Ray-aligned Occupancy Map Array (ROMA) for Fast Approximate Ray Tracing

Zheng Zeng[1], Zilin Xu[2], Lu Wang[2], Lifan Wu[3], Ling-Qi Yan[1]

[1]University of California, Santa Barbara

[2]Shandong University

[3]NVIDIA

Background

Real-time ray tracing

Ray traced results

How to make it real-time?

Hardware Ray Tracing (HWRT) is good, but …

Ray traced results

Make it "real-time"

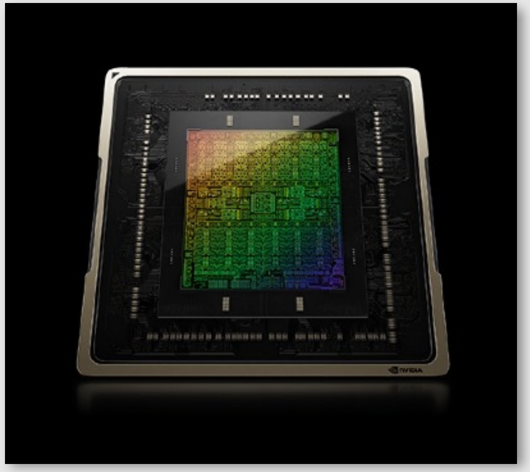

Ray tracing hardware
[NVIDIA Ada Lovelace]

Hardware Ray Tracing (HWRT)

Only high-end platforms support it!

Hardware Ray Tracing (HWRT) is good, but …

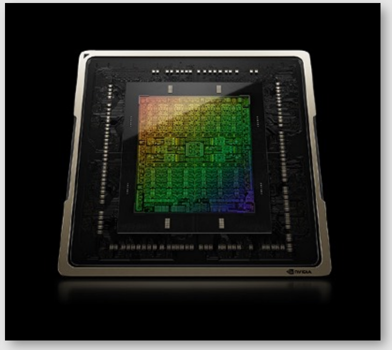

Hardware Ray Tracing (HWRT)

Approximate

Software Ray Tracing (SWRT)

SWRT: Screen-space Ray Tracing

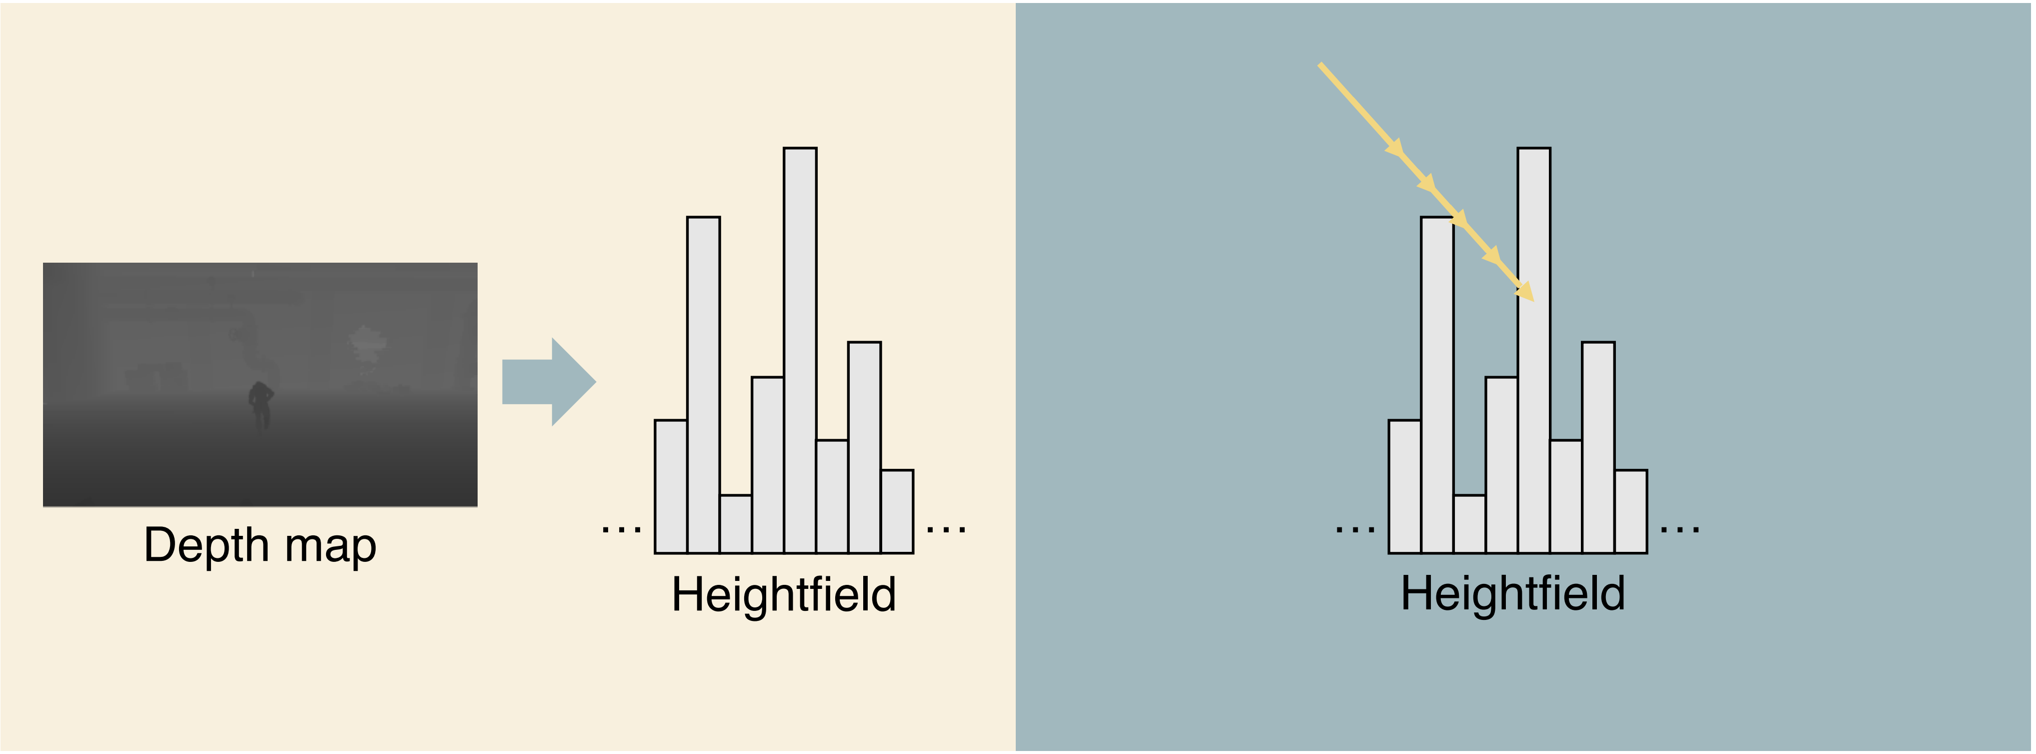

Generation

Tracing

SWRT: Screen-space Ray Tracing

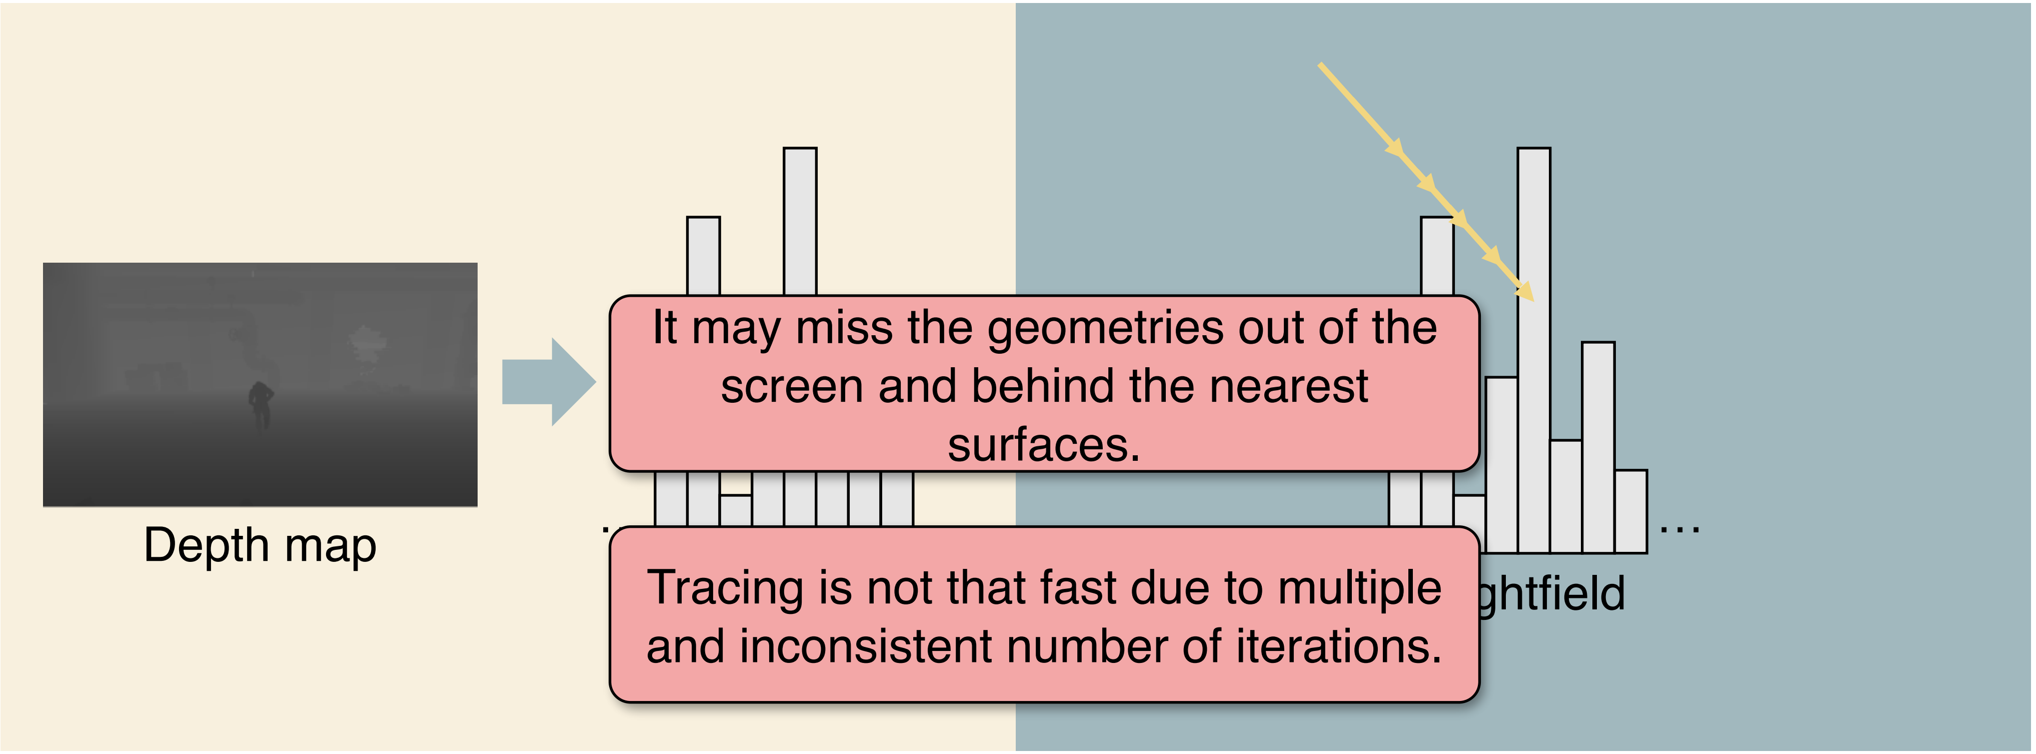

Generation

Tracing

SWRT: Distance Fields (DFs)

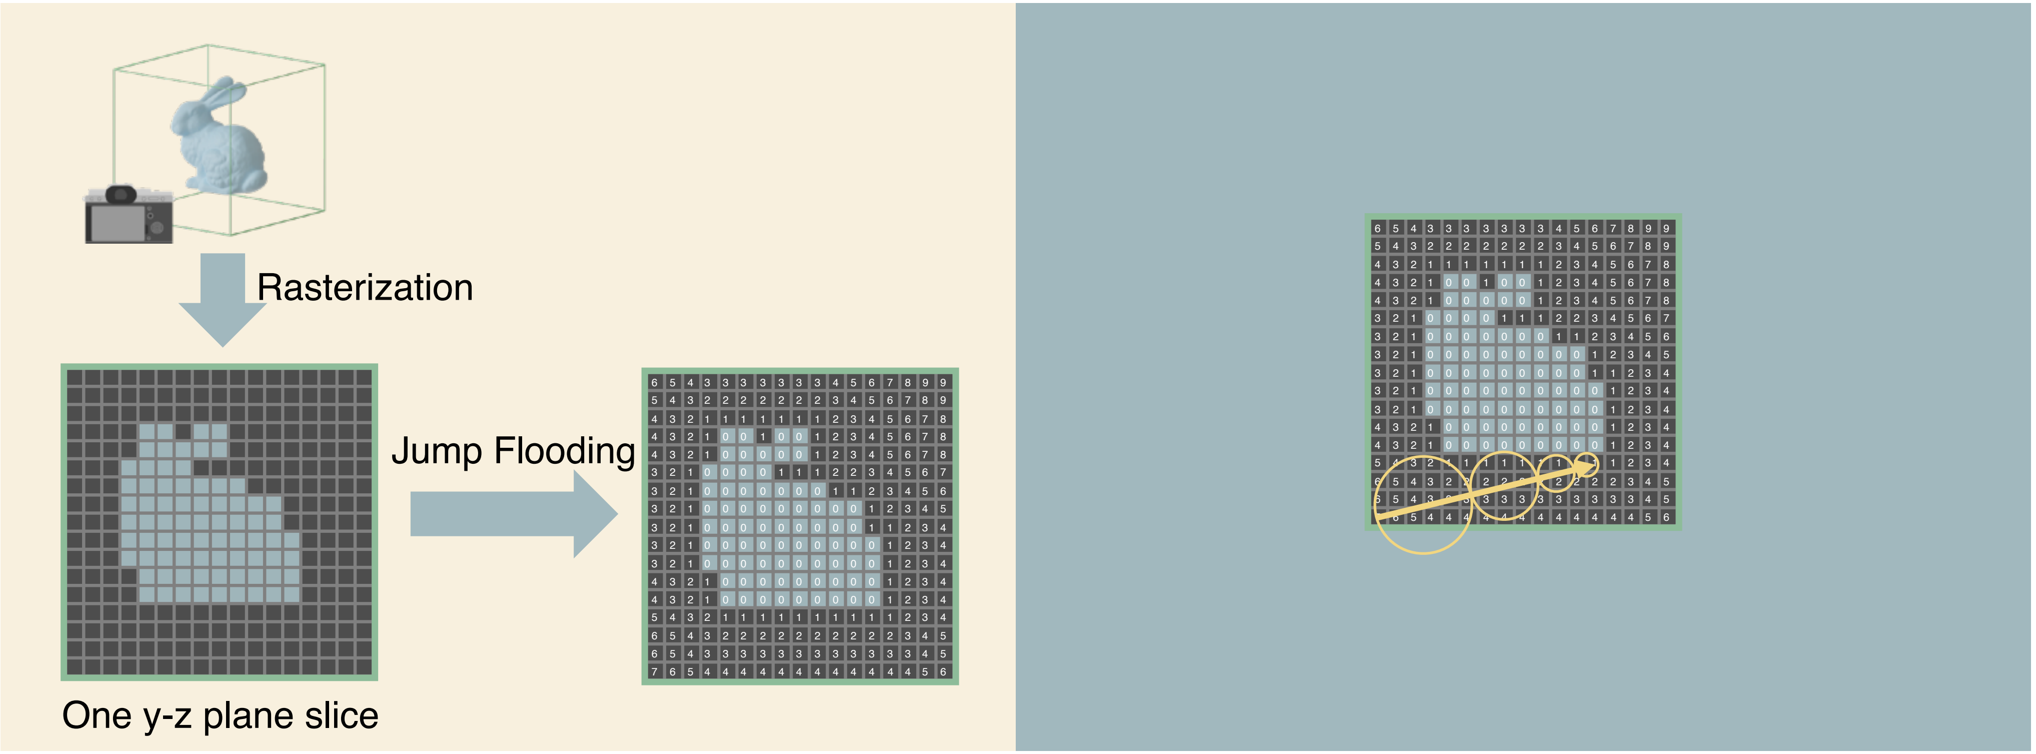

Generation

Tracing

SWRT: Distance Fields (DFs)

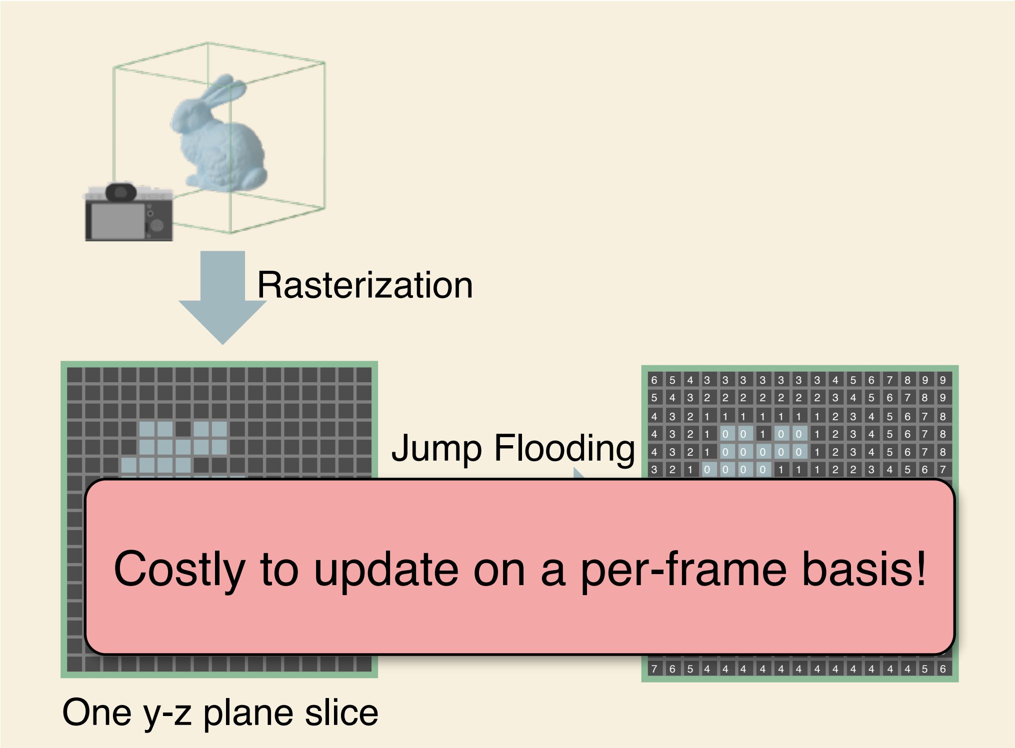

Generation

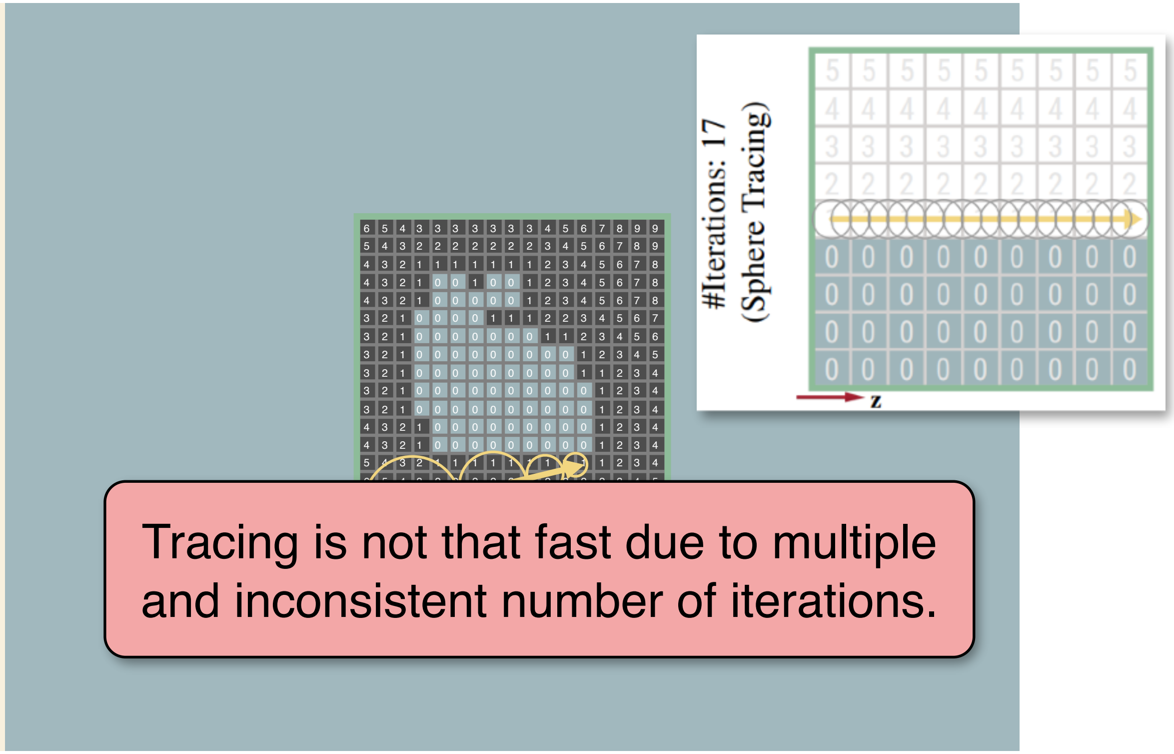

Tracing

SWRT: Occupancy Maps (OMs)

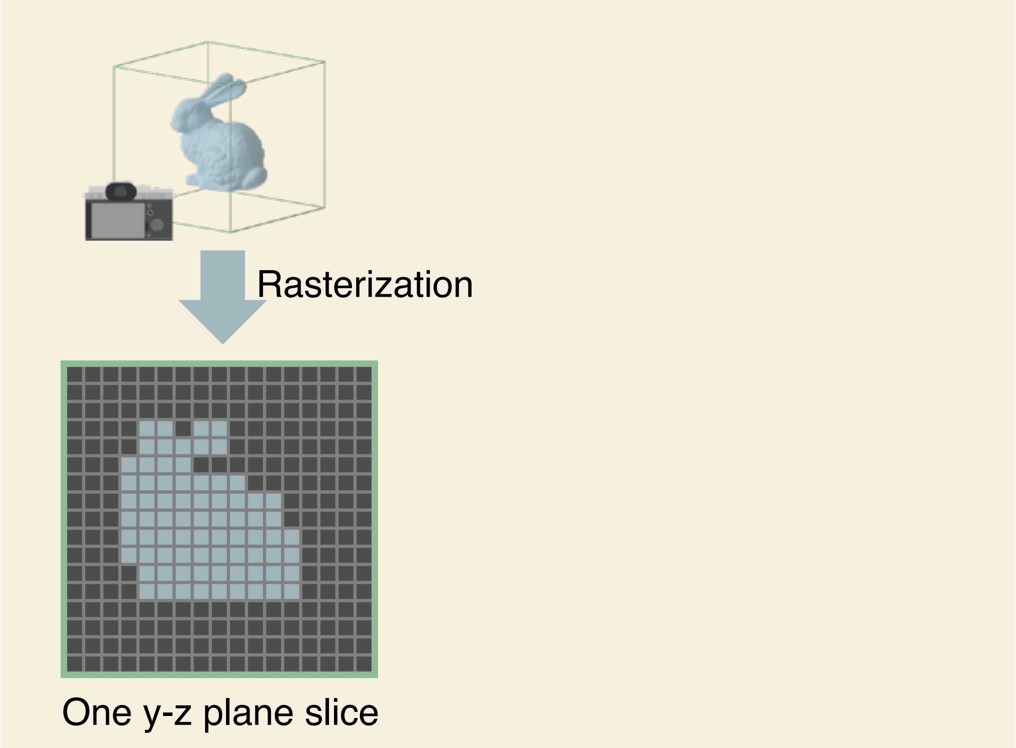

Generation

SWRT: Occupancy Maps (OMs)

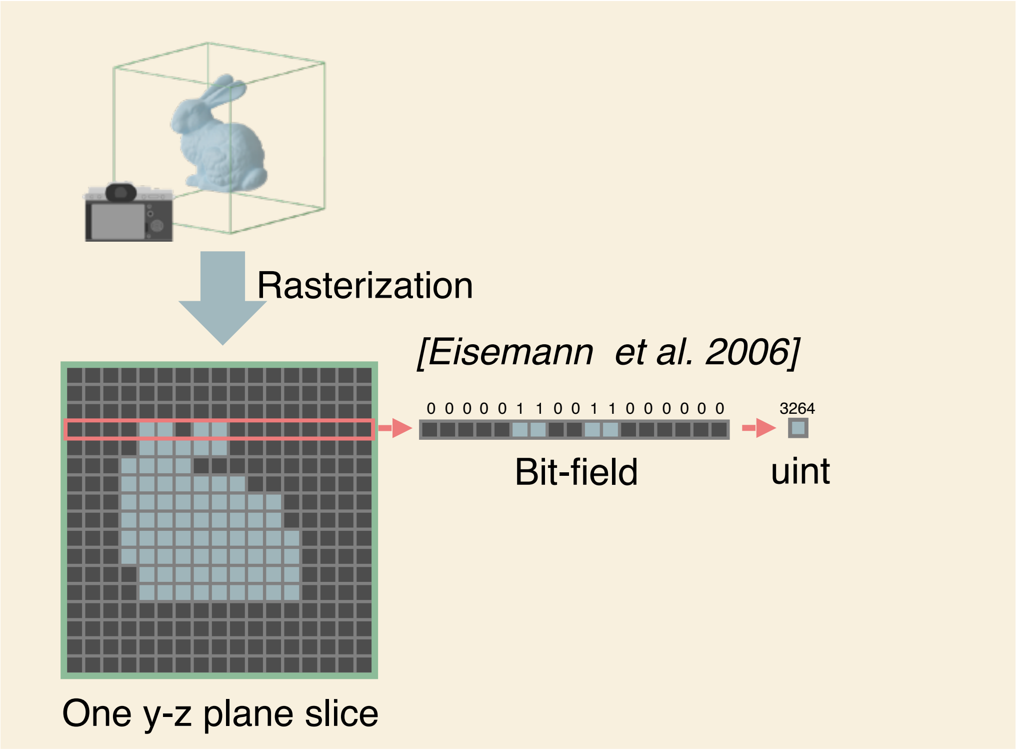

Generation

SWRT: Occupancy Maps (OMs)

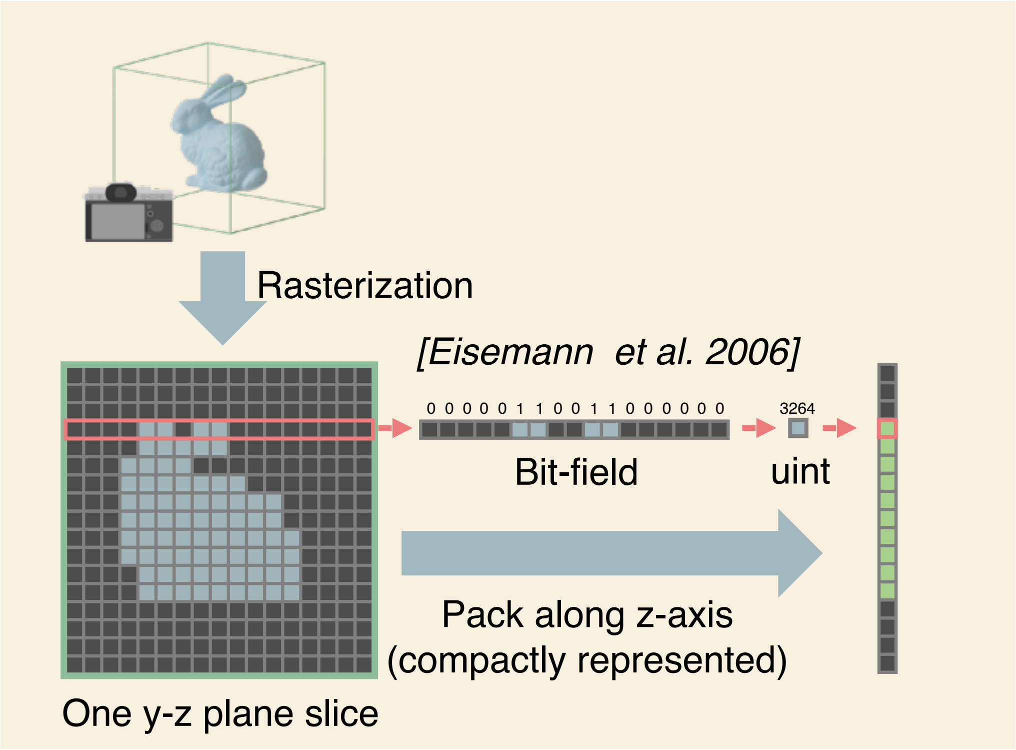

Generation

SWRT: Occupancy Maps (OMs)

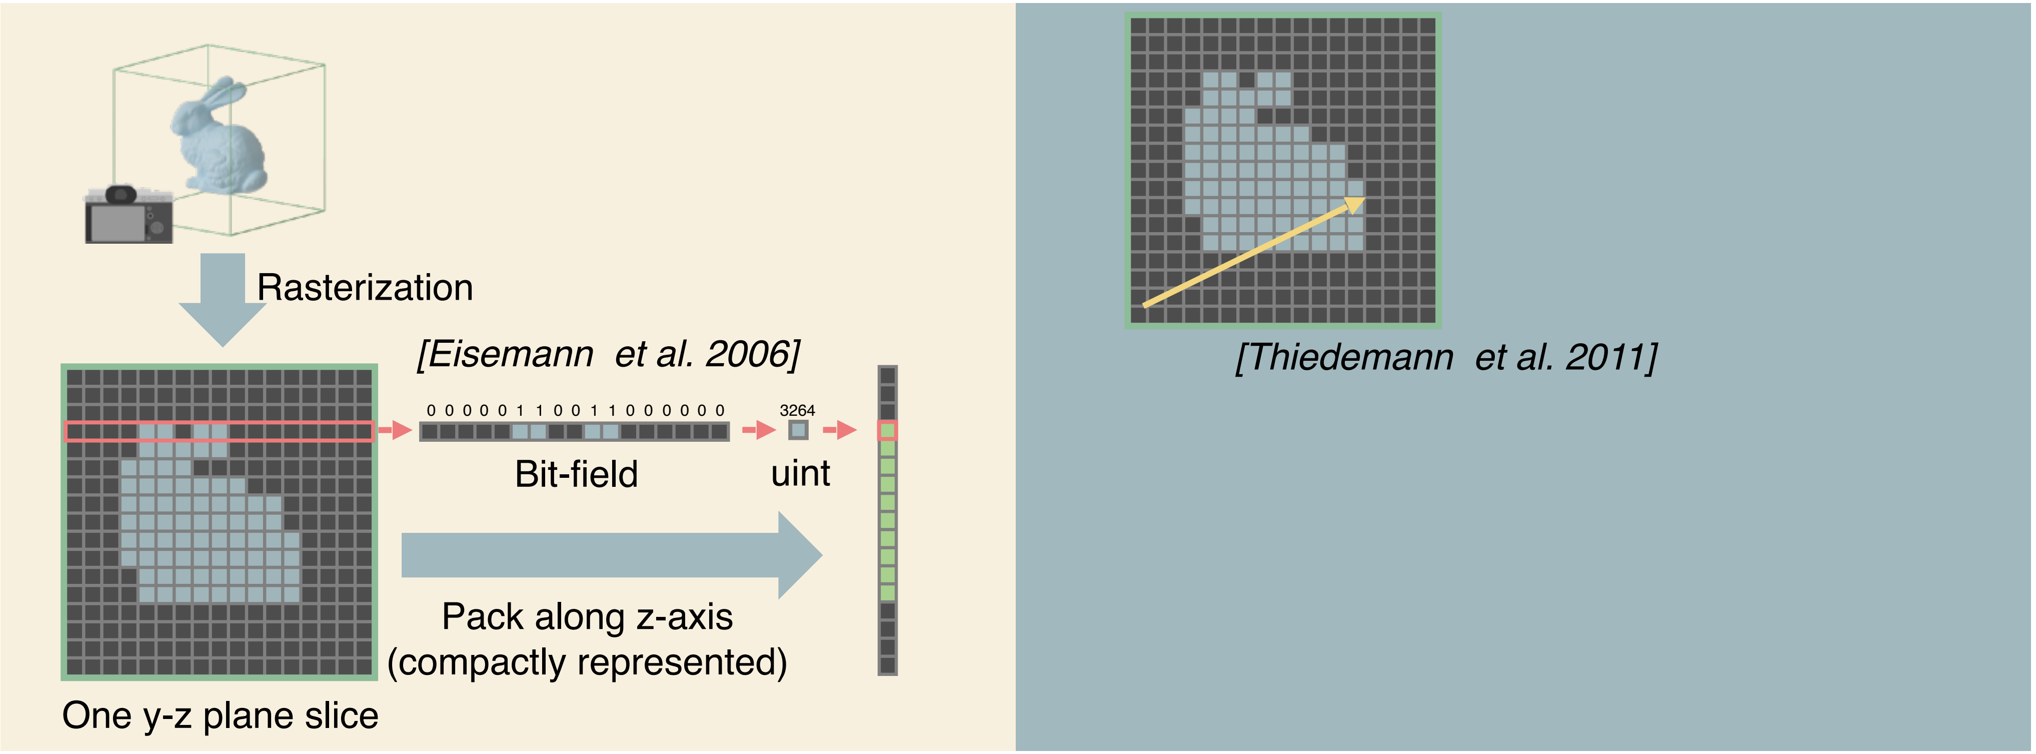

Generation Tracing

SWRT: Occupancy Maps (OMs)

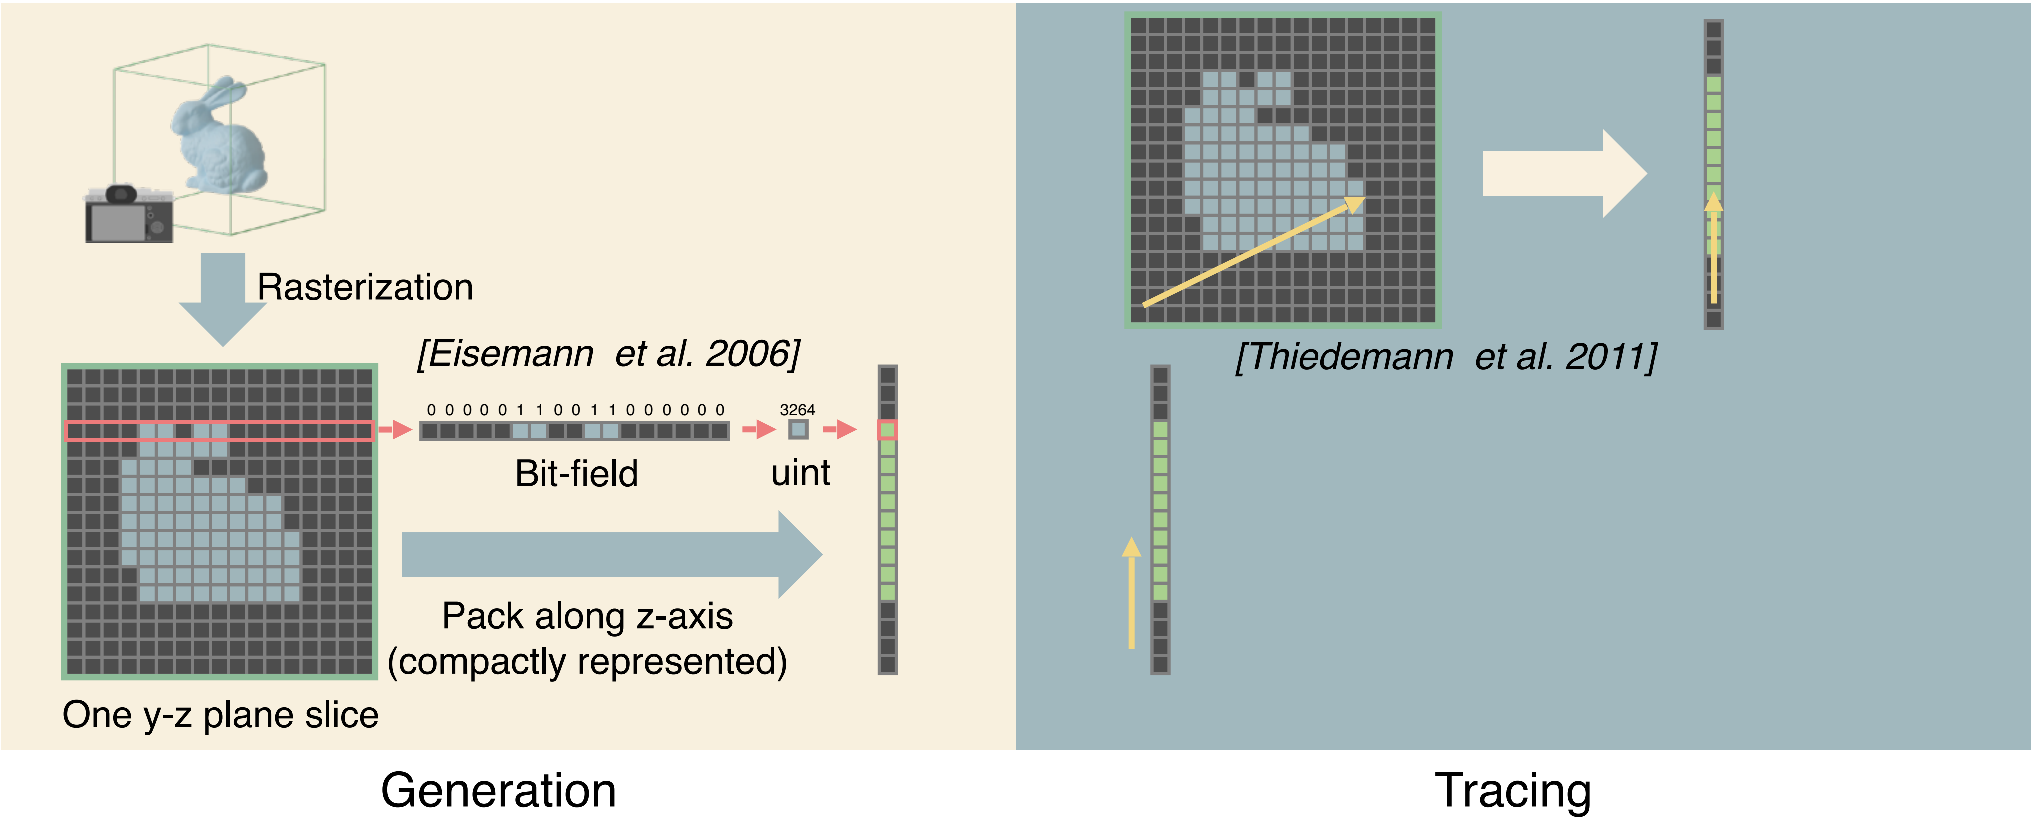

SWRT: Occupancy Maps (OMs)

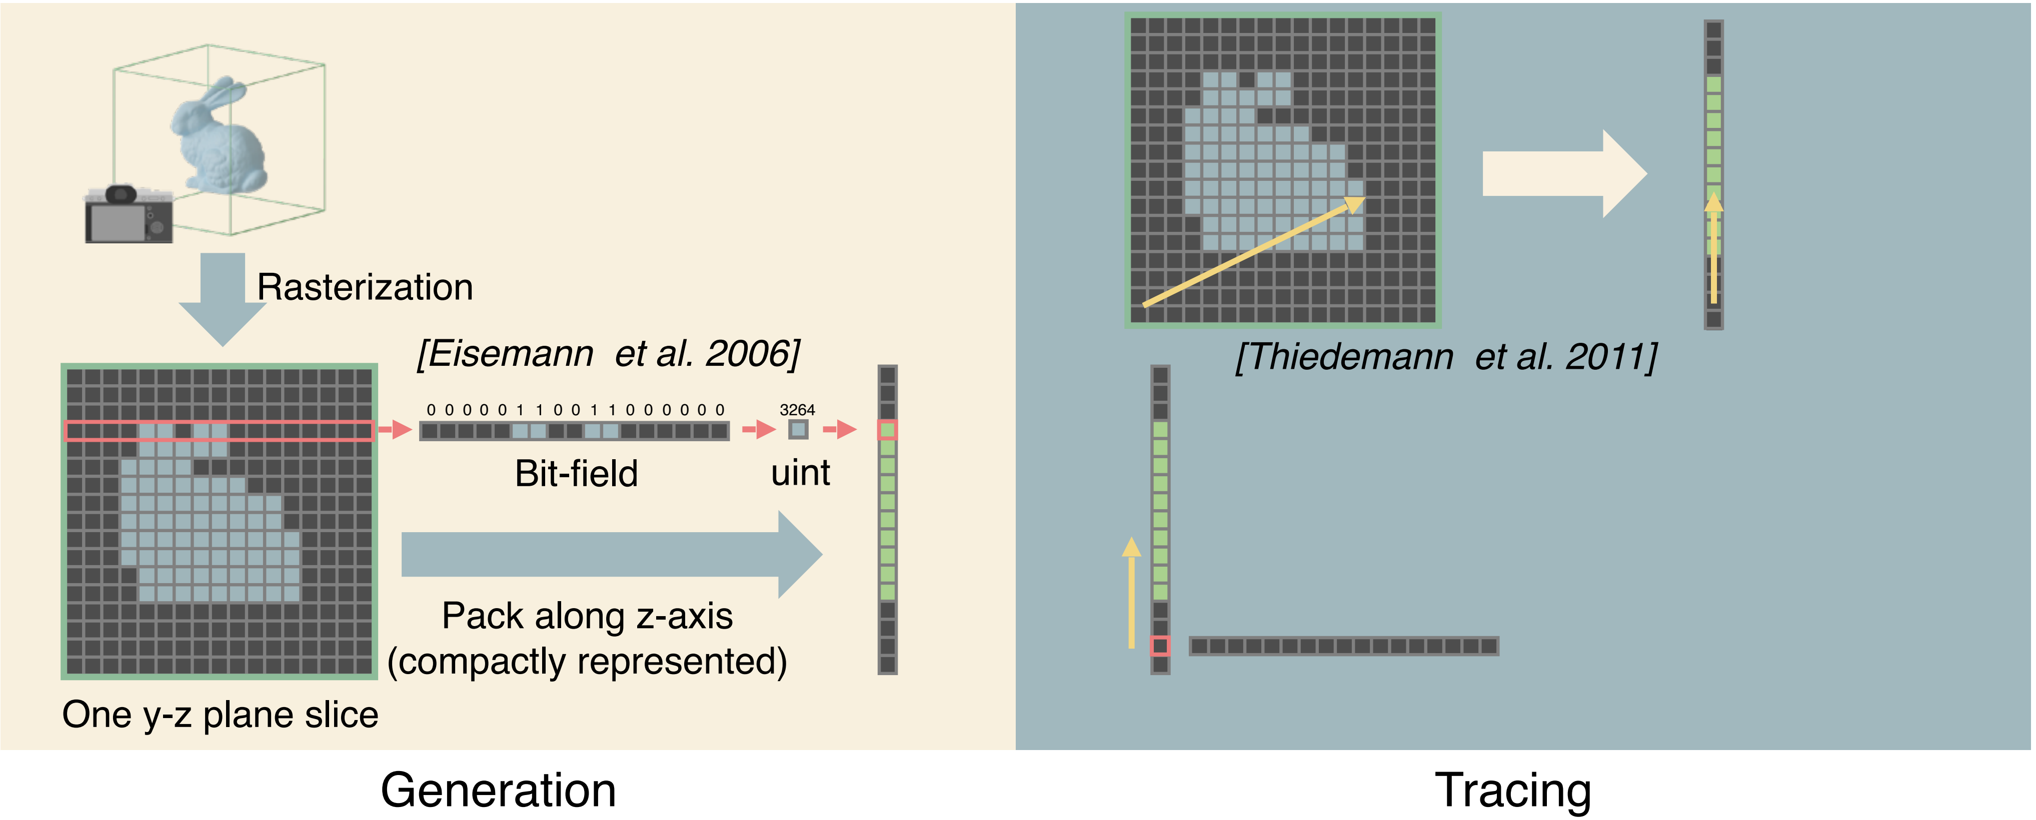

SWRT: Occupancy Maps (OMs)

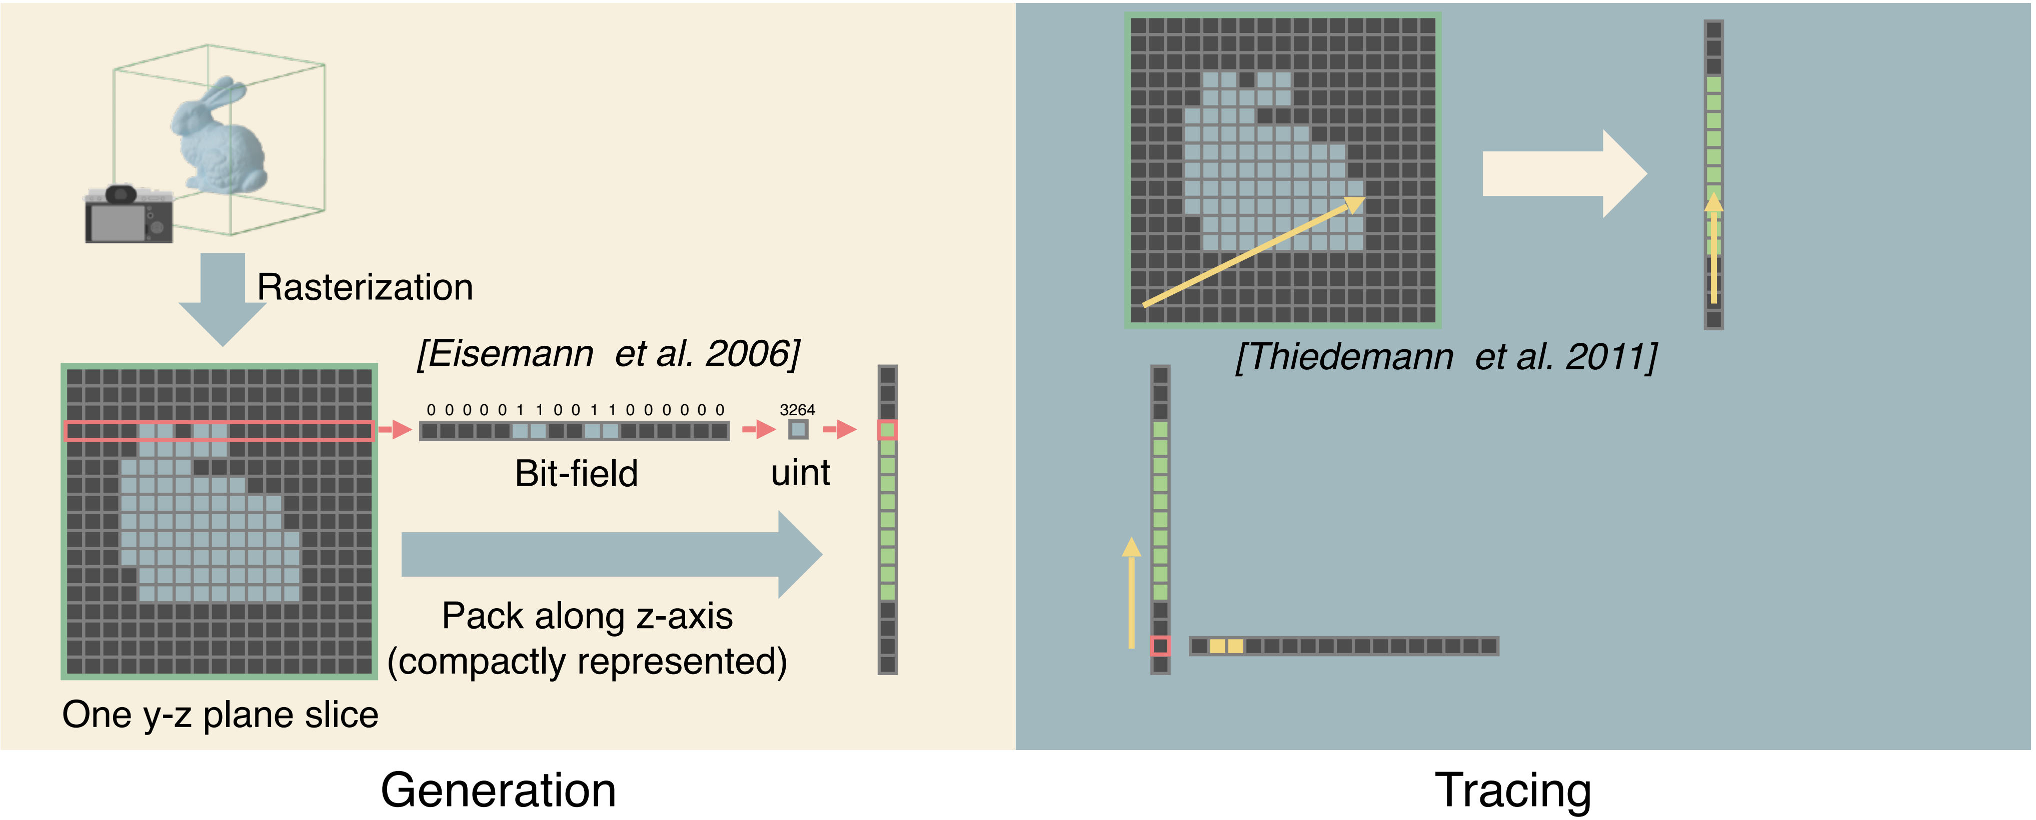

SWRT: Occupancy Maps (OMs)

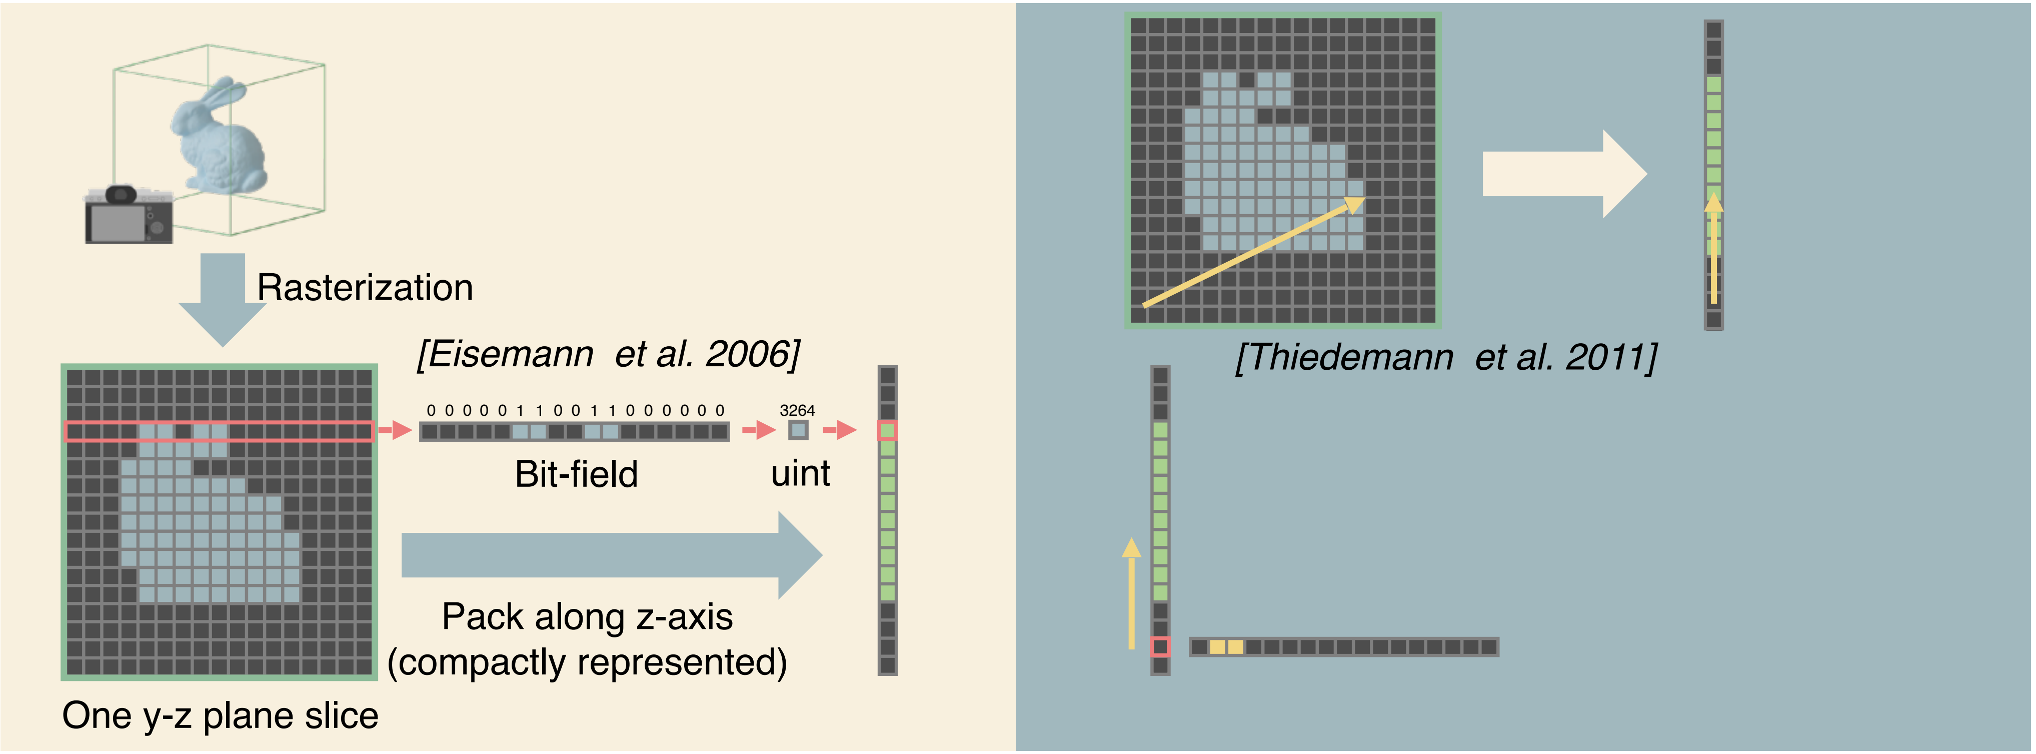

Generation

Tracing

SWRT: Occupancy Maps (OMs)

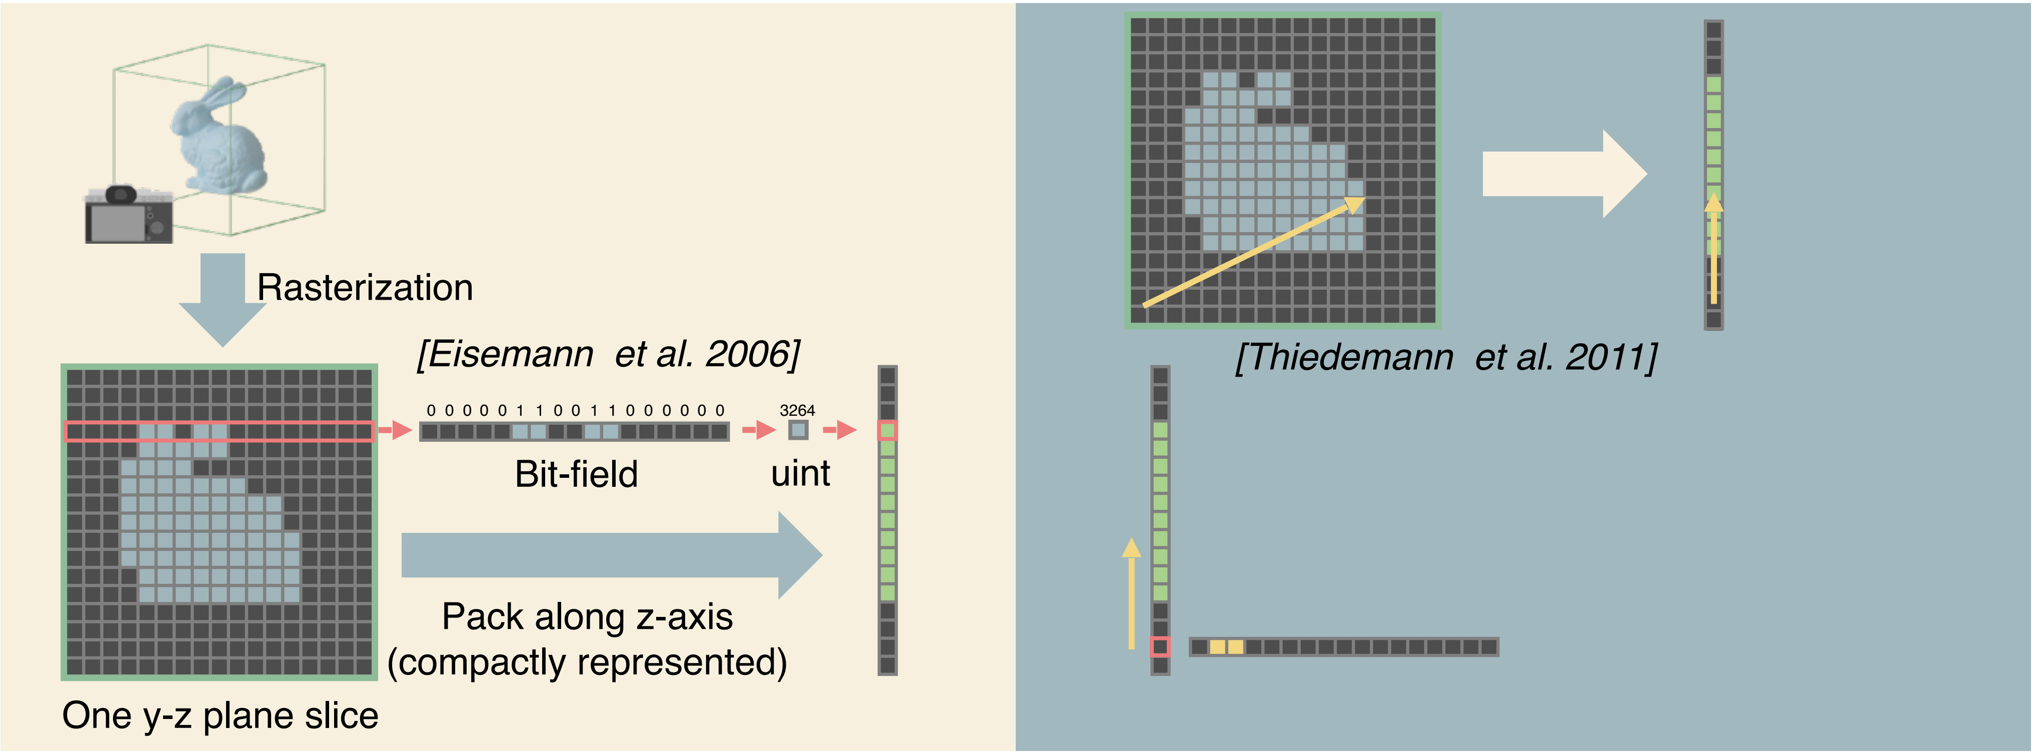

Generation

Tracing

SWRT: Occupancy Maps (OMs)

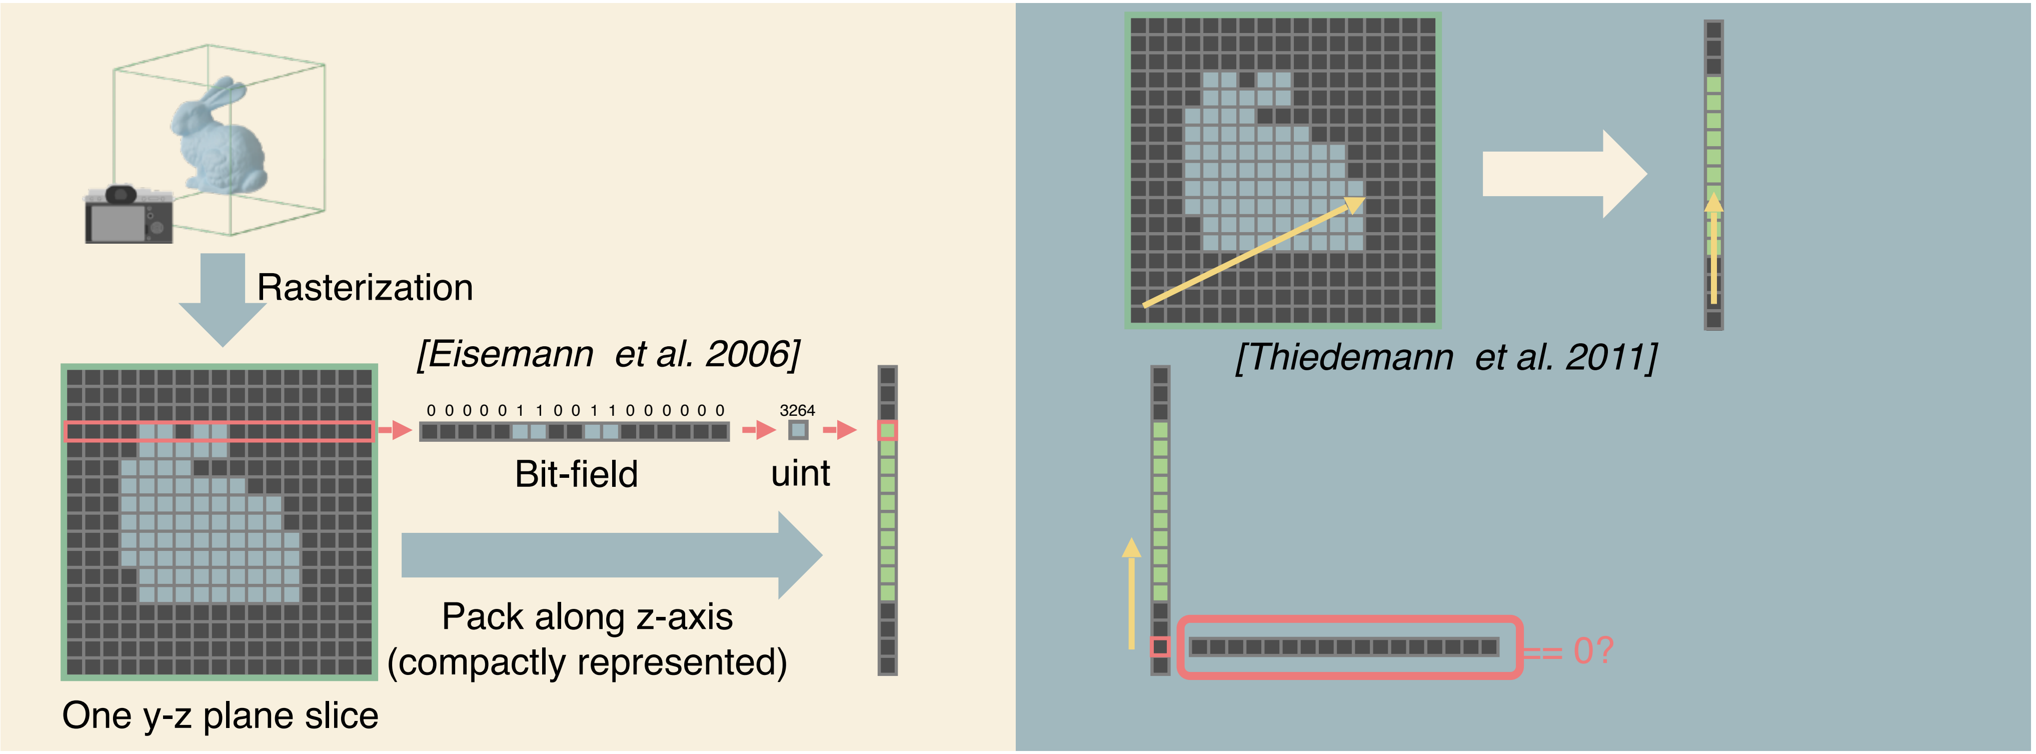

Generation

Tracing

SWRT: Occupancy Maps (OMs)

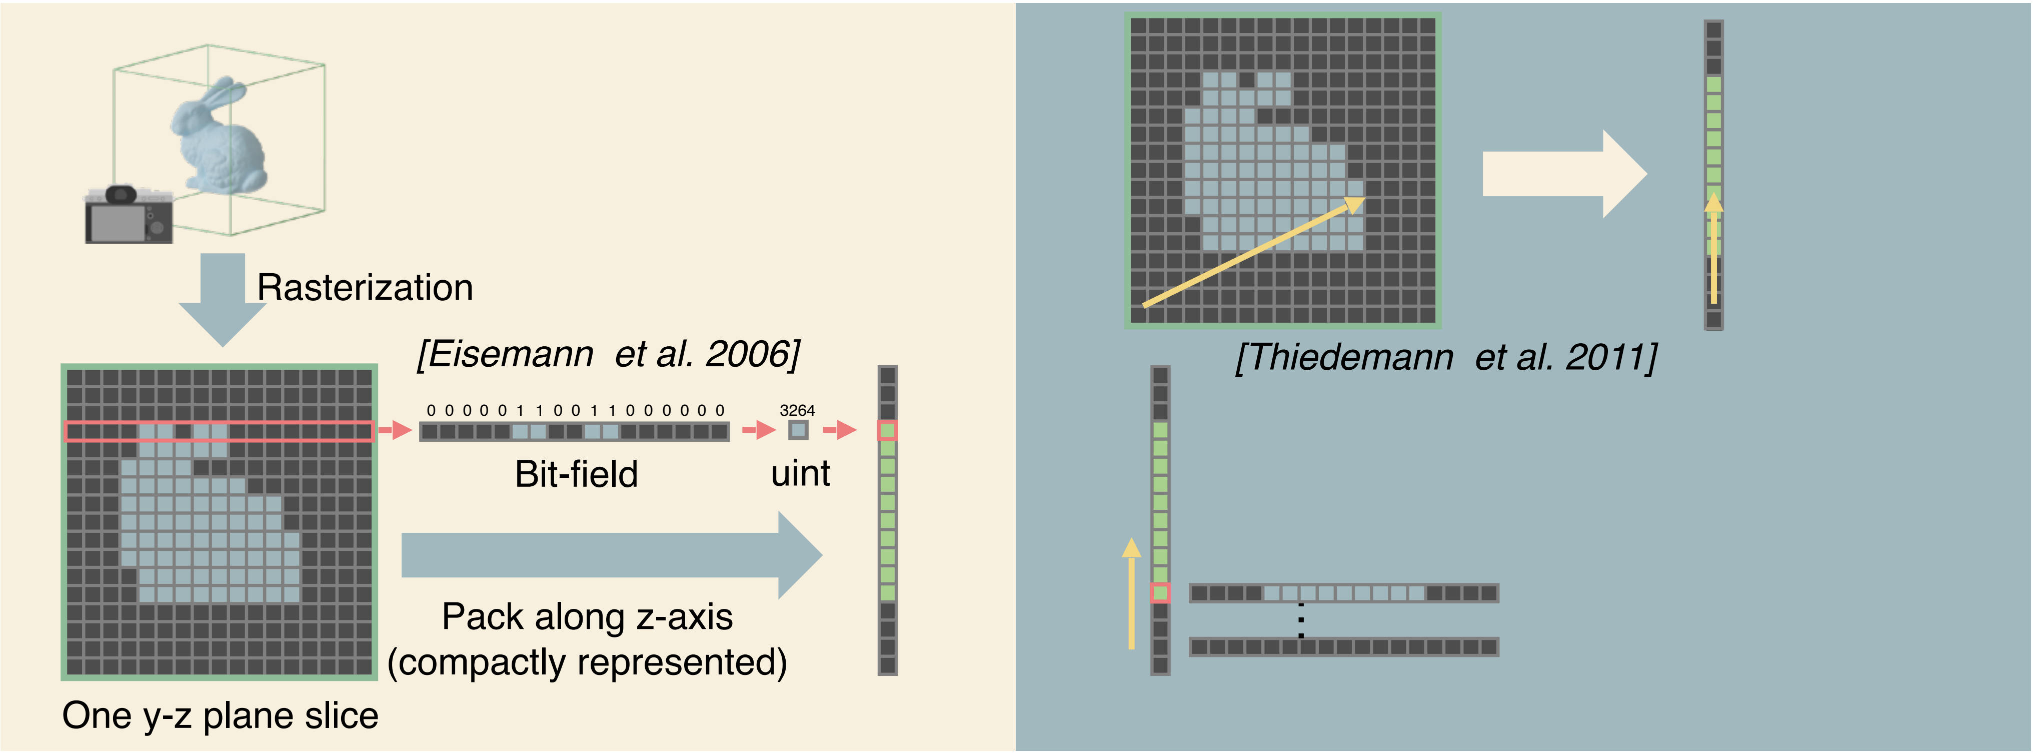

Generation

Tracing

SWRT: Occupancy Maps (OMs)

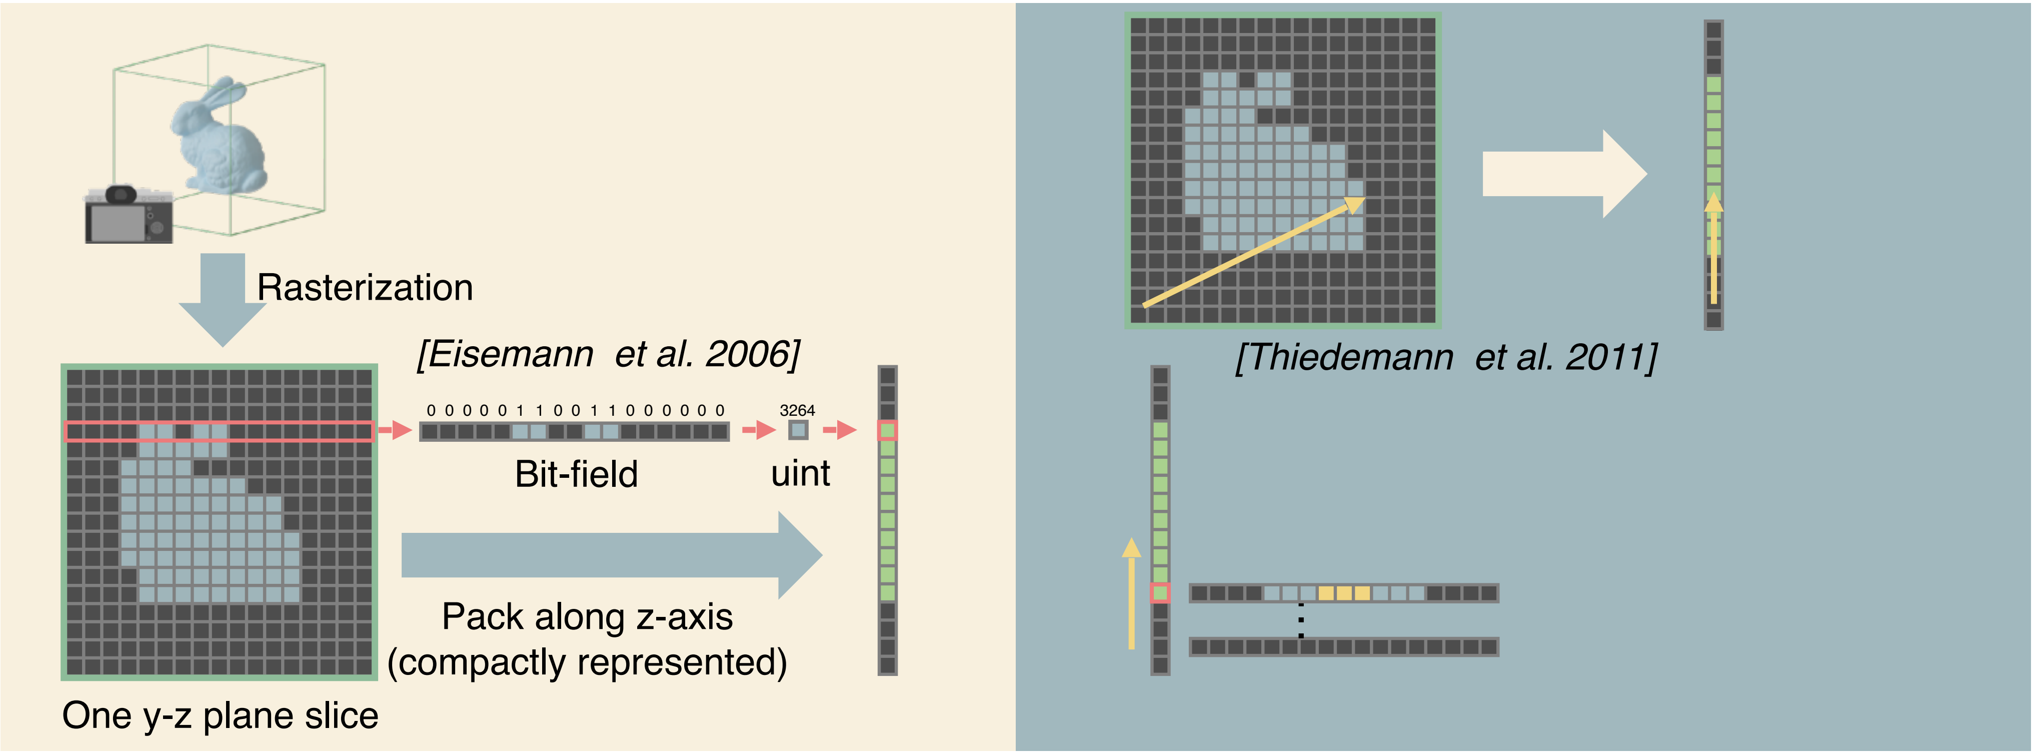

Generation

Tracing

SWRT: Occupancy Maps (OMs)

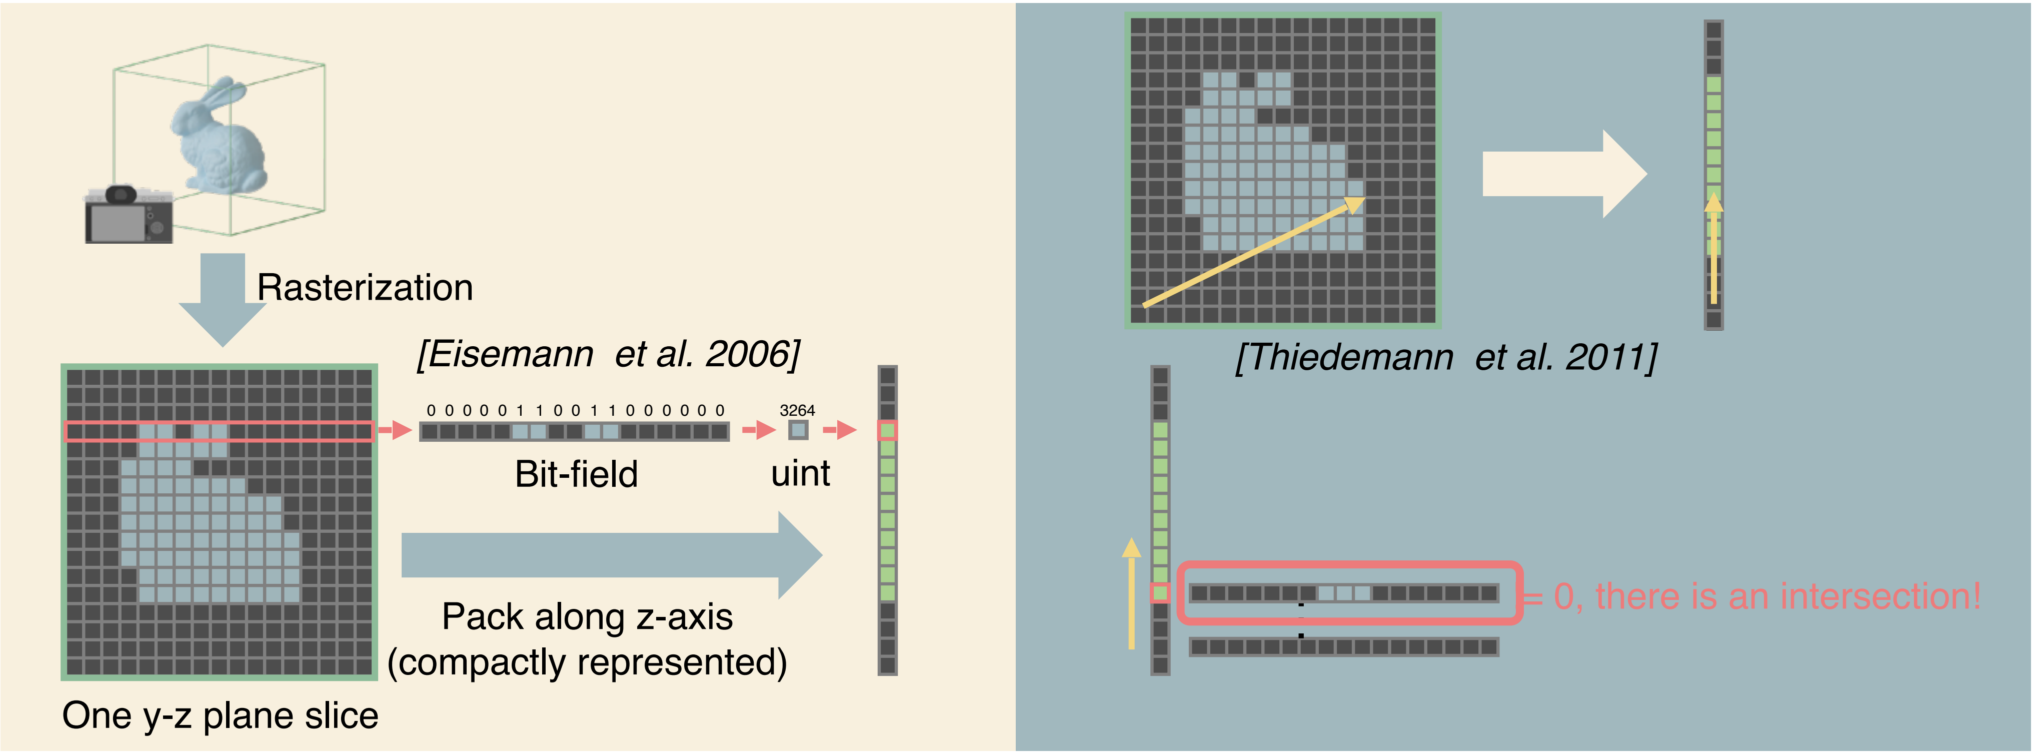

Generation

Tracing

Our goal

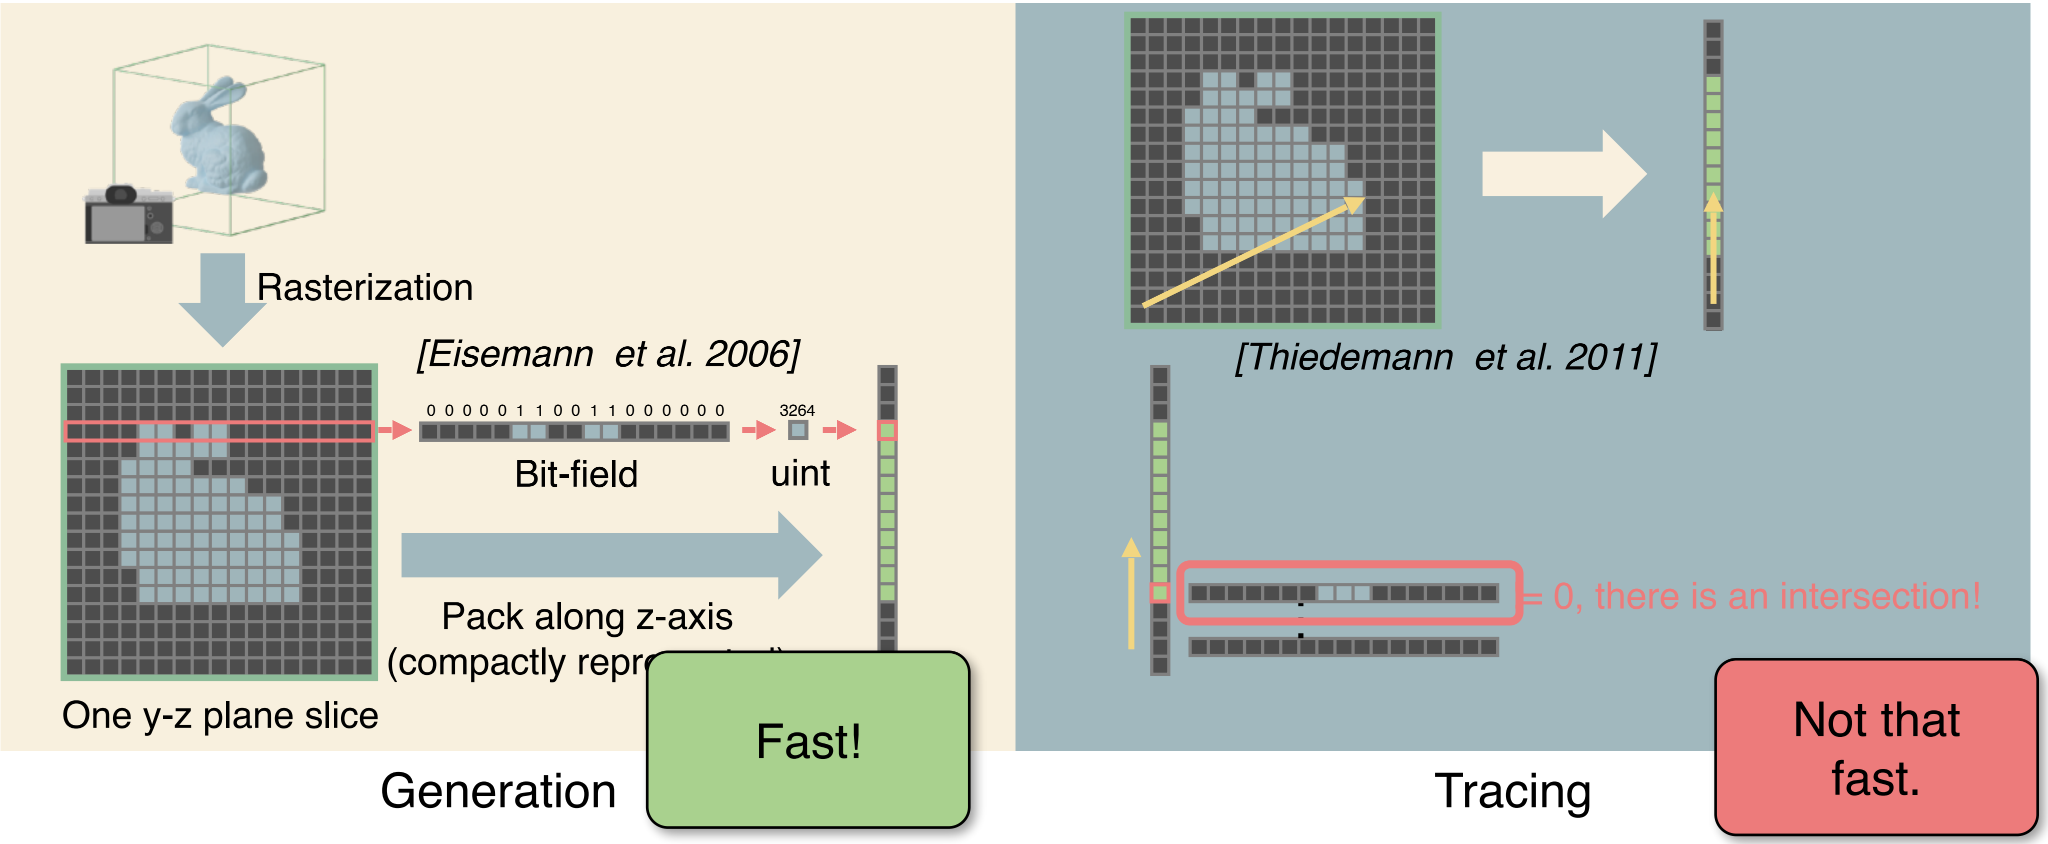
Rasterization
[Eisemann et al. 2006]
0 0 0 0 0 1 1 0 0 1 1 0 0 0 0 0 0
3264
Bit-field
uint
Pack along z-axis
(compactly repr
One y-z plane slice
Generation
Fast!
[Thiedemann et al. 2011]
0, there is an intersection!
Tracing
Not that fast.

Our goal

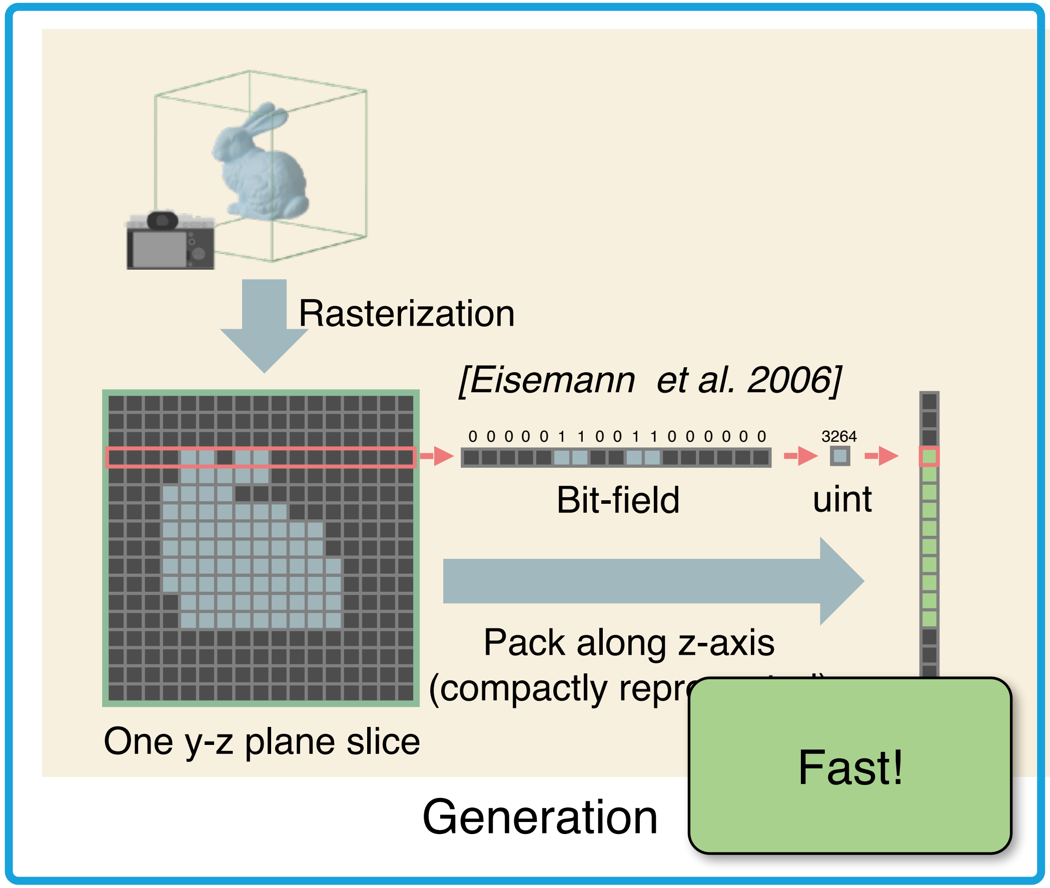

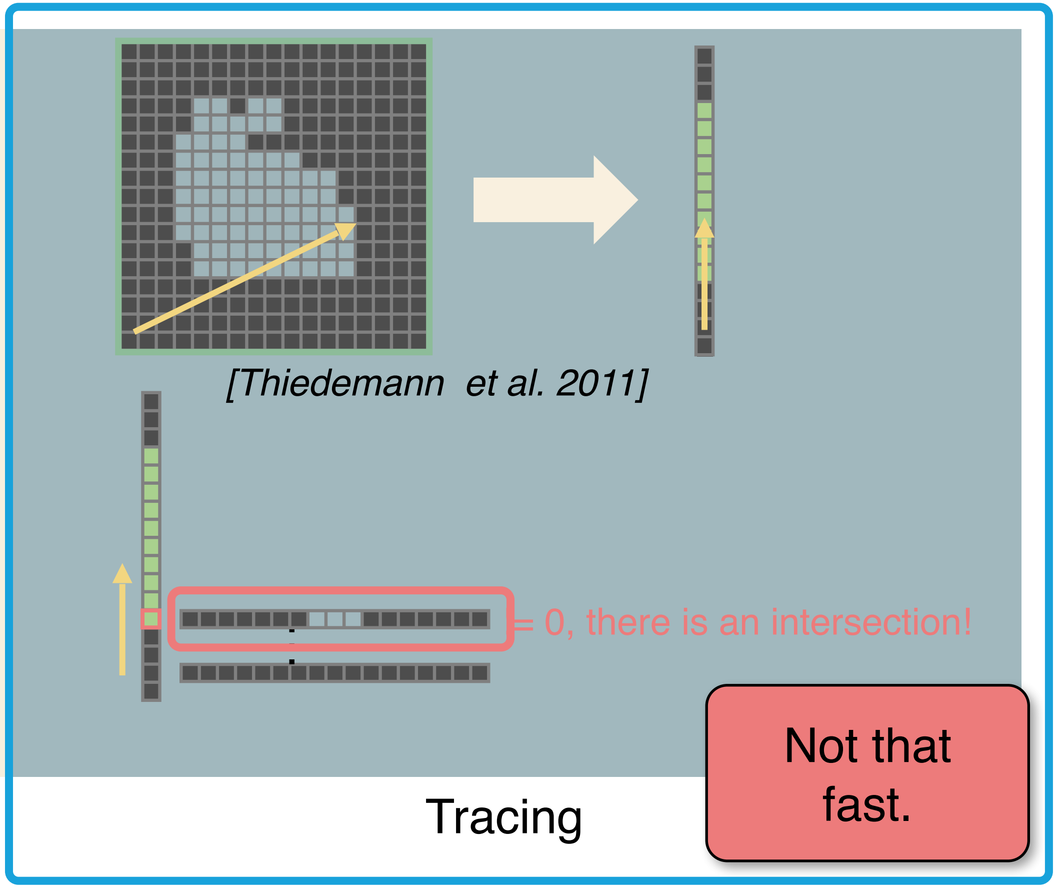

Better utilize the trait of fast generation

Further speed up ray tracing against OMs

Method

Key observation

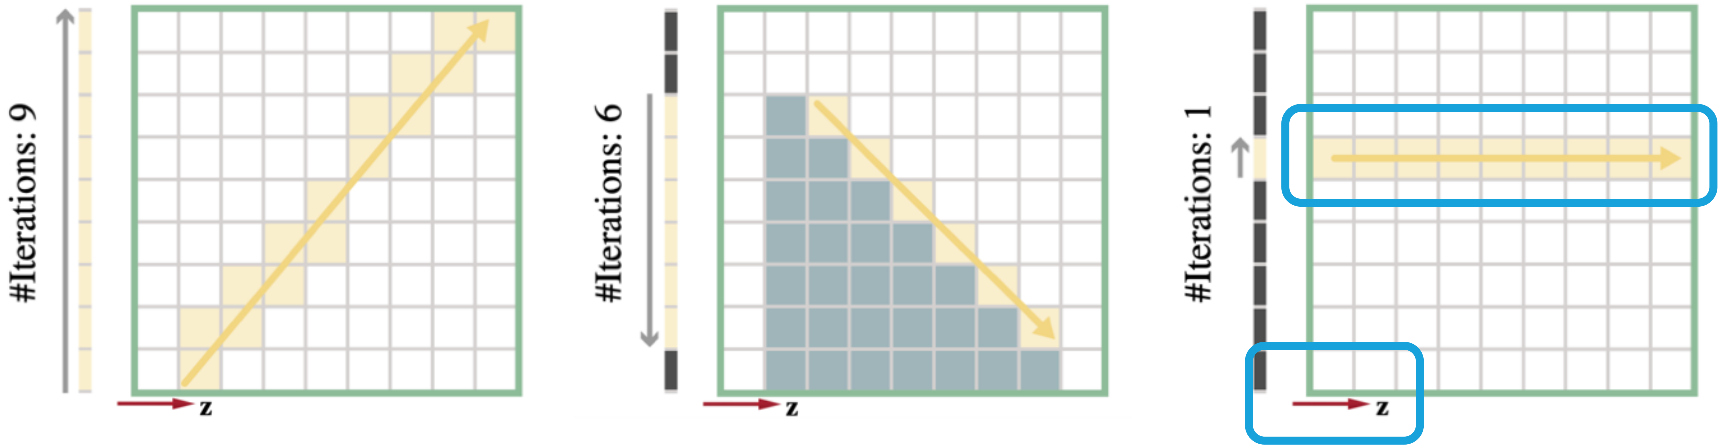

When ray is tracing along the z-axis, only **a single lookup** is required!

Key observation

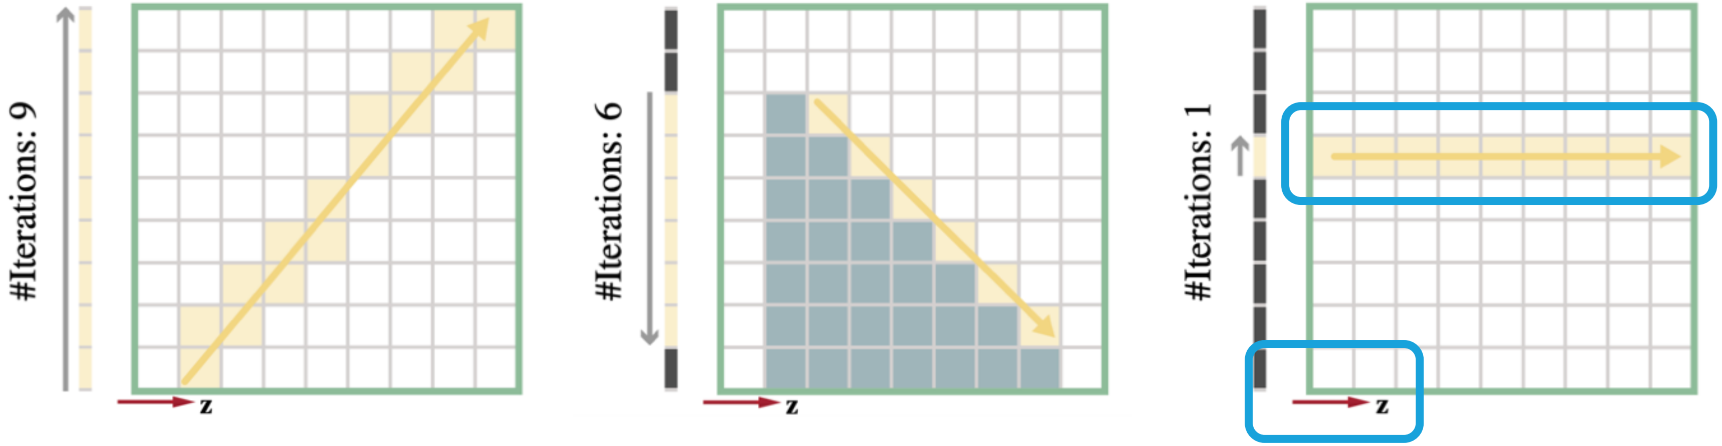

When ray is tracing along the z-axis, only **a single lookup** is required!

If all our rays are tracing along the z-axis
(OMs are aligned with rays),
we will achieve O(1) tracing performance.

Key observation

If all our rays are tracing along the z-axis
(OMs are aligned with rays),
we will achieve O(1) tracing performance.

↓ An infinite number of ray-aligned OMs is required.

↓ Each ray-aligned OM invokes one pass of rasterization.

Key idea

An infinite number of ray-aligned OMs is required.

Each ray-aligned OM invokes one pass of rasterization.

Key idea

An infinite number of ray-aligned OMs is required.

Each ray-aligned OM invokes one pass of rasterization.

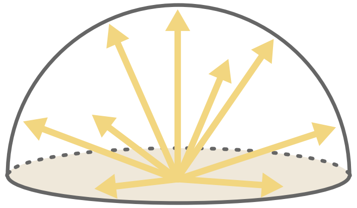

Candidate directions

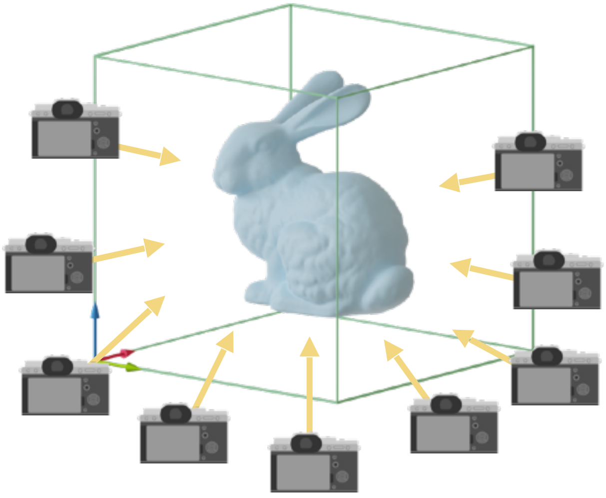

Rasterization

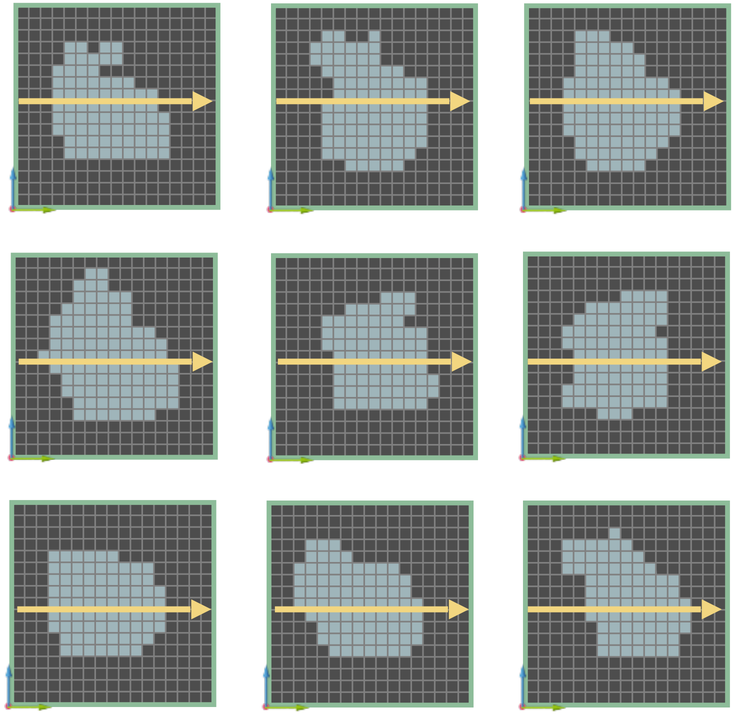

Key idea

Only a subset of ray-aligned OMs are required.

Each ray-aligned OM invokes one pass of rasterization.

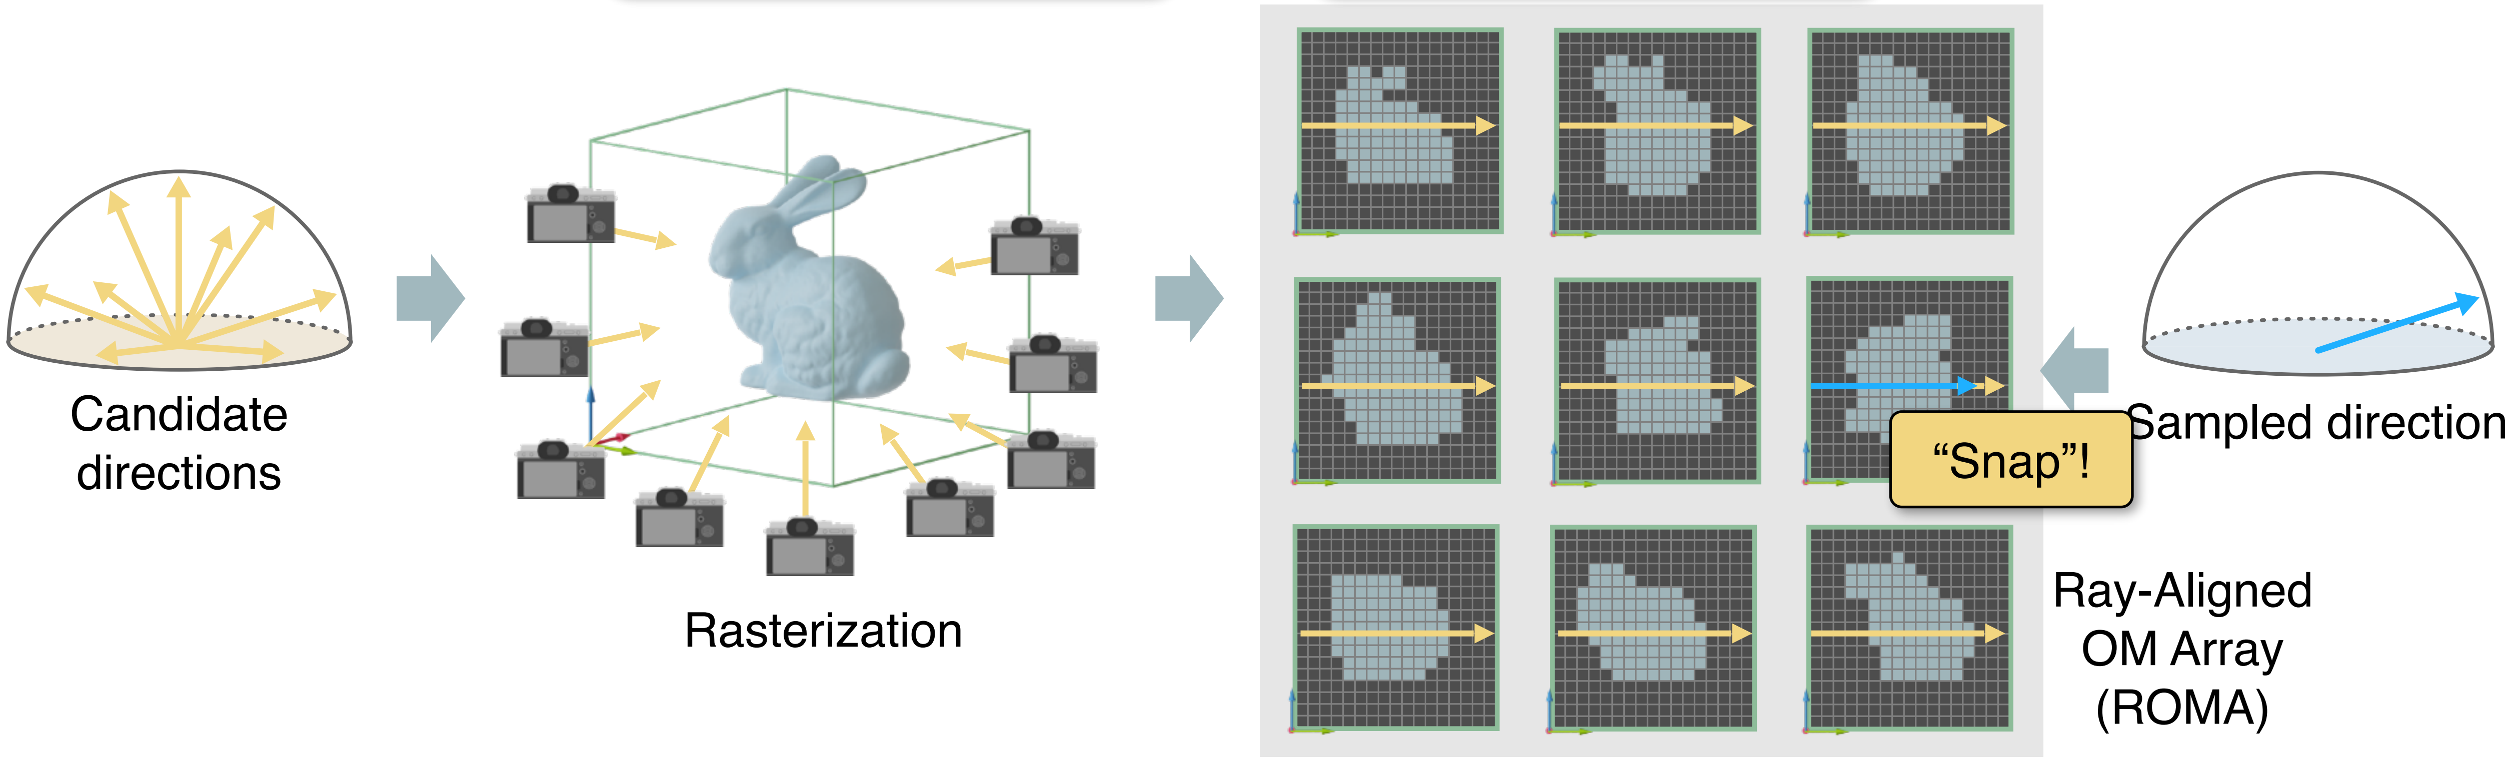

Key idea

Only a subset of ray-aligned OMs are required.

Only one pass of rasterization is required.

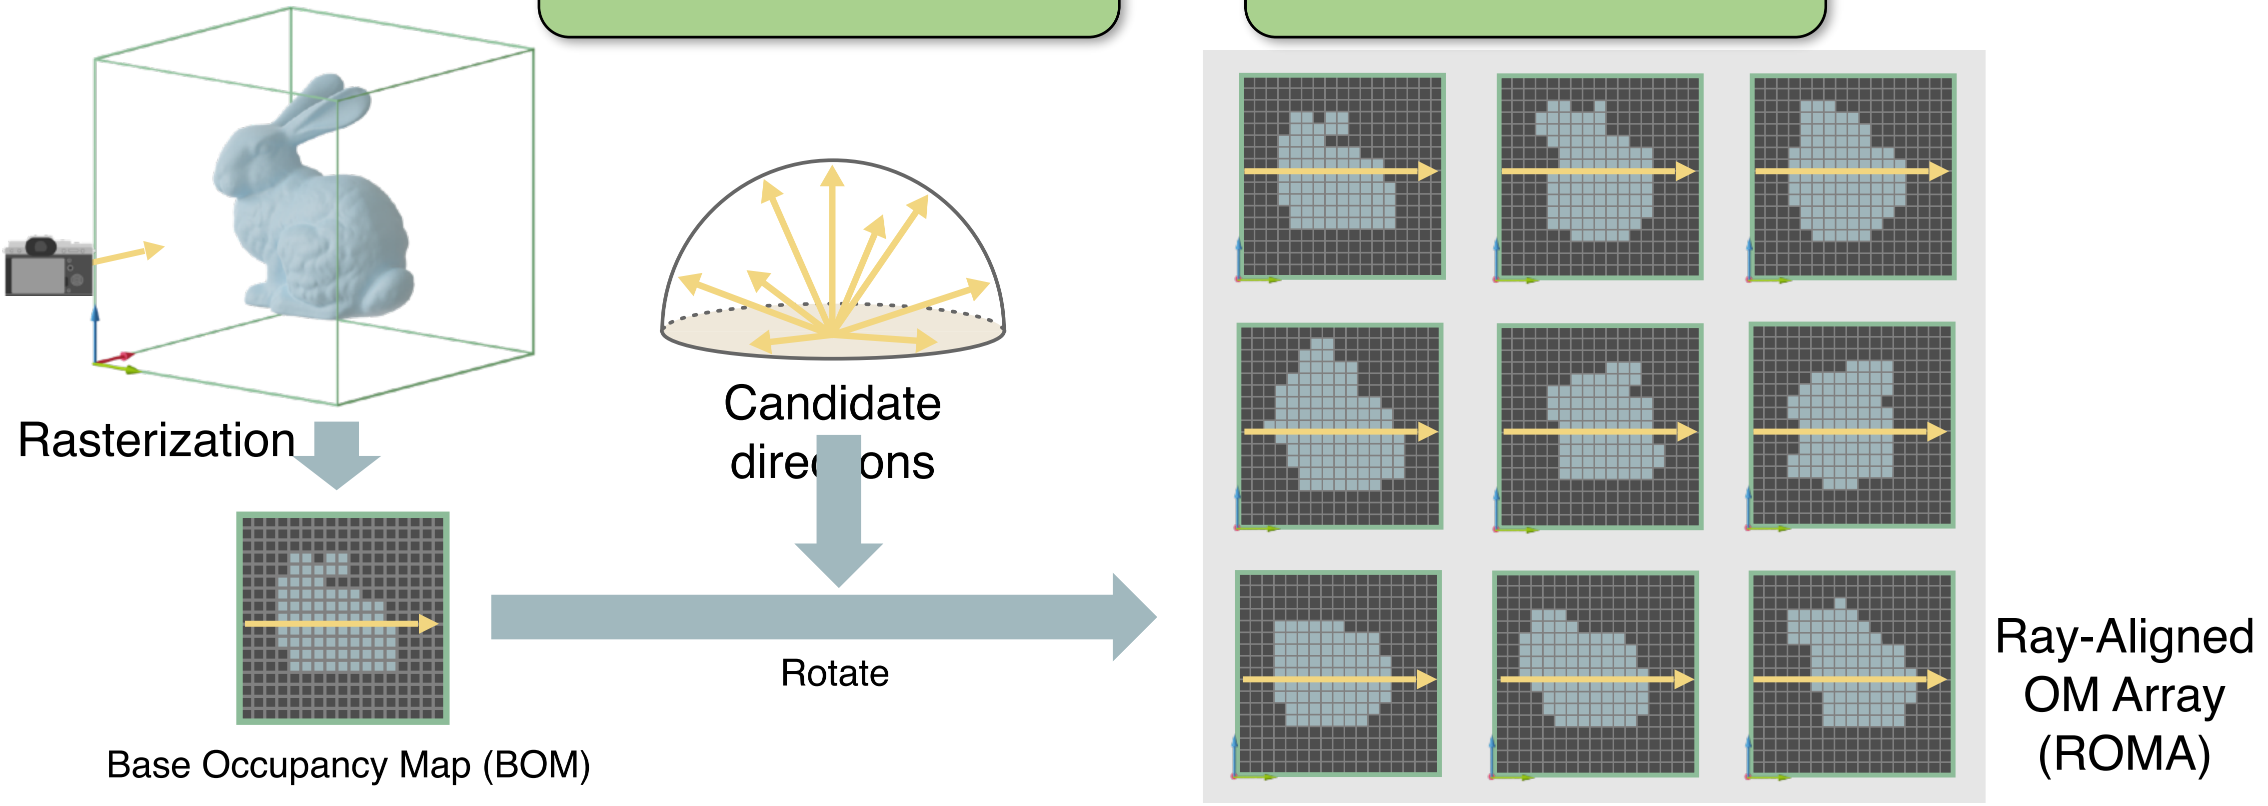

Overview

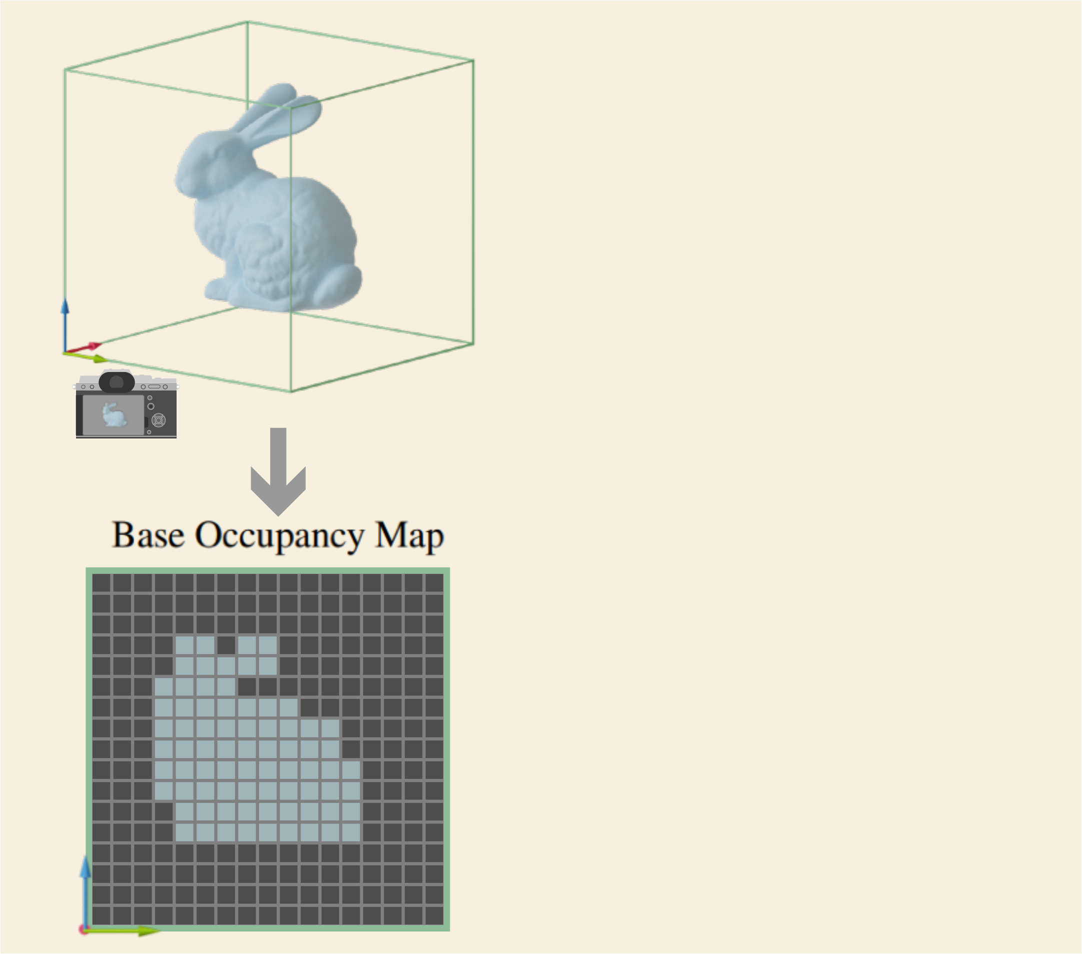

GENERATION STAGE

Overview

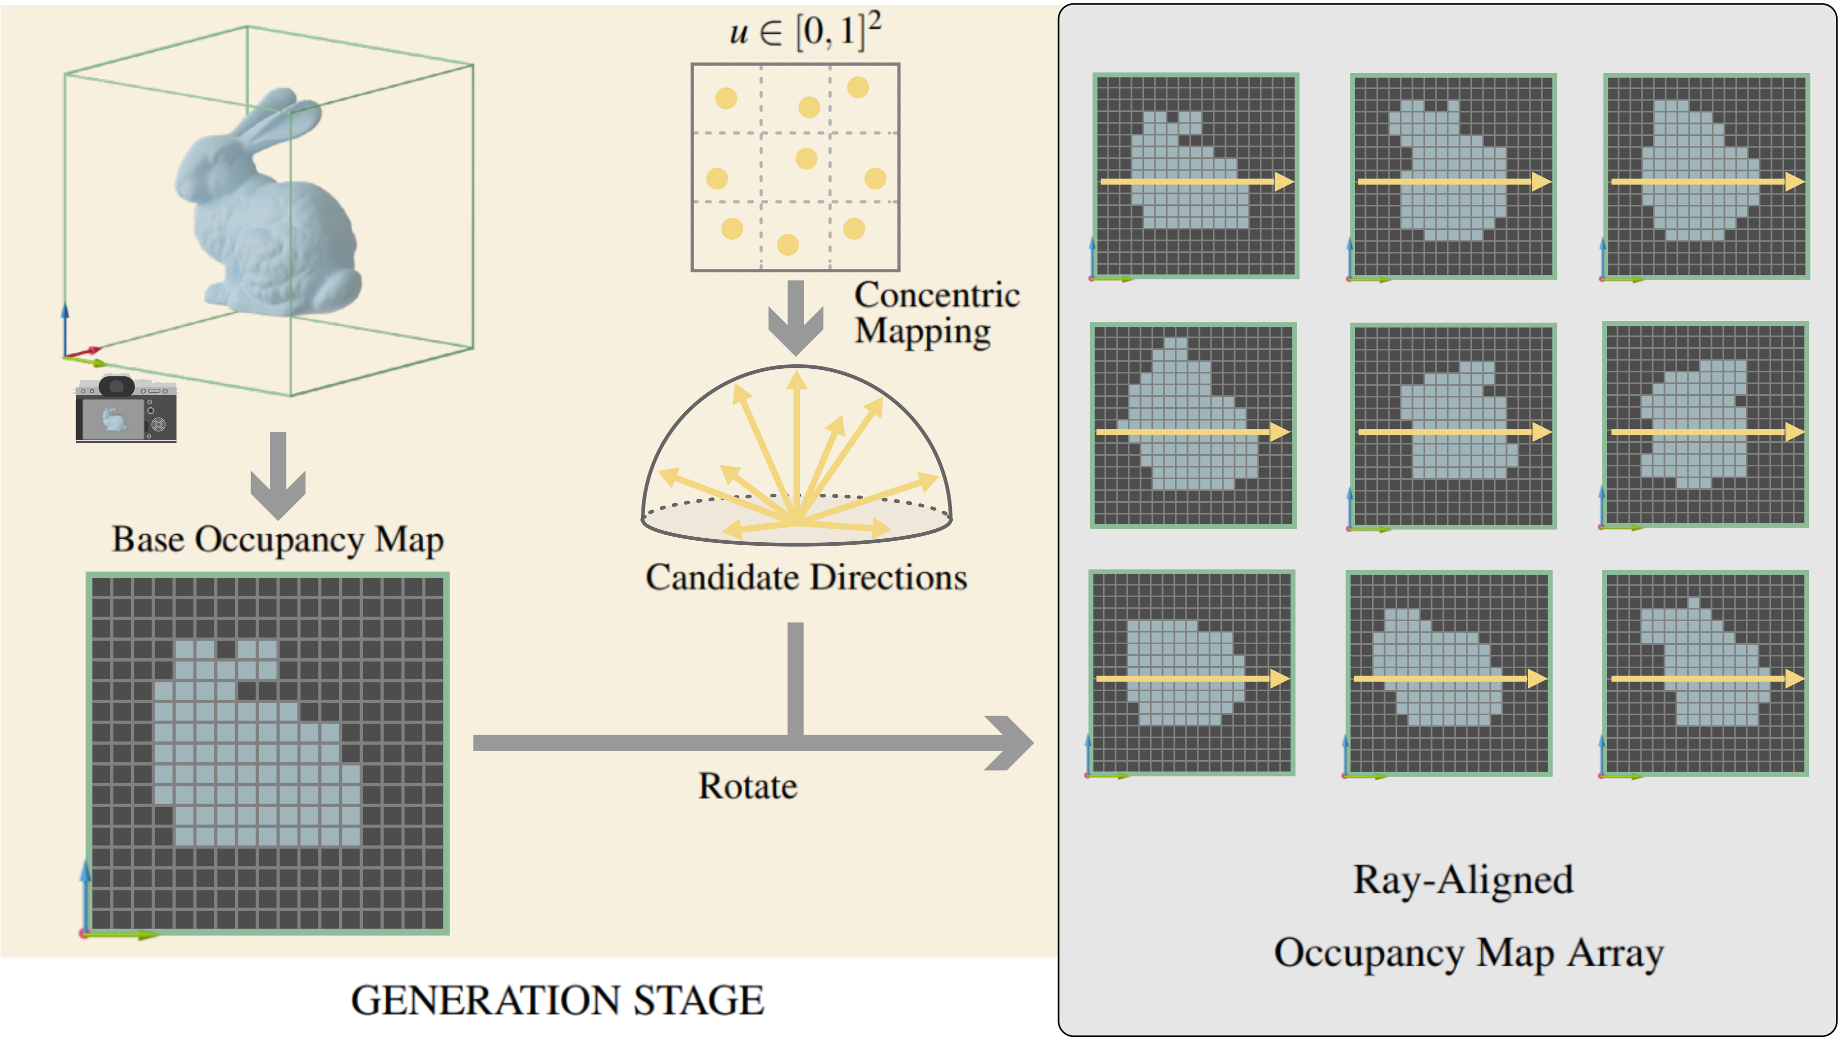
$u \in [0,1]^2$
Concentric Mapping
Base Occupancy Map
Candidate Directions
Rotate
Ray-Aligned
Occupancy Map Array
GENERATION STAGE

Overview

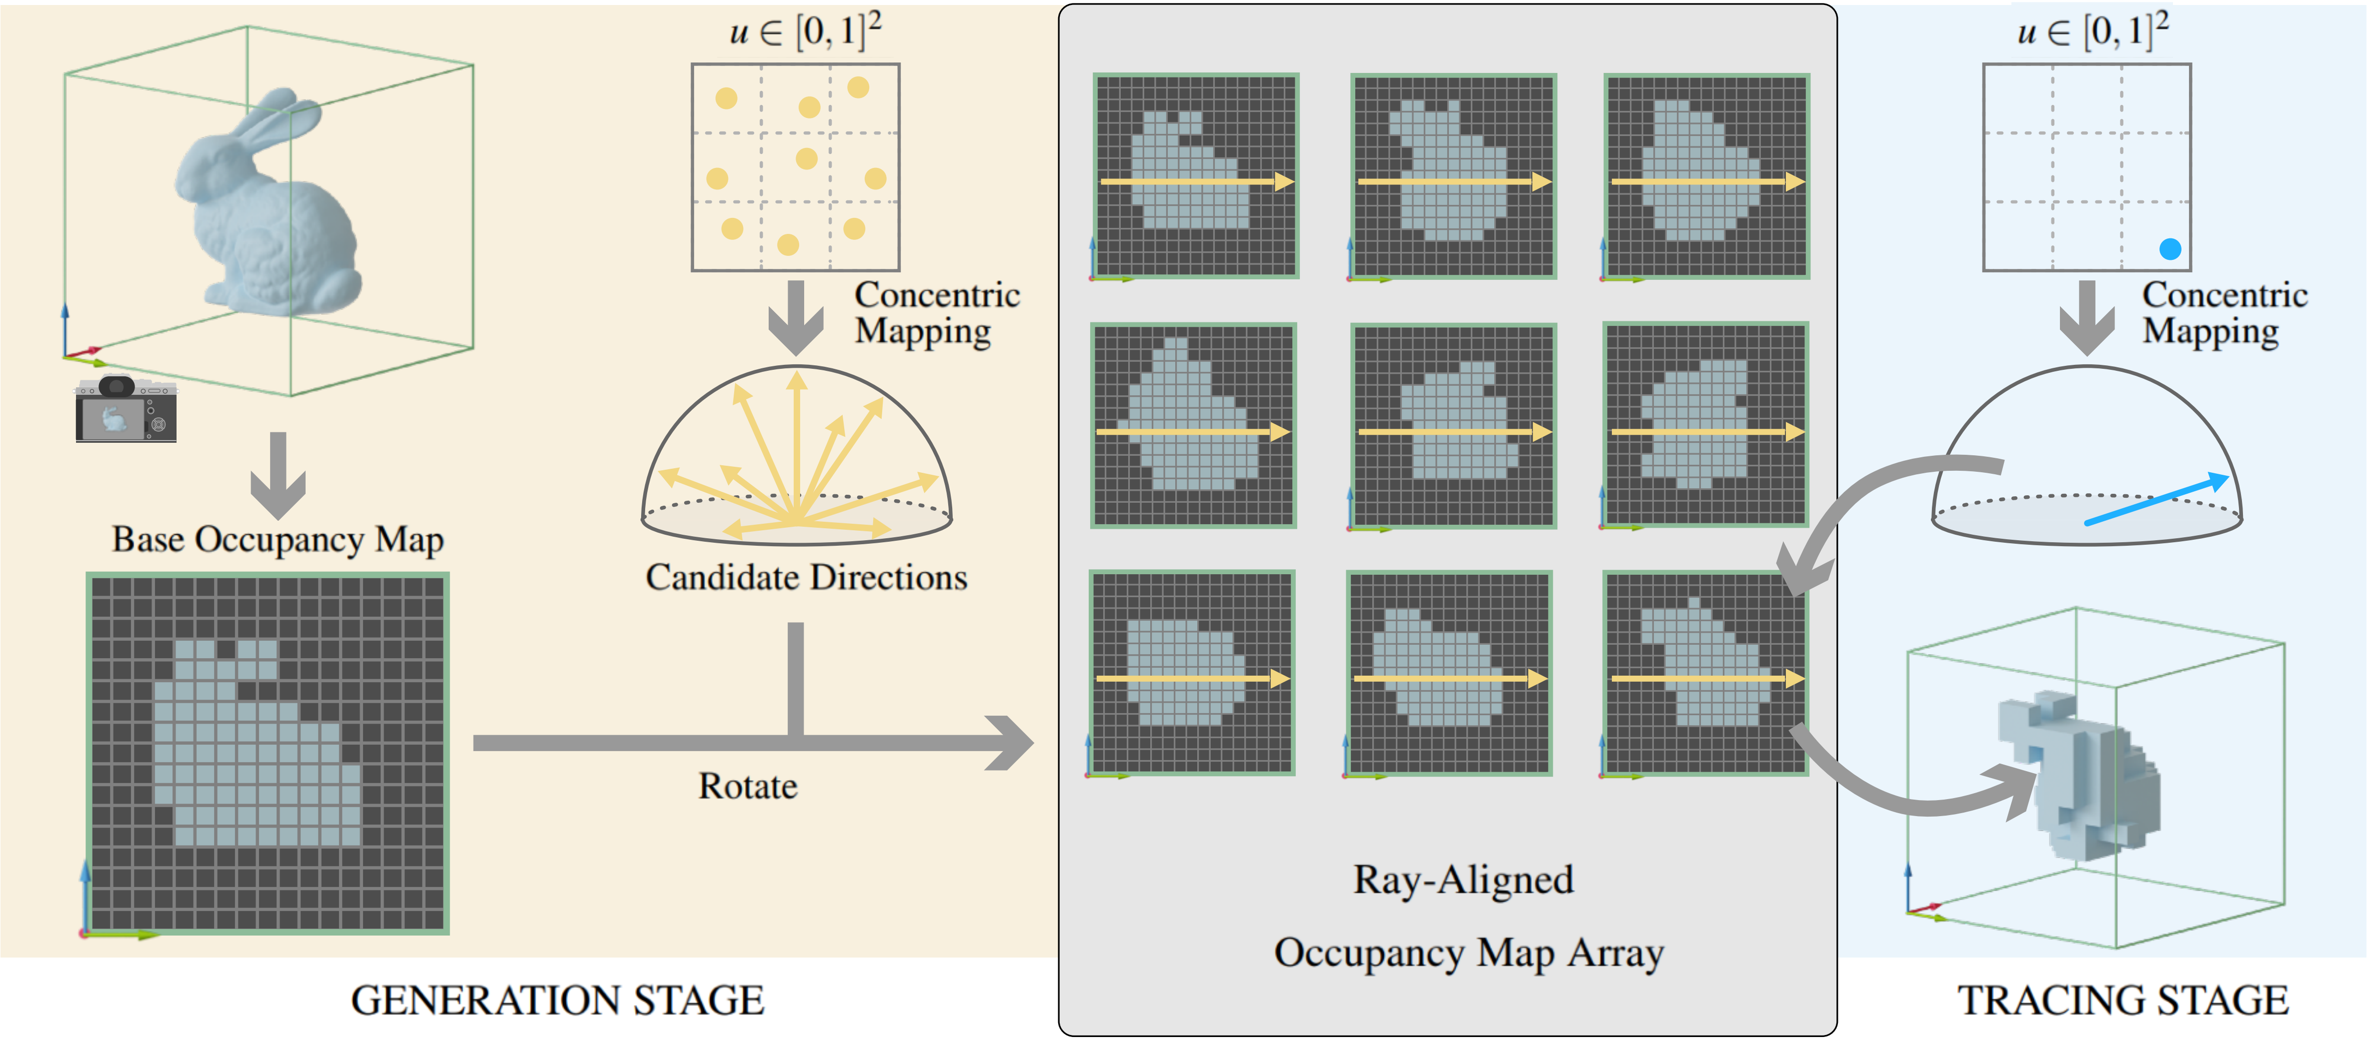

Base Occupancy Map
$u \in [0, 1]^2$
Concentric Mapping
Candidate Directions
Rotate
GENERATION STAGE
Ray-Aligned
Occupancy Map Array
$u \in [0, 1]^2$
Concentric Mapping
TRACING STAGE

Step 1: Generate BOM

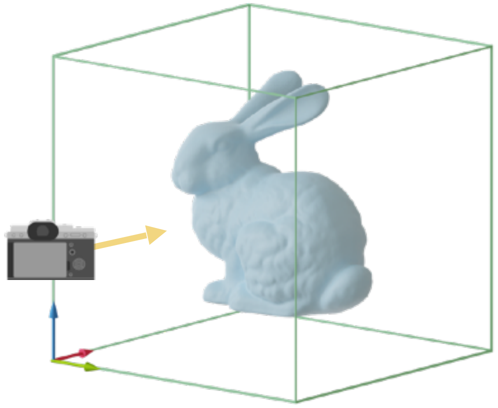

Rasterization

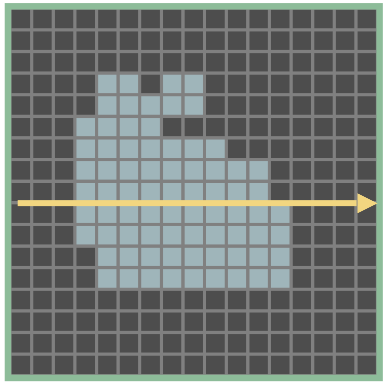

Base Occupancy Map (BOM)

Step 1: Generate BOM

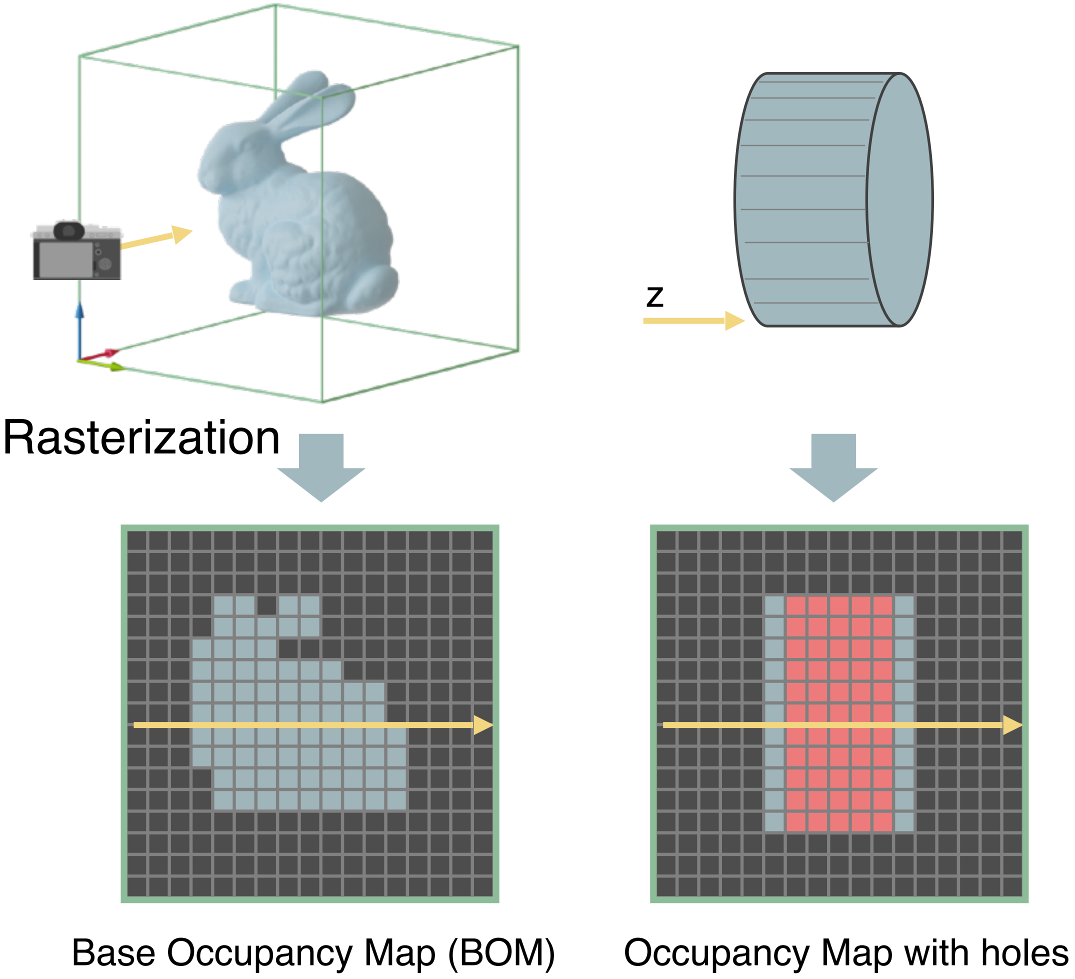

Step 1: Generate BOM

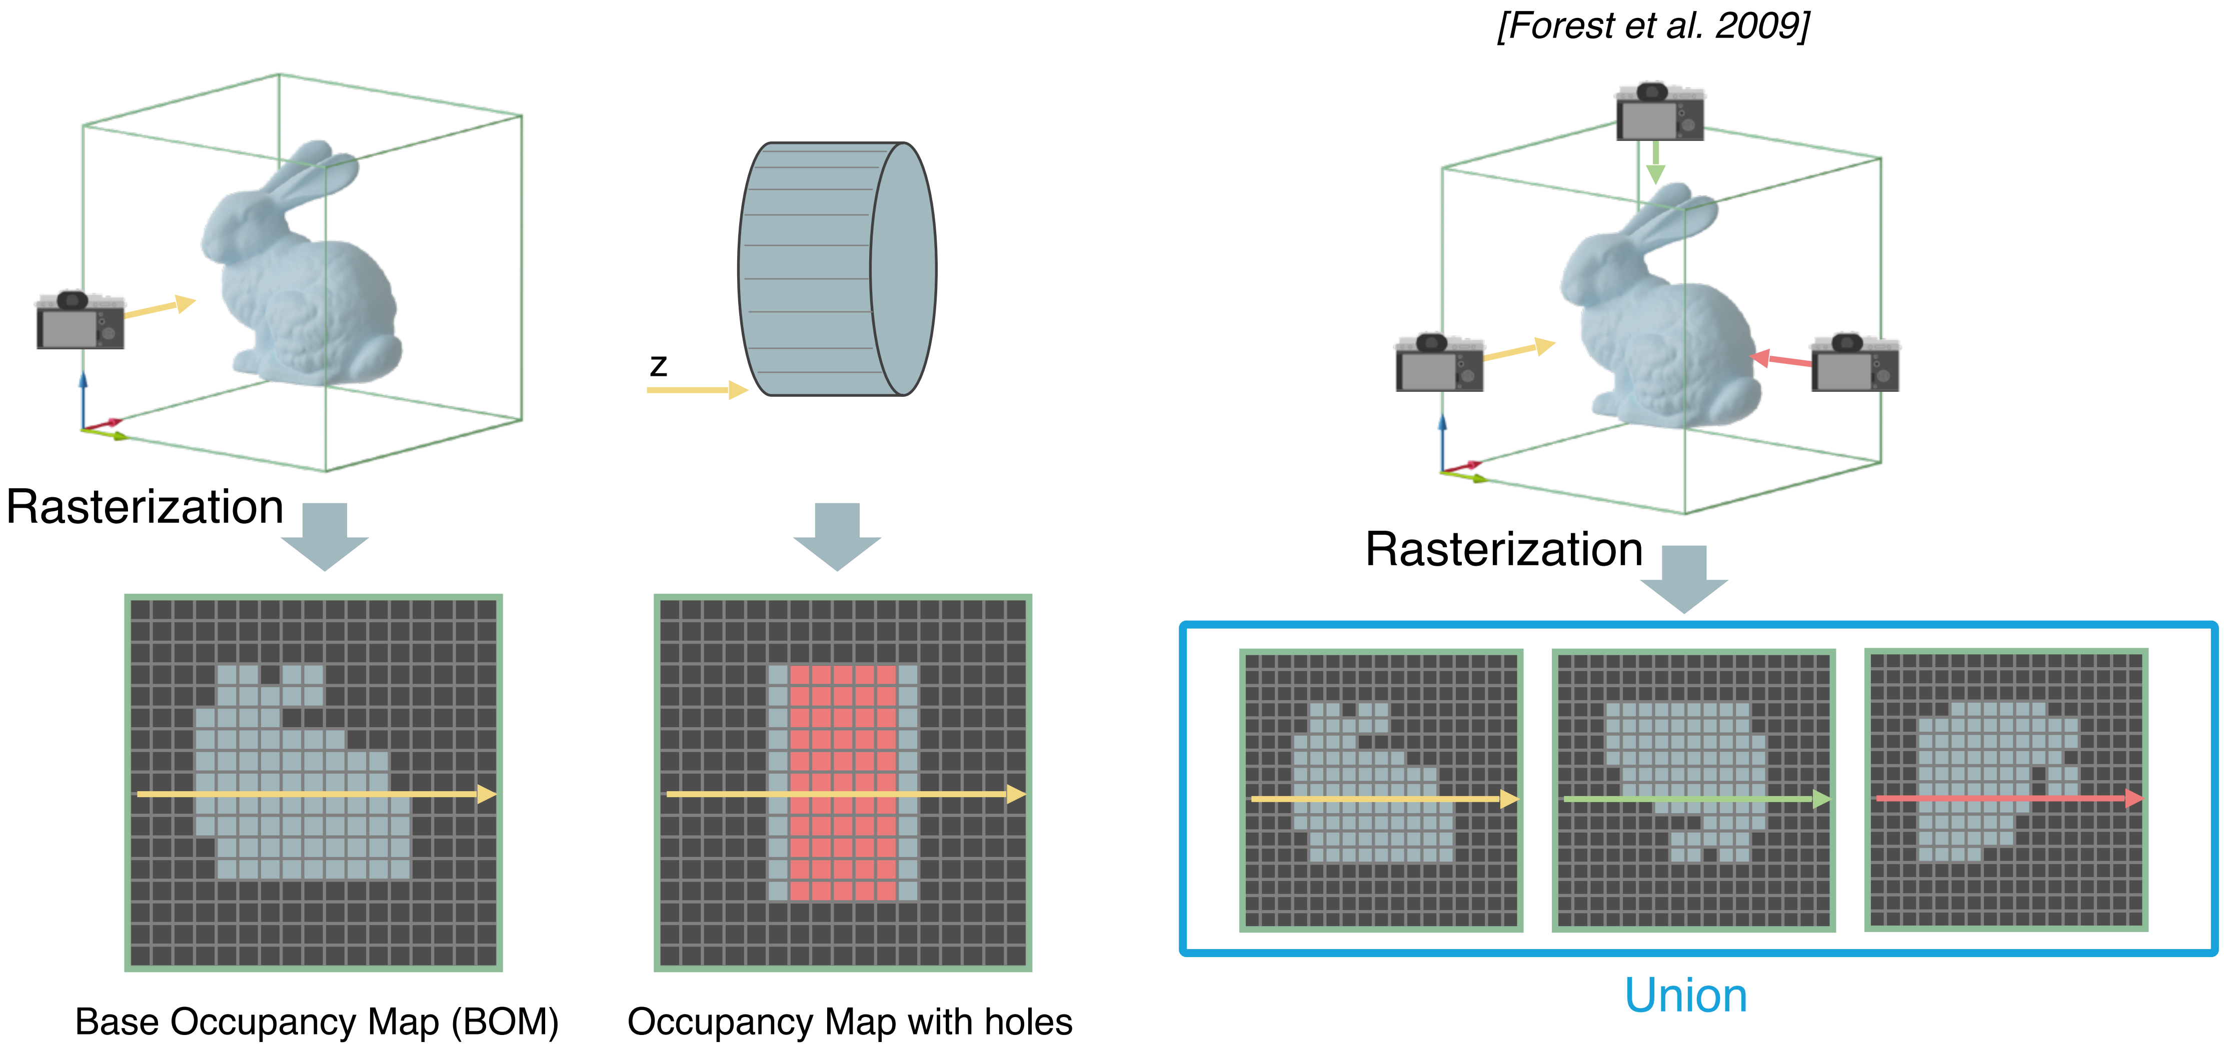

Step 2: “Rotate” BOM to ROMA

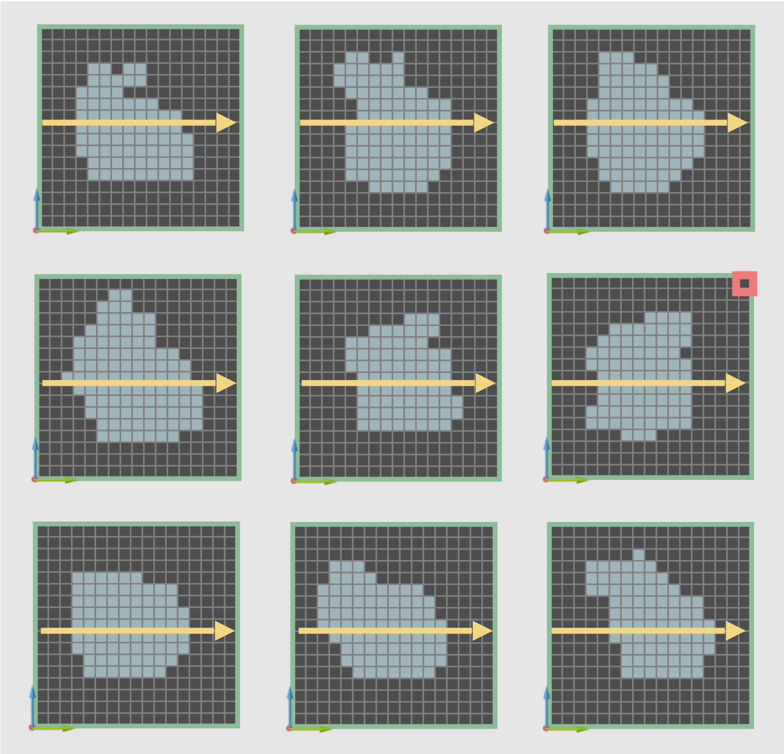

Ray-Aligned OM Array
(ROMA)

Step 2: “Rotate” BOM to ROMA

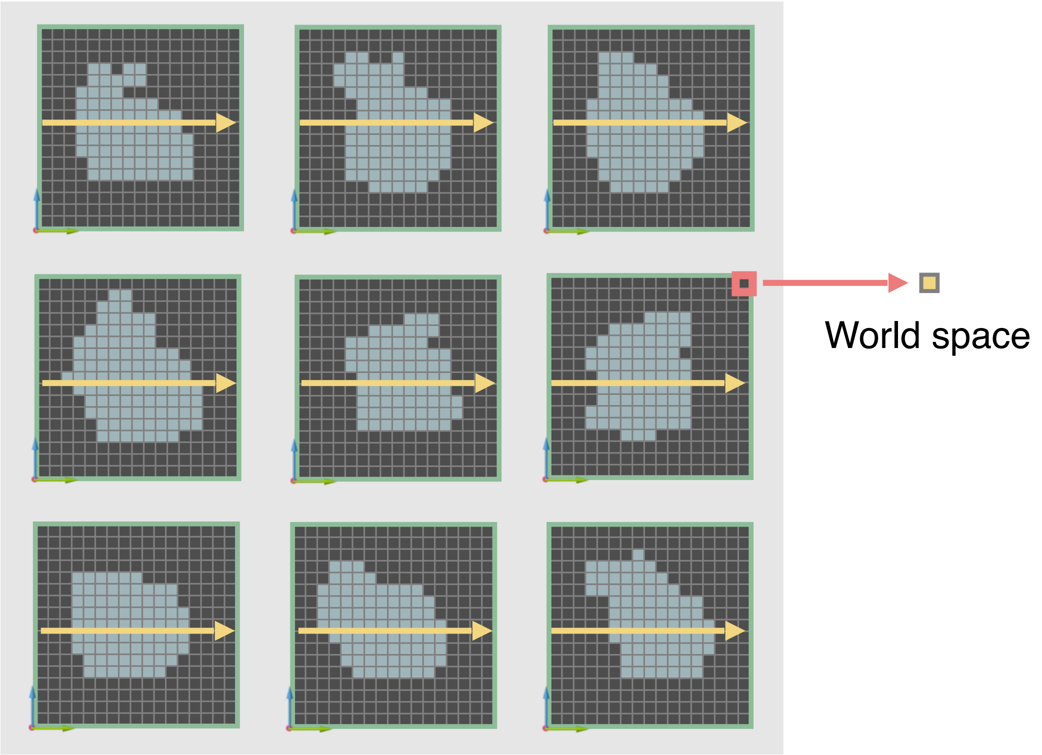

Ray-Aligned OM Array
(ROMA)

Step 2: “Rotate” BOM to ROMA

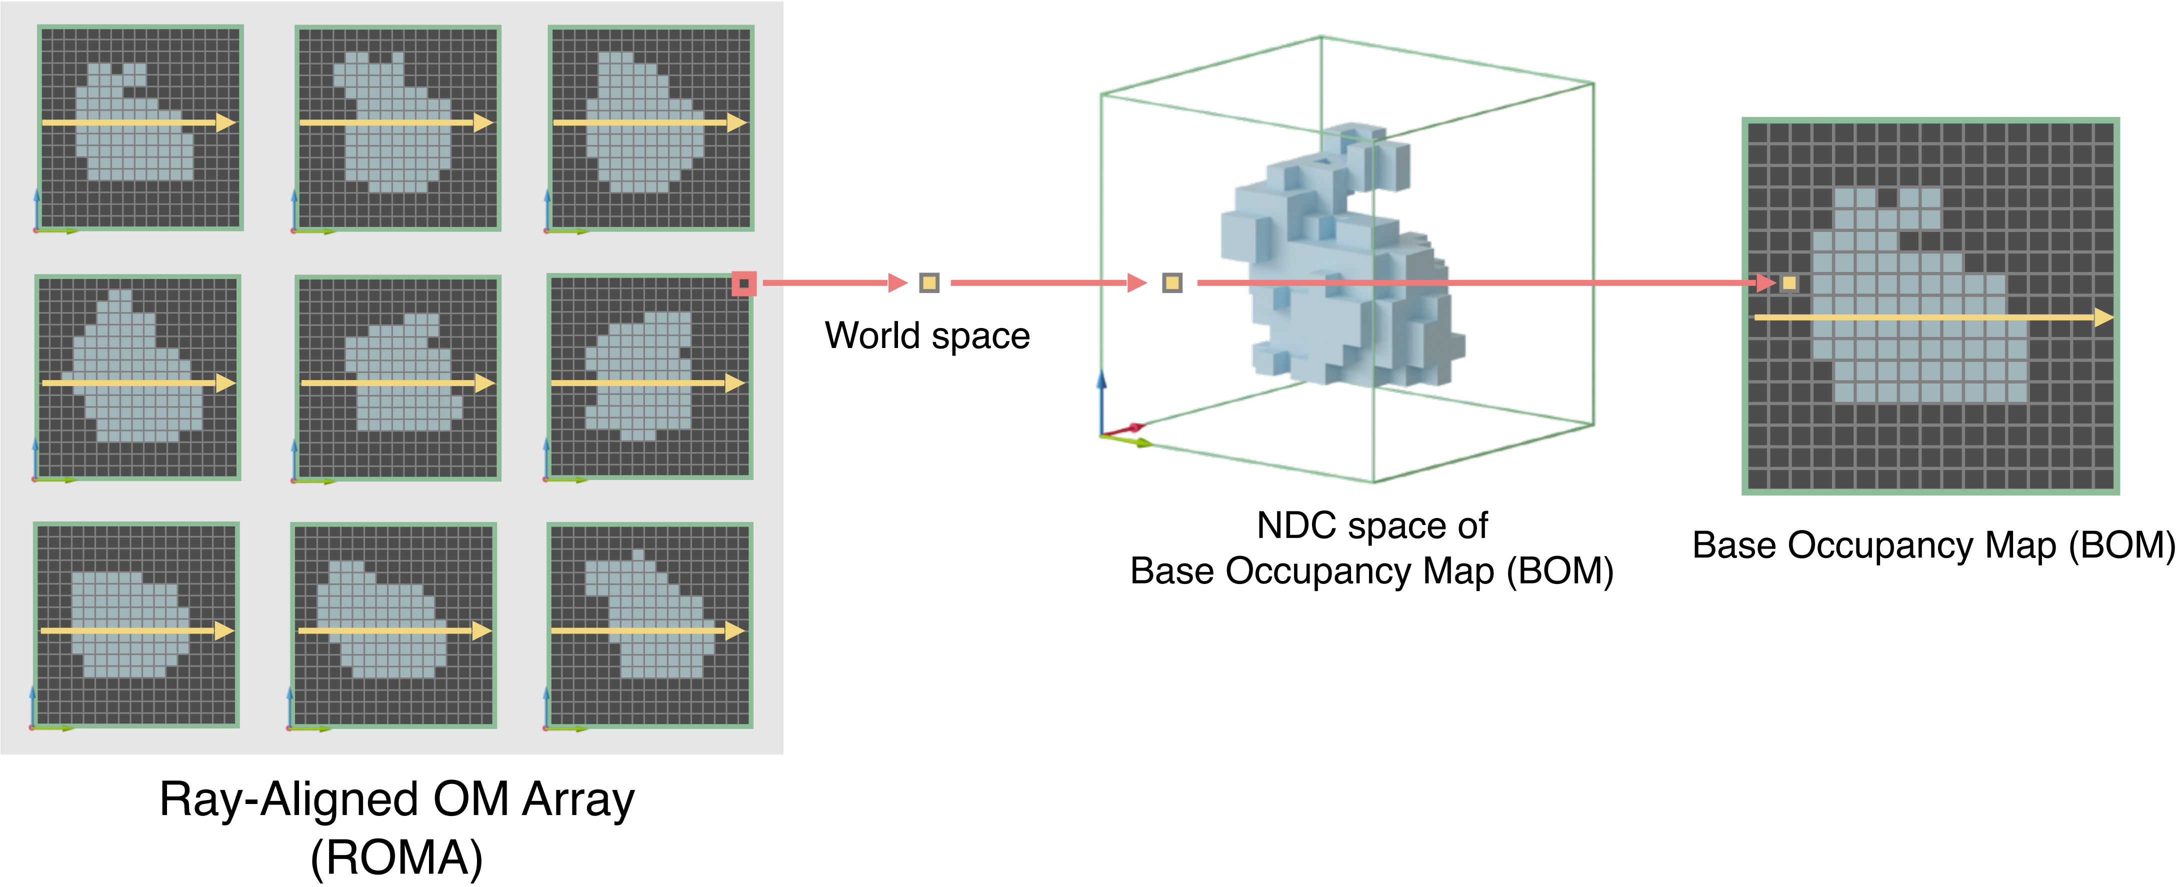

Step 3: “Snap” and Trace

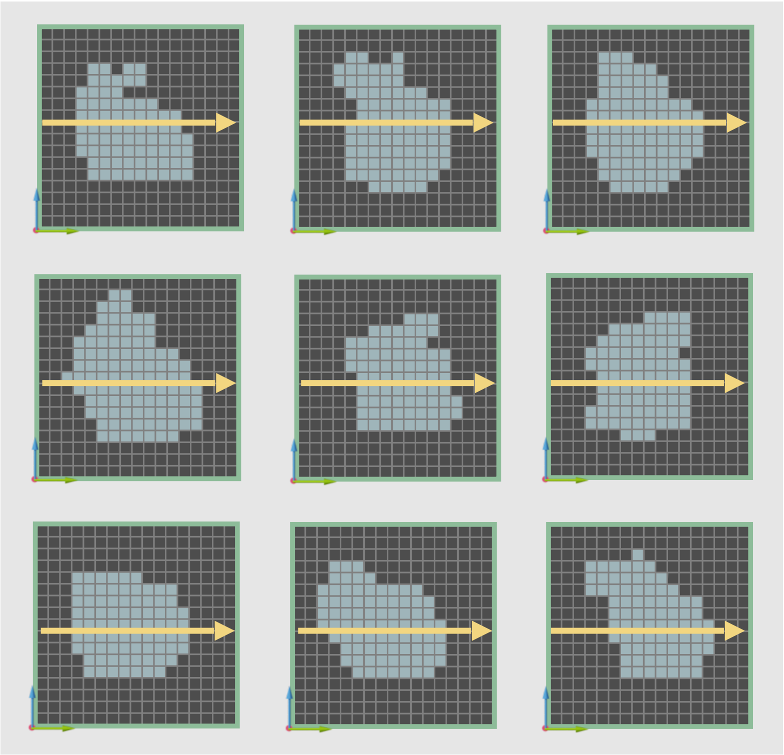

Ray-Aligned OM Array
(ROMA)

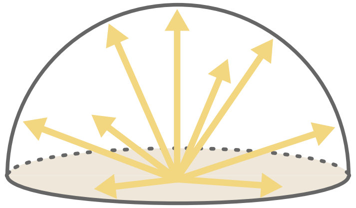

Candidate
directions

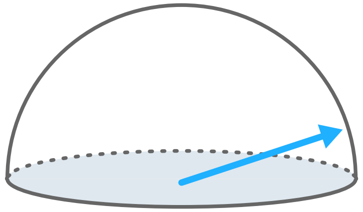

Sampled direction

Step 3: “Snap” and Trace

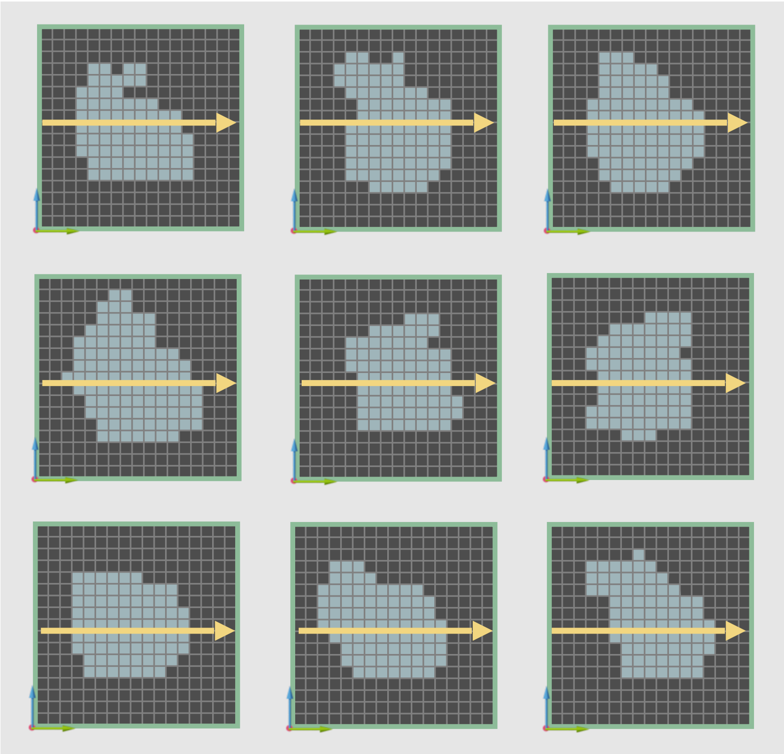

Ray-Aligned OM Array (ROMA)

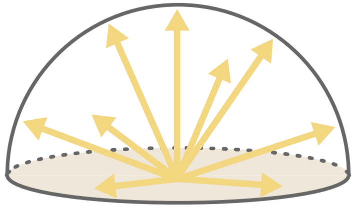

Candidate directions

Which is the closest candidate direction?

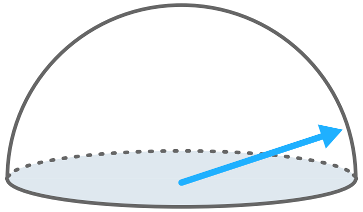

Sampled direction

Step 3: "Snap" and Trace

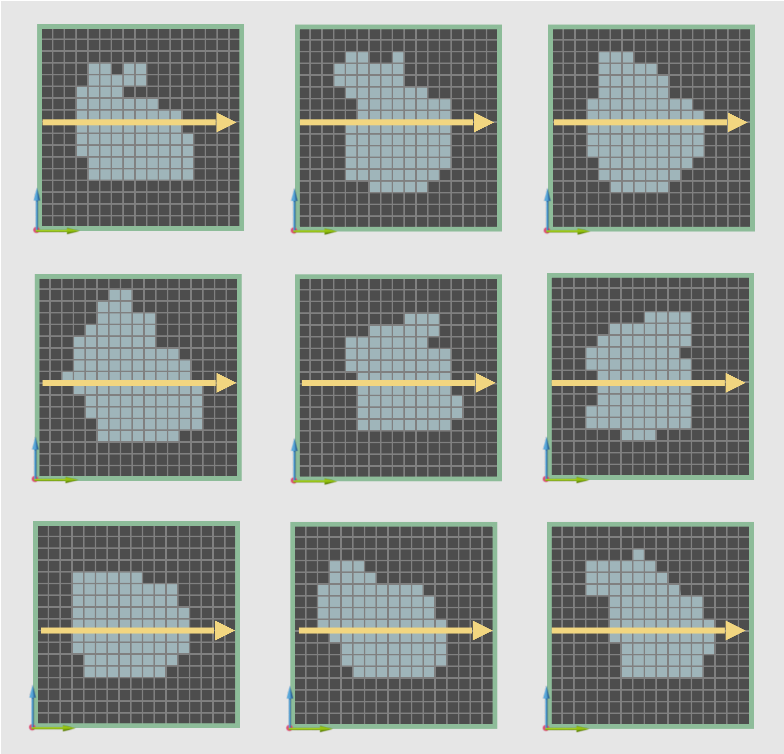

Ray-Aligned OM Array (ROMA)

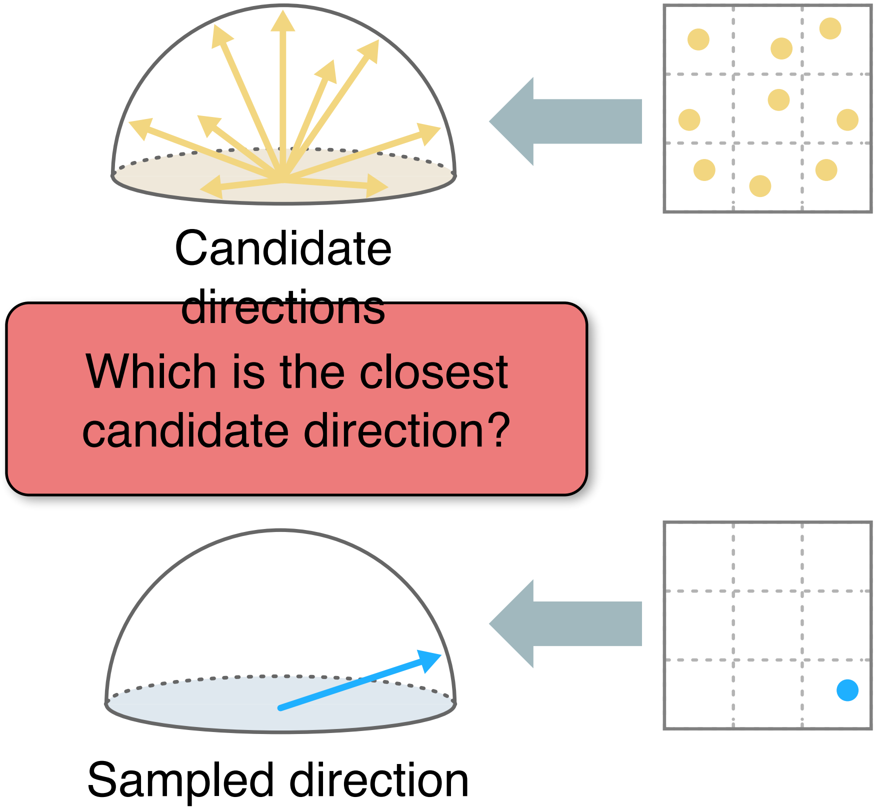

Step 3: "Snap" and Trace

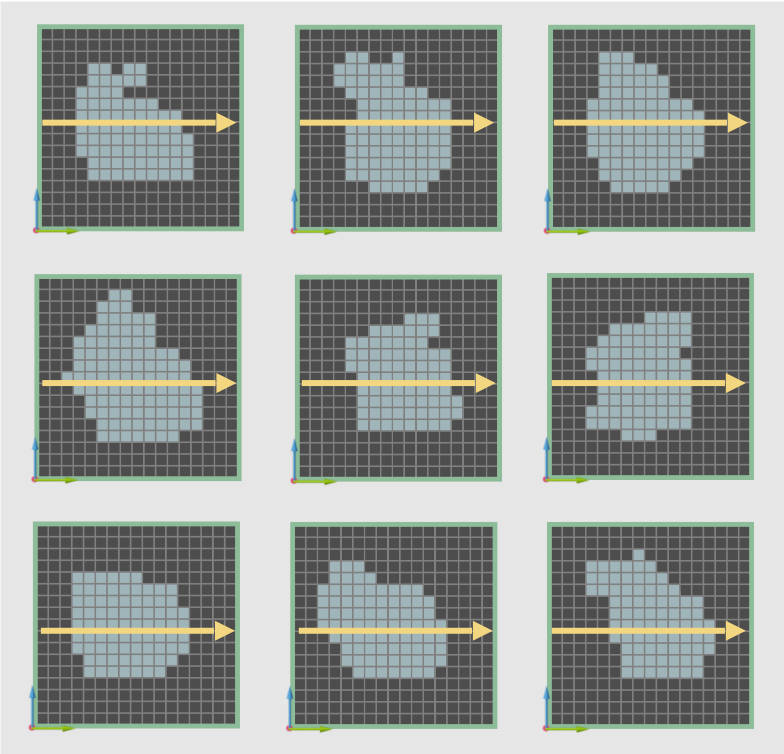

Ray-Aligned OM Array (ROMA)

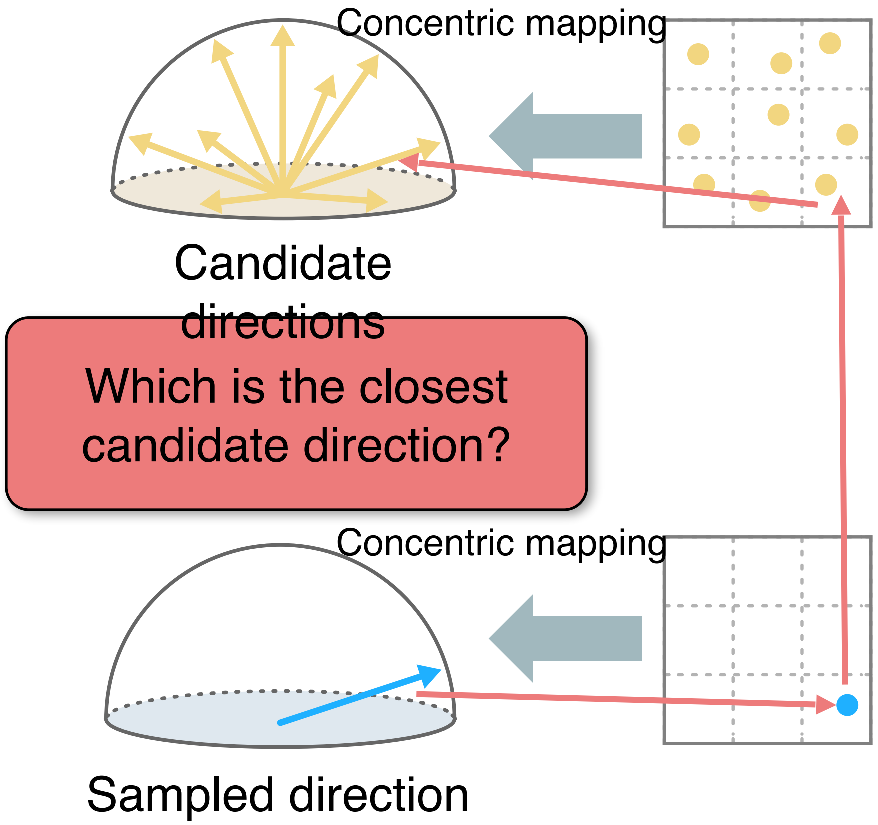

Step 3: “Snap” and Trace

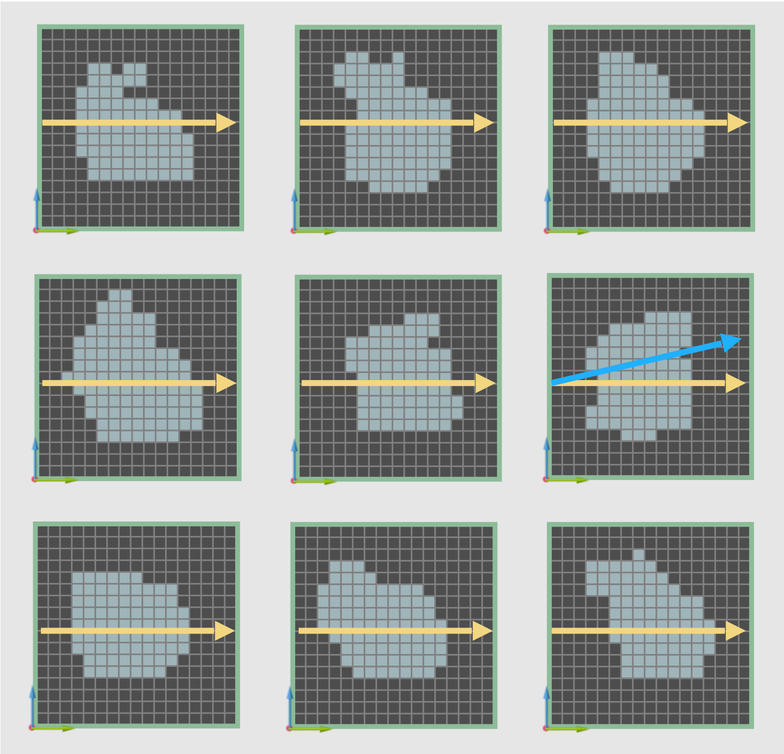

Ray-Aligned OM Array
(ROMA)

Step 3: “Snap” and Trace

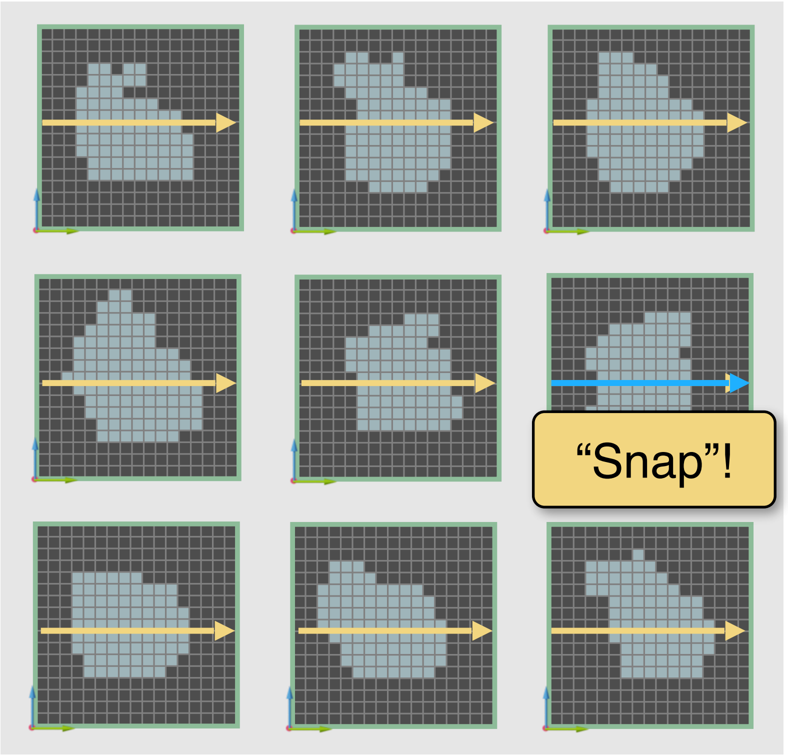

Ray-Aligned OM Array
(ROMA)

Step 3: “Snap” and Trace

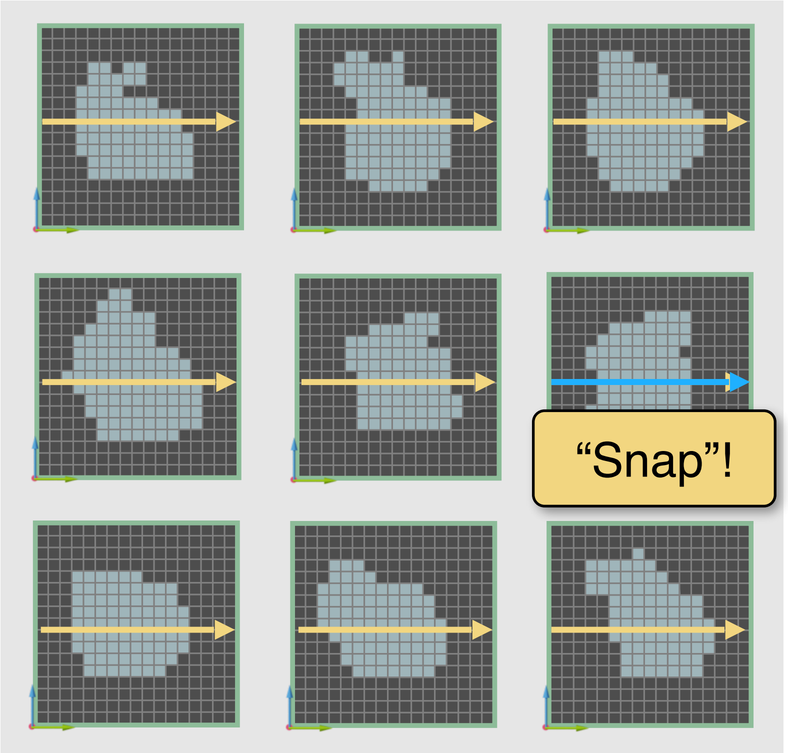

Ray-Aligned OM Array
(ROMA)

any hit query

Step 3: “Snap” and Trace

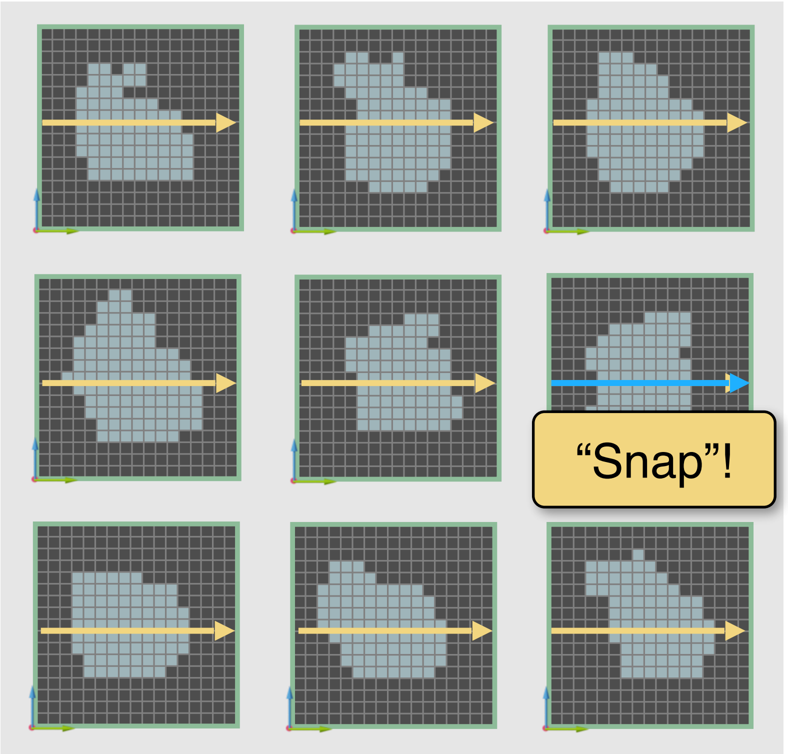

Ray-Aligned OM Array
(ROMA)

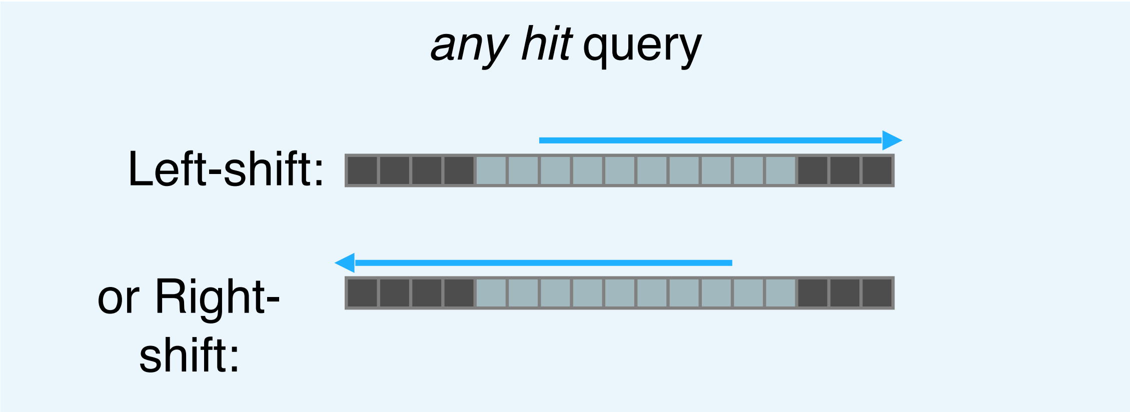

Step 3: “Snap” and Trace

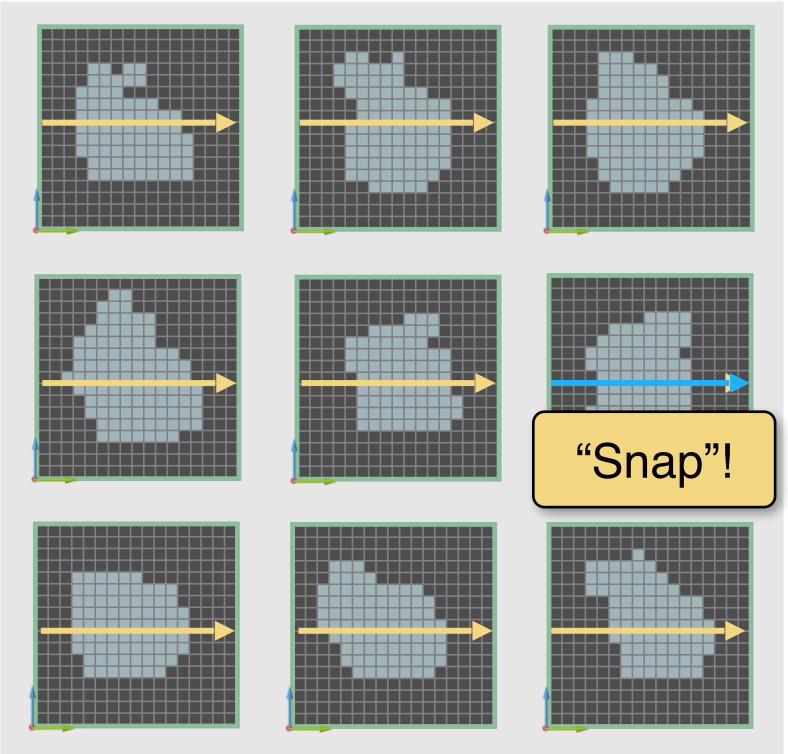

Ray-Aligned OM Array (ROMA)

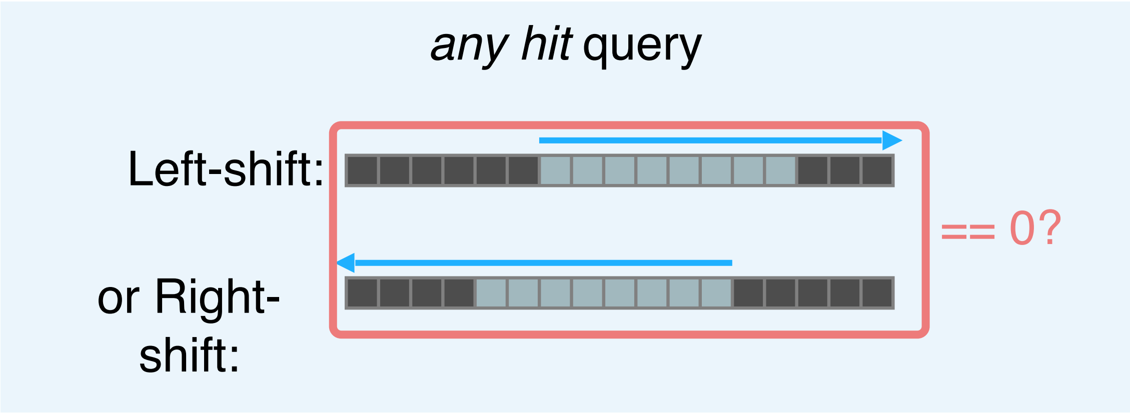

Step 3: “Snap” and Trace

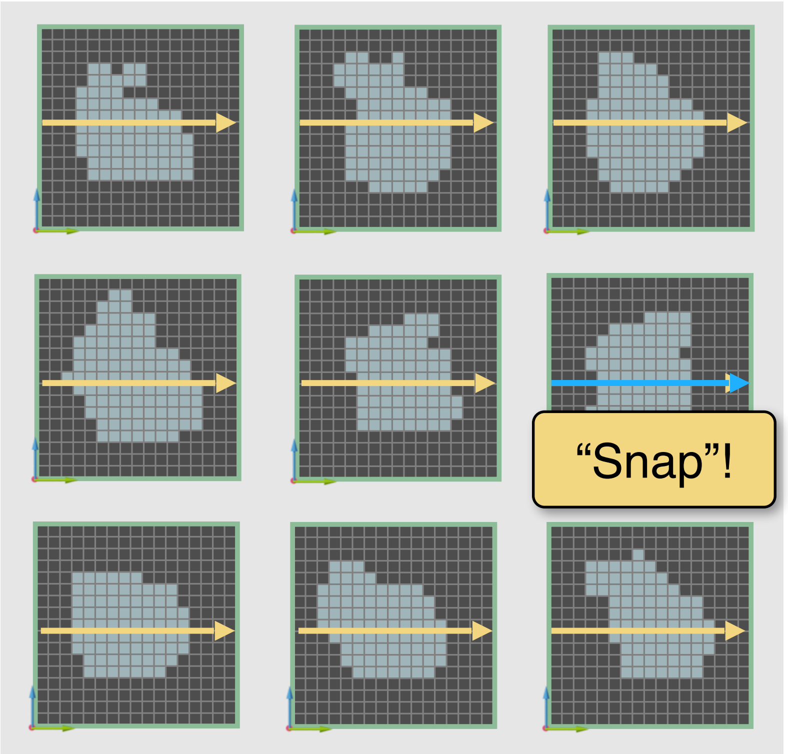

Ray-Aligned OM Array
(ROMA)

any hit query

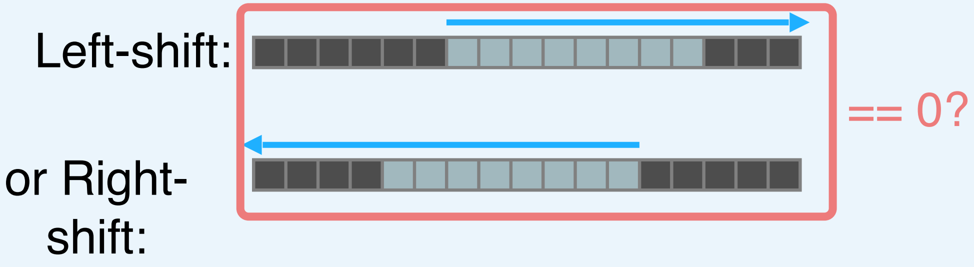

closest hit query

Step 3: “Snap” and Trace

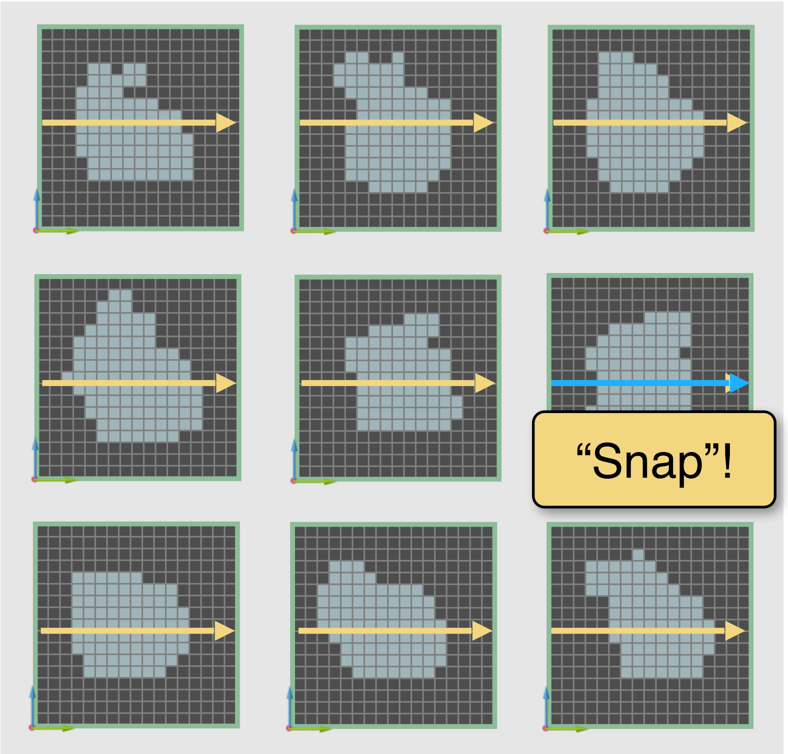

Ray-Aligned OM Array
(ROMA)

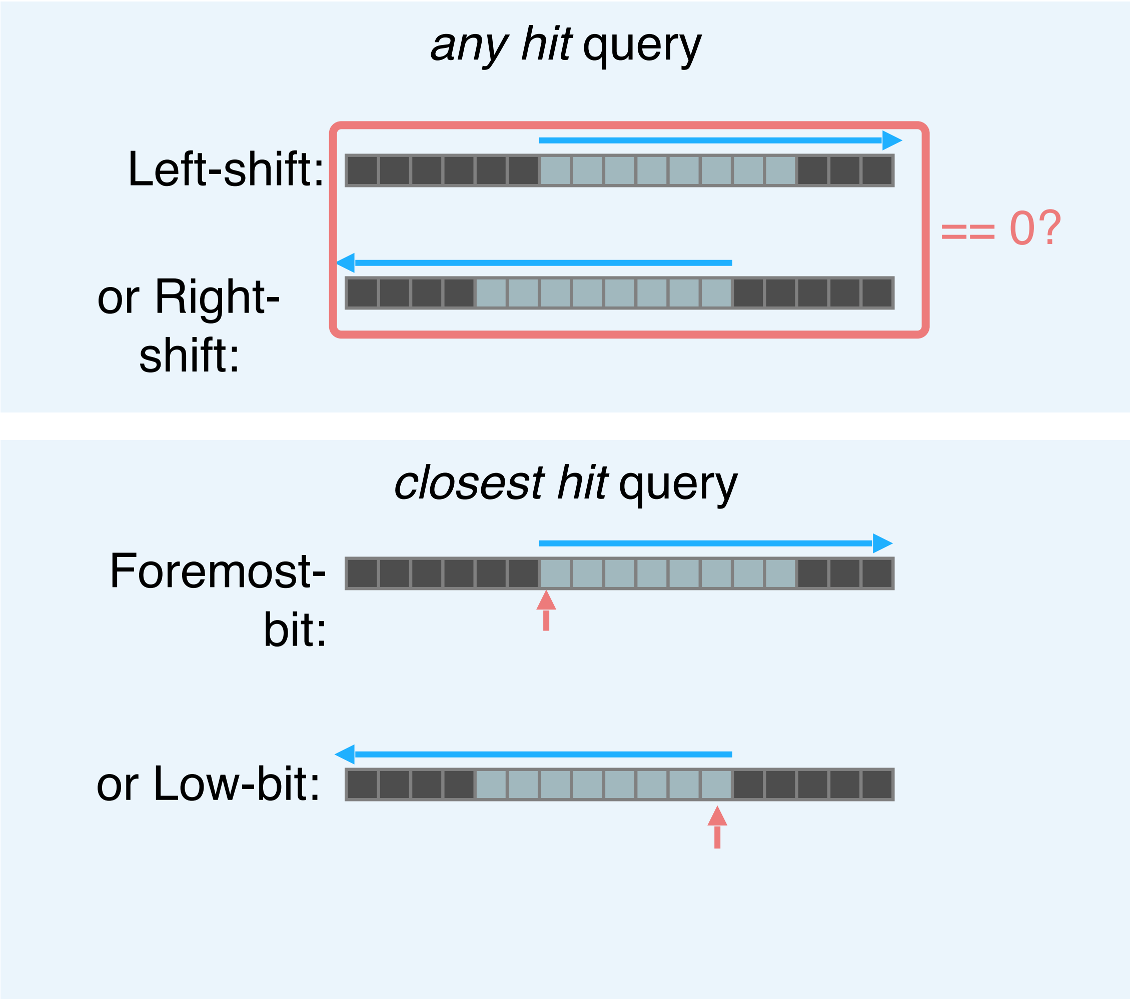

Step 3: “Snap” and Trace

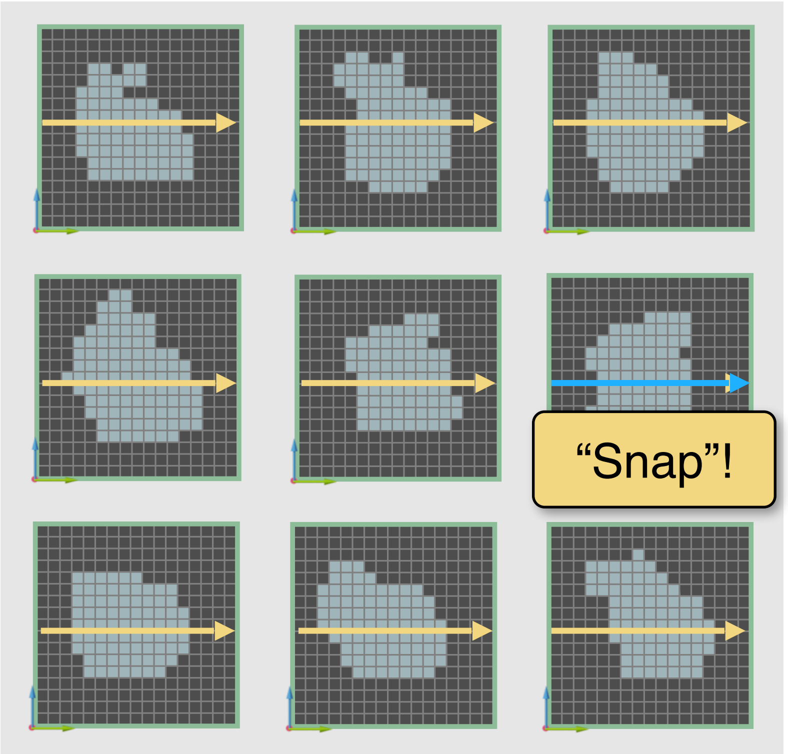

Ray-Aligned OM Array
(ROMA)

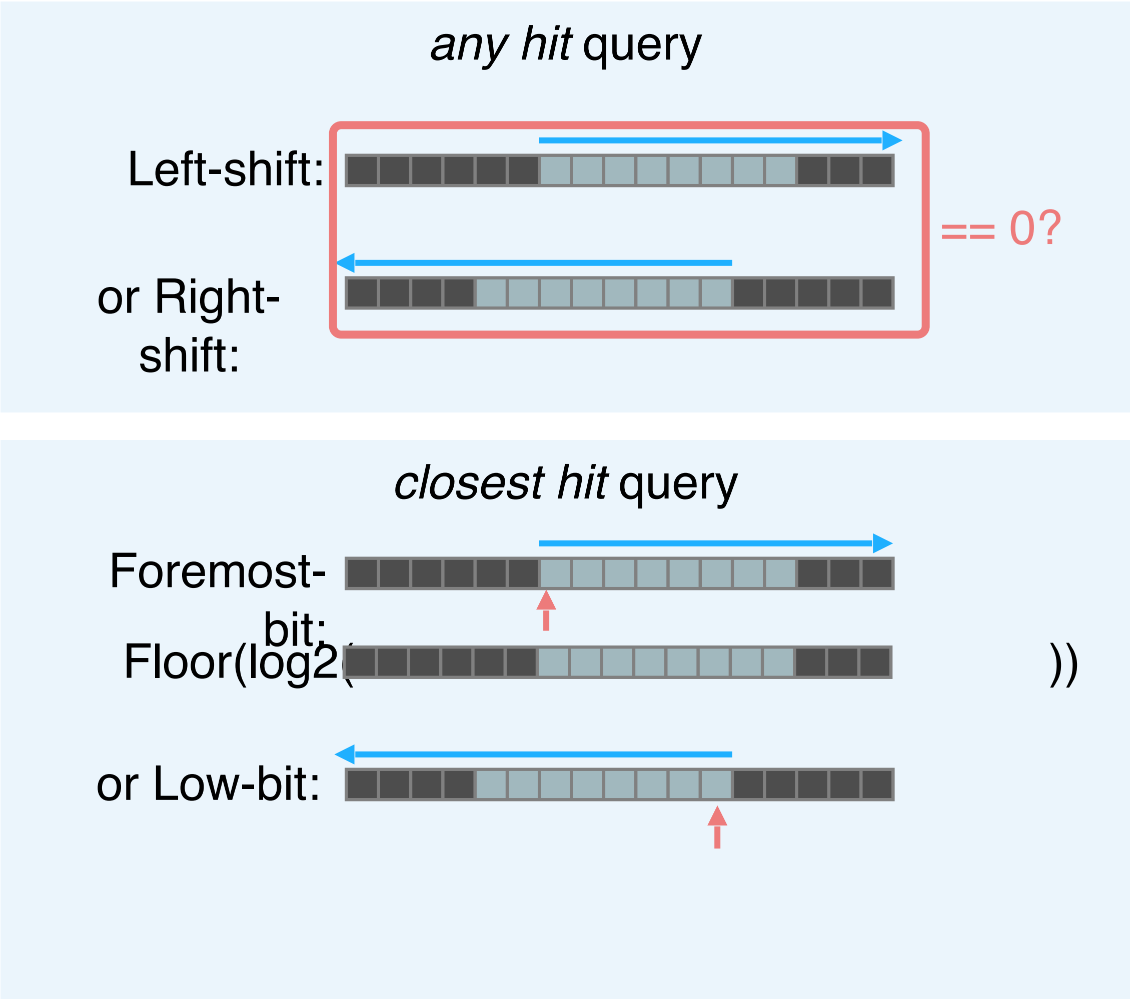

Step 3: “Snap” and Trace

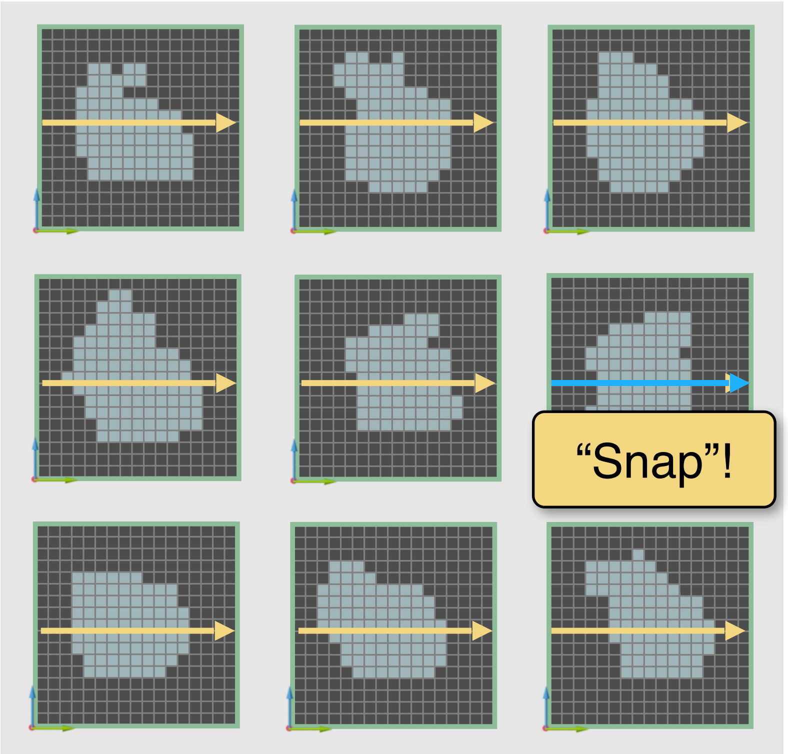

Ray-Aligned OM Array
(ROMA)

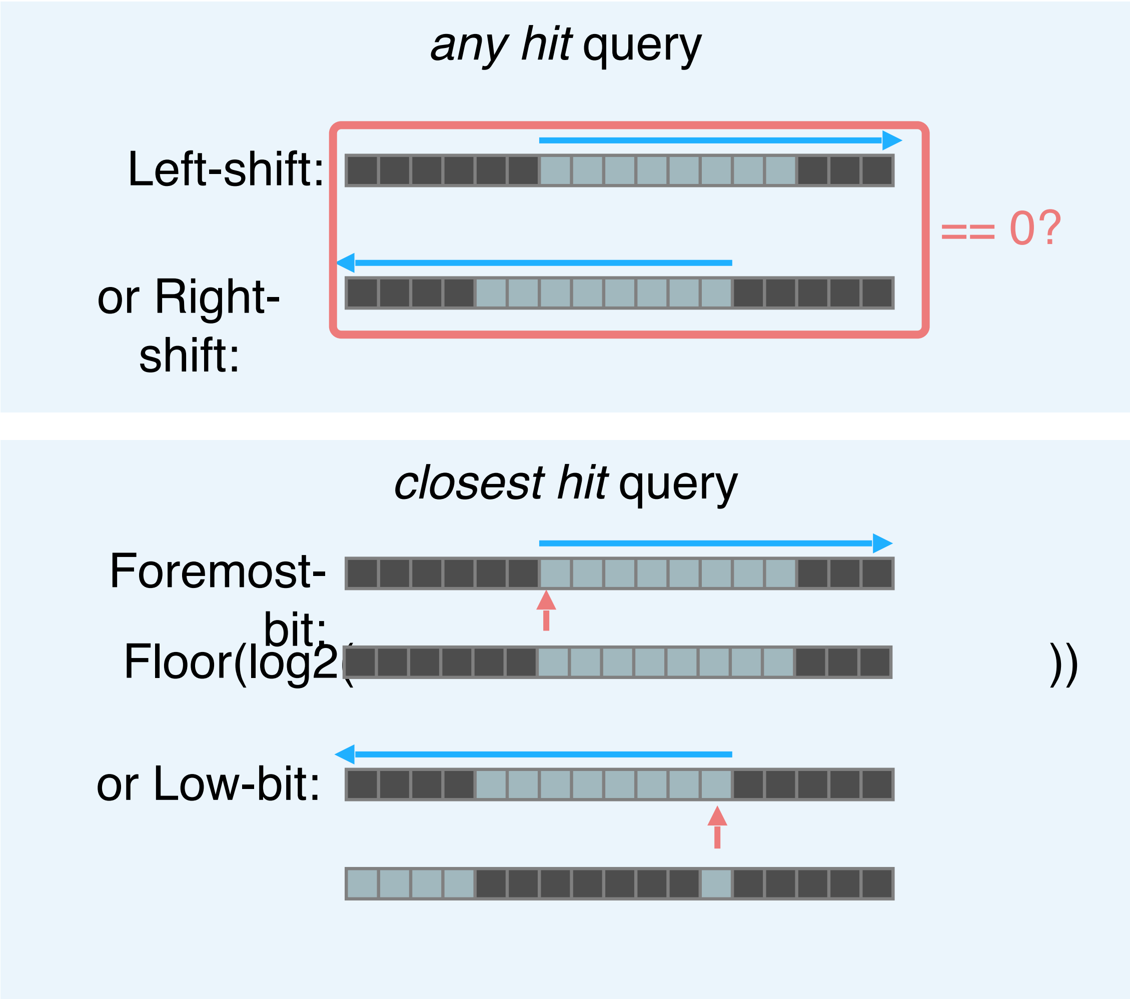

Step 3: “Snap” and Trace

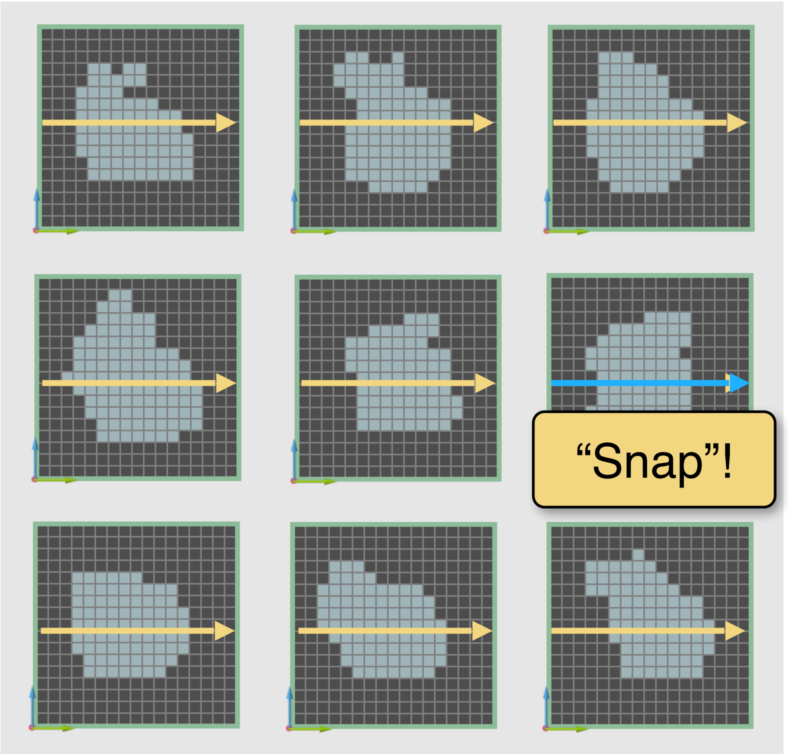

Ray-Aligned OM Array (ROMA)

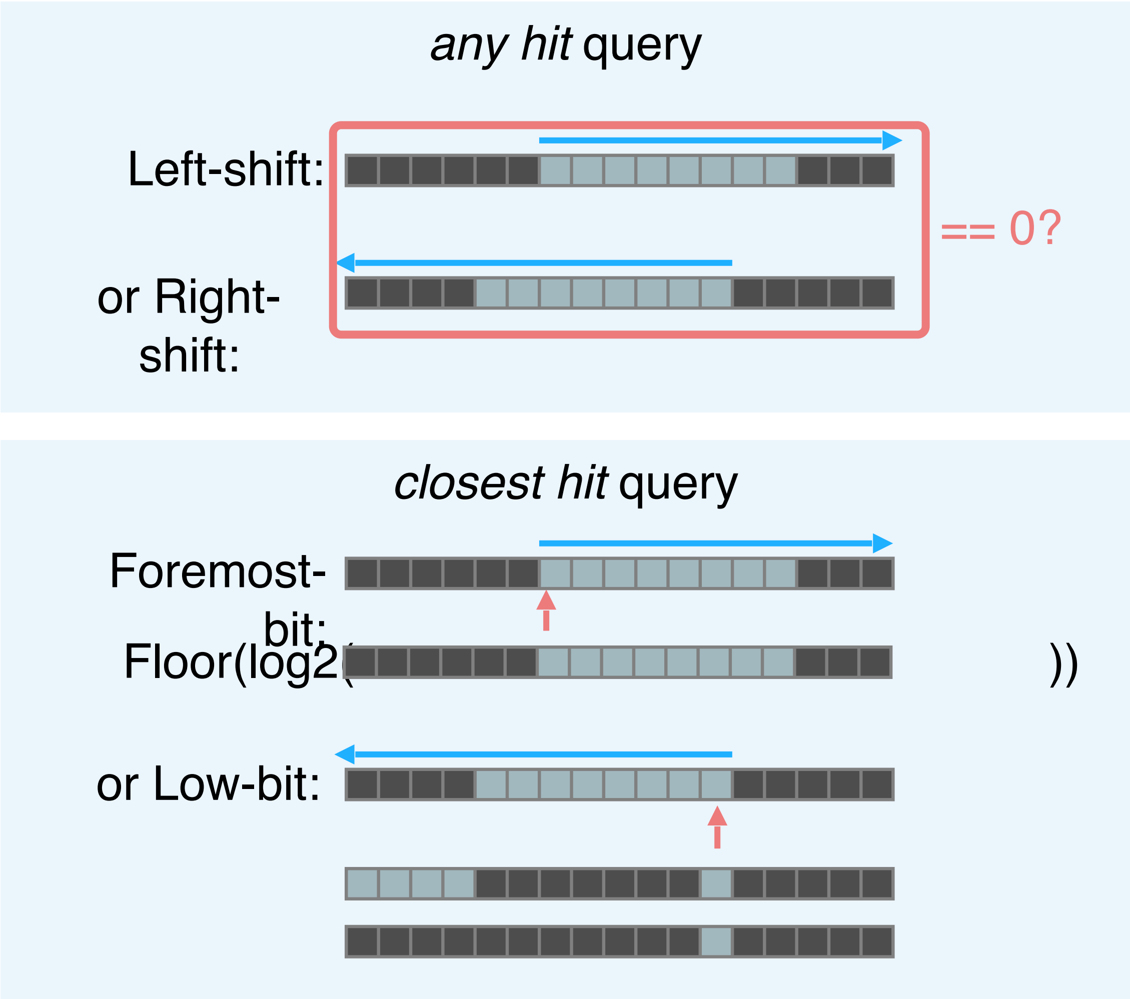

Step 3: “Snap” and Trace

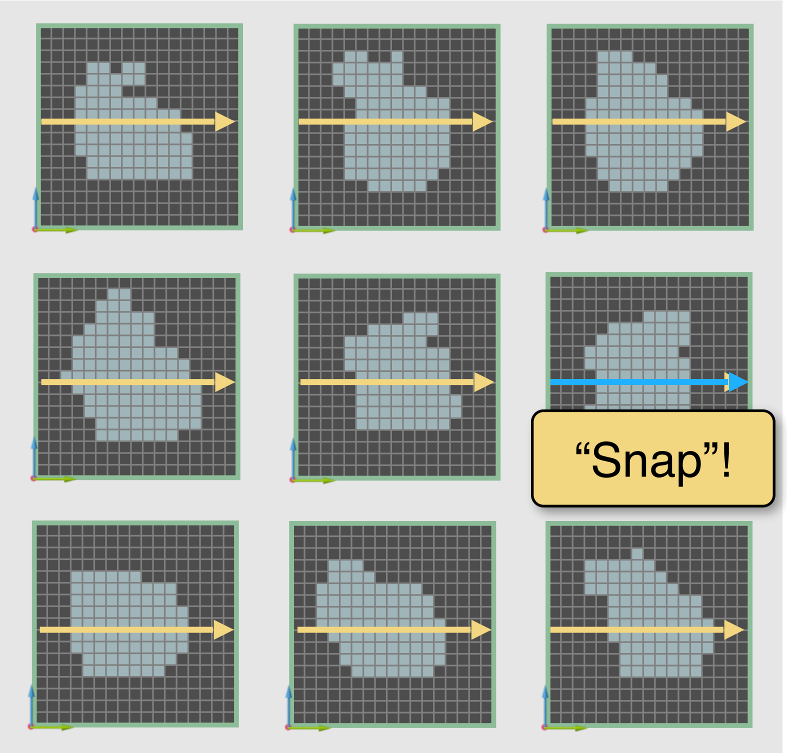

Ray-Aligned OM Array
(ROMA)

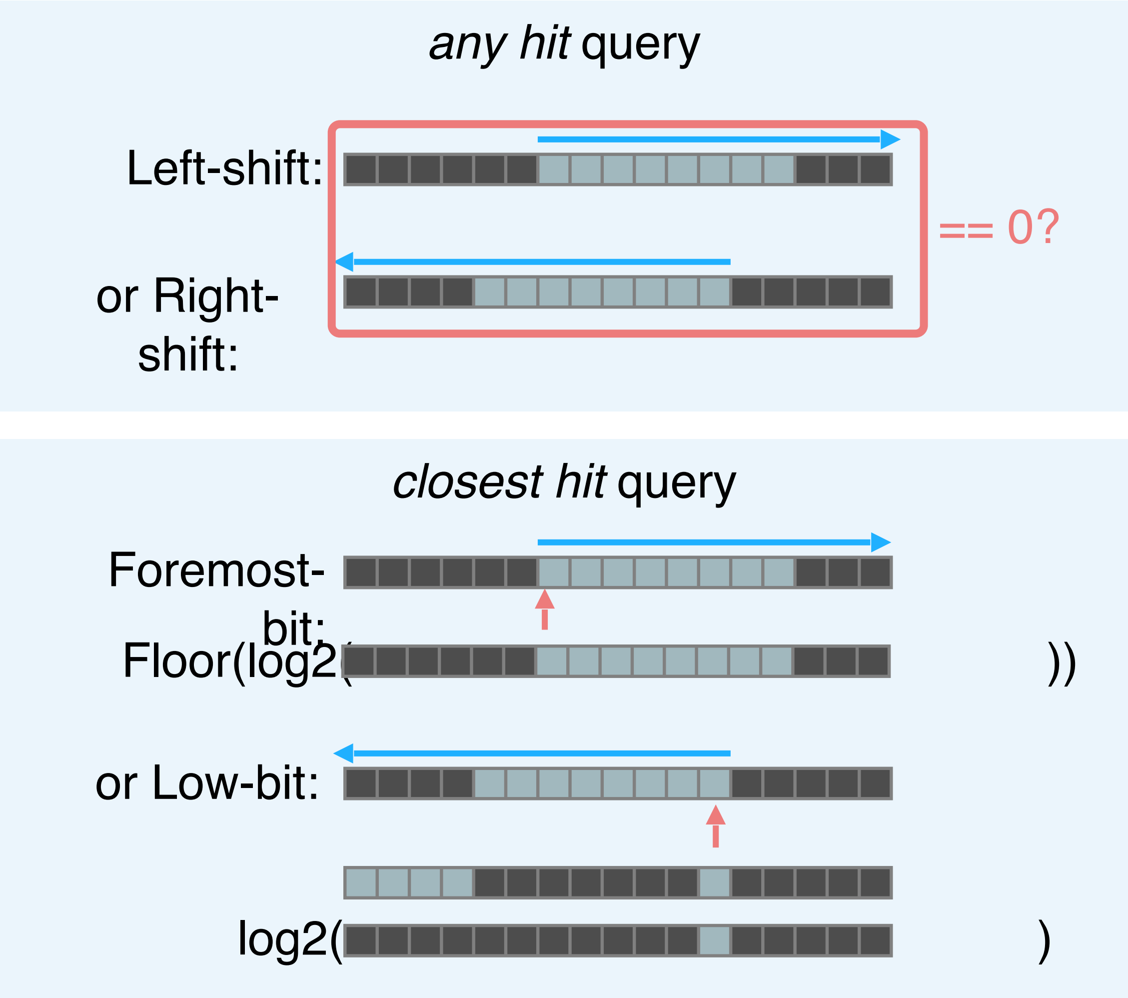

Spatiotemporal scheme

- Higher positional and angular resolutions are critical for ROMA.
- However, the extra time and space costs are always unacceptable.

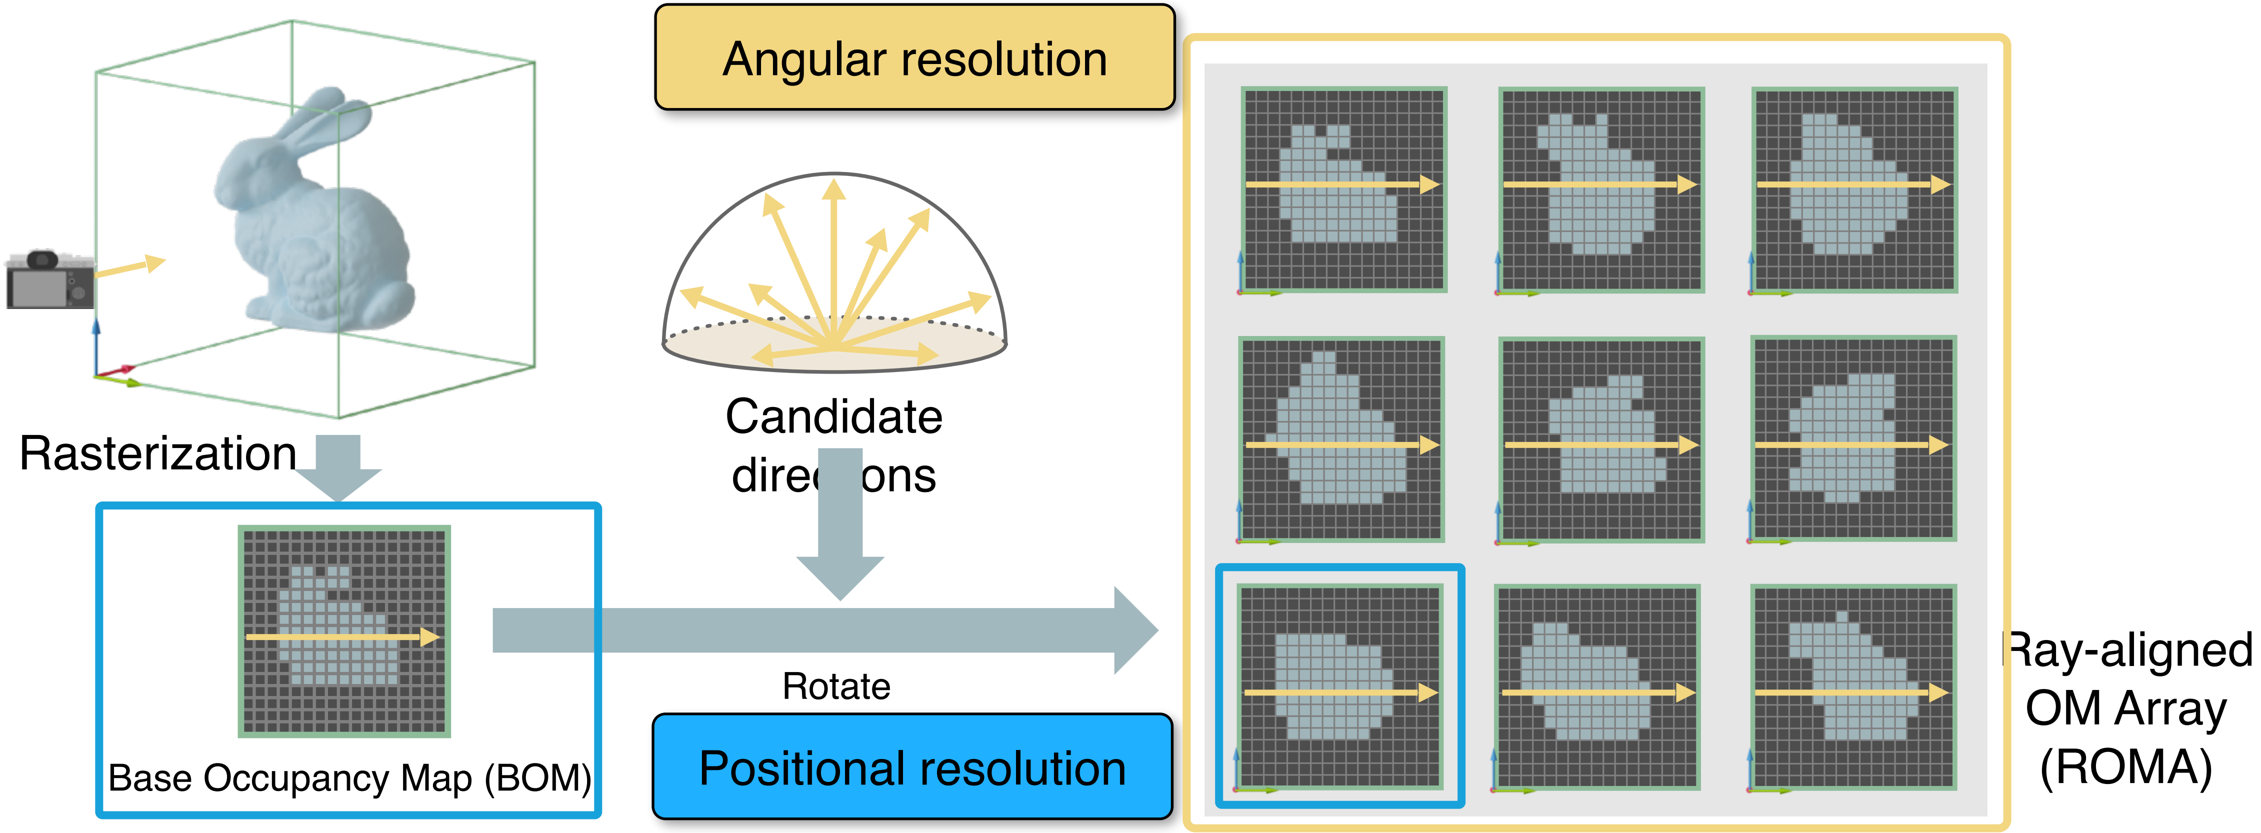

Spatiotemporal scheme

- A spatiotemporal scheme to further boost performance while alleviating aliasing!

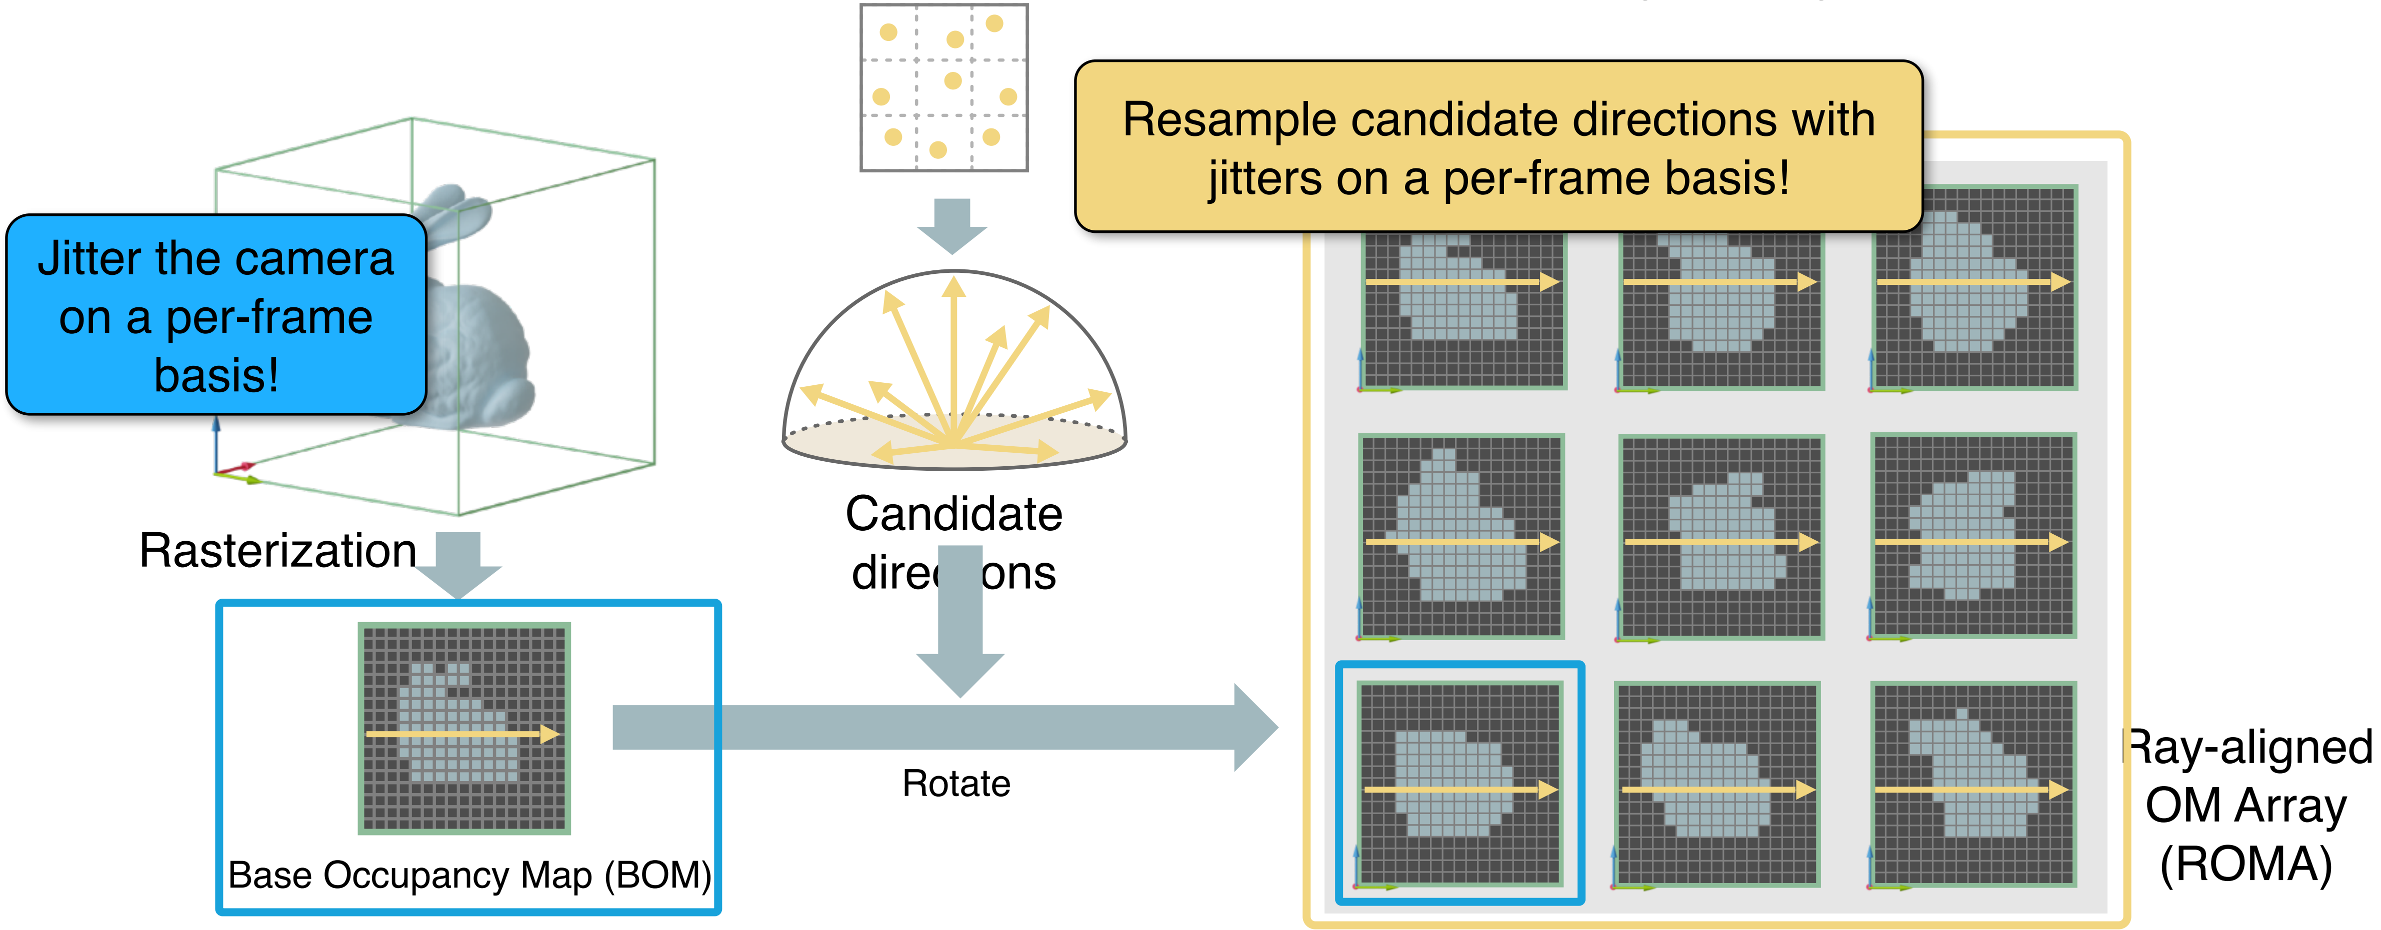

Spatiotemporal scheme

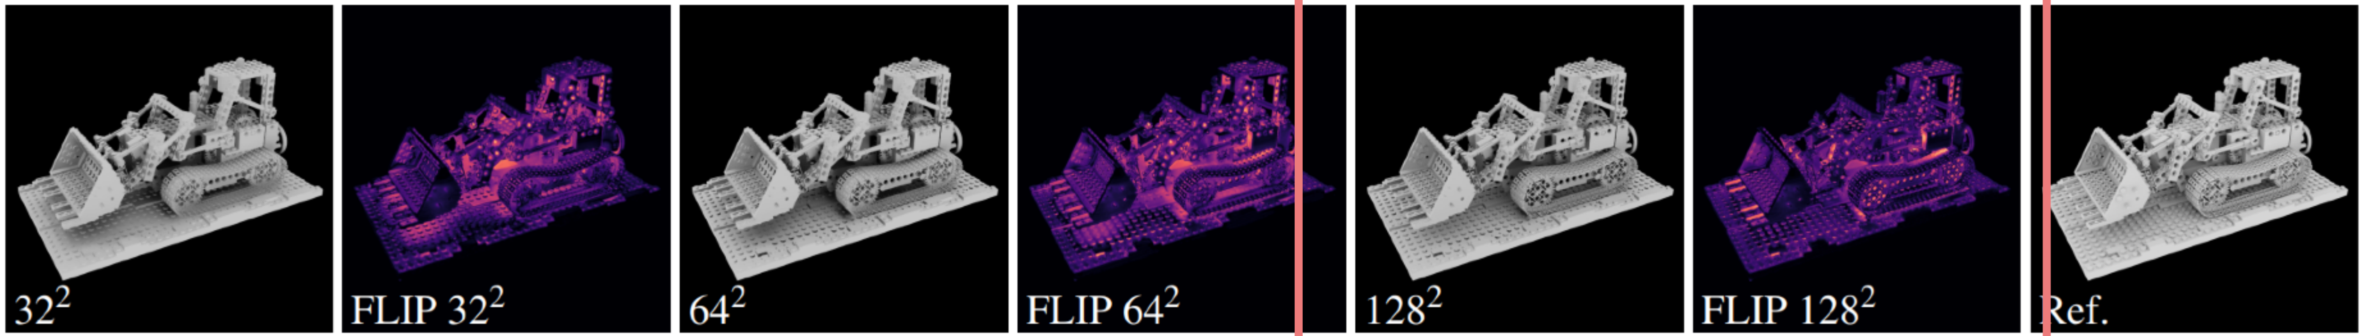

(a) *Different positional resolution (with a same angular resolution of* 4^2*).*

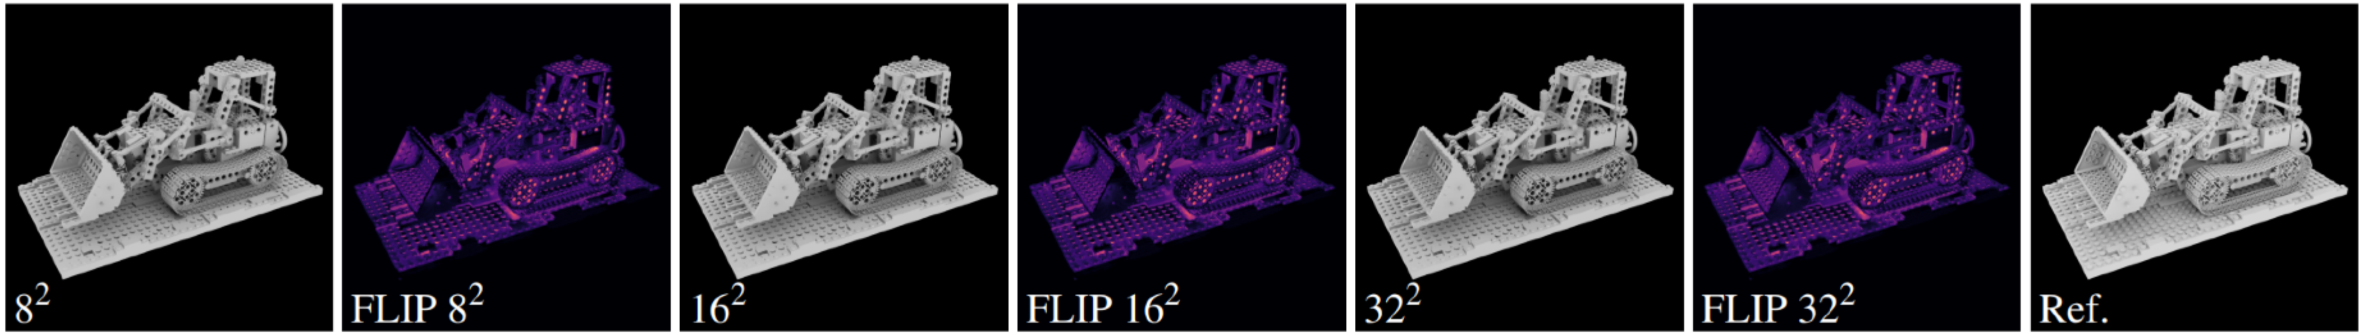

(b) *Different angular resolution (with a same positional resolution of* 128^2*).*

“Snap” or not

- Sometimes, we cannot simply snap the ray

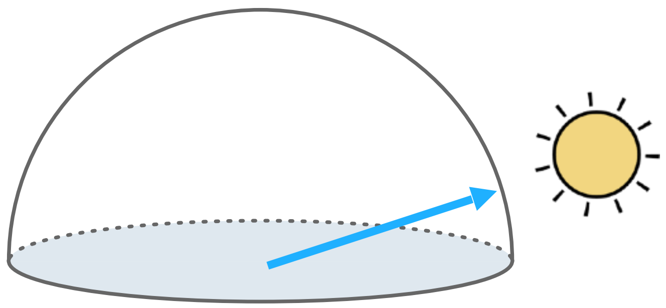

Sampled direction

“Snap” or not

- Sometimes, we cannot simply snap the ray

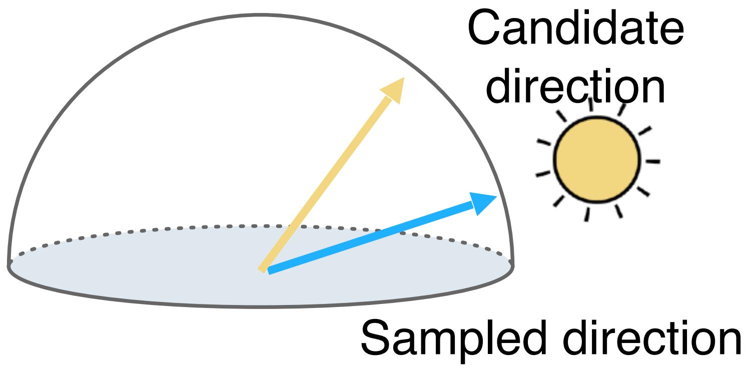

“Snap” or not

- Sometimes, we cannot simply snap the ray

Sampled rays’ direction	Applications	Support	“Snap”
Entire hemisphere	Ambient occlusion Diffuse reflection	✓	✓

“Snap” or not

- Sometimes, we cannot simply snap the ray

Sampled rays’ direction	Applications	Support	“Snap”
Entire hemisphere	Ambient occlusion Diffuse reflection	✓	✓

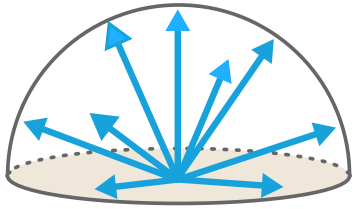

Entire hemisphere

“Snap” or not

- Sometimes, we cannot simply snap the ray

Sampled rays’ direction	Applications	Support	“Snap”
Entire hemisphere	Ambient occlusion Diffuse reflection	✓	✓
Specific directions	Soft shadows Glossy reflection	✓	✗

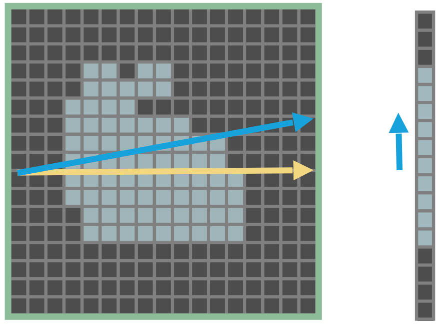

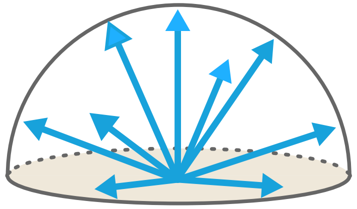

Entire hemisphere

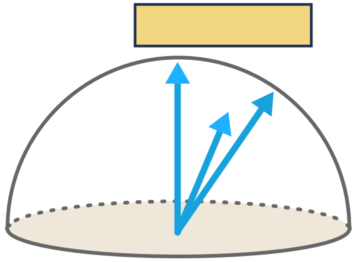

Specific directions

“Snap” or not

- Sometimes, we cannot simply snap the ray

Sampled rays’ direction	Applications	Support	“Snap”
Entire hemisphere	Ambient occlusion Diffuse reflection	✓	✓
Specific directions	Soft shadows Glossy reflection	✓	✗
One specific direction	Hard shadows Pure specular reflection	✗	N/A

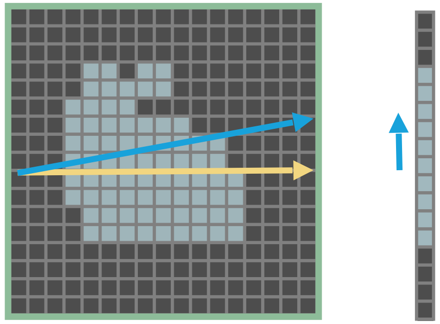

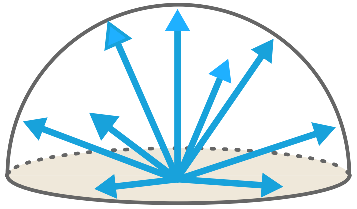

Entire hemisphere

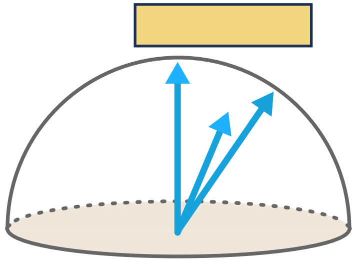

Specific directions

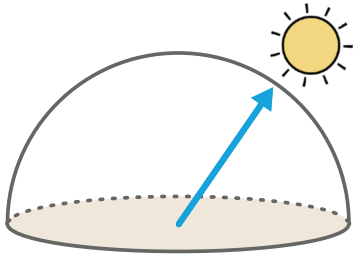

One Specific direction

“Snap” or not

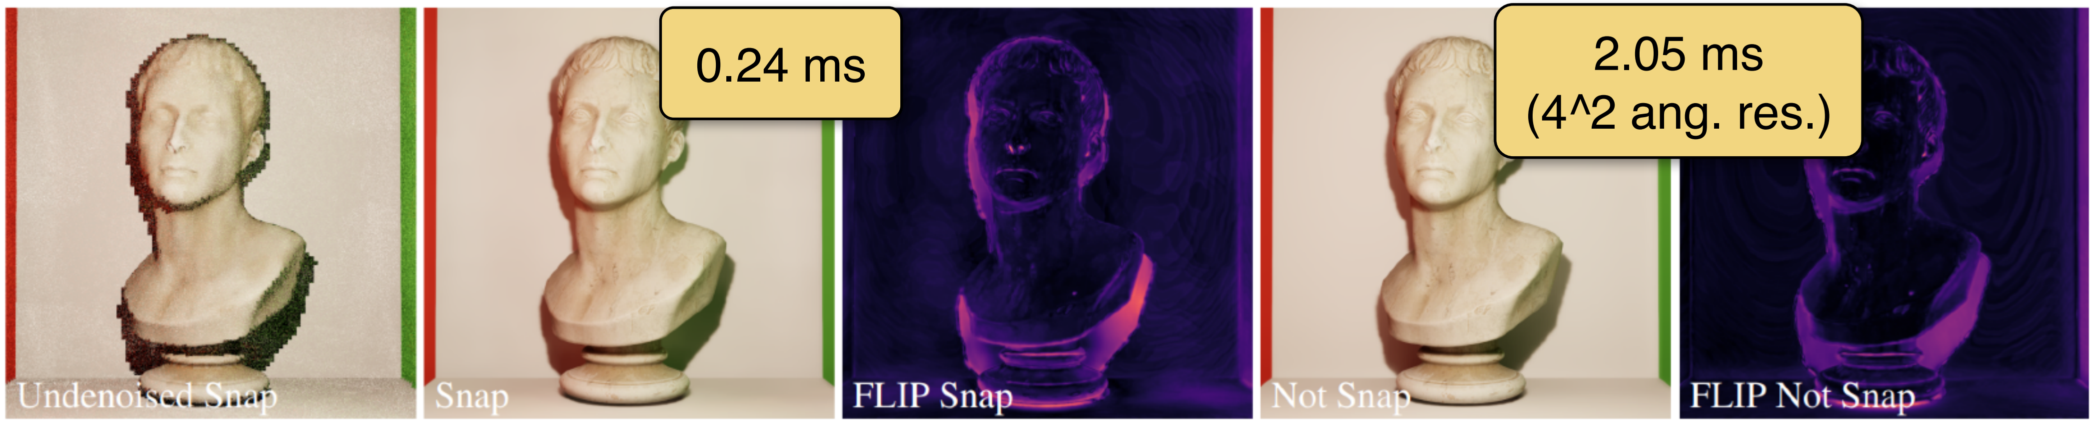

(a) MARBLE *scene. We show both direct illumination from RSM and one-bounce diffuse indirect illumination traced using ROMA.*

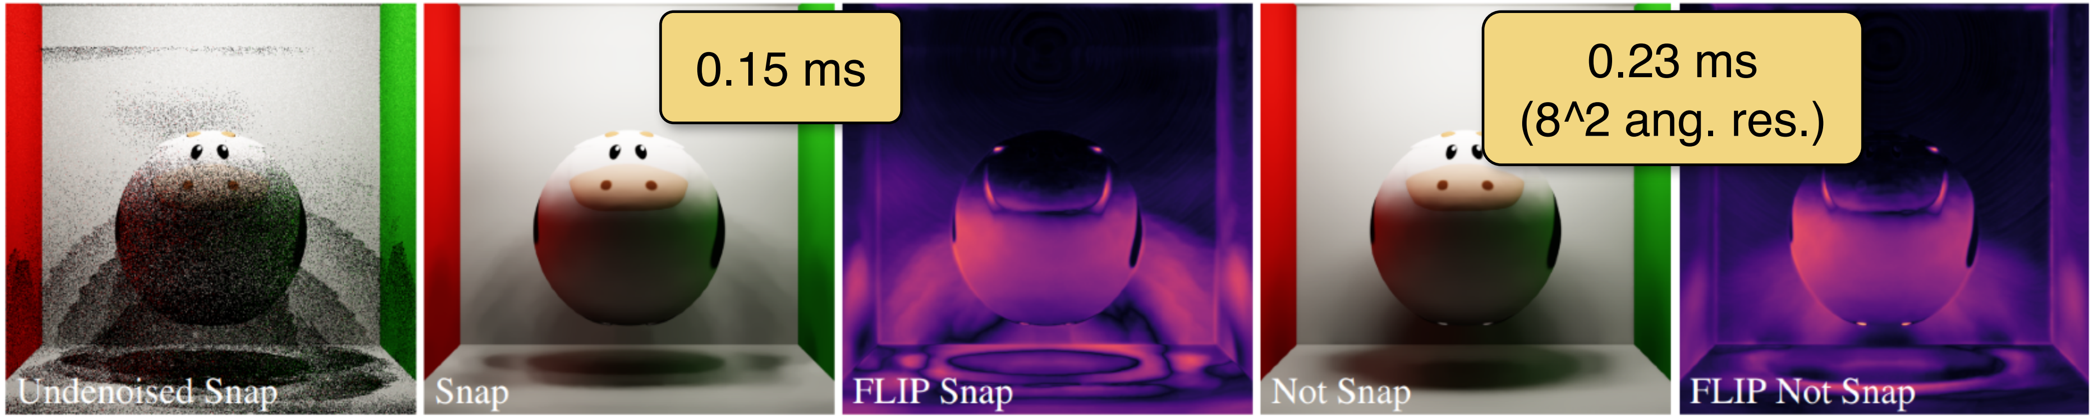

(b) MORPHING SPOT *scene. We show both direct illumination with soft shadow and one-bounce diffuse indirect illumination traced using ROMA.*

Result

Ambient Occlusion (AO)

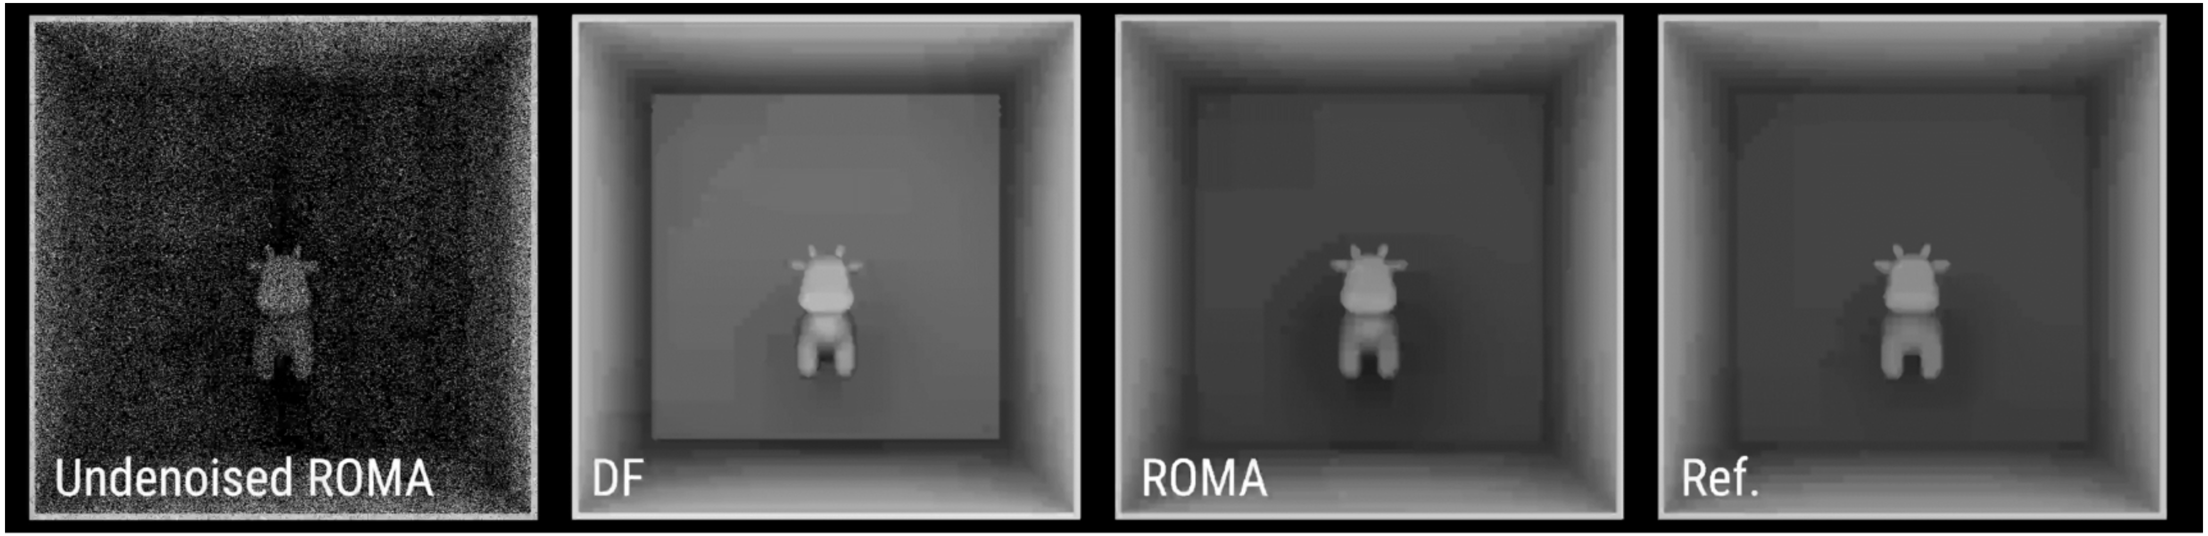

Performance

- The times are measured on 1080P for ambient occlusion (snapped rays).
- **Pos. resolution**: resolution of OMs; **ang. resolution**: the number of candidate directions.

Pos. Res.	32^2		64^2		128^2	
Ang. Res.	4^2	8^2	4^2	8^2	4^2	8^2
GENERATION						
ROMA	0.14 ms	0.18 ms	0.20 ms	0.31 ms	0.30 ms	0.69 ms
Distance field	0.31 ms		0.55 ms		3.31 ms	
(Speed-up)	(2.2×)	(1.7×)	(2.7×)	(1.7×)	(11.0×)	(4.8×)
(Storage)	(1×)	(4×)	(1×)	(4×)	(1×)	(4×)
HWRT	0.08 ms					

Performance

- The times are measured on 1080P for ambient occlusion (snapped rays).
- **Pos. resolution**: resolution of OMs; **ang. resolution**: the number of candidate directions.

Pos. Res.	32^2		64^2		128^2	
Ang. Res.	4^2	8^2	4^2	8^2	4^2	8^2
GENERATION						
ROMA	0.14 ms	0.18 ms	0.20 ms	0.31 ms	0.30 ms	0.69 ms
Distance field	0.31 ms		0.55 ms		3.31 ms	
(Speed-up)	(2.2×)	(1.7×)	(2.7×)	(1.7×)	(11.0×)	(4.8×)
(Storage)	(1×)	(4×)	(1×)	(4×)	(1×)	(4×)
HWRT	0.08 ms					

Always < 1ms!

Performance

- The times are measured on 1080P for ambient occlusion (snapped rays).
- **Pos. resolution**: resolution of OMs; **ang. resolution**: the number of candidate directions.

Pos. Res.	32^2		64^2		128^2	
Ang. Res.	4^2	8^2	4^2	8^2	4^2	8^2
GENERATION						
ROMA	0.14 ms	0.18 ms	0.20 ms	0.31 ms	0.30 ms	0.69 ms
Distance field	0.31 ms		0.55 ms		3.31 ms	
(Speed-up)	(2.2×)	(1.7×)	(2.7×)	(1.7×)	(11.0×)	(4.8×)
(Storage)	(1×)	(4×)	(1×)	(4×)	(1×)	(4×)
HWRT	0.08 ms					

Always < 1ms!

11x faster

Performance

- The times are measured on 1080P for ambient occlusion (snapped rays).
- **Pos. resolution**: resolution of OMs; **ang. resolution**: the number of candidate directions.

Pos. Res.	32^2		64^2		128^2	
Ang. Res.	4^2	8^2	4^2	8^2	4^2	8^2
GENERATION						
ROMA	0.14 ms	0.18 ms	0.20 ms	0.31 ms	0.30 ms	0.69 ms
Distance field	0.31 ms		0.55 ms		3.31 ms	
(Speed-up)	(2.2×)	(1.7×)	(2.7×)	(1.7×)	(11.0×)	(4.8×)
(Storage)	(1×)	(4×)	(1×)	(4×)	(1×)	(4×)
HWRT	0.08 ms					

Always < 1ms!

1.7x-11.0x faster

11x faster

Performance

- The times are measured on 1080P for ambient occlusion (snapped rays).
- **Pos. resolution**: resolution of OMs; **ang. resolution**: the number of candidate directions.

Pos. Res.	32^2		64^2		128^2	
Ang. Res.	4^2	8^2	4^2	8^2	4^2	8^2
GENERATION						
ROMA	0.14 ms	0.18 ms	0.20 ms	0.31 ms	0.30 ms	0.69 ms
Distance field	0.31 ms		0.55 ms		3.31 ms	
(Speed-up)	(2.2×)	(1.7×)	(2.7×)	(1.7×)	(11.0×)	(4.8×)
(Storage)	(1×)	(4×)	(1×)	(4×)	(1×)	(4×)
HWRT	0.08 ms					
TRACING (1 sample per pixel)						
ROMA	0.16 ms	0.16 ms	0.16 ms	0.16 ms	0.16 ms	0.16 ms
Distance field	0.55 ms		0.84 ms		1.30 ms	
(Speed-up)	(3.4×)		(5.25×)		(8.1×)	
HWRT	0.31 ms					
(Speed-up)	(1.9×)					

Always < 1ms!

1.7x-11.0x faster

11x faster

Aways 0.16ms!

Performance

- The times are measured on 1080P for ambient occlusion (snapped rays).
- **Pos. resolution**: resolution of OMs; **ang. resolution**: the number of candidate directions.

Pos. Res.	32^2		64^2		128^2	
Ang. Res.	4^2	8^2	4^2	8^2	4^2	8^2
GENERATION						
ROMA	0.14 ms	0.18 ms	0.20 ms	0.31 ms	0.30 ms	0.69 ms
Distance field	0.31 ms		0.55 ms		3.31 ms	
(Speed-up)	(2.2×)	(1.7×)	(2.7×)	(1.7×)	(11.0×)	(4.8×)
(Storage)	(1×)	(4×)	(1×)	(4×)	(1×)	(4×)
HWRT	0.08 ms					
TRACING (1 sample per pixel)						
ROMA	0.16 ms	0.16 ms	0.16 ms	0.16 ms	0.16 ms	0.16 ms
Distance field	0.55 ms		0.84 ms		1.30 ms	
(Speed-up)	(3.4×)		(5.25×)		(8.1×)	
HWRT	0.31 ms					
(Speed-up)	(1.9×)					

Always < 1ms!

1.7x-11.0x faster

11x faster

Aways 0.16ms!

3.4x-8.1x faster

Performance

- The times are measured on 1080P for ambient occlusion (snapped rays).
- **Pos. resolution**: resolution of OMs; **ang. resolution**: the number of candidate directions.

Pos. Res. / Ang. Res.	32^2 / 4^2	32^2 / 8^2	64^2 / 4^2	64^2 / 8^2	128^2 / 4^2	128^2 / 8^2
GENERATION						
ROMA	0.14 ms	0.18 ms	0.20 ms	0.31 ms	0.30 ms	0.69 ms
Distance field	0.31 ms		0.55 ms		3.31 ms	
(Speed-up)	(2.2×)	(1.7×)	(2.7×)	(1.7×)	(11.0×)	(4.8×)
(Storage)	(1×)	(4×)	(1×)	(4×)	(1×)	(4×)
HWRT	0.08 ms					
TRACING (1 sample per pixel)						
ROMA	0.16 ms	0.16 ms	0.16 ms	0.16 ms	0.16 ms	0.16 ms
Distance field	0.55 ms		0.84 ms		1.30 ms	
(Speed-up)	(3.4×)		(5.25×)		(8.1×)	
HWRT	0.31 ms					
(Speed-up)	(1.9×)					
TOTAL						
ROMA	0.30 ms	0.34 ms	0.36 ms	0.47 ms	0.46 ms	0.85 ms
Distance field	0.86 ms		1.39 ms		4.61 ms	
(Speed-up)	(2.9×)	(2.5×)	(3.9×)	(3.0×)	(10.0×)	(5.4×)

Always < 1ms!

11x faster

1.7x-11.0x faster

Aways 0.16ms!

3.4x-8.1x faster

2.5x-10x faster

Performance

- The times are measured on 1080P for ambient occlusion (snapped rays).
- **Pos. resolution**: resolution of OMs; **ang. resolution**: the number of candidate directions.

Pos. Res.	32^2		64^2		128^2	
Ang. Res.	4^2	8^2	4^2	8^2	4^2	8^2
GENERATION						
ROMA	0.14 ms	0.18 ms	0.20 ms	0.31 ms	0.30 ms	0.69 ms
Distance field	0.31 ms		0.55 ms		3.31 ms	
(Speed-up)	(2.2×)	(1.7×)	(2.7×)	(1.7×)	(11.0×)	(4.8×)
(Storage)	(1×)	(4×)	(1×)	(4×)	(1×)	(4×)
HWRT	0.08 ms					
TRACING (1 sample per pixel)						
ROMA	0.16 ms	0.16 ms	0.16 ms	0.16 ms	0.16 ms	0.16 ms
Distance field	0.55 ms		0.84 ms		1.30 ms	
(Speed-up)	(3.4×)		(5.25×)		(8.1×)	
HWRT	0.31 ms					
(Speed-up)	(1.9×)					
TOTAL						
ROMA	0.30 ms	0.34 ms	0.36 ms	0.47 ms	0.46 ms	0.85 ms
Distance field	0.86 ms		1.39 ms		4.61 ms	
(Speed-up)	(2.9×)	(2.5×)	(3.9×)	(3.0×)	(10.0×)	(5.4×)

Always < 1ms!
1.7x-11.0x faster
11x faster
Aways 0.16ms!
3.4x-8.1x faster
1.9x faster
2.5x-10x faster

One-bounce Diffuse GI & Soft Shadows

- We want to show:
  - Soft shadows from direct illumination (un-snapped rays)
  - Color bleedings from indirect illumination (snapped rays)

Generation: ~2.86 ms (3.2x)
Tracing: ~0.90 ms (2.0x)

Generation: ~0.89 ms
Tracing: ~0.45 ms

One-bounce Diffuse GI & Soft Shadows

- We want to show:
  - Soft shadows from direct illumination (un-snapped rays)
  - Color bleedings from indirect illumination (snapped rays)

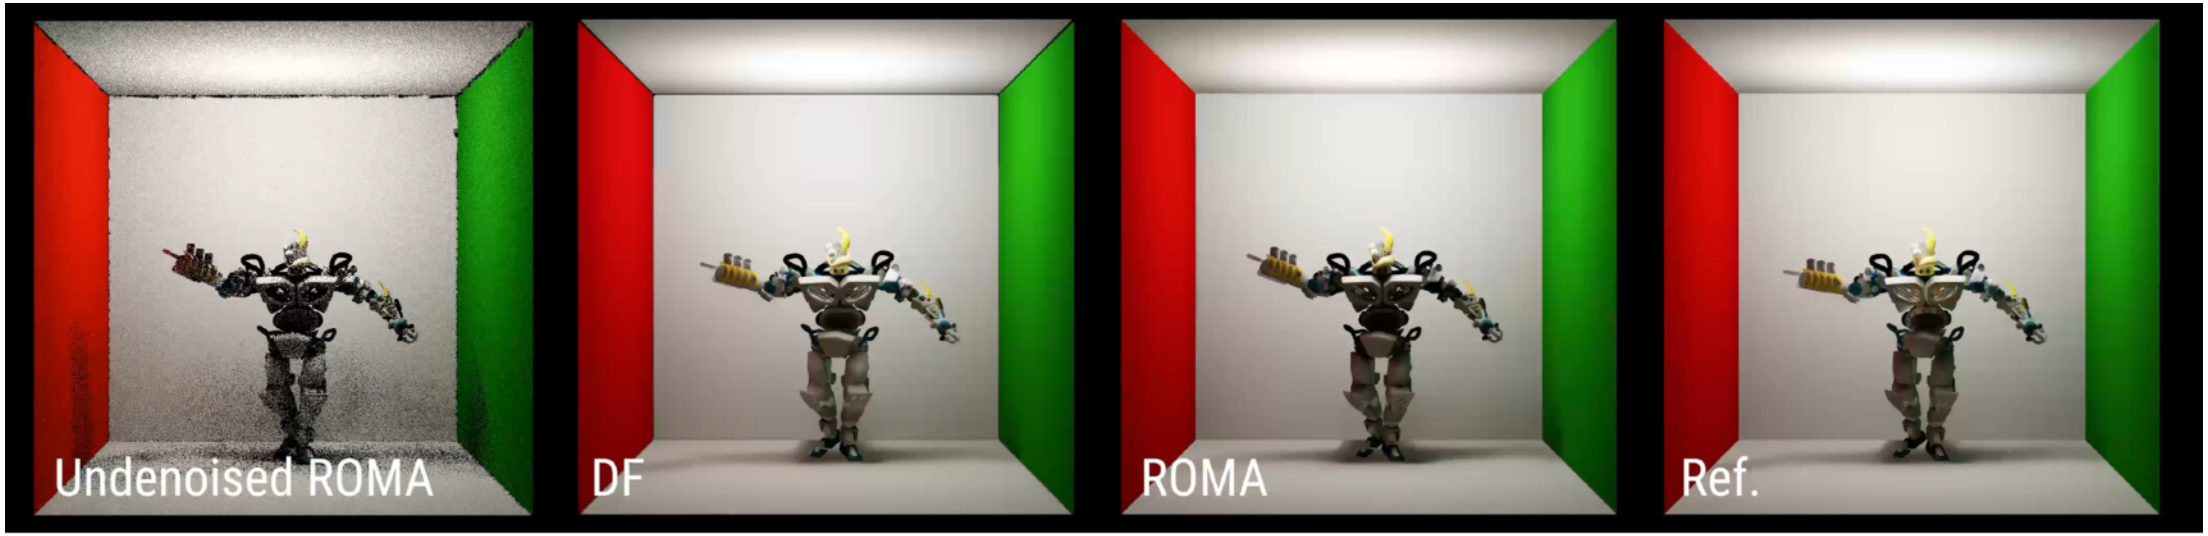

Generation: ~2.86 ms (3.2x)
Tracing: ~0.90 ms (2.0x)

Generation: ~0.89 ms
Tracing: ~0.45 ms

One-bounce Diffuse GI & Soft Shadows

- We want to show:
    - Soft shadows from direct illumination (un-snapped rays)
    - Color bleedings from indirect illumination (snapped rays)

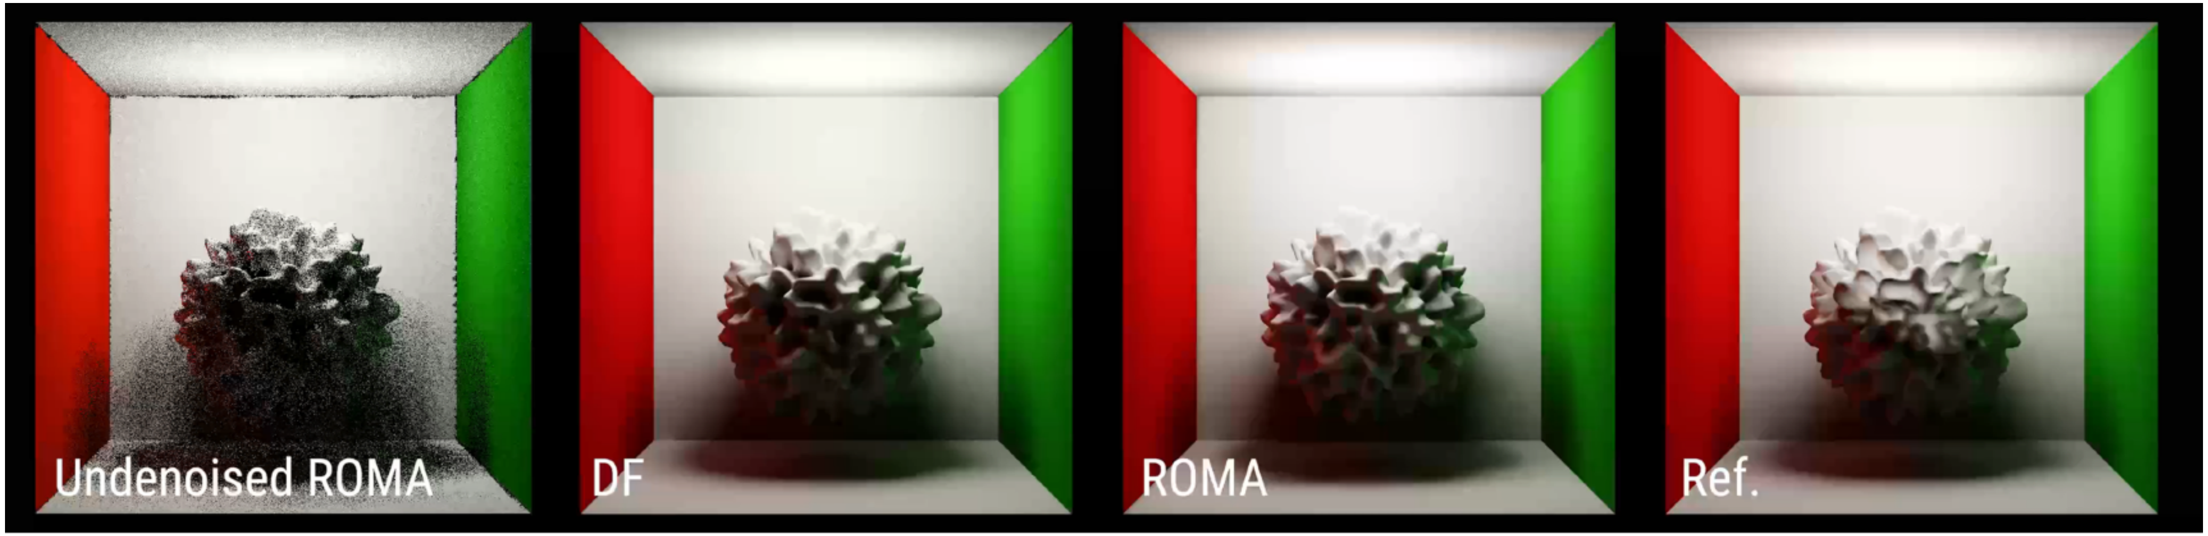

Generation: ~2.86 ms (3.2x)
Tracing: ~0.90 ms (2.0x)

Generation: ~0.89 ms
Tracing: ~0.45 ms

Conclusion

Conclusion

- A new SWRT solution, ROMA, that enables fast approximate ray tracing.

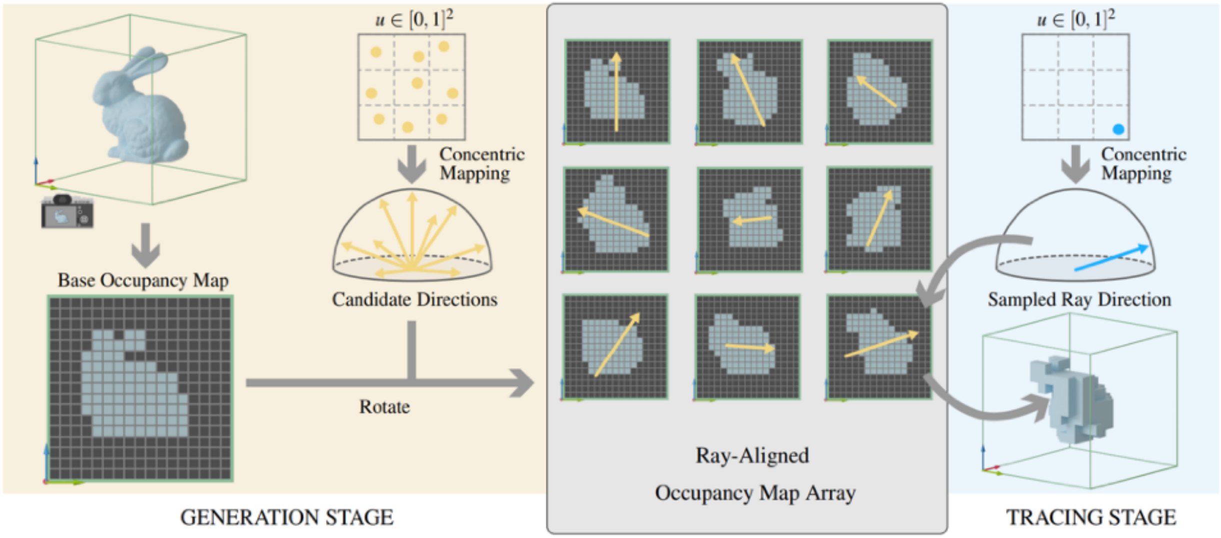

- Fast: generation & tracing.
- Approximate: scene voxelization & tracing snapped rays.
- Fully scalable: balancing the performance and quality in a spatiotemporal way.

Limitations

- ROMA generation is slower than the BVH construction/update for HWRT on small-scale scenes.
  - ROMA generation does not have specially optimized refitting method and is not handled by drivers.
- ROMA does not have these good properties as DFs have.
  - Filtered to get coarser levels.
  - Differentiable rendering.
- Shadow acne / light leaking.

Future work

- Extend ROMA to support larger scenes.
    - Cascades.
    - Local ROMAs for near-field tracing and global ROMA for far-field tracing. (Lumen)
    - Local mesh ROMAs coupled with AABB tree traversal. (AMD’s Brixelizer)
- Hardware support for ROMA to boost its performance.
- Explore hybrid solutions
    - Combine screen-space ray tracing for near-field tracing and using ROMA only for far-field tracing.

Ray-aligned Occupancy Map Array for Fast Approximate Ray Tracing

Thank you!

ROMA wasn't built in a day, but in <1 millisecond :)