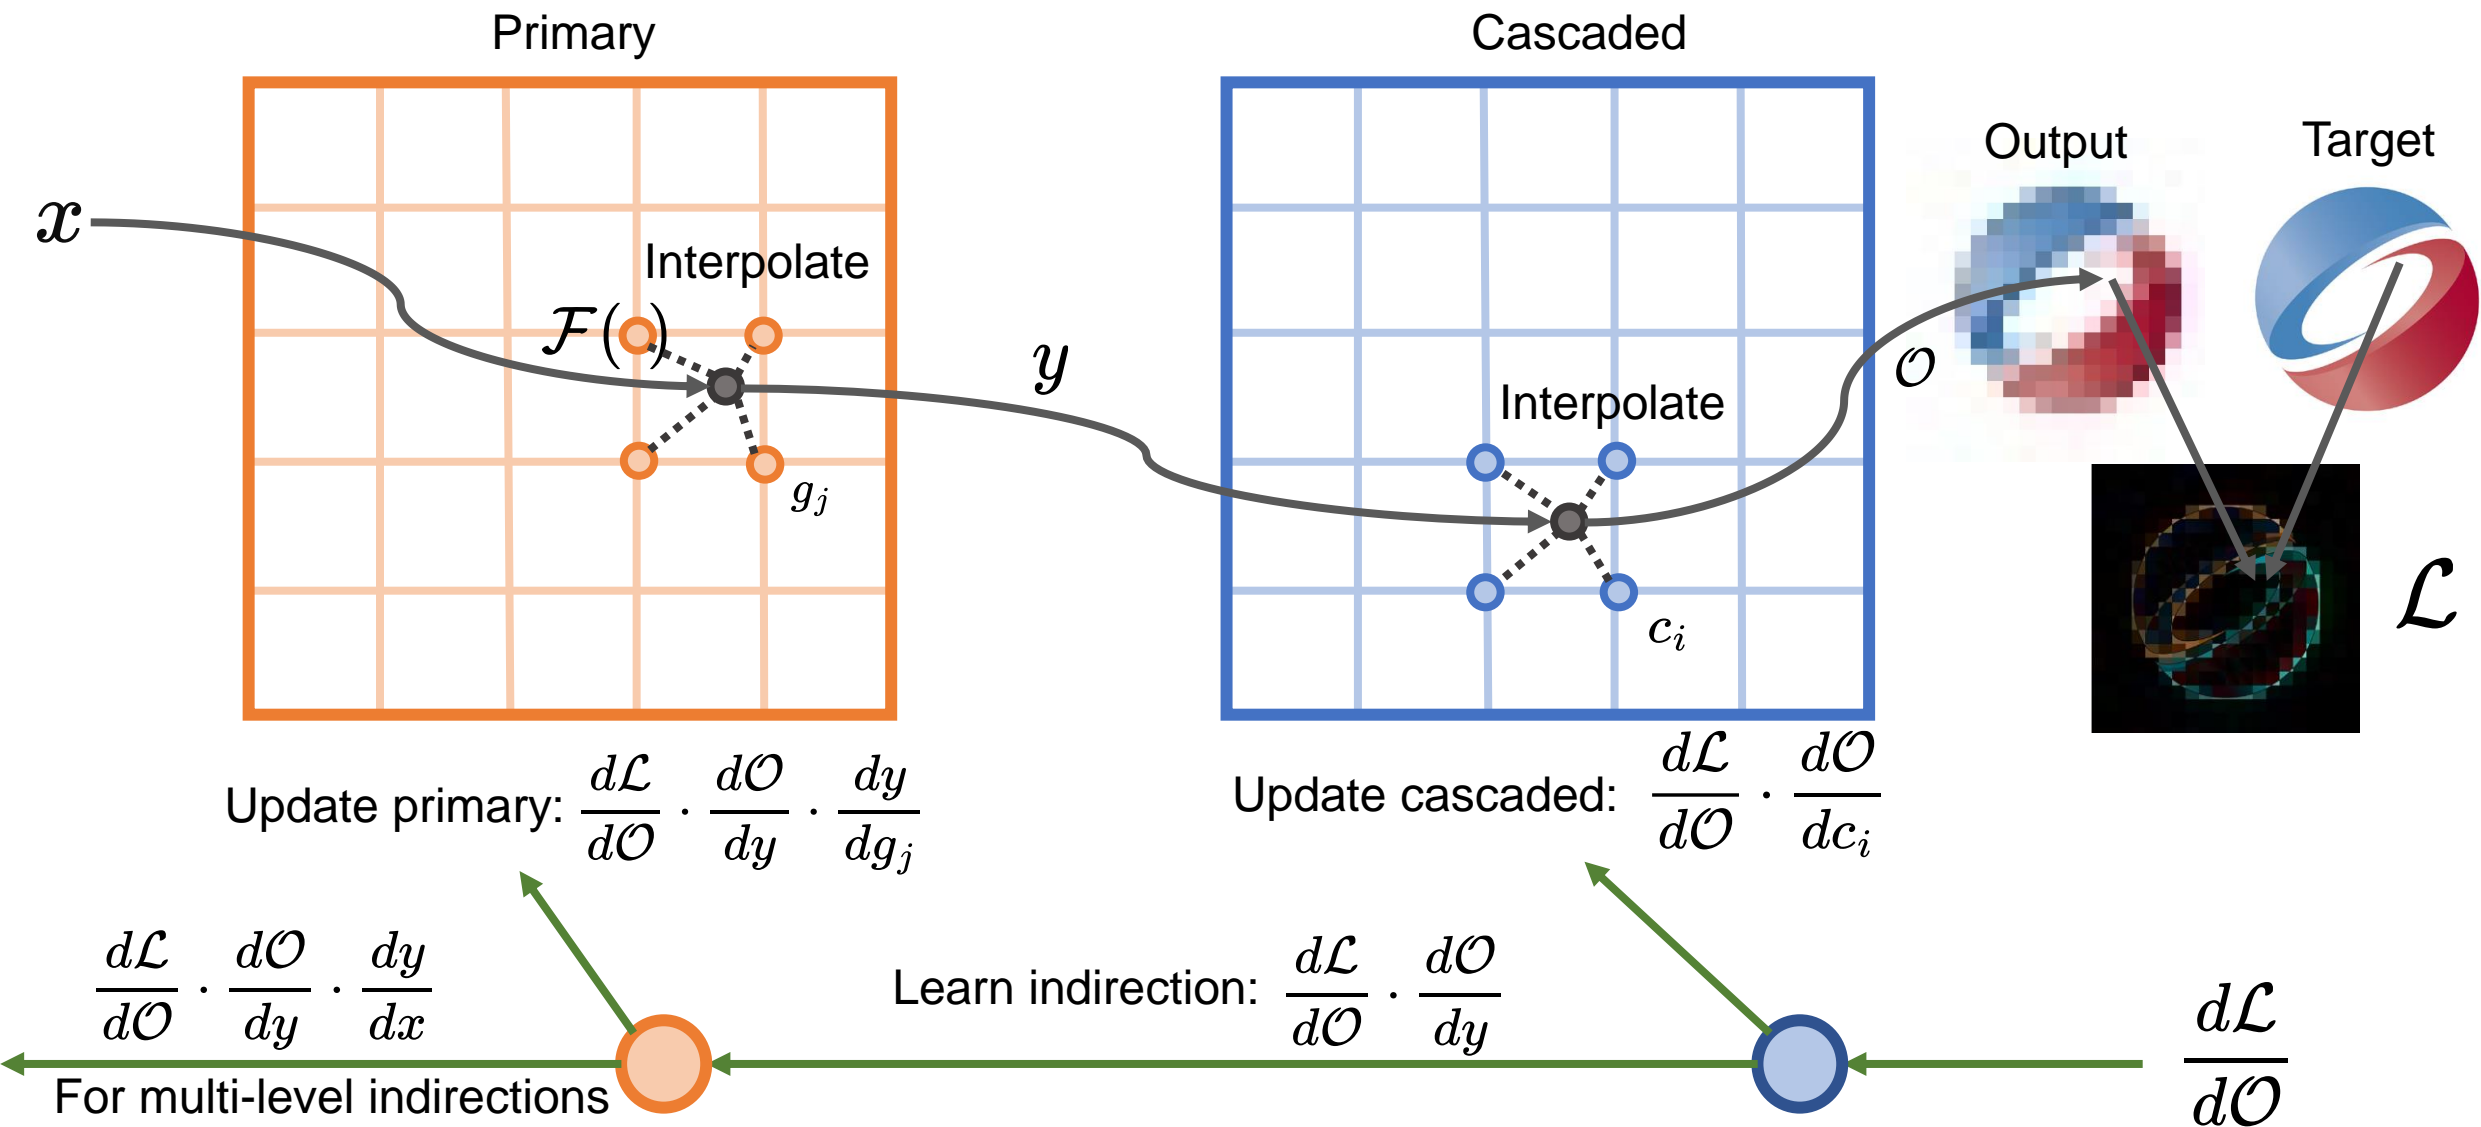
Update cascaded:  $\frac{d\mathcal{L}}{d\mathcal{O}} \cdot \frac{d\mathcal{O}}{dc_i}$

Learn indirection:  $\frac{d\mathcal{L}}{d\mathcal{O}} \cdot \frac{d\mathcal{O}}{dy}$

$\frac{d\mathcal{L}}{d\mathcal{O}}$




# Technical Summary : Backward pass



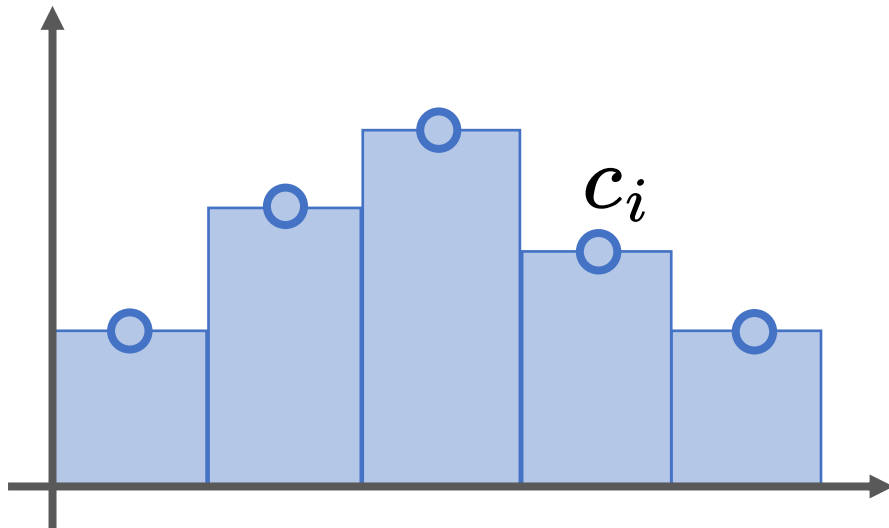

# Intuition





# Intuition



# Intuition




# Intuition



# Intuition


1D-array



$C_i$

Two gradients:

# Intuition




Two gradients:

$$\frac{dO}{dc_i}$$

↓

Gradient of output w.r.t. array cell values.

# Intuition



Two gradients:

$$\frac{dO}{dc_i}$$

↓


Gradient of output w.r.t. array cell values.

$$\frac{dO}{dy}$$

↓

Gradient of output w.r.t. array input coordinates.

# Intuition



Two gradients:

$$\frac{dO}{dc_i}$$

↓


Gradient of output w.r.t. array cell values.  
Purpose: Update array cells.

$$\frac{dO}{dy}$$

↓

Gradient of output w.r.t. array input coordinates.

# Intuition



Two gradients:

$$\frac{dO}{dc_i}$$

↓

Gradient of output w.r.t. array cell values.  
Purpose: Update array cells.

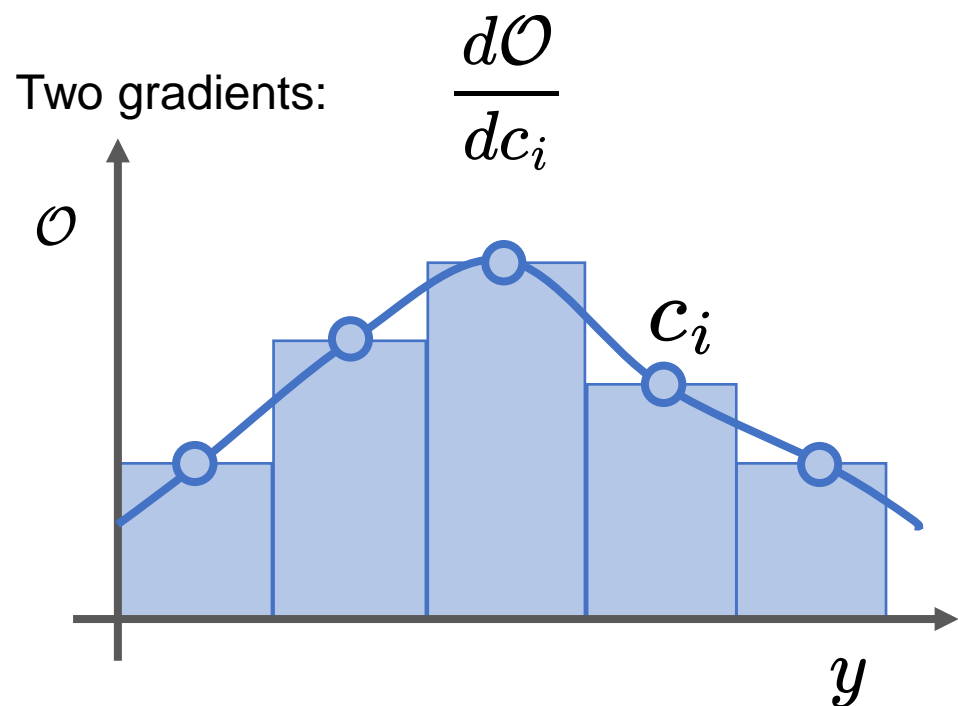

$$\frac{dO}{dy}$$

↓

Gradient of output w.r.t. array input coordinates.  
Purpose: Learn indirection.





# Intuition





$$\frac{dO}{dy}$$

# Intuition




Muller et. al 2022, Kanwar et. al 2022...

# Intuition





Muller et. al 2022, Kanwar et. al 2022...




and additionally.

# Intuition





Muller et. al 2022, Kanwar et. al 2022...




and additionally.

# Intuition




Muller et. al 2022, Kanwar et. al 2022...




and additionally.


# Intuition : Adaptive resolution



Adjust cell boundaries

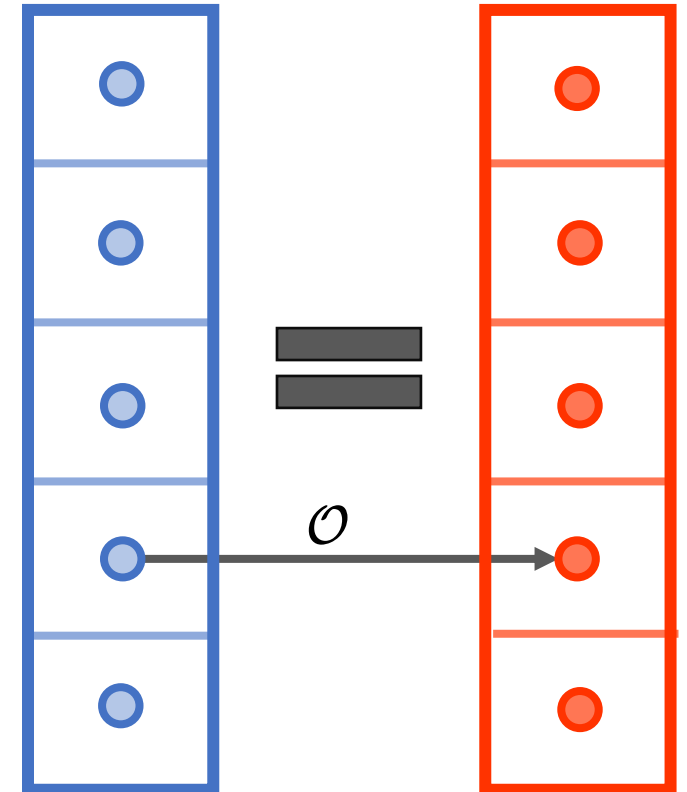


# Intuition : Adaptive resolution




Adjust cell boundaries

Hypothetical array adaptive resolution




# Intuition : Adaptive resolution



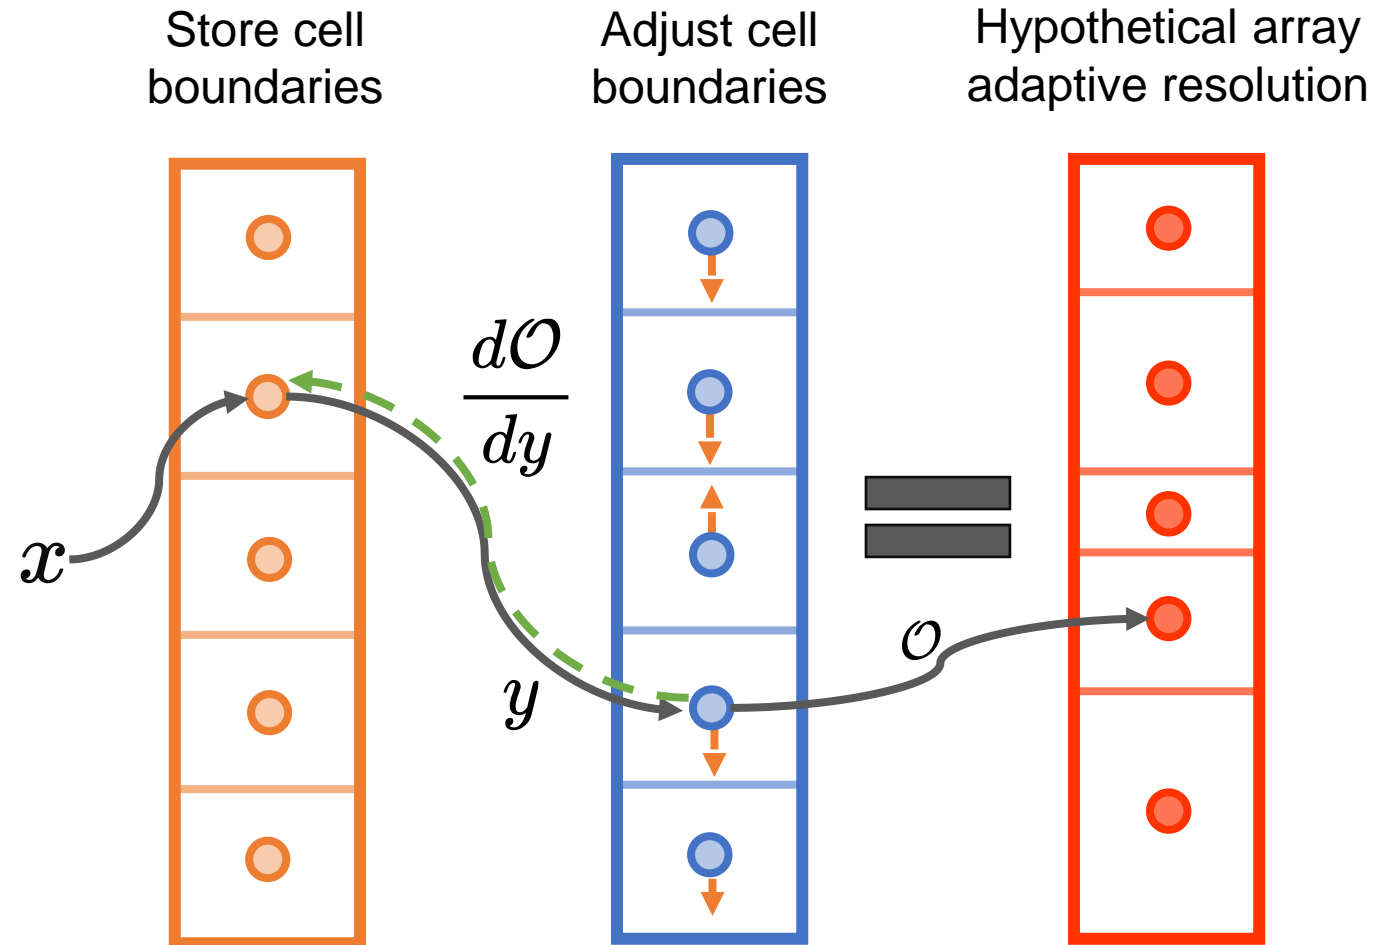

Adjust cell boundaries

Hypothetical array adaptive resolution



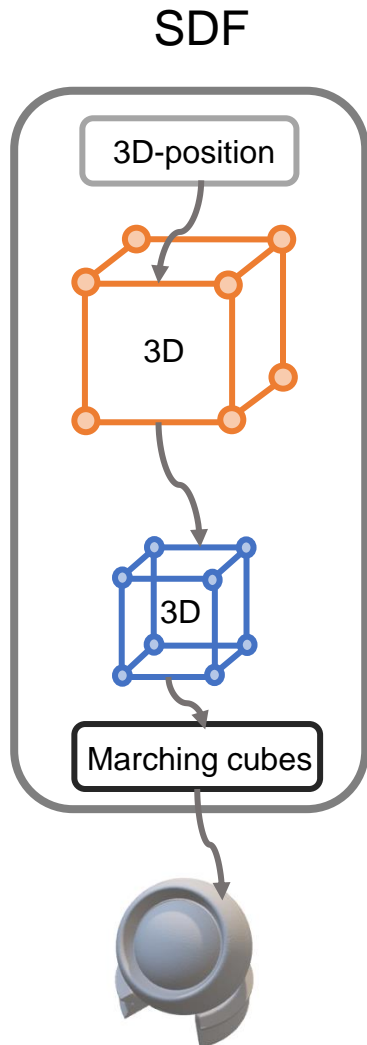


# Intuition : Adaptive resolution

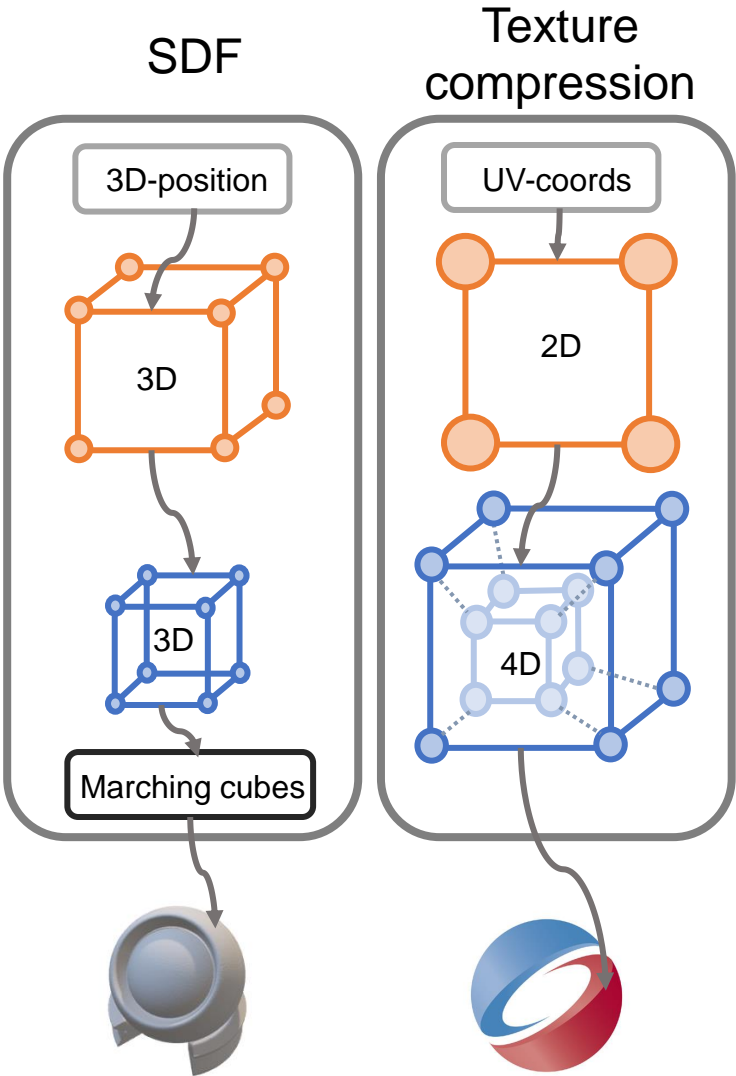


# Applications : Network structure

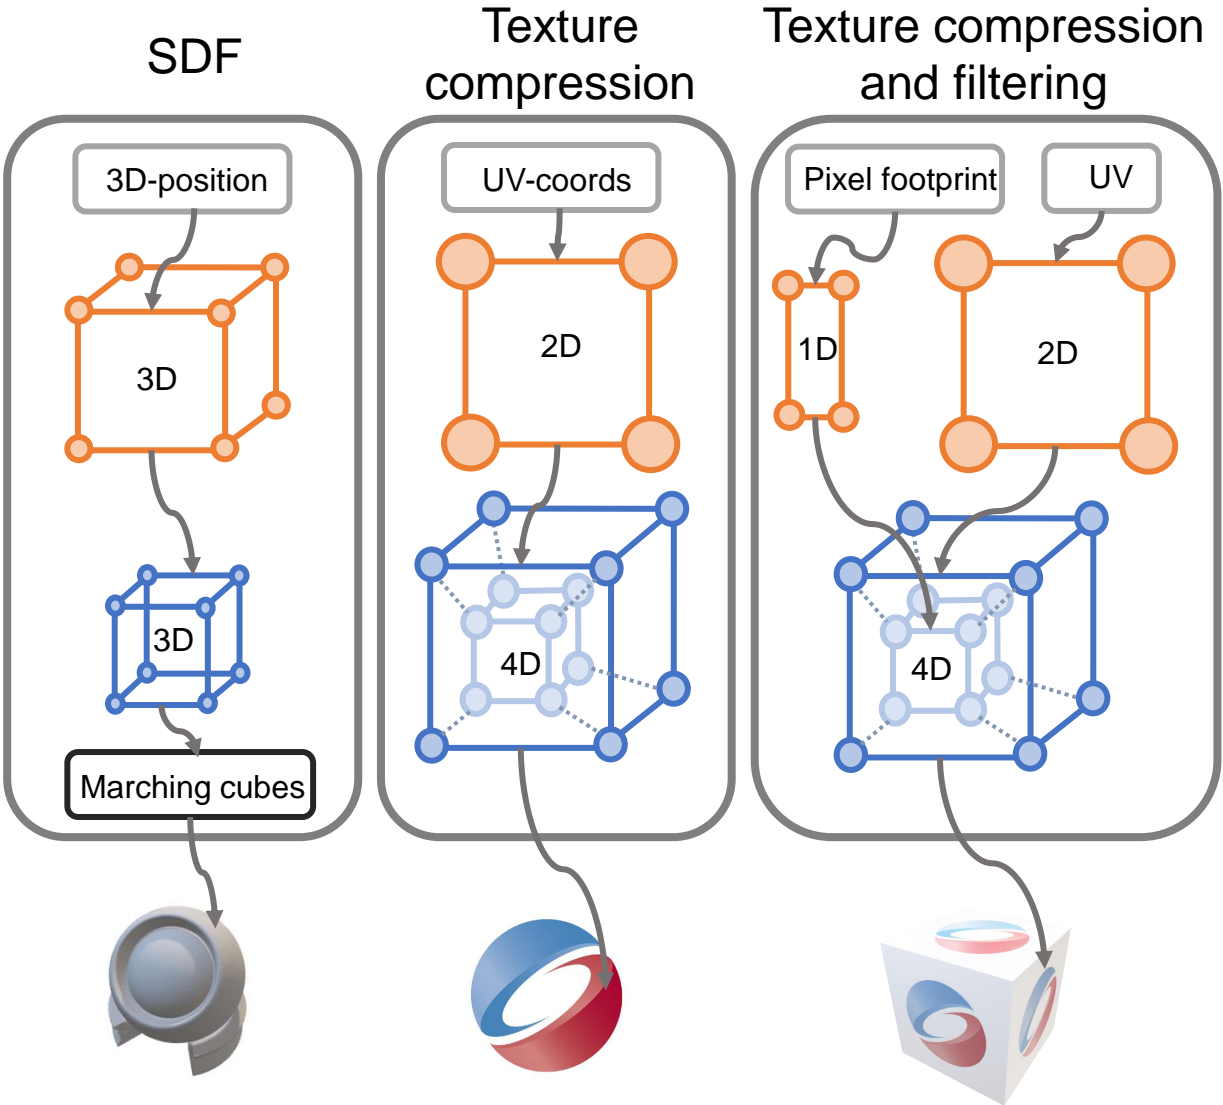
# Applications : Network structure



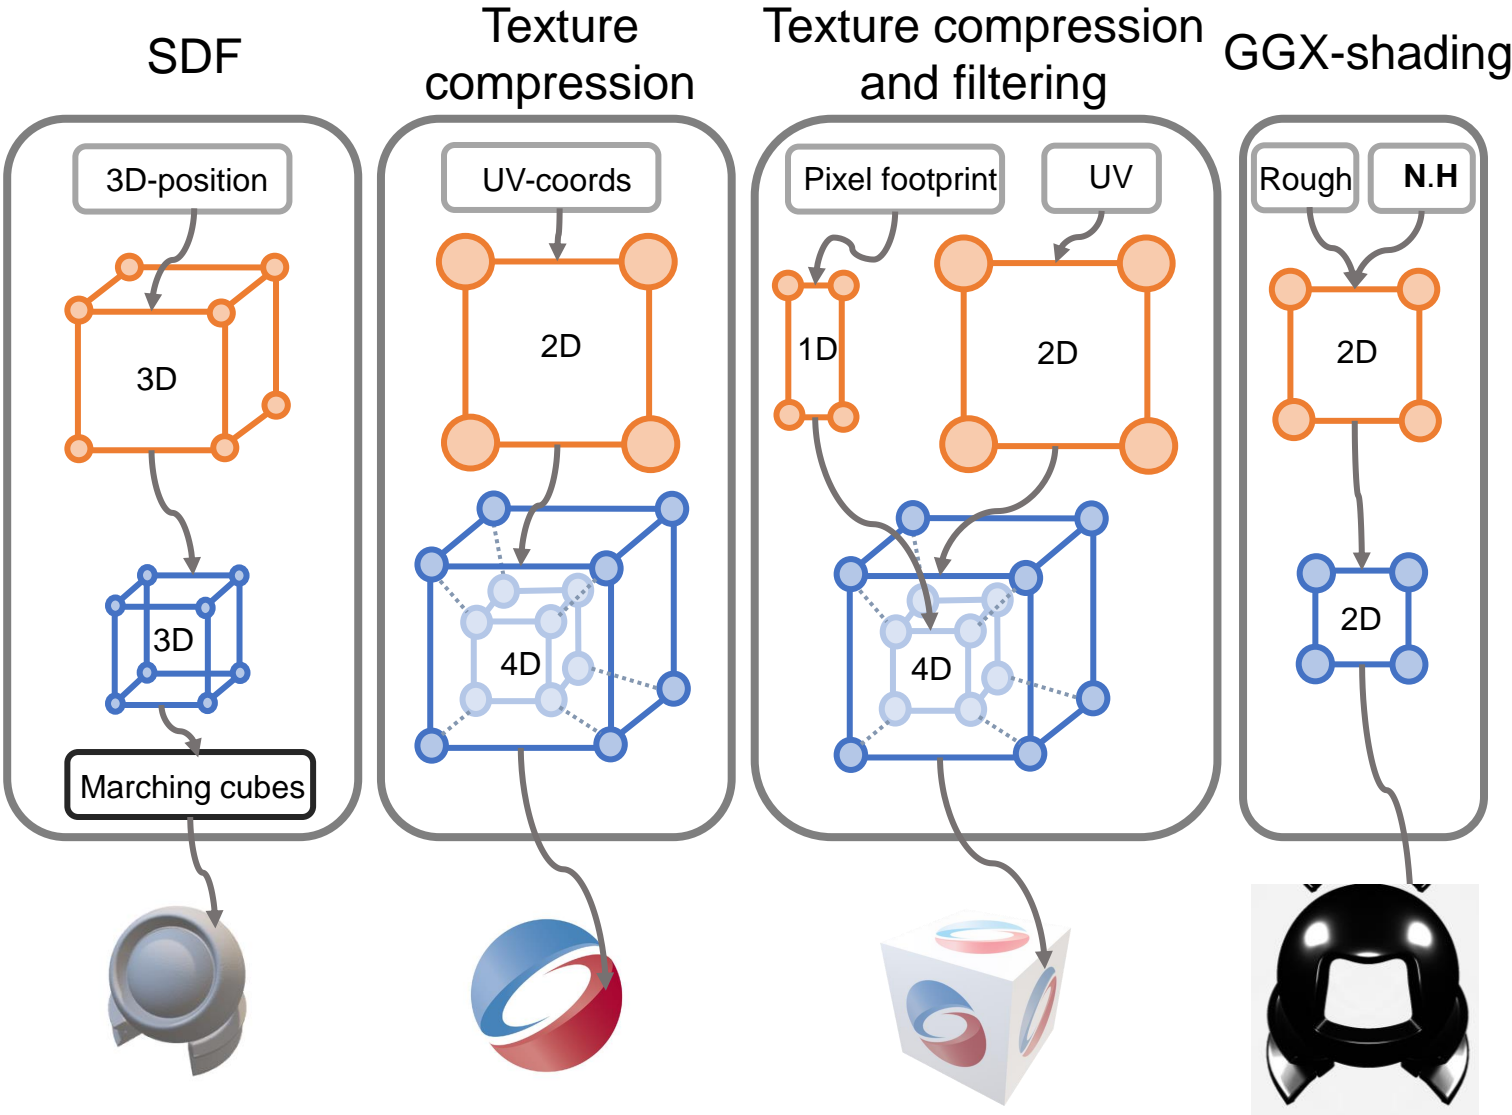
# Applications : Network structure



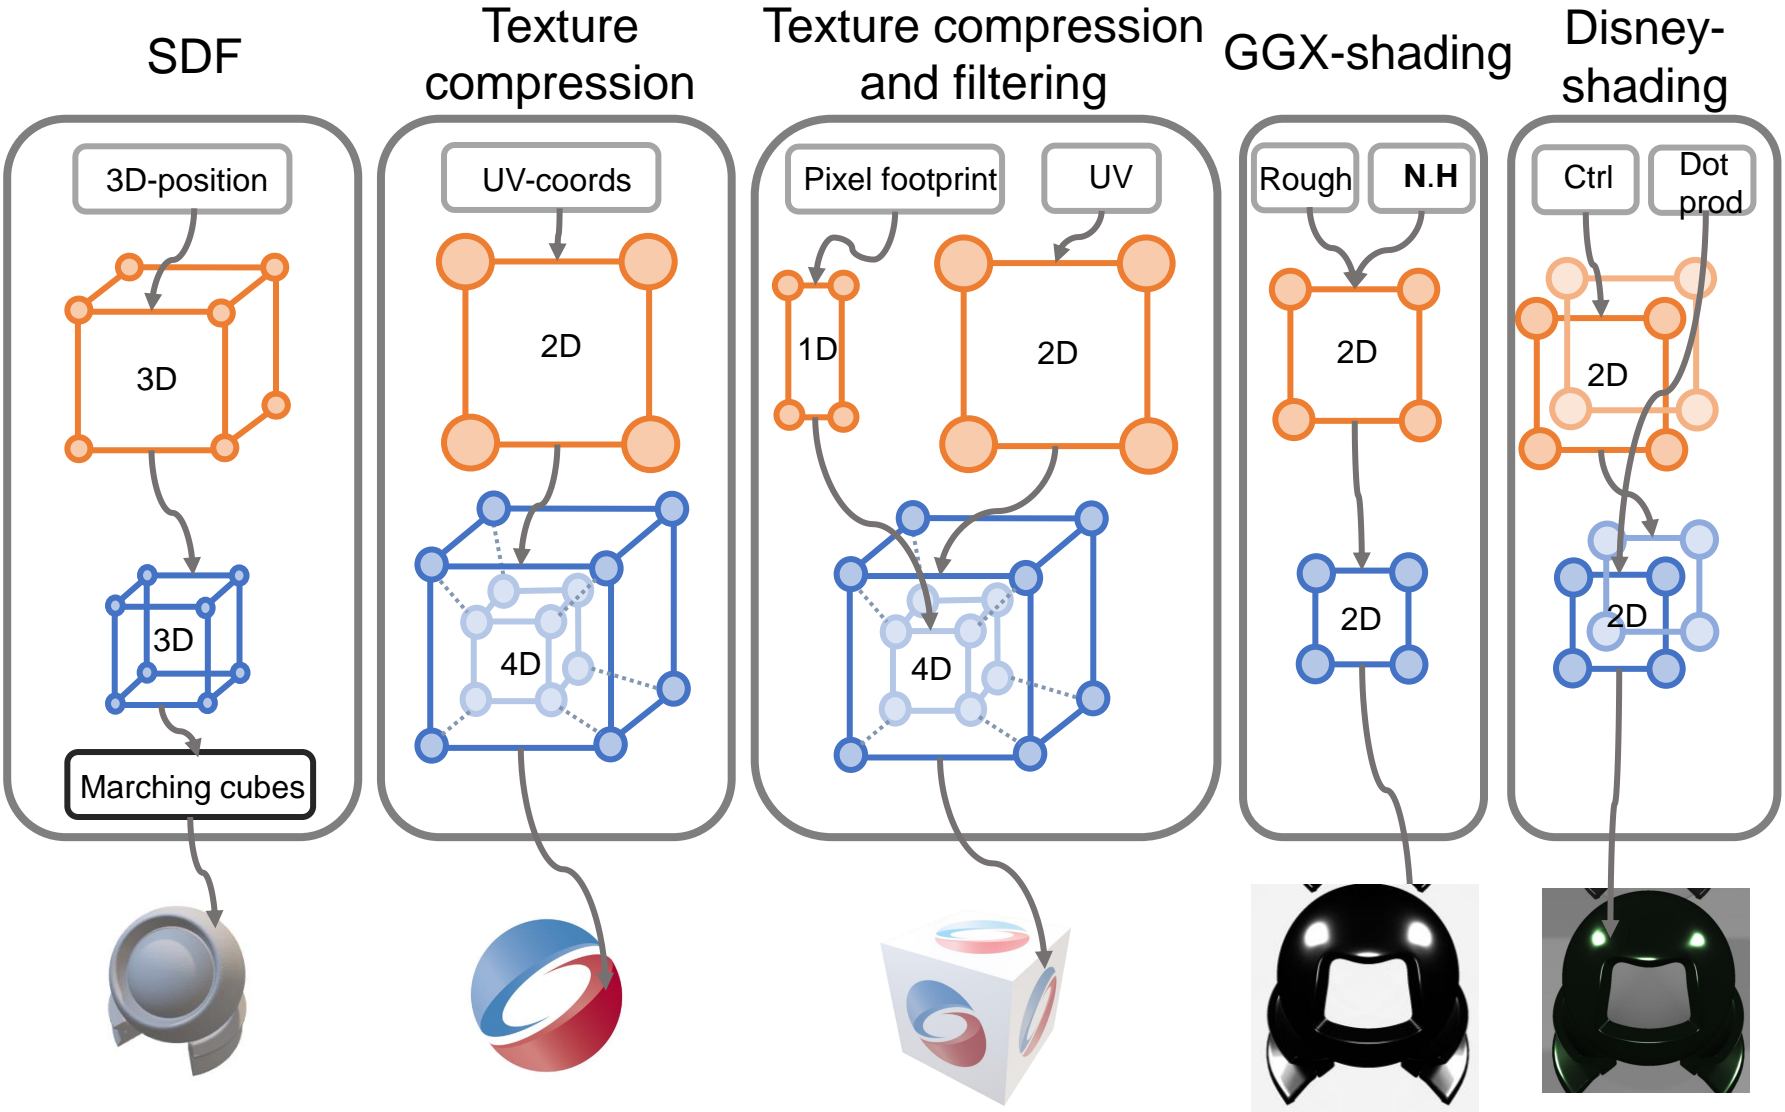
# Applications : Network structure



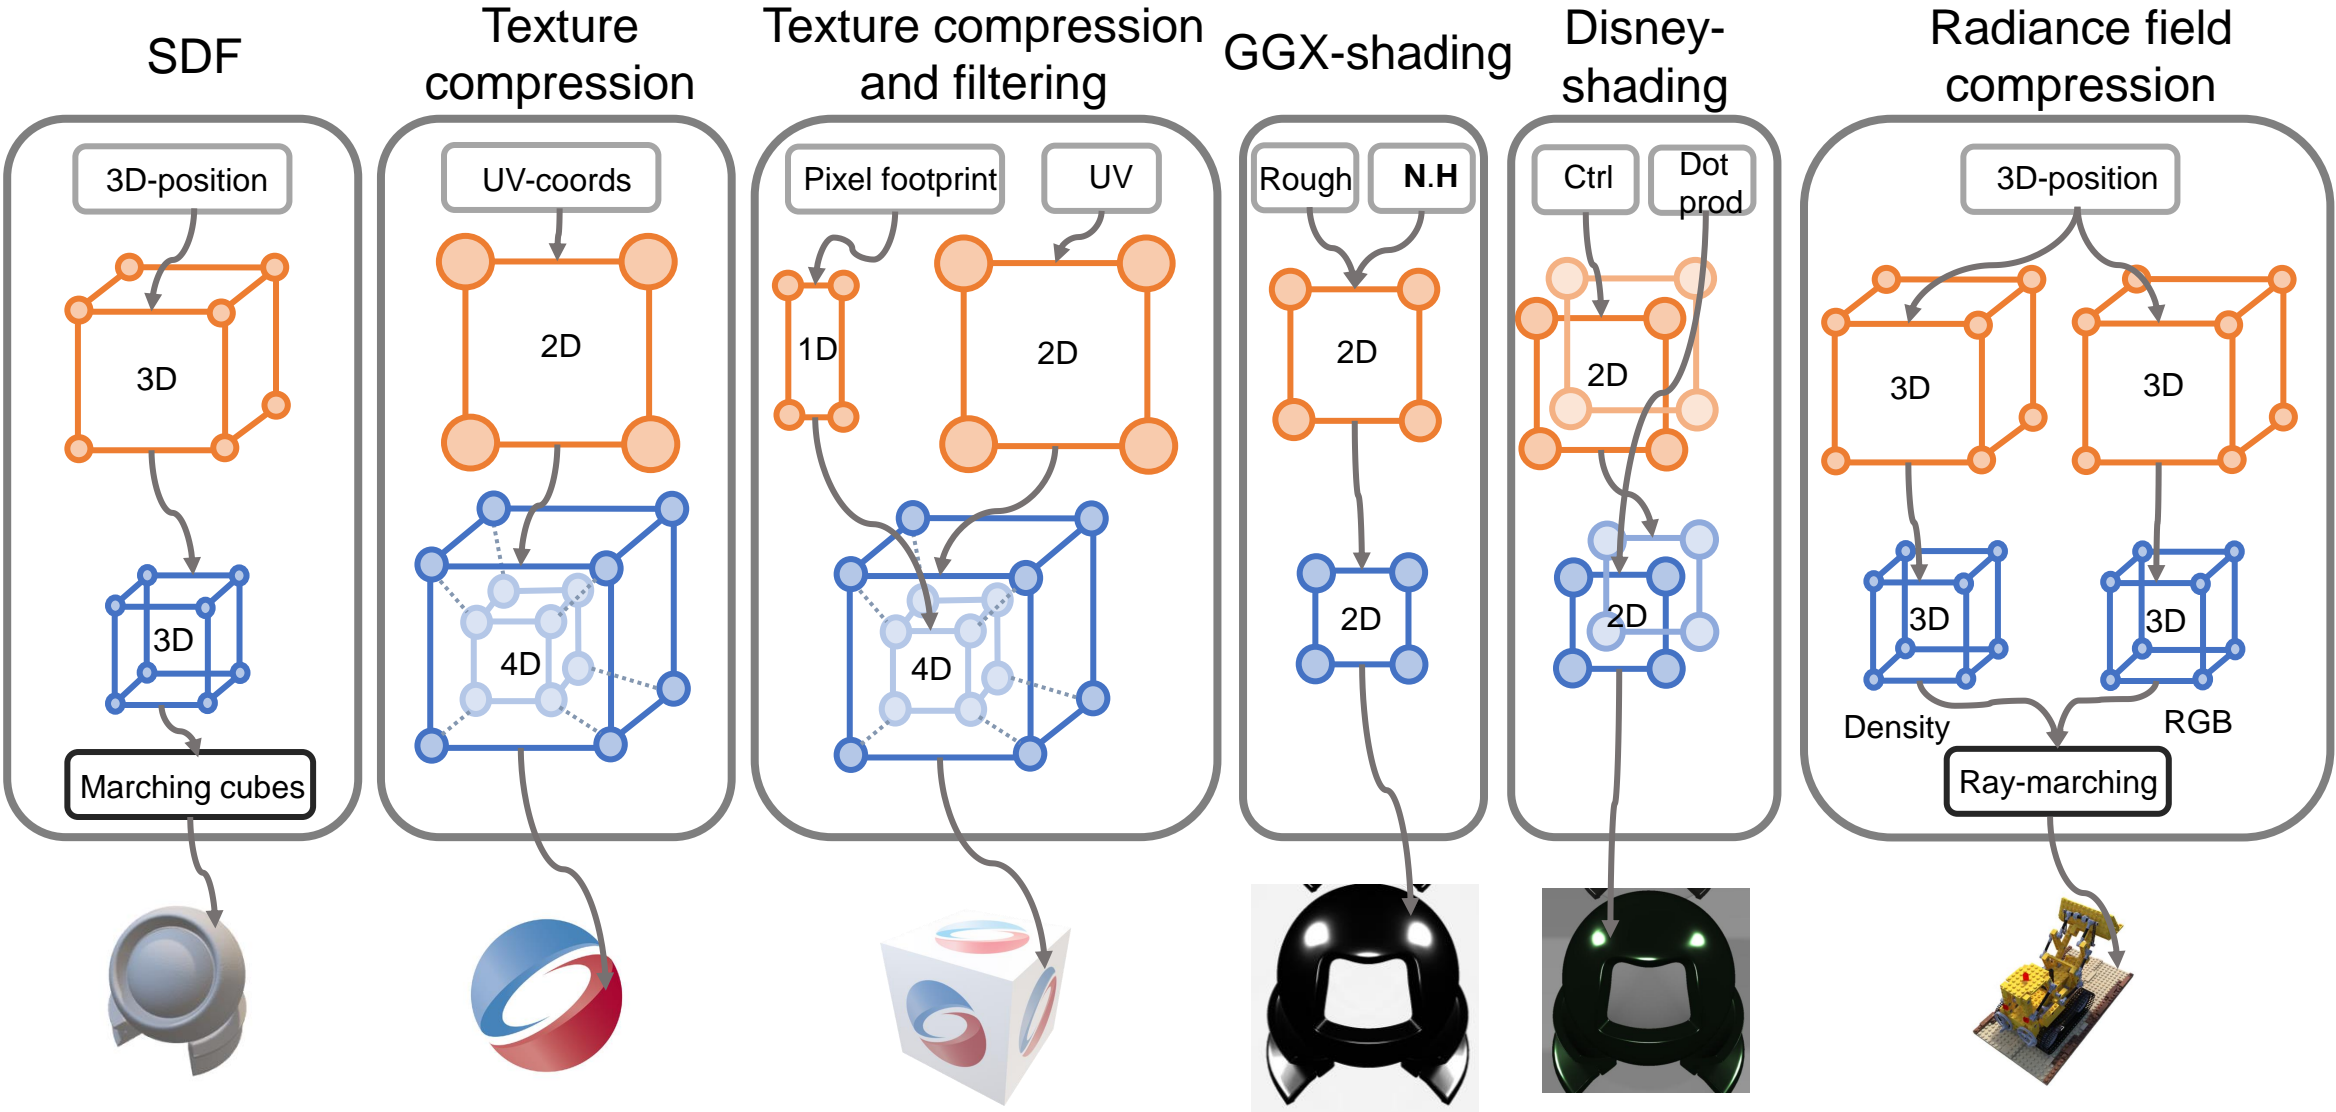
# Applications : Network structure



# Applications : Network structure



# Applications : Network structure





# Efficiency

# Efficiency

Illustrated through Isotropic GGX approximation.


- A 2D function with highly non-linear scalar output.

# Efficiency

Illustrated through Isotropic GGX approximation.

- A 2D function with highly non-linear scalar output.

Iso-GGX function as 2D texture




# Efficiency

Illustrated through Isotropic GGX approximation.

- A 2D function with highly non-linear scalar output.


Iso-GGX function as 2D texture



Shade



Scene shaded with Isotropic-GGX function




# Efficiency

Illustrated through Isotropic GGX approximation.

- A 2D function with highly non-linear scalar output.

Iso-GGX function as 2D texture



Approximate




Approximate Isotropic-GGX with:

# Efficiency

Illustrated through Isotropic GGX approximation.

- A 2D function with highly non-linear scalar output.


Iso-GGX function as 2D texture



Approximate  
→

MLP only:

Approximate Isotropic-GGX with:




# Efficiency

Illustrated through Isotropic GGX approximation.

- A 2D function with highly non-linear scalar output.


Iso-GGX function as 2D texture




Approximate  
→

Approximate Isotropic-GGX with:

MLP only:



1-level lookup:




# Efficiency

Illustrated through Isotropic GGX approximation.

- A 2D function with highly non-linear scalar output.

Iso-GGX function as 2D texture




Approximate




Approximate Isotropic-GGX with:


MLP only:



1-level lookup:



MLP + lookup:






# Efficiency

Illustrated through Isotropic GGX approximation.

- A 2D function with highly non-linear scalar output.


Iso-GGX function as 2D texture




Approximate  
→

Approximate Isotropic-GGX with:


MLP only:




1-level lookup:



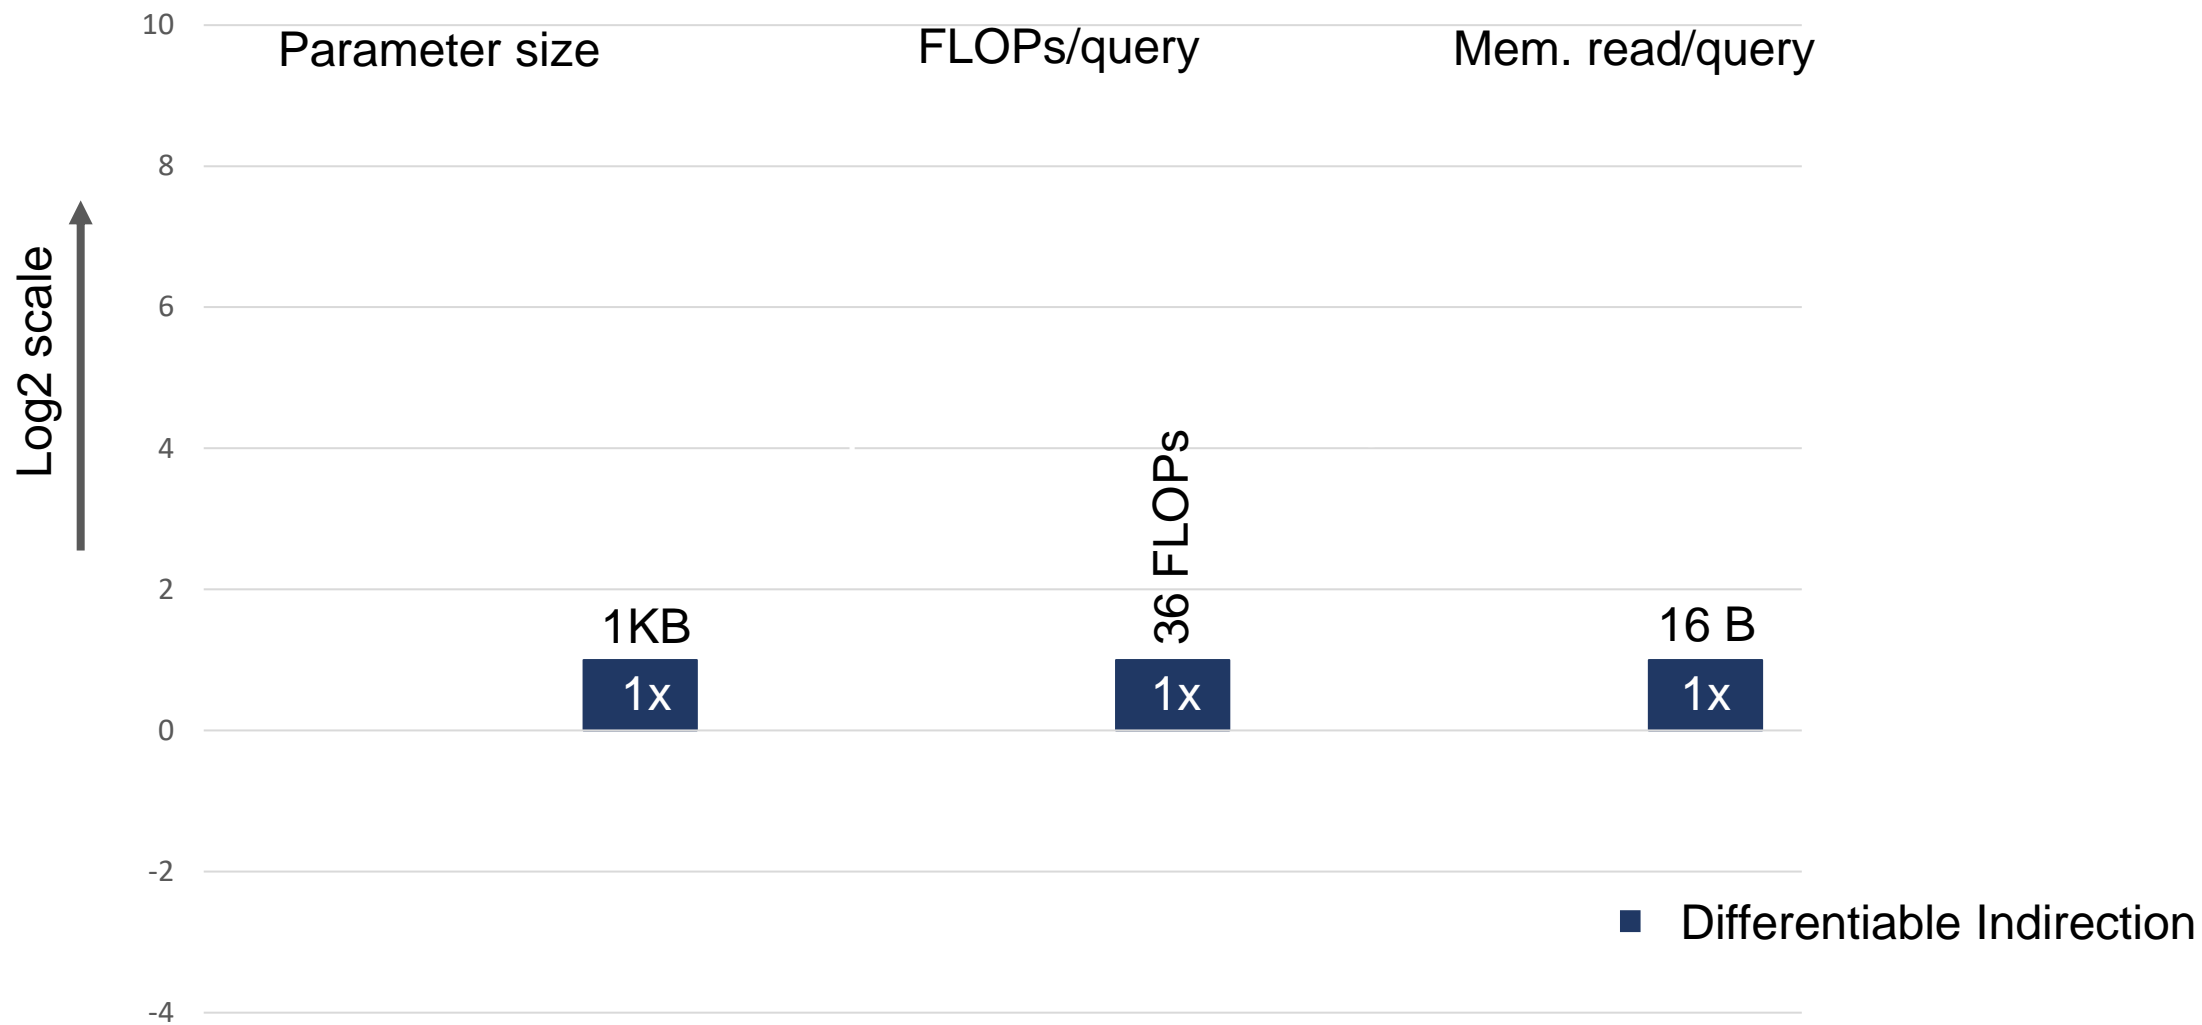
MLP + lookup:




Ours:



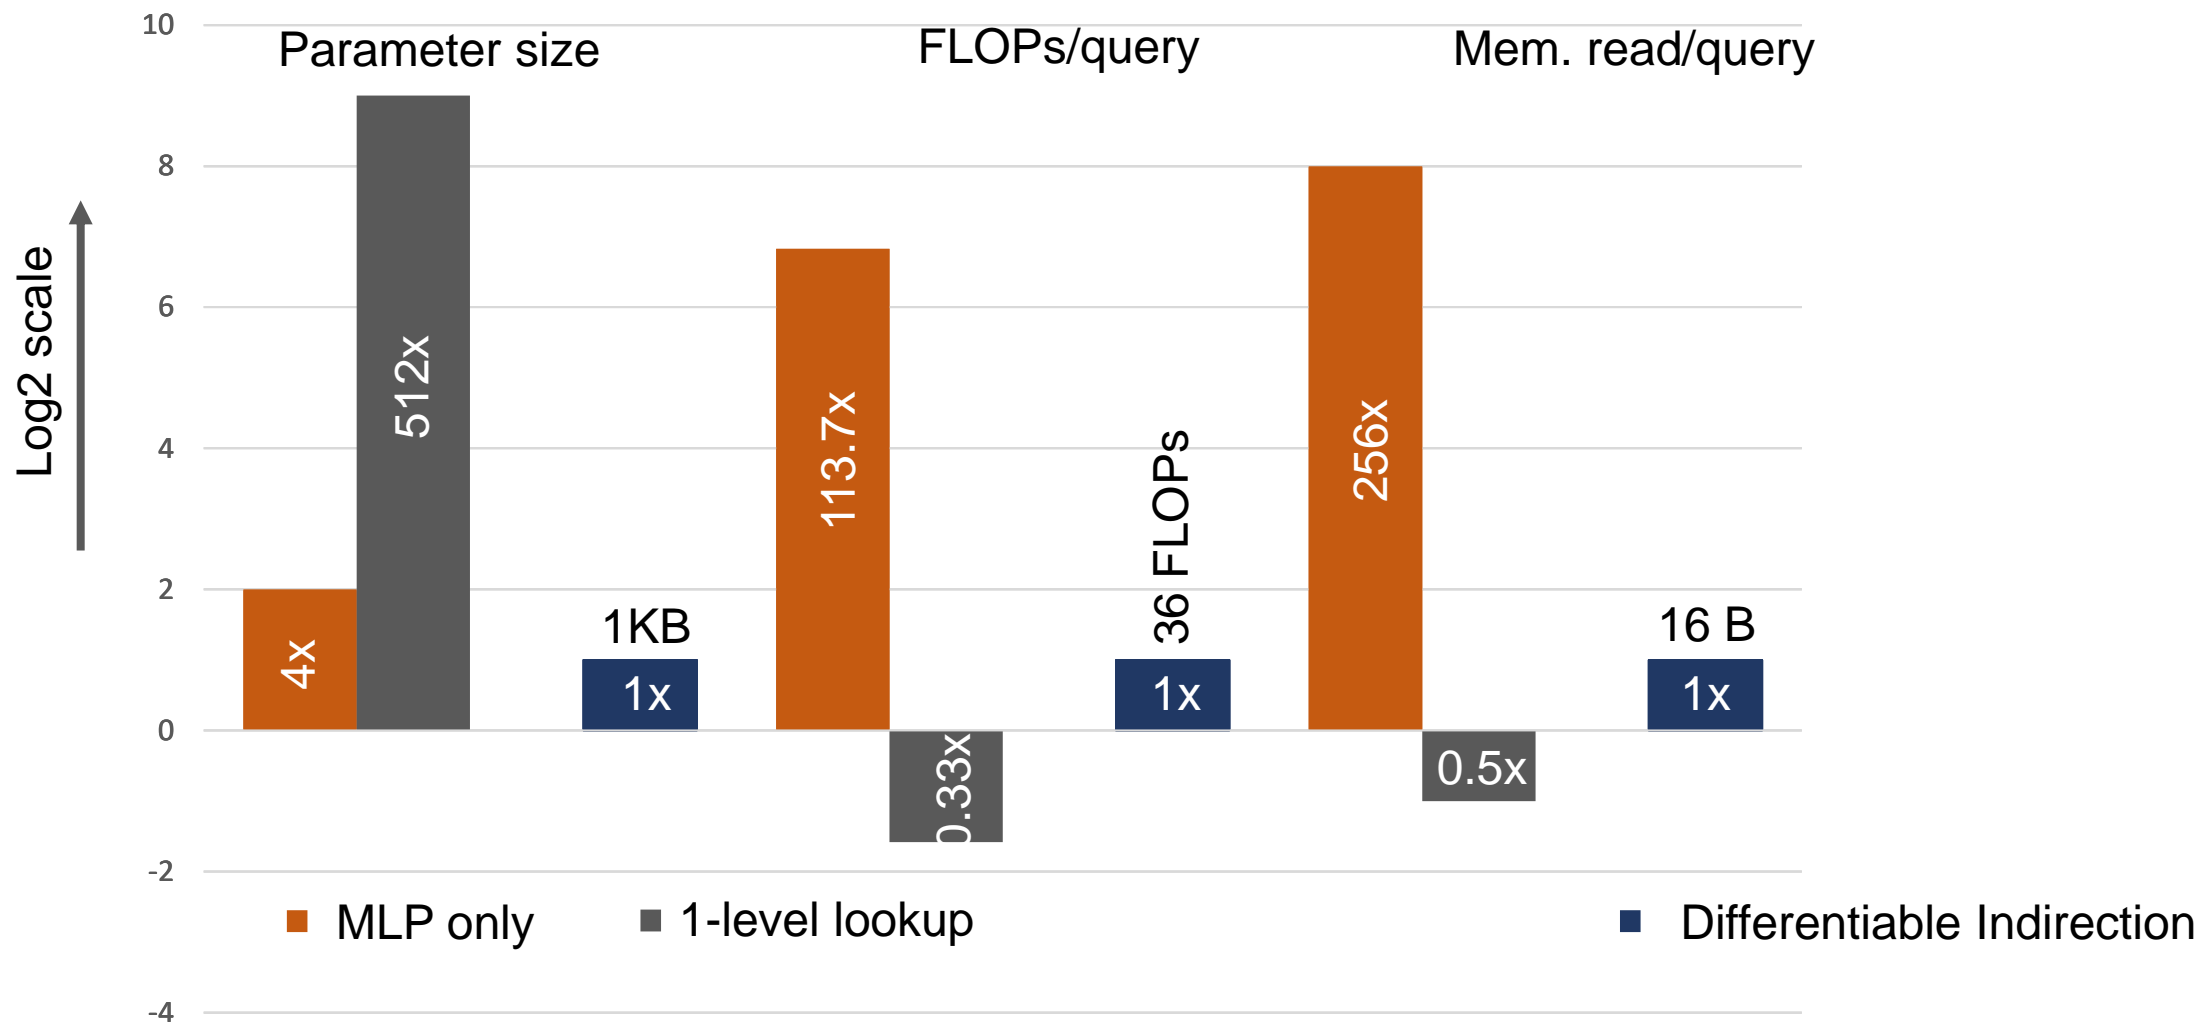
# Efficiency




# Efficiency



# Efficiency




# Efficiency



# Efficiency : more examples

# Efficiency : more examples

SDF



12 GB (Surface samples)



24 MB (IoU: 0.982)

# Efficiency : more examples

SDF

Texture  
Compression



12 GB (Surface samples)



24 MB (IoU: 0.982)

6x-48x



# Efficiency : more examples

SDF



12 GB (Surface samples)




24 MB (IoU: 0.982)

Texture  
Compression



6x-48x

Compression  
and Filtering



Built-in  
sampler

# Efficiency : more examples

SDF



12 GB (Surface samples)




24 MB (IoU: 0.982)

Texture  
Compression



6x-48x

Compression  
and Filtering



Built-in  
sampler

Disney  
Shading



240+ FLOPs



5 lookups + 41 FLOPs

# Efficiency : more examples

SDF



12 GB (Surface samples)




24 MB (IoU: 0.982)

Texture  
Compression



6x-48x

Compression  
and Filtering



Built-in  
sampler

Disney  
Shading



240+ FLOPs



5 lookups + 41 FLOPs

Radiance Field  
Compression



860 MB (PSNR: 33)



89 MB (PSNR: 31)




# Results

Texture  
Compression



SDF



Shading and  
Filtering



Radiance Field  
Compression



## Conclusion

- Efficient – minimal parameter size, FLOPs, and bandwidth.
- Flexible – several applications in graphics pipeline and potentially beyond.
- Adaptive – supports adaptive spatial resolution.





**SIGGRAPH  
ASIA 2023  
SYDNEY**

# Efficient Graphics Representation with Differentiable Indirection

Scan for details:



Thank you.

