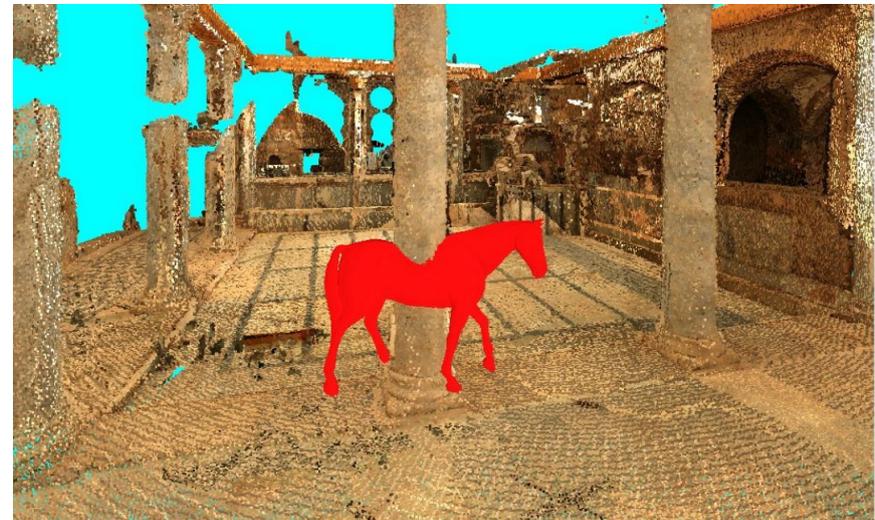
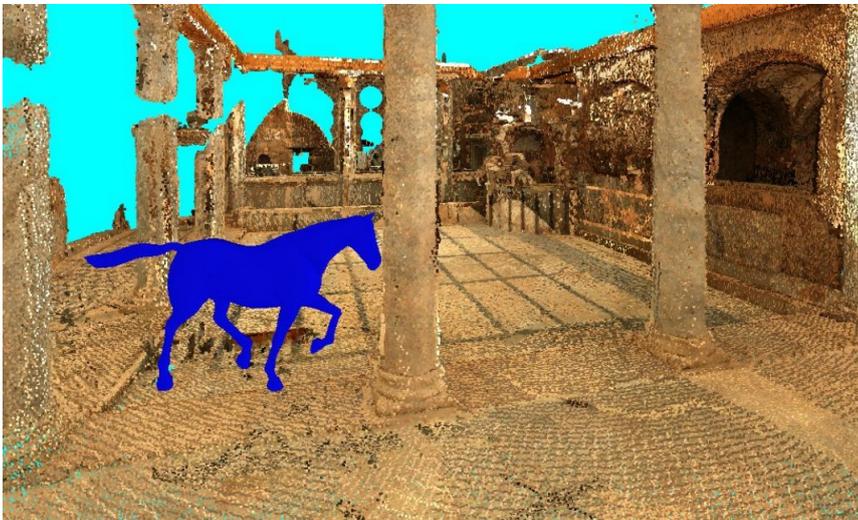


Efficient Collision Detection While Rendering Dynamic Point Clouds

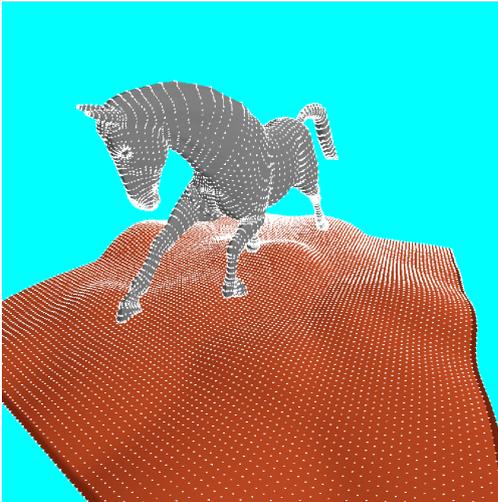


M. Radwan, S. Ohrhallinger and M. Wimmer

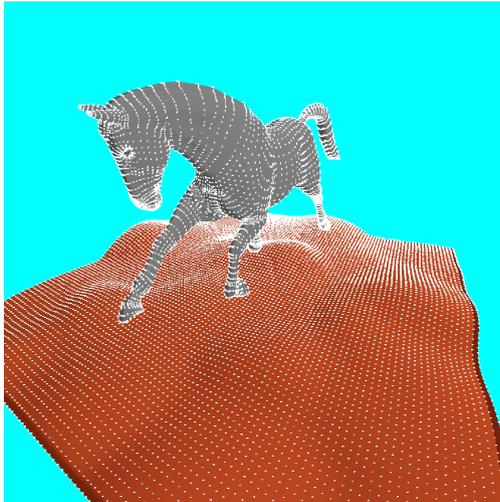
Vienna University of Technology, Austria



- Point clouds are queried using bounding hierarchy
- Construction: $O(N \log N)$, query time: $O(\log N)$

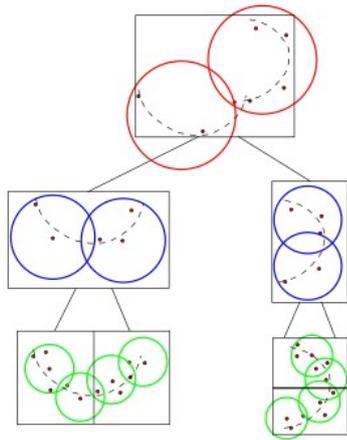


- Point clouds are queried using bounding hierarchy
- Construction: $O(N \log N)$, query time: $O(\log N)$



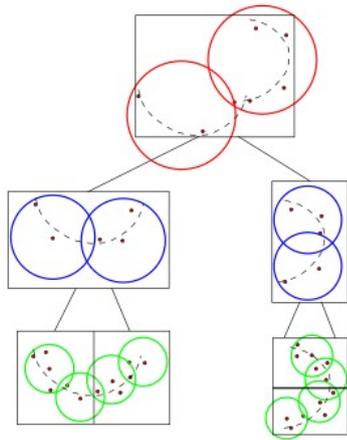
- Dynamic points without any time coherency:
per-frame construction → too ~~slow~~ for $N=10000000$



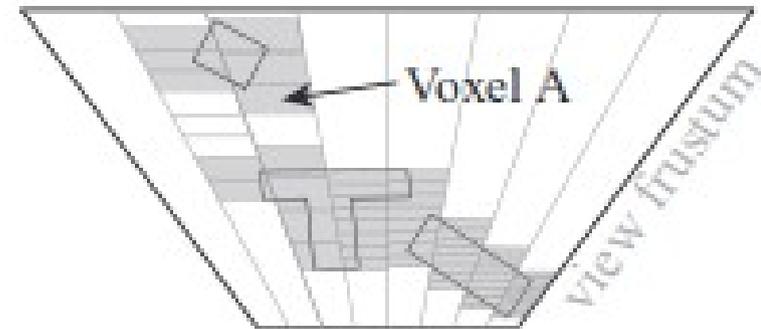


BVH [Klein et al '04]



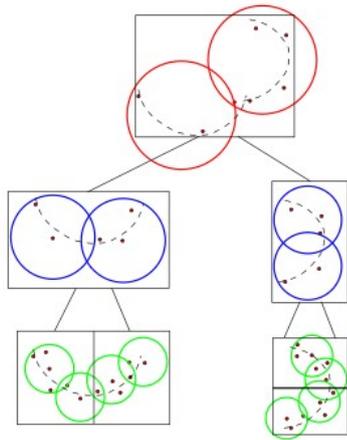


BVH [Klein et al '04]

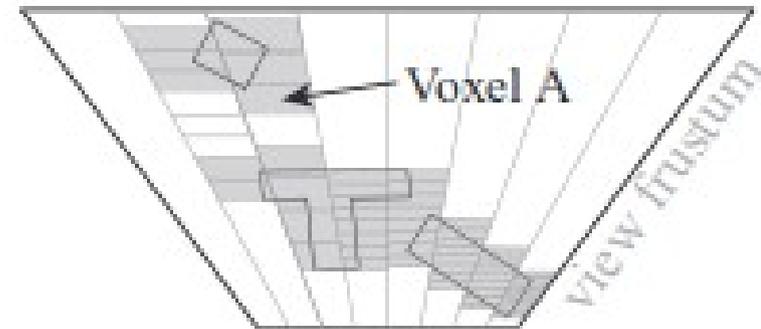


Voxels [Eisemann et al '06]

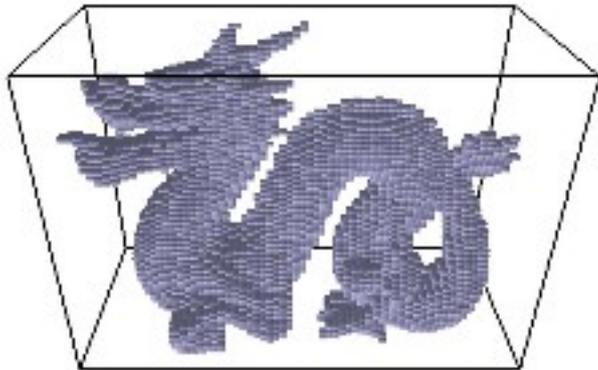




BVH [Klein et al '04]

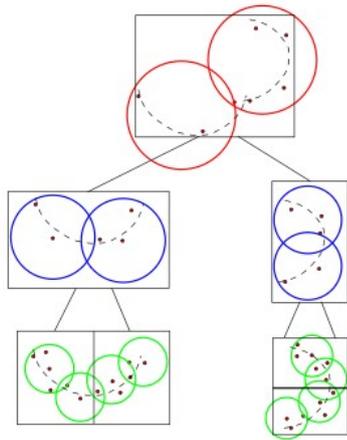


Voxels [Eisemann et al '06]

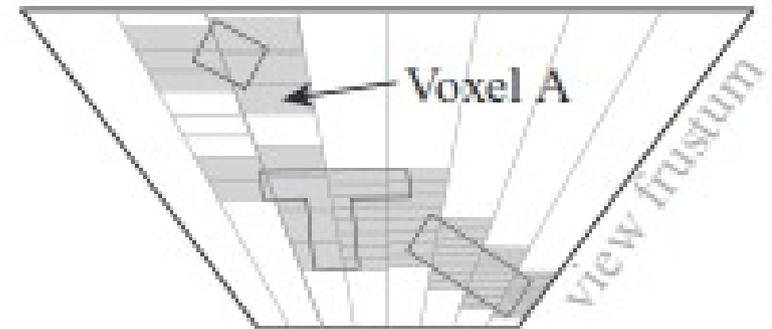


LDI [Heidelberger et al '04]

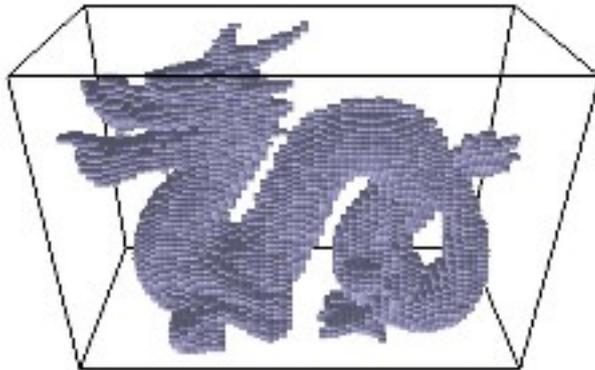




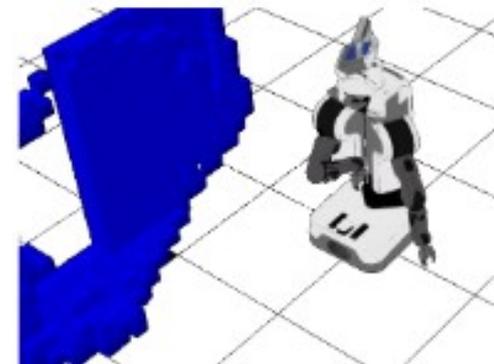
BVH [Klein et al '04]



Voxels [Eisemann et al '06]



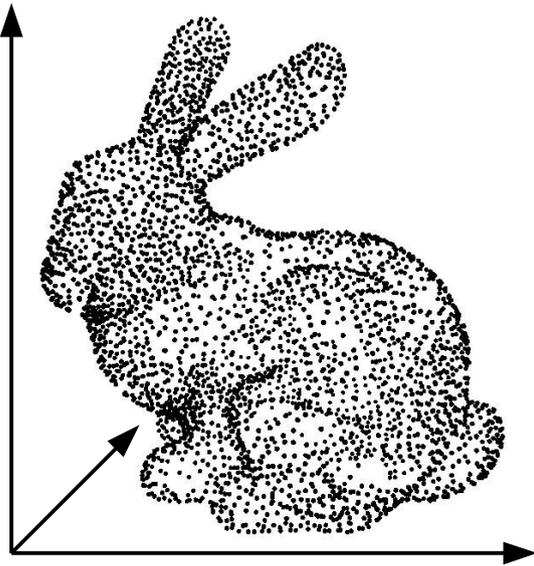
LDI [Heidelberger et al '04]



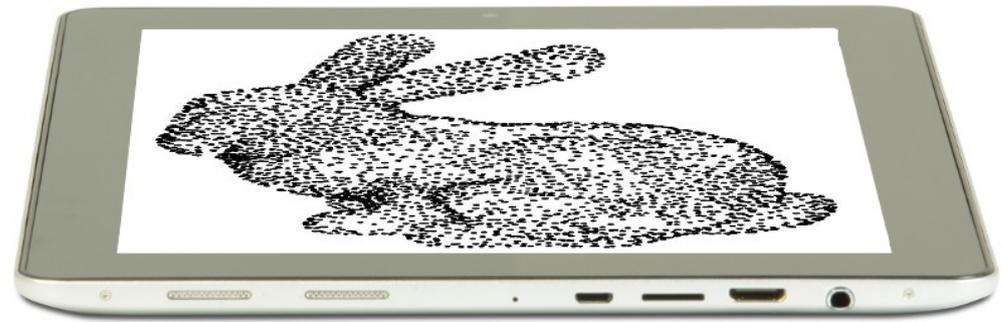
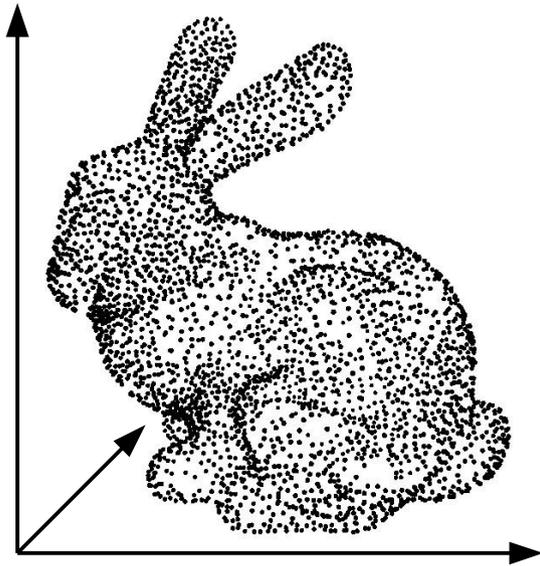
Dynamic [Pan et al '13]



- 3D point cloud is really sampled on 2D surface



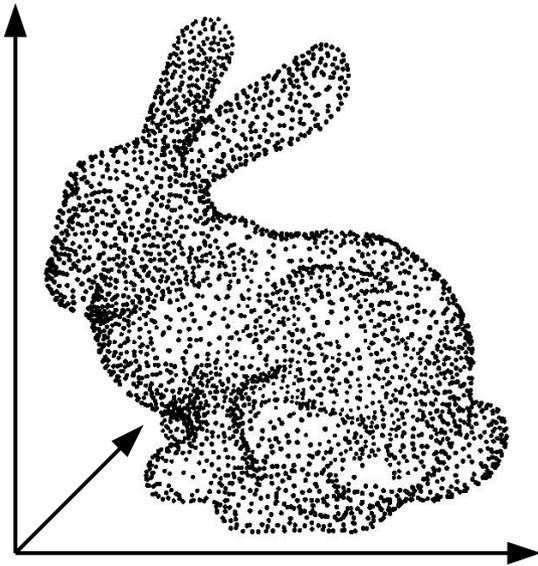
- 3D point cloud is really sampled on 2D surface



→ flatten to depth images (in screen space): $O(N)$



- 3D point cloud is really sampled on 2D surface



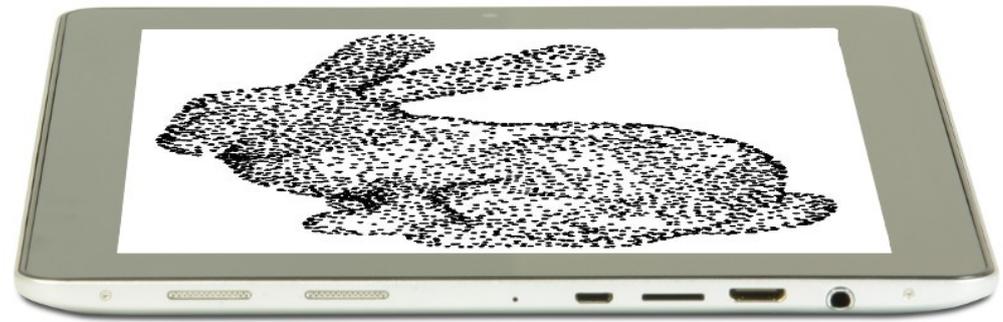
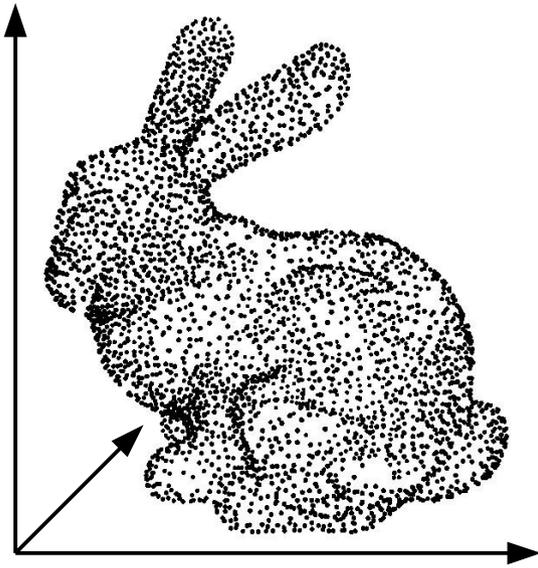
→ flatten to depth images (in screen space): $O(N)$



- Incidental benefits of our method:



- 3D point cloud is really sampled on 2D surface

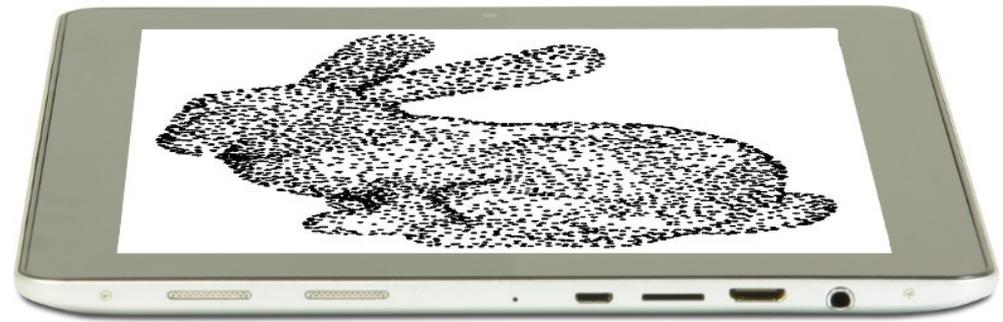
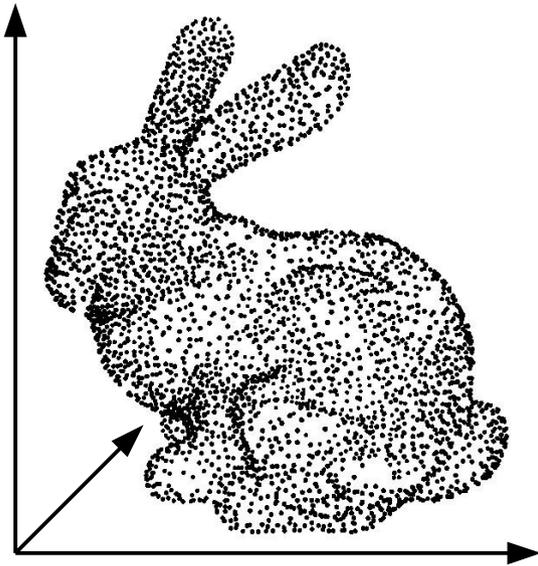


→ flatten to depth images (in screen space): $O(N)$ ✓

- Incidental benefits of our method:
Superior accuracy ✓



- 3D point cloud is really sampled on 2D surface



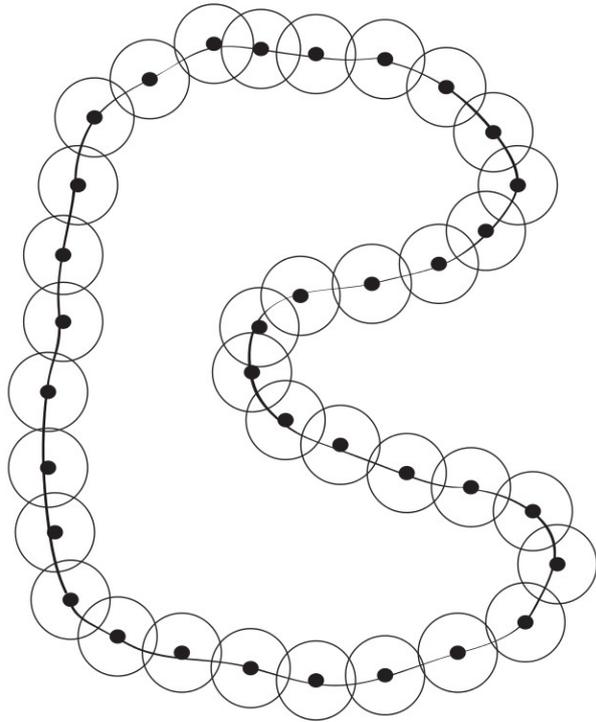
→ flatten to depth images (in screen space): $O(N)$

- Incidental benefits of our method:

Superior accuracy

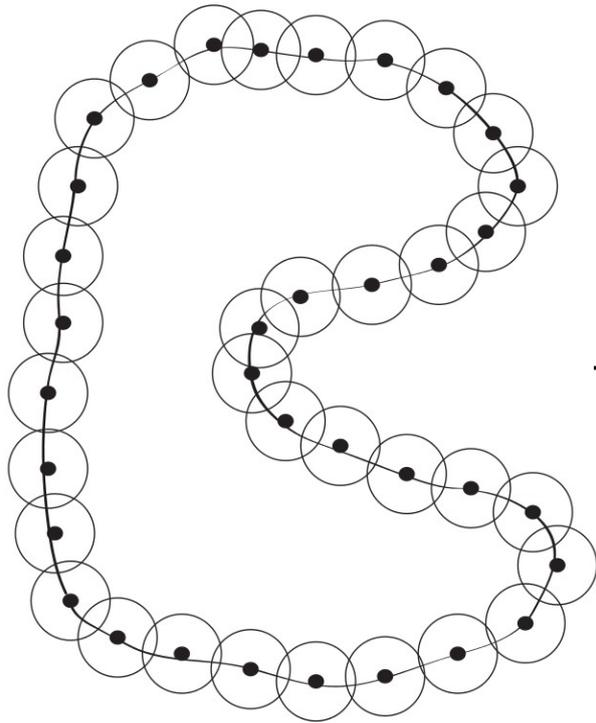
Robustness to sensor noise



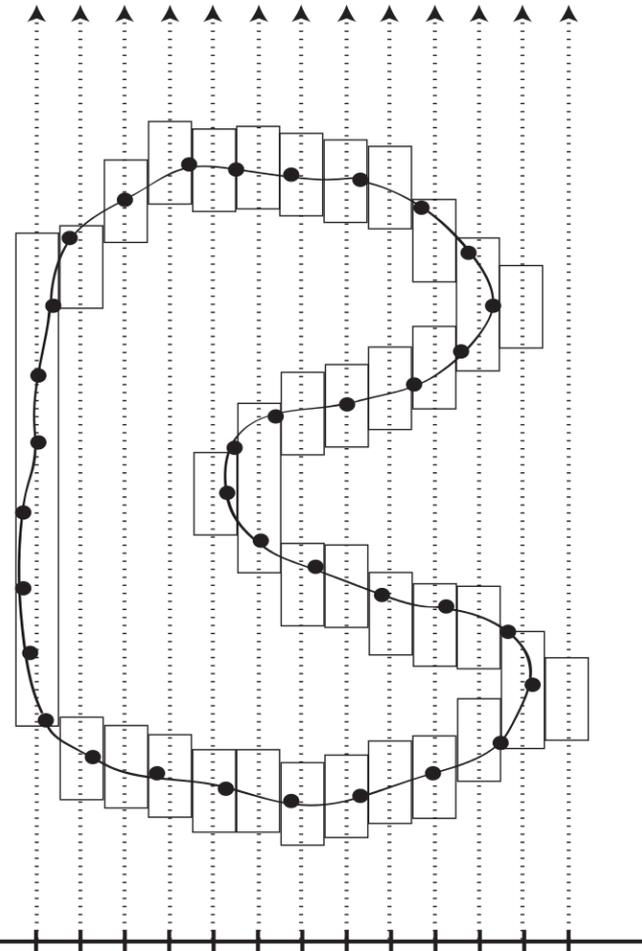
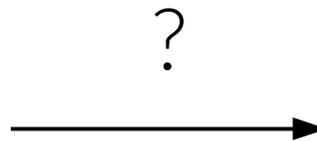


\mathbb{R}^3 : spherical cover





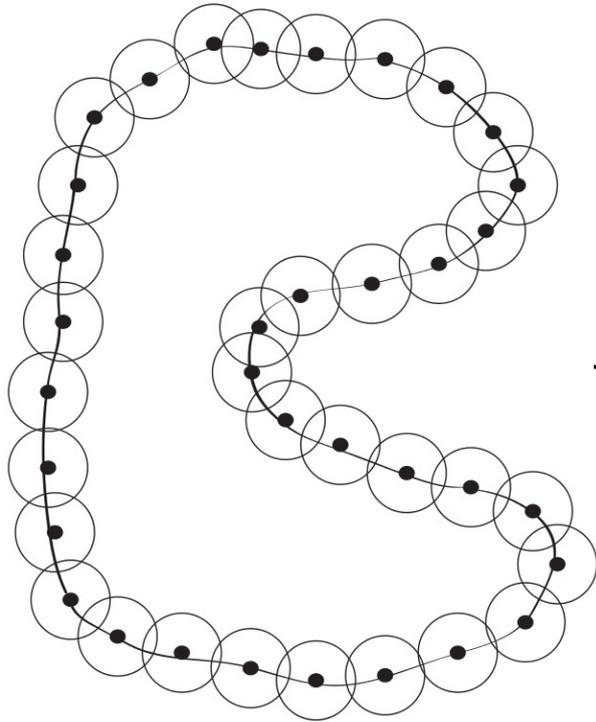
\mathbb{R}^3 : spherical cover



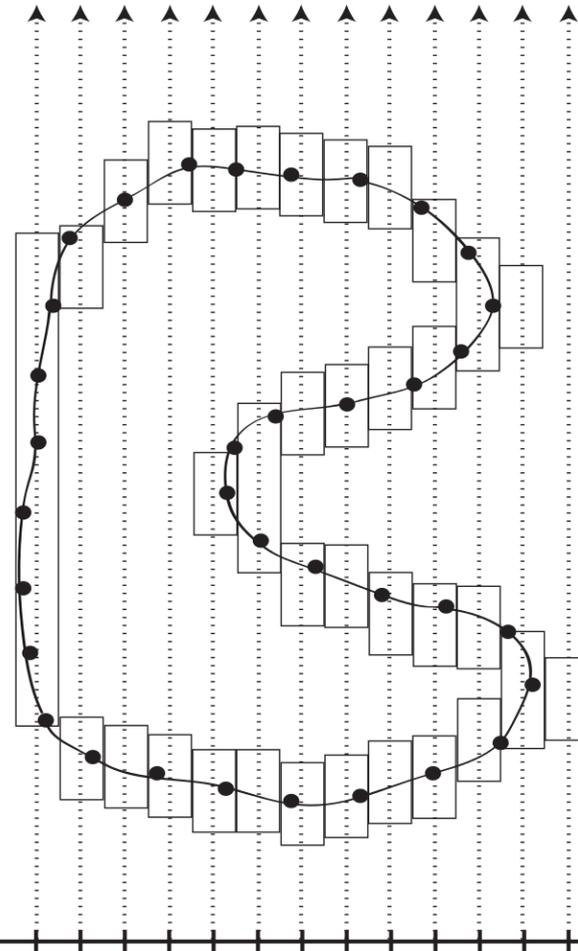
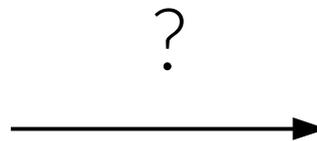
view ray depth intervals



inequal ~~boundary~~ thickness

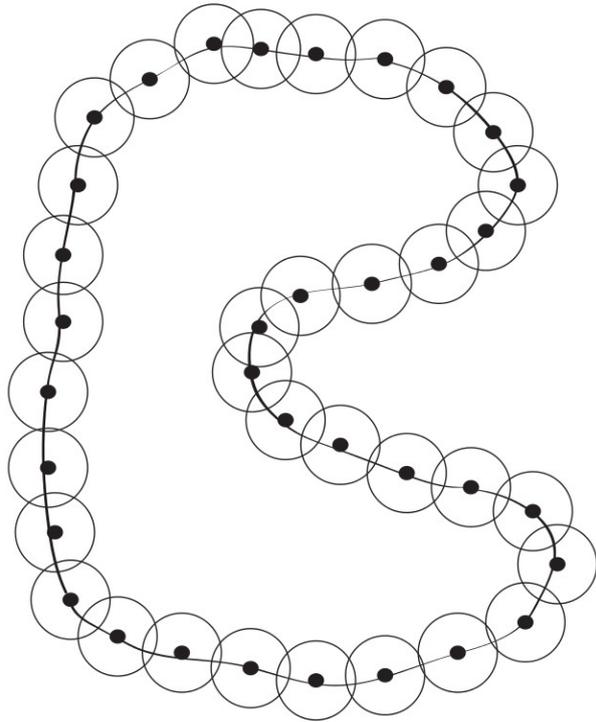


\mathbb{R}^3 : spherical cover

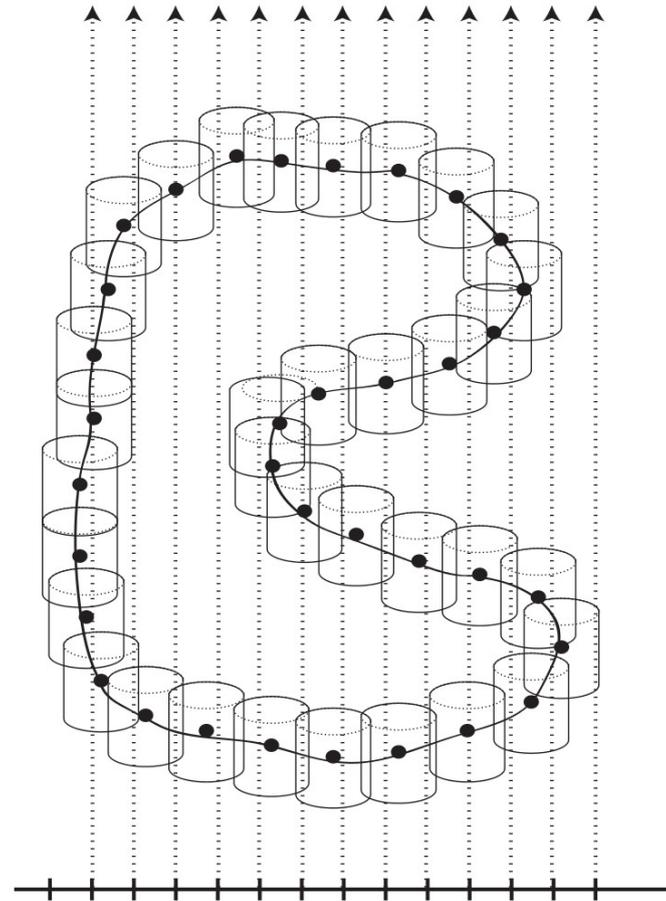


view ray depth intervals



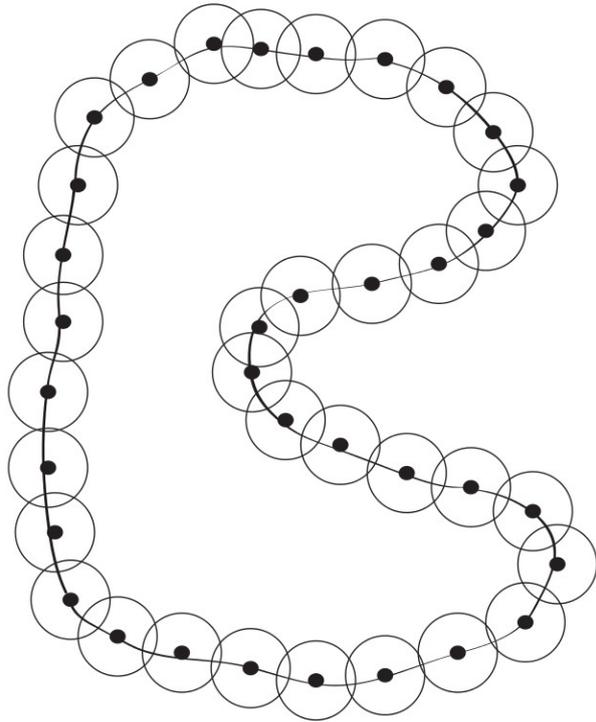


\mathbb{R}^3 : spherical cover

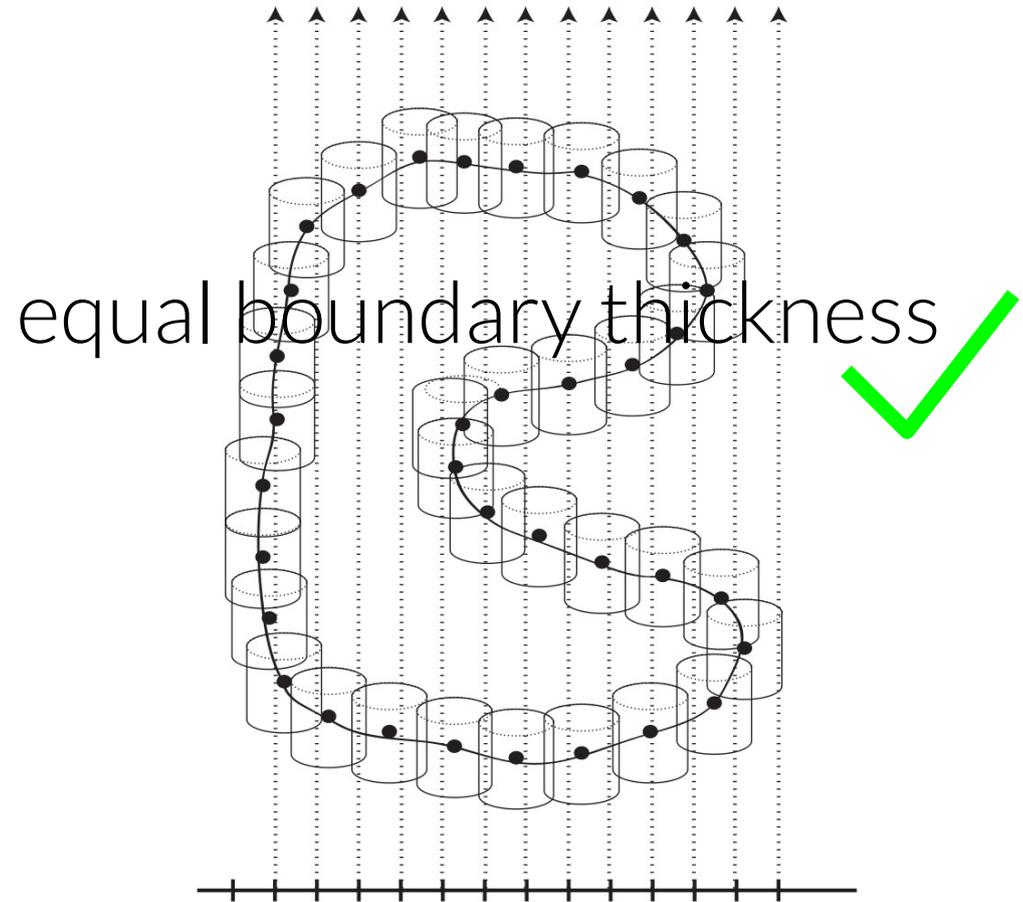


\mathbb{R}^3 : cylindrical cover



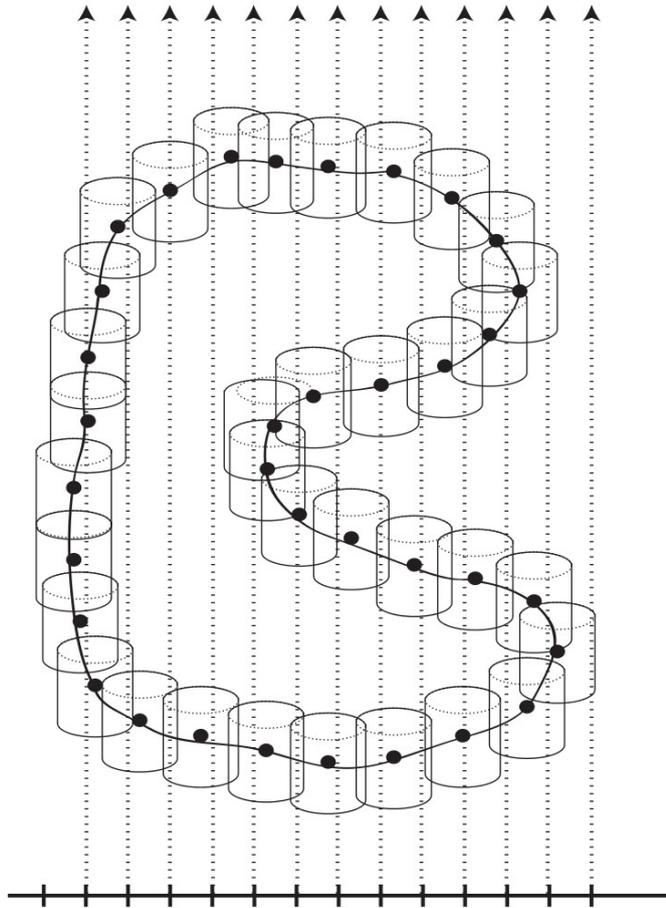


\mathbb{R}^3 : spherical cover



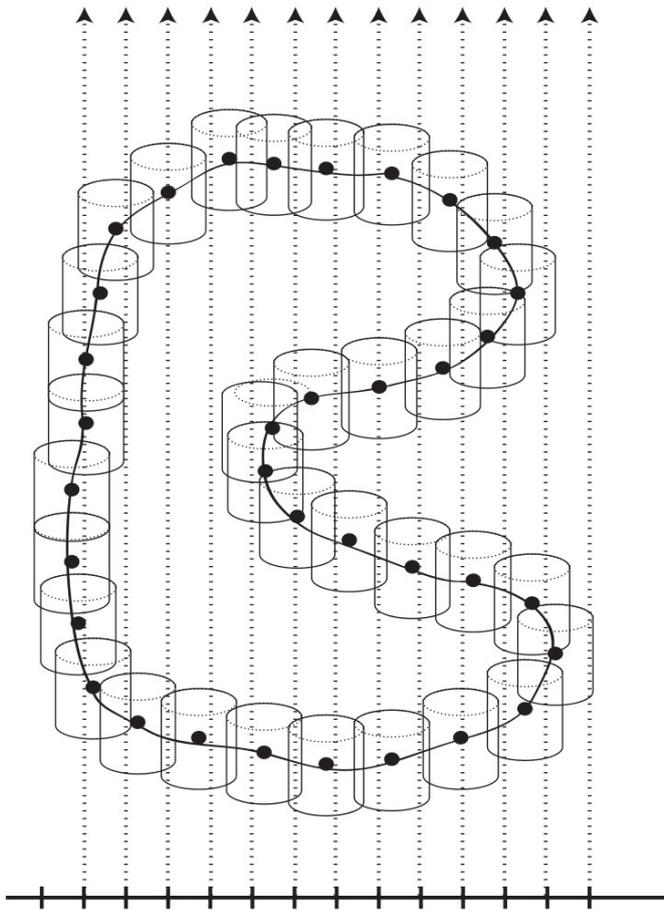
\mathbb{R}^3 : cylindrical cover



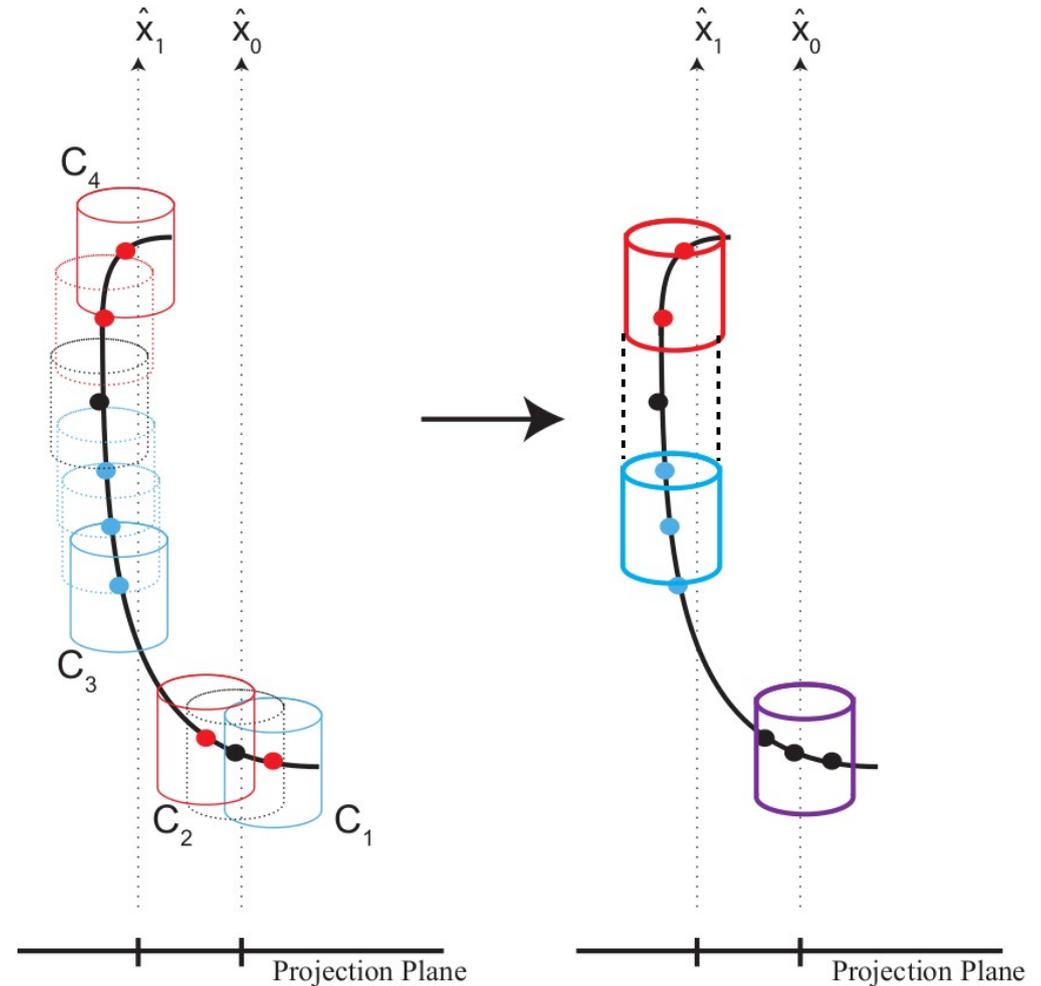


\mathbb{R}^3 : cylindrical cover



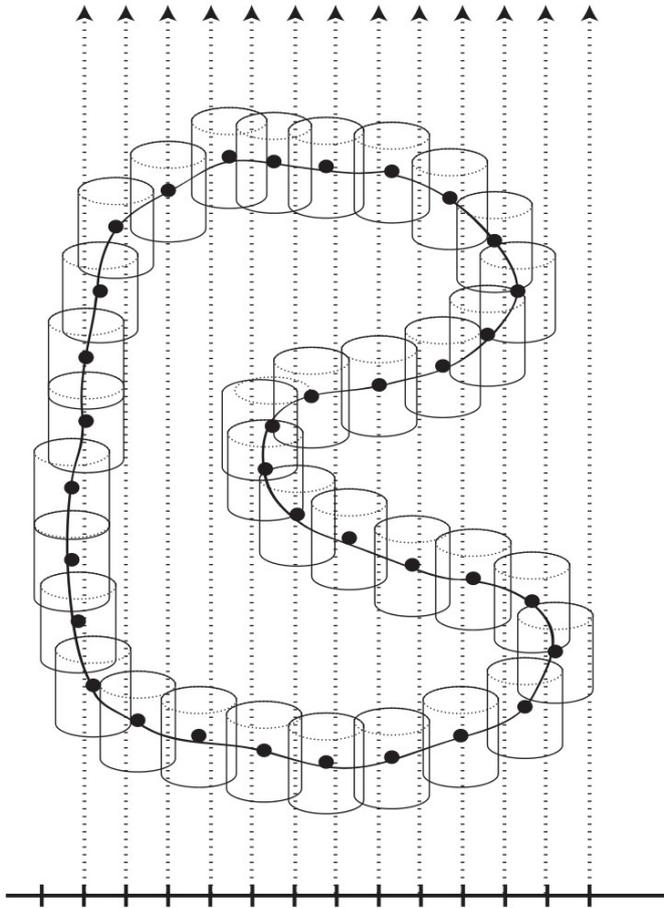


\mathbb{R}^3 : cylindrical cover

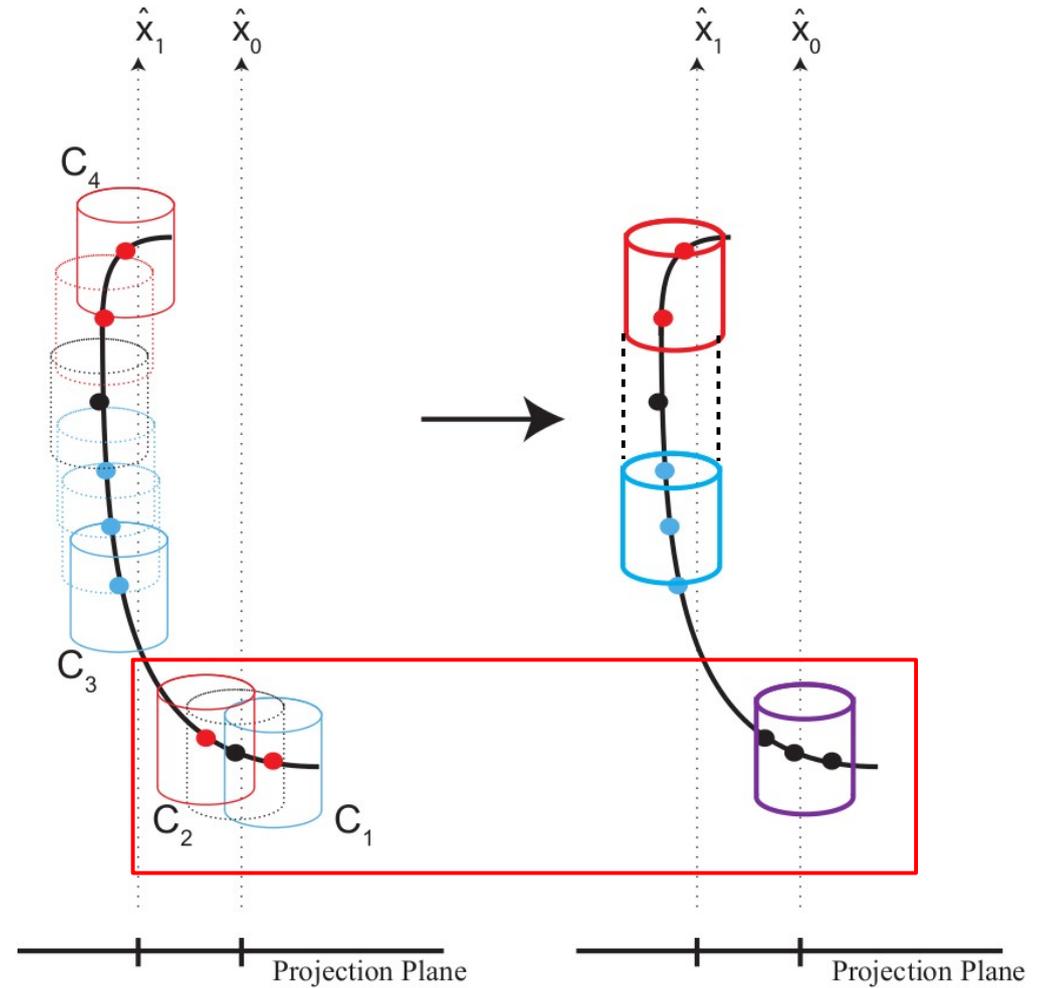


cylinders \rightarrow blended view rays



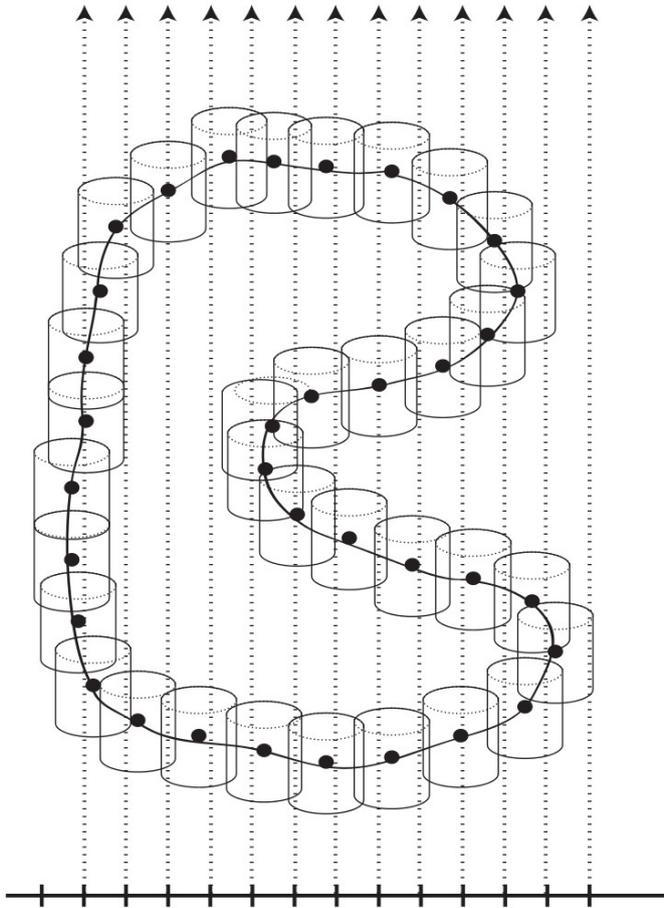


\mathbb{R}^3 : cylindrical cover

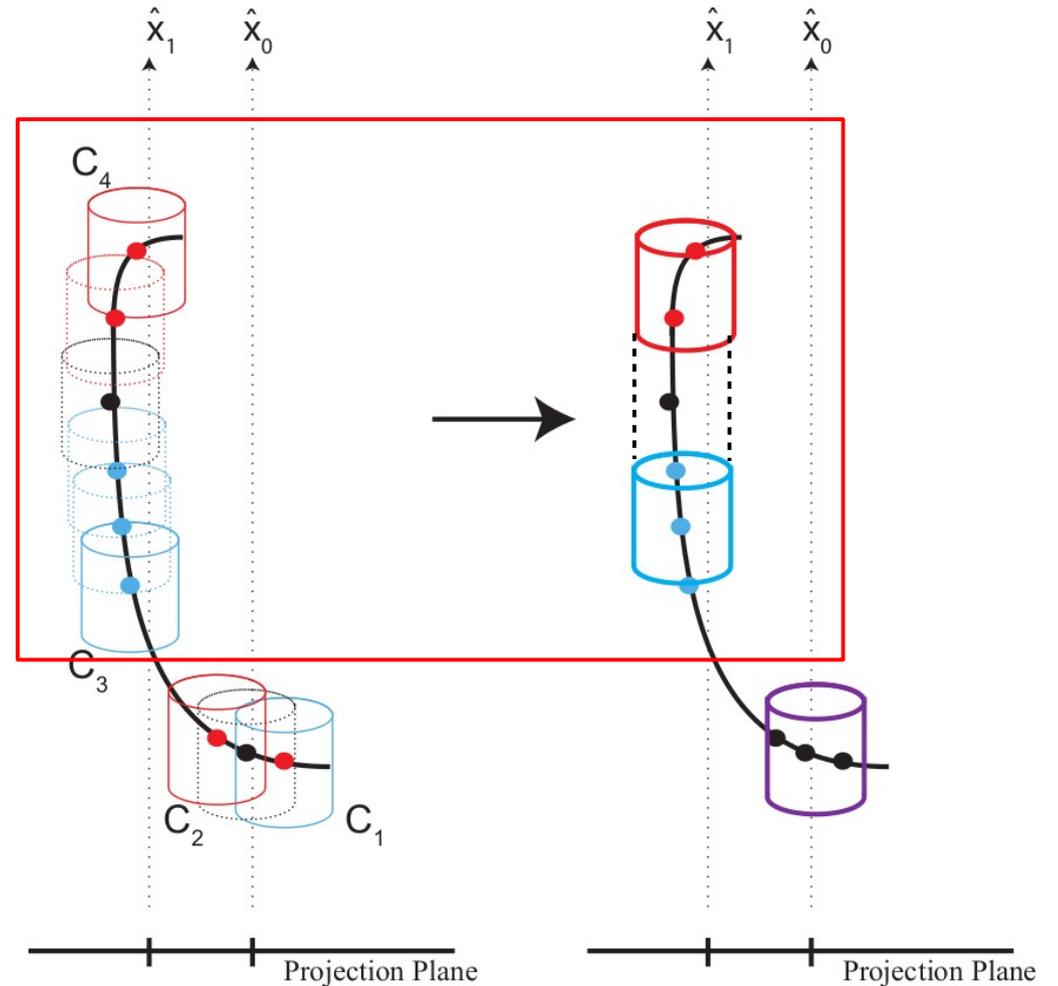


cylinders \rightarrow blended view rays



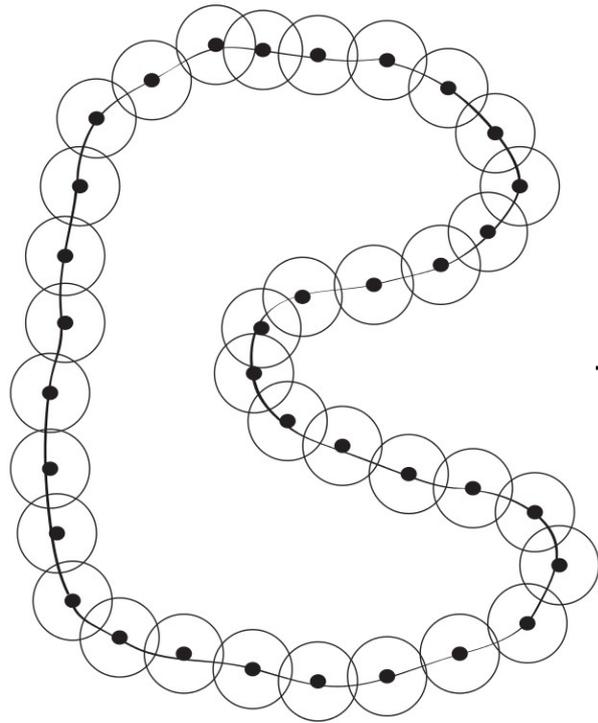


\mathbb{R}^3 : cylindrical cover



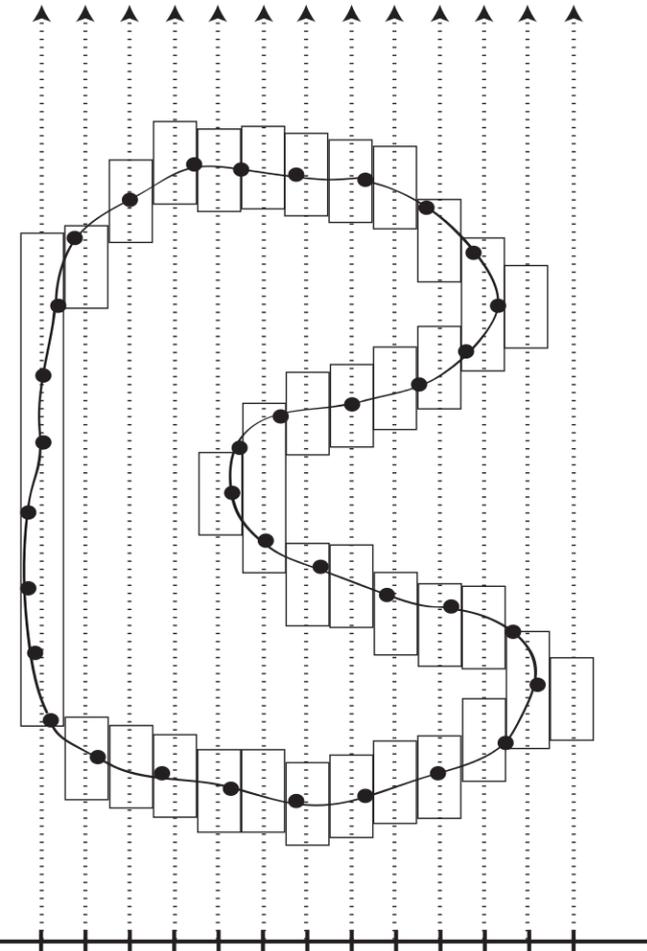
cylinders \rightarrow blended view rays





\mathbb{R}^3 : spherical cover

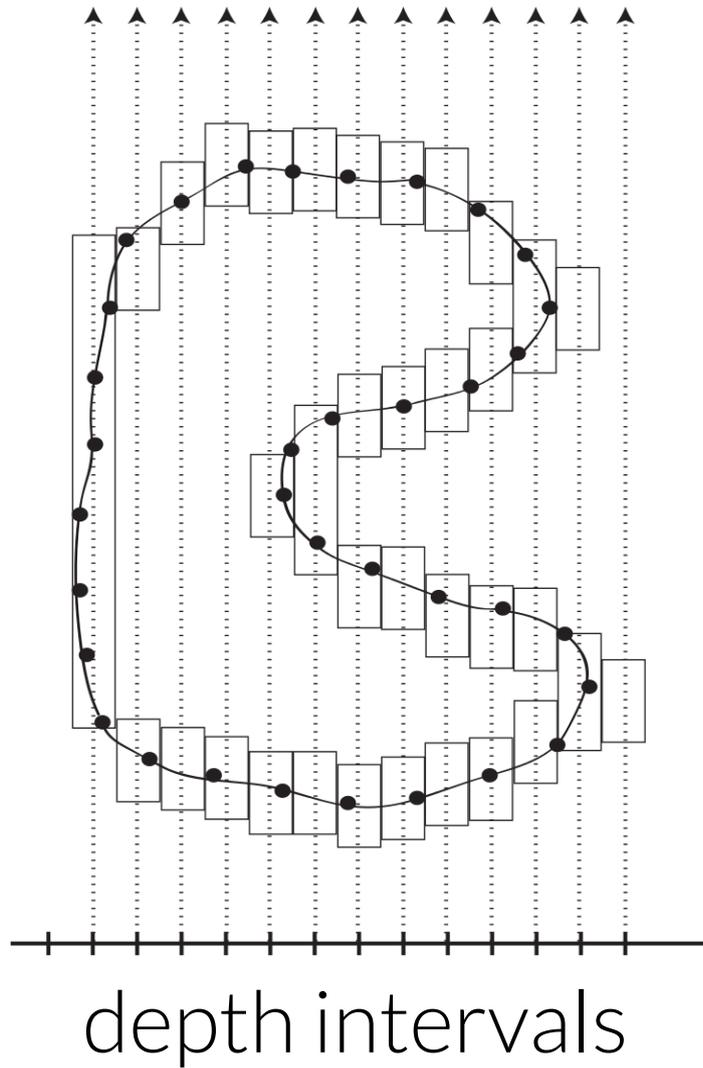
blended



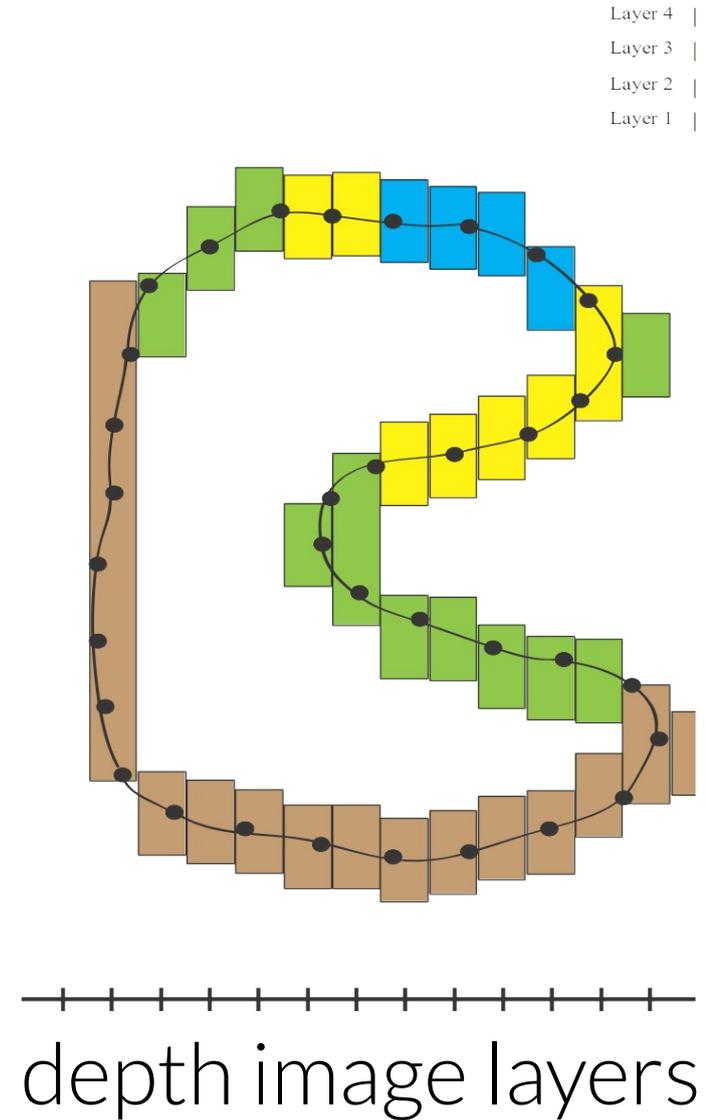
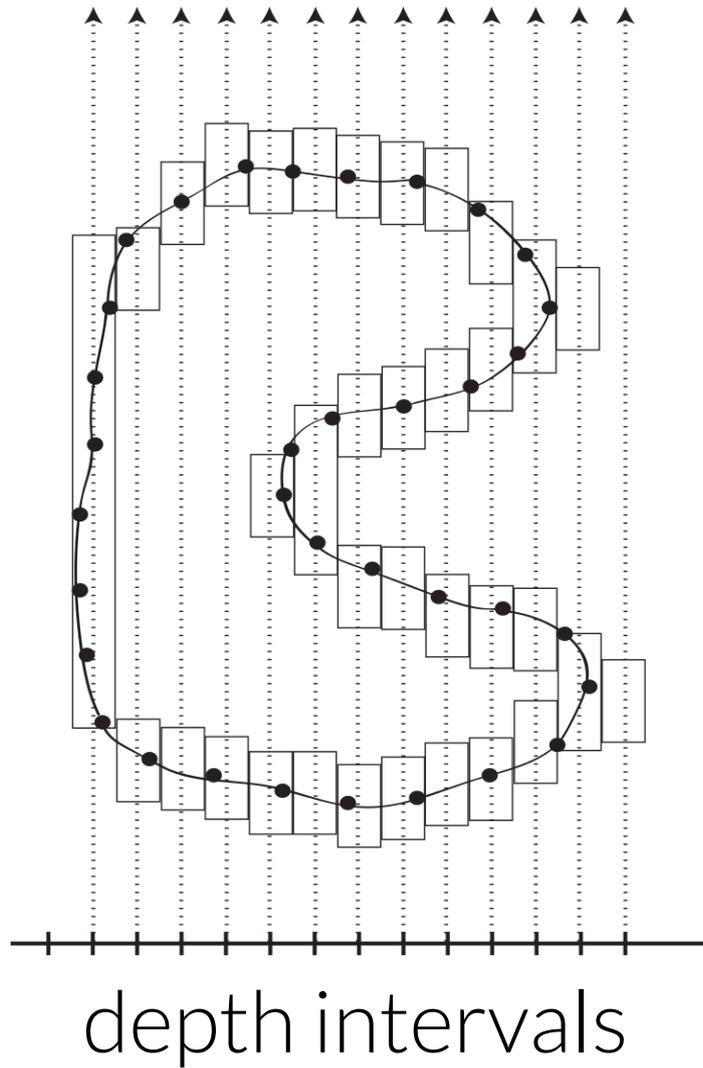
view ray depth intervals



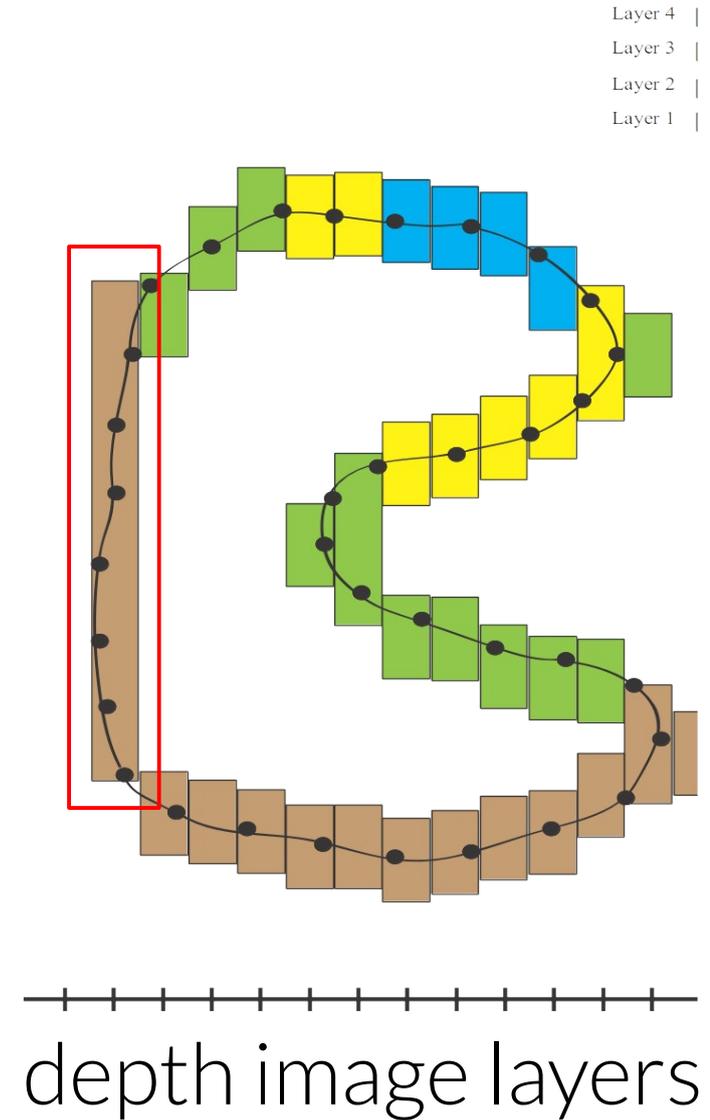
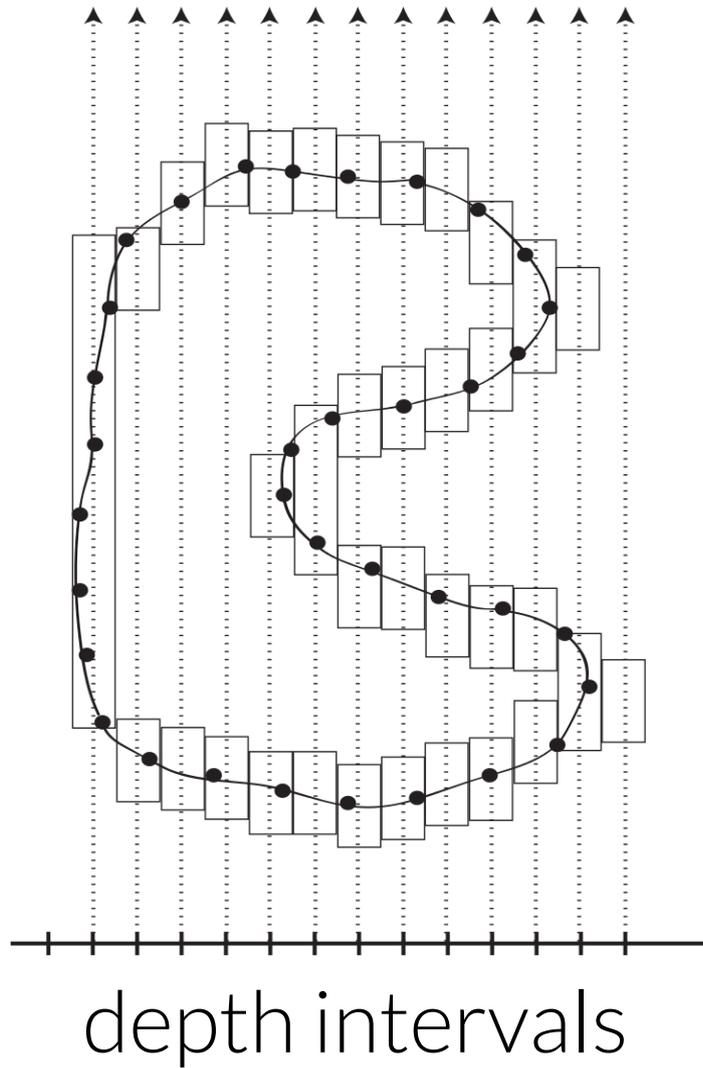
Thickened LDI (layered depth images)



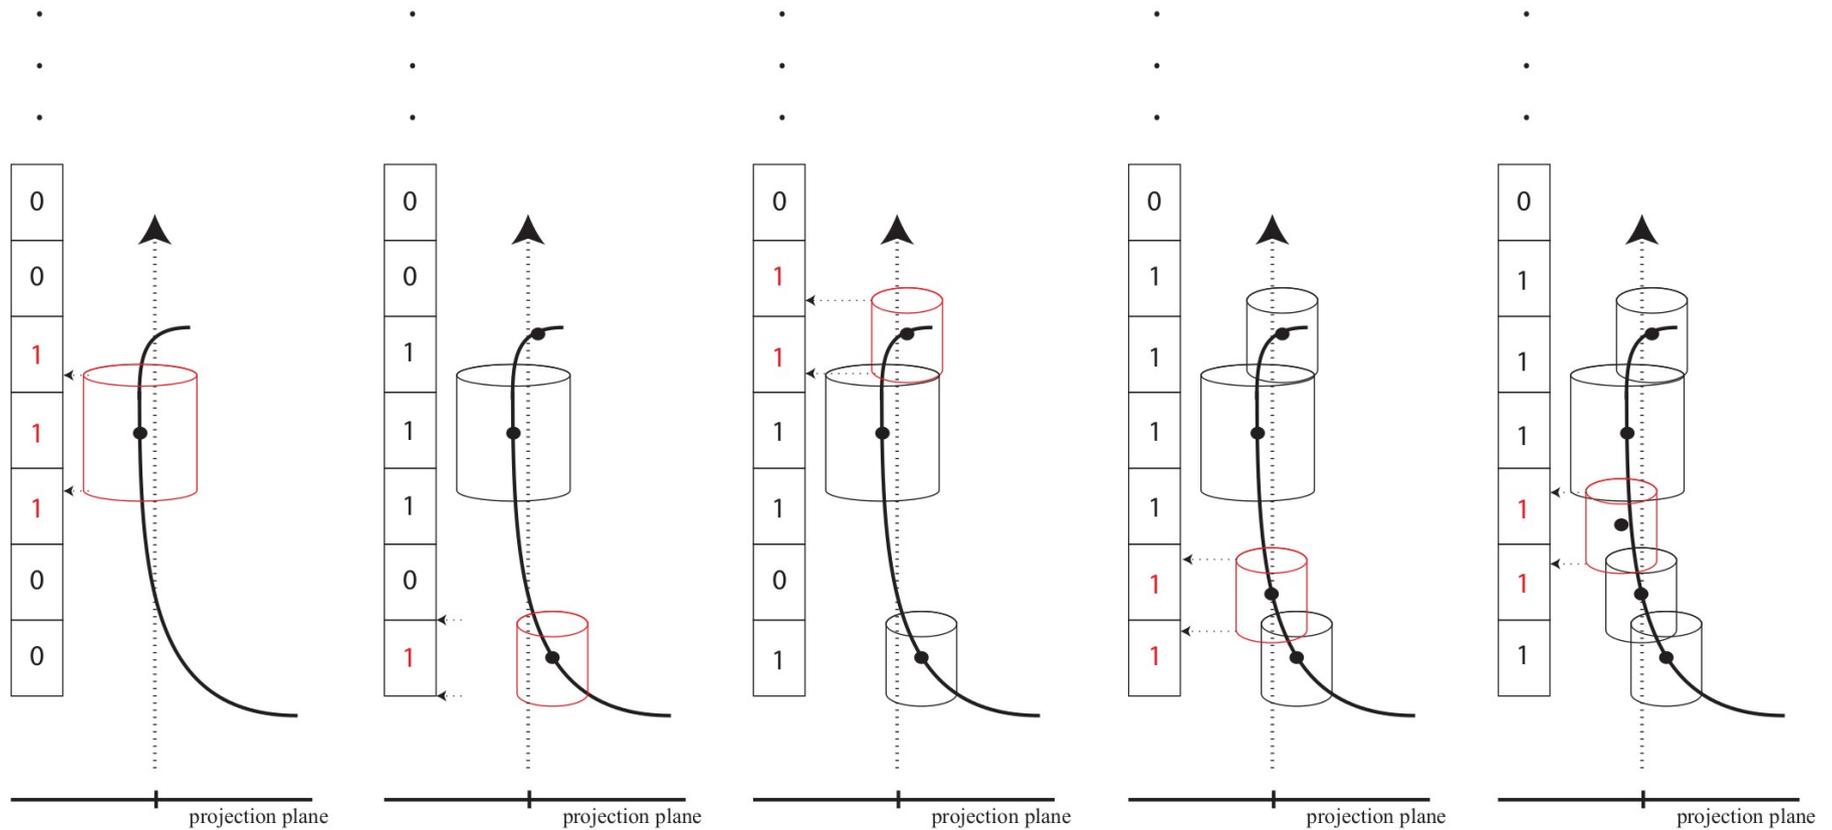
Thickened LDI (layered depth images)



Thickened LDI (layered depth images)

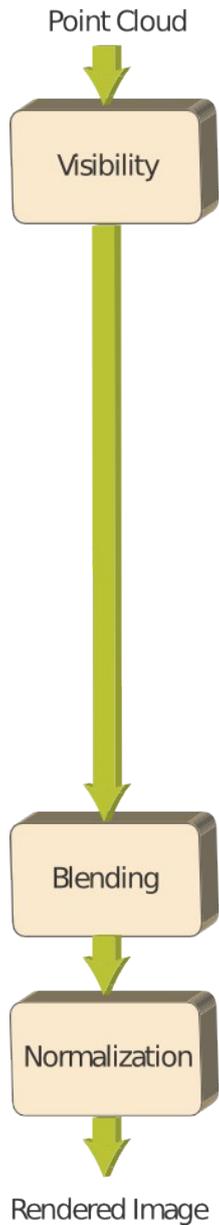


Stacking cylinders into a depth layer

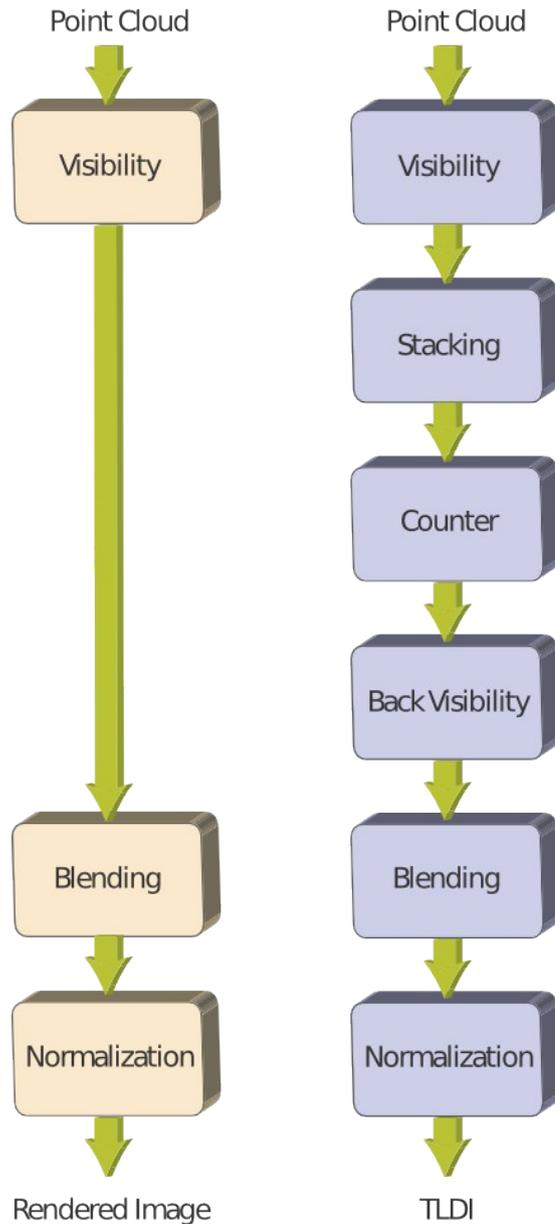


Detect layer connectivity with bit array occupancy

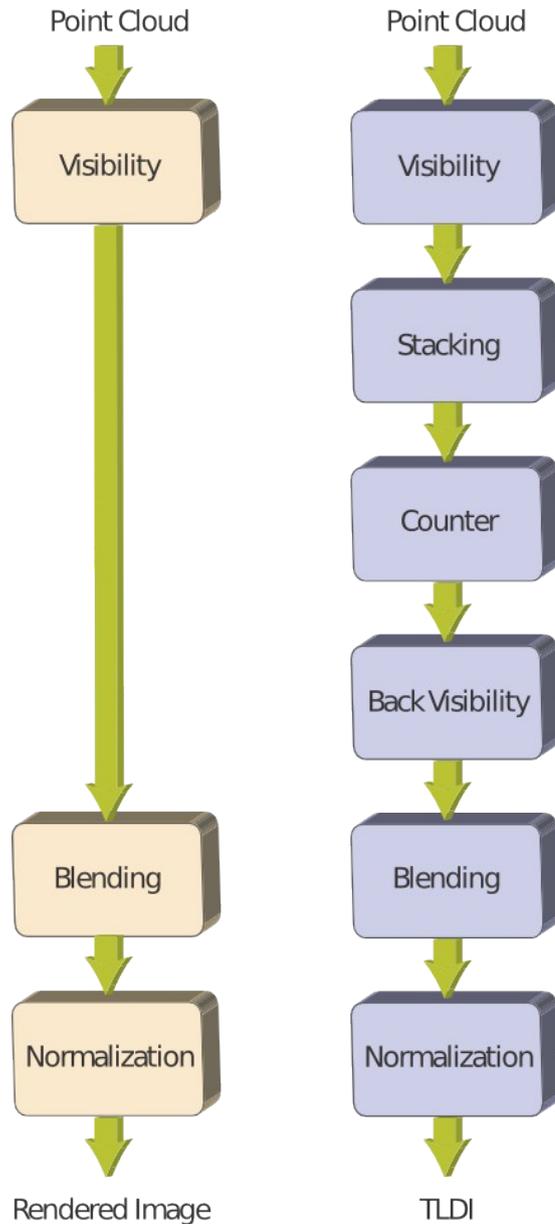




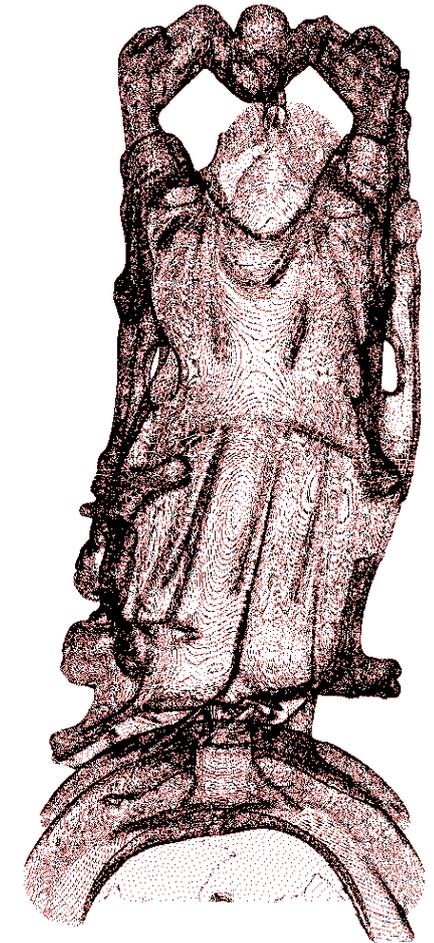
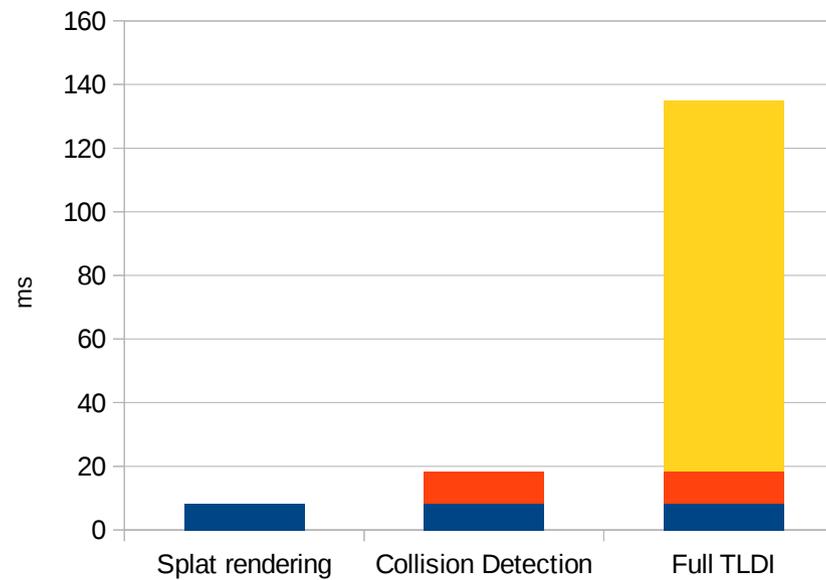
Reuse of point-based rendering pipeline



Reuse of point-based rendering pipeline

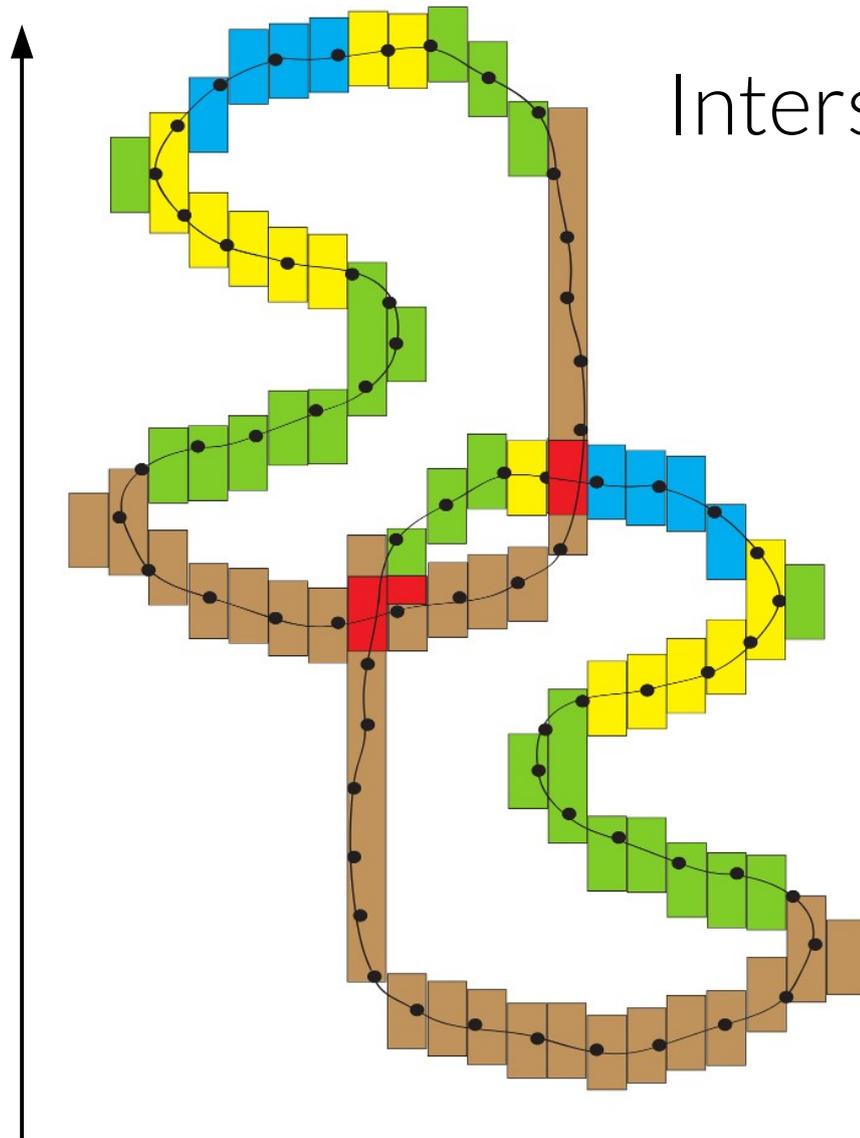


Happy buddha, 500k points



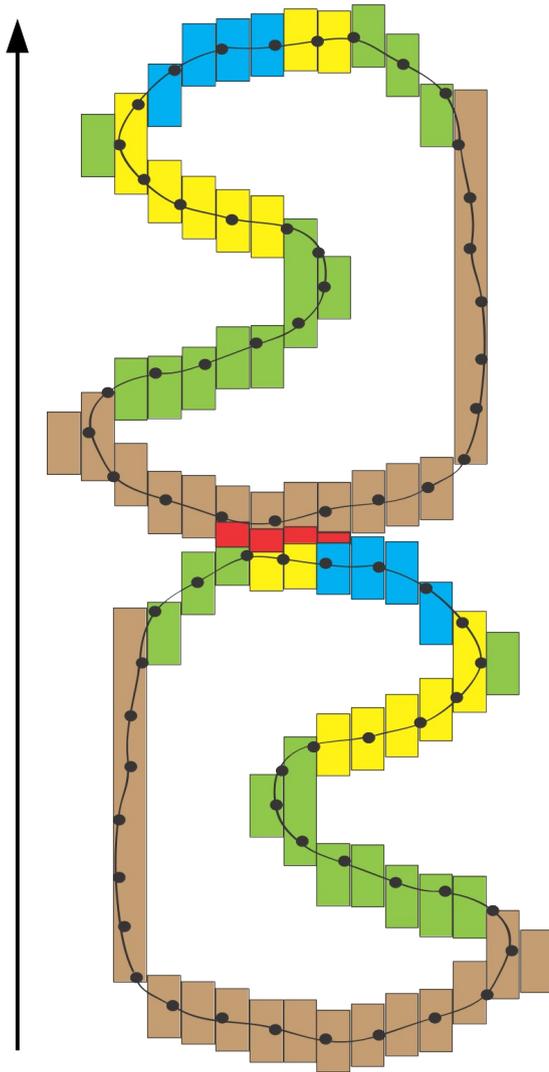
Almost half of run time reused

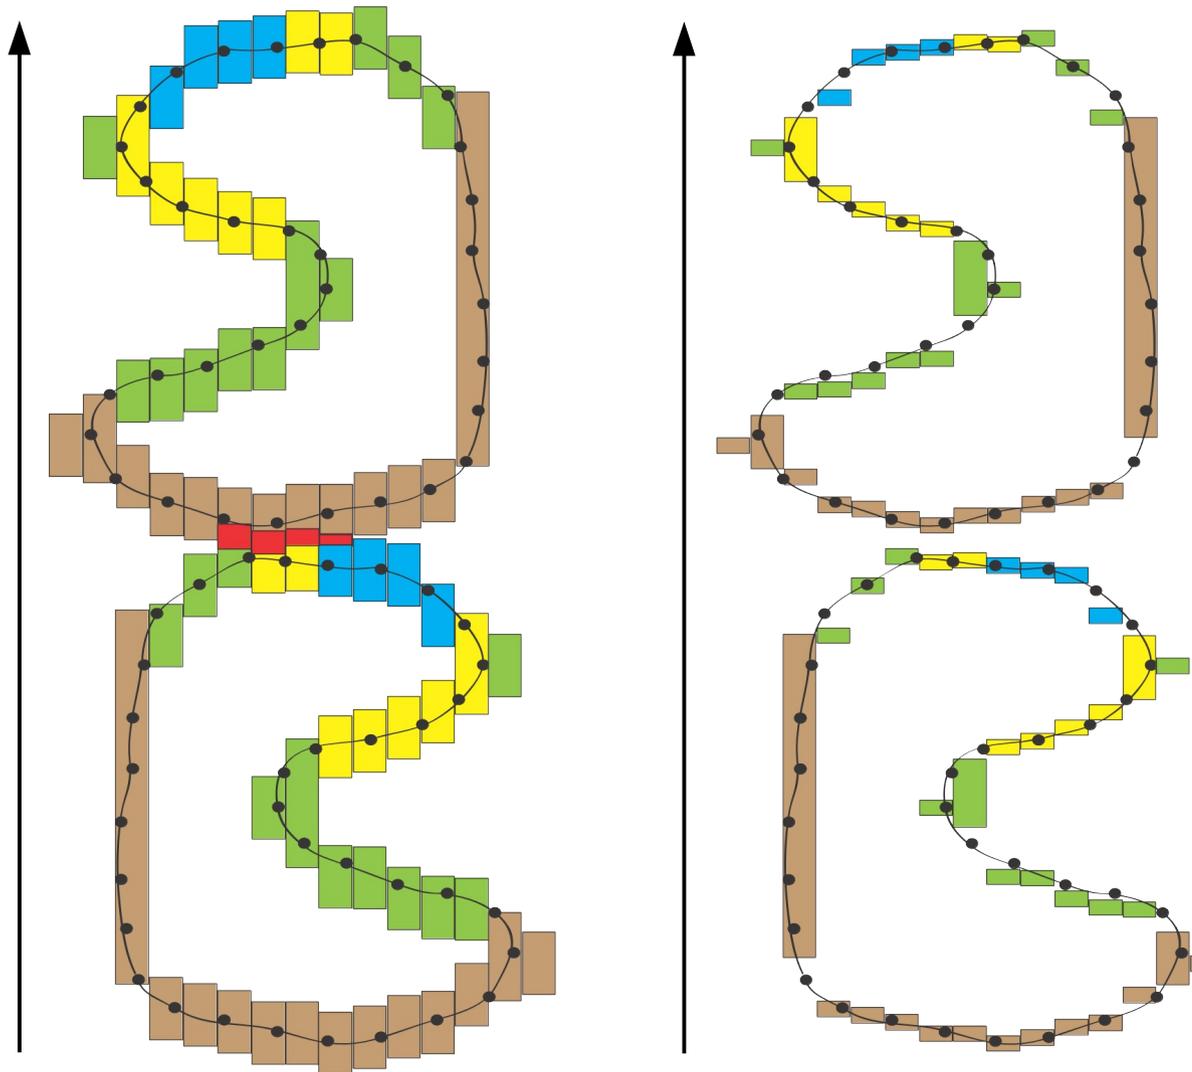




Intersection of view ray intervals

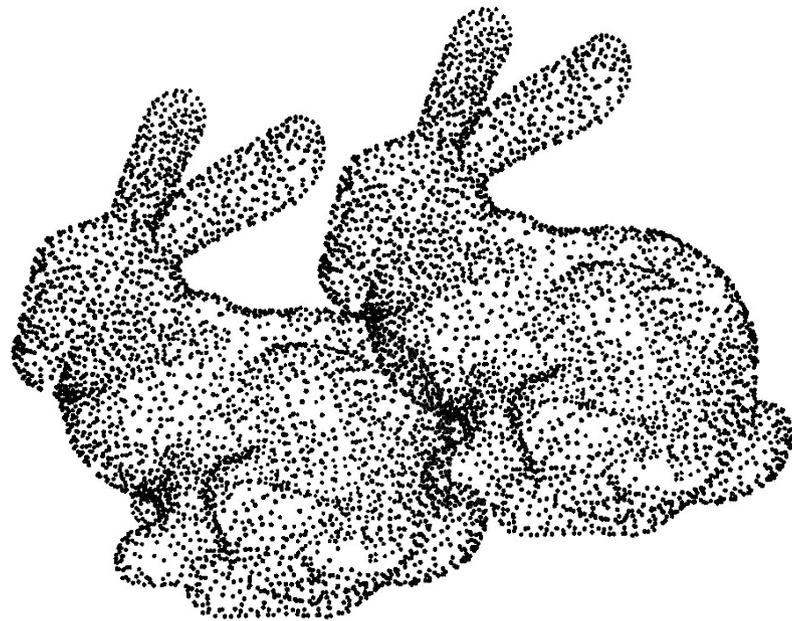






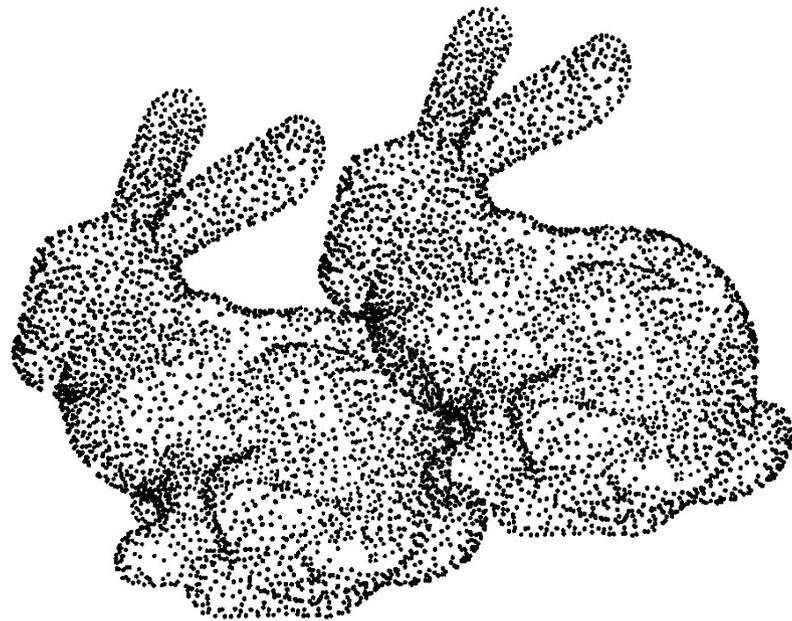
Squash in view direction – but how much?



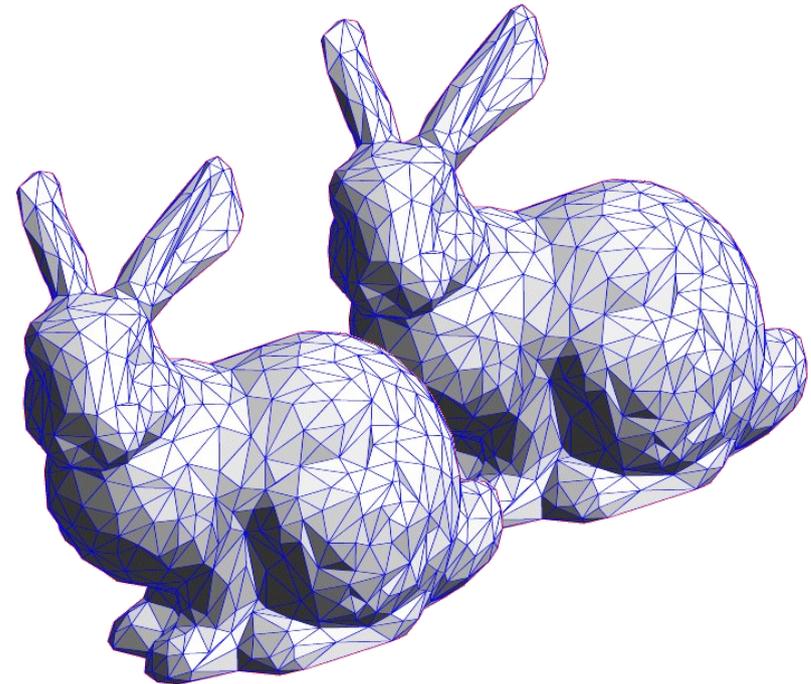


point clouds





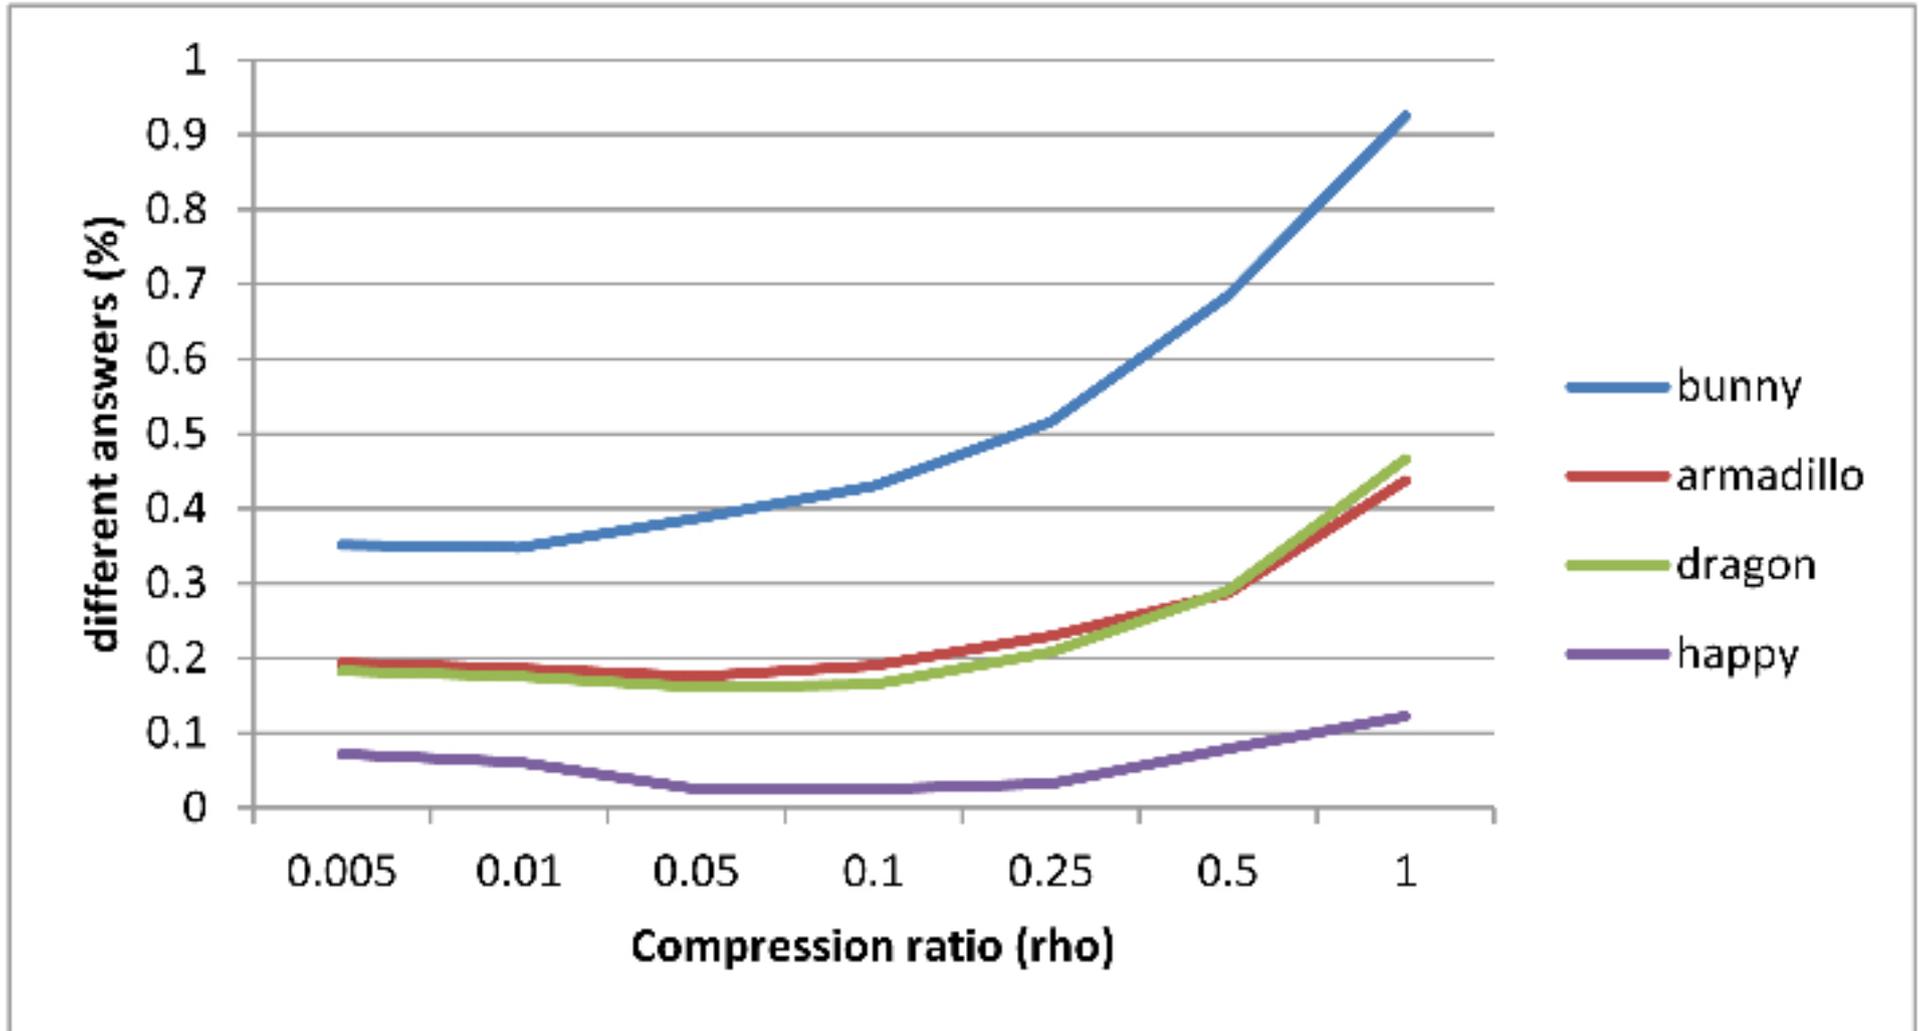
point clouds

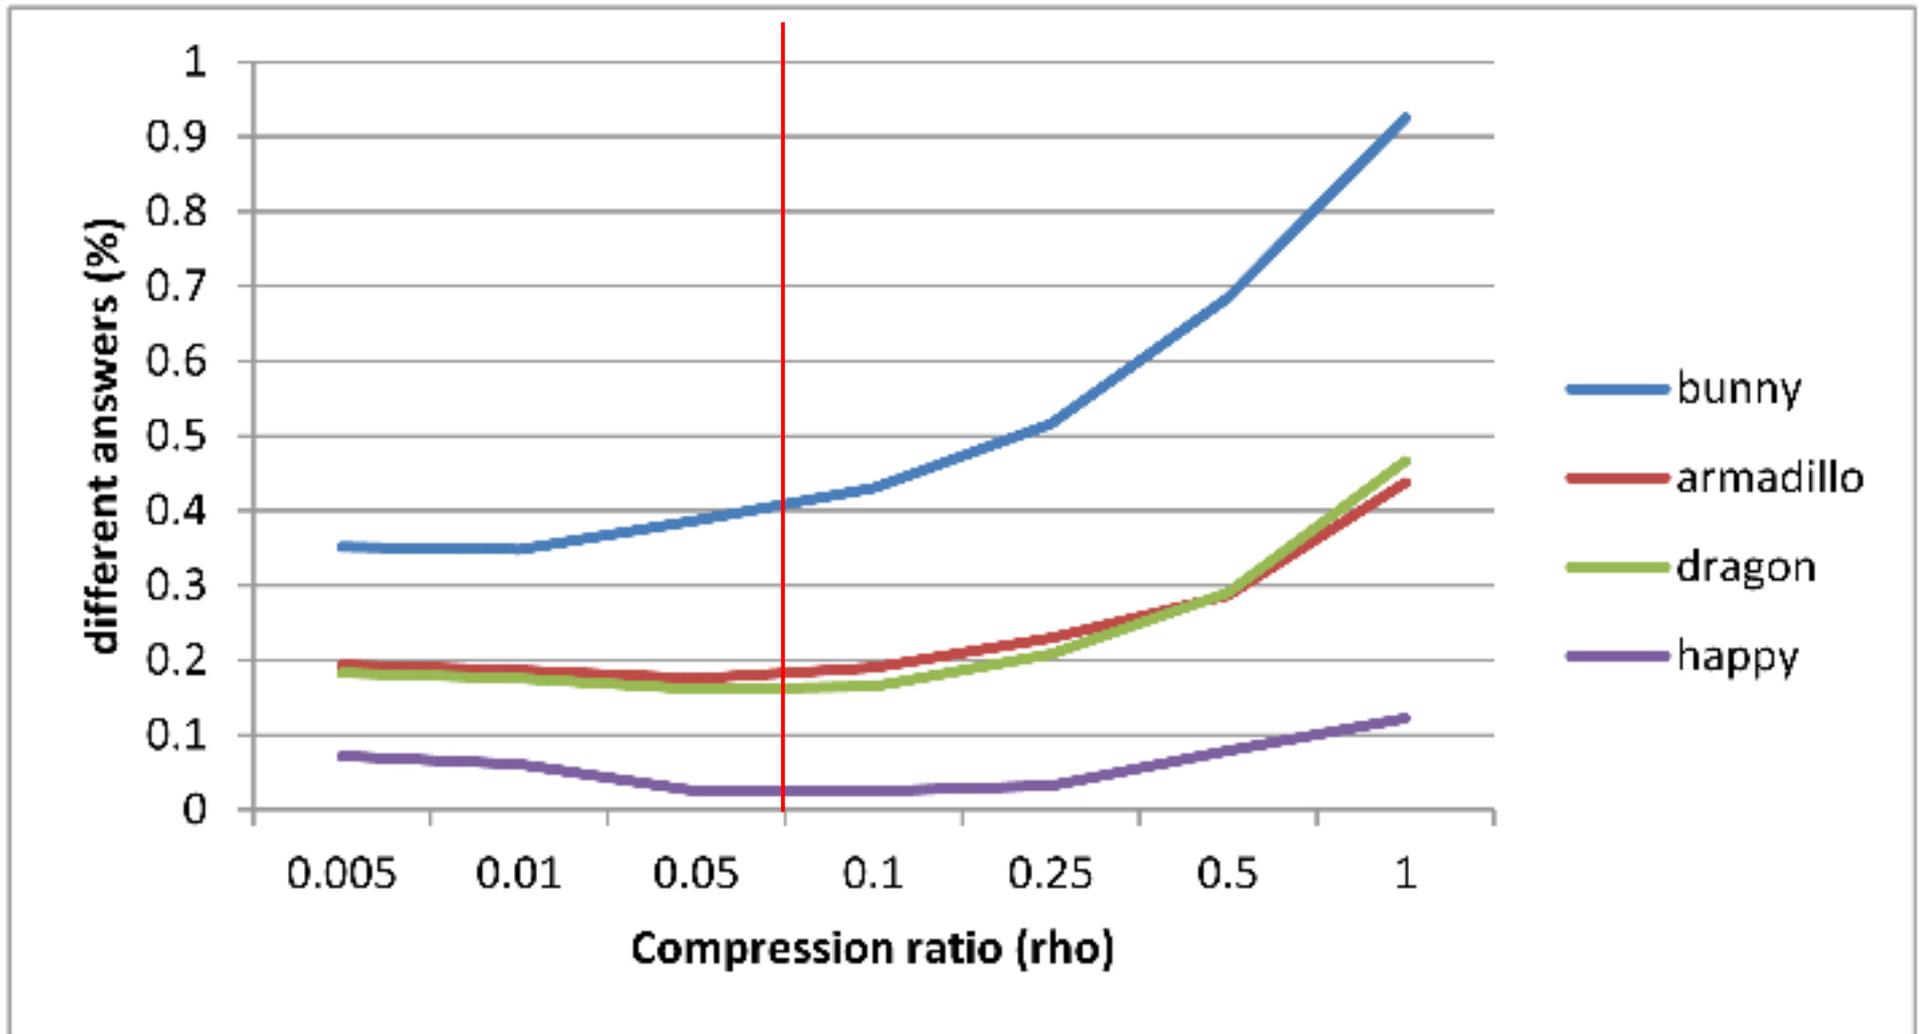


mesh = reference



Squashing the bounding volume

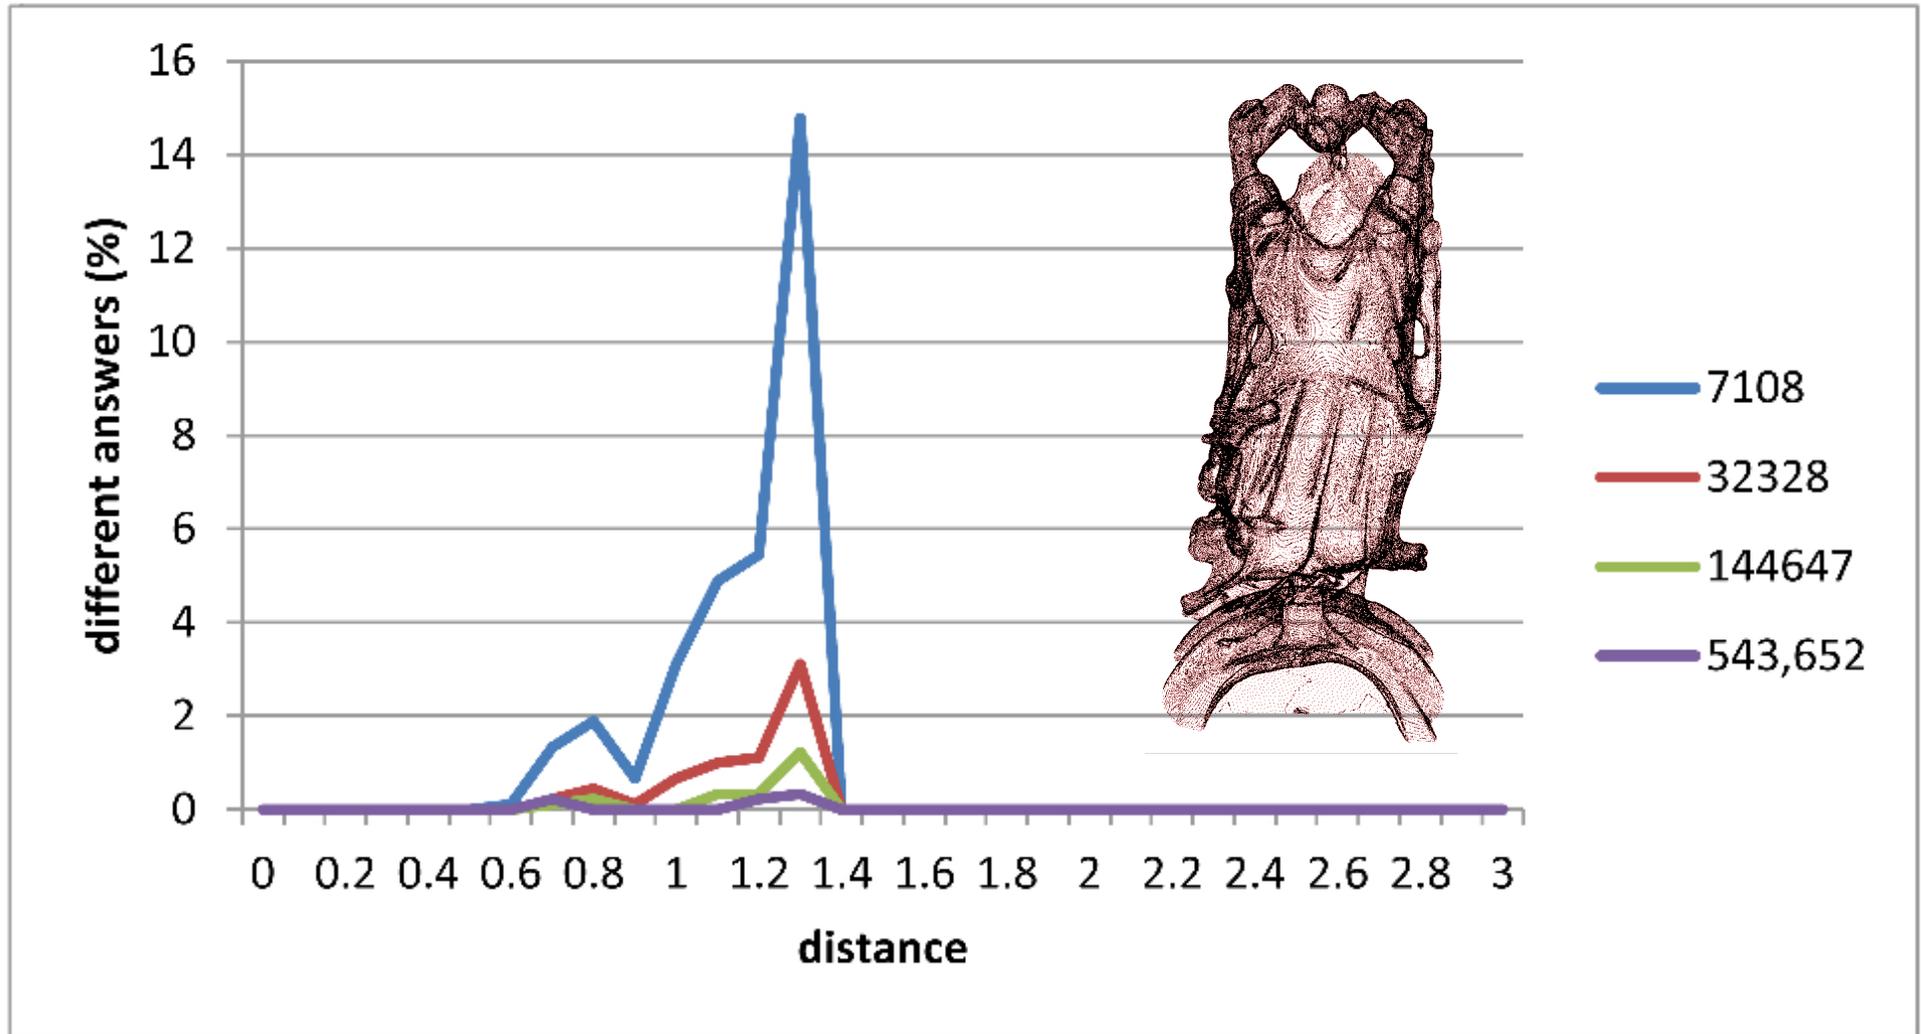




$\rho=0.05$ is good \rightarrow squashed by factor 20

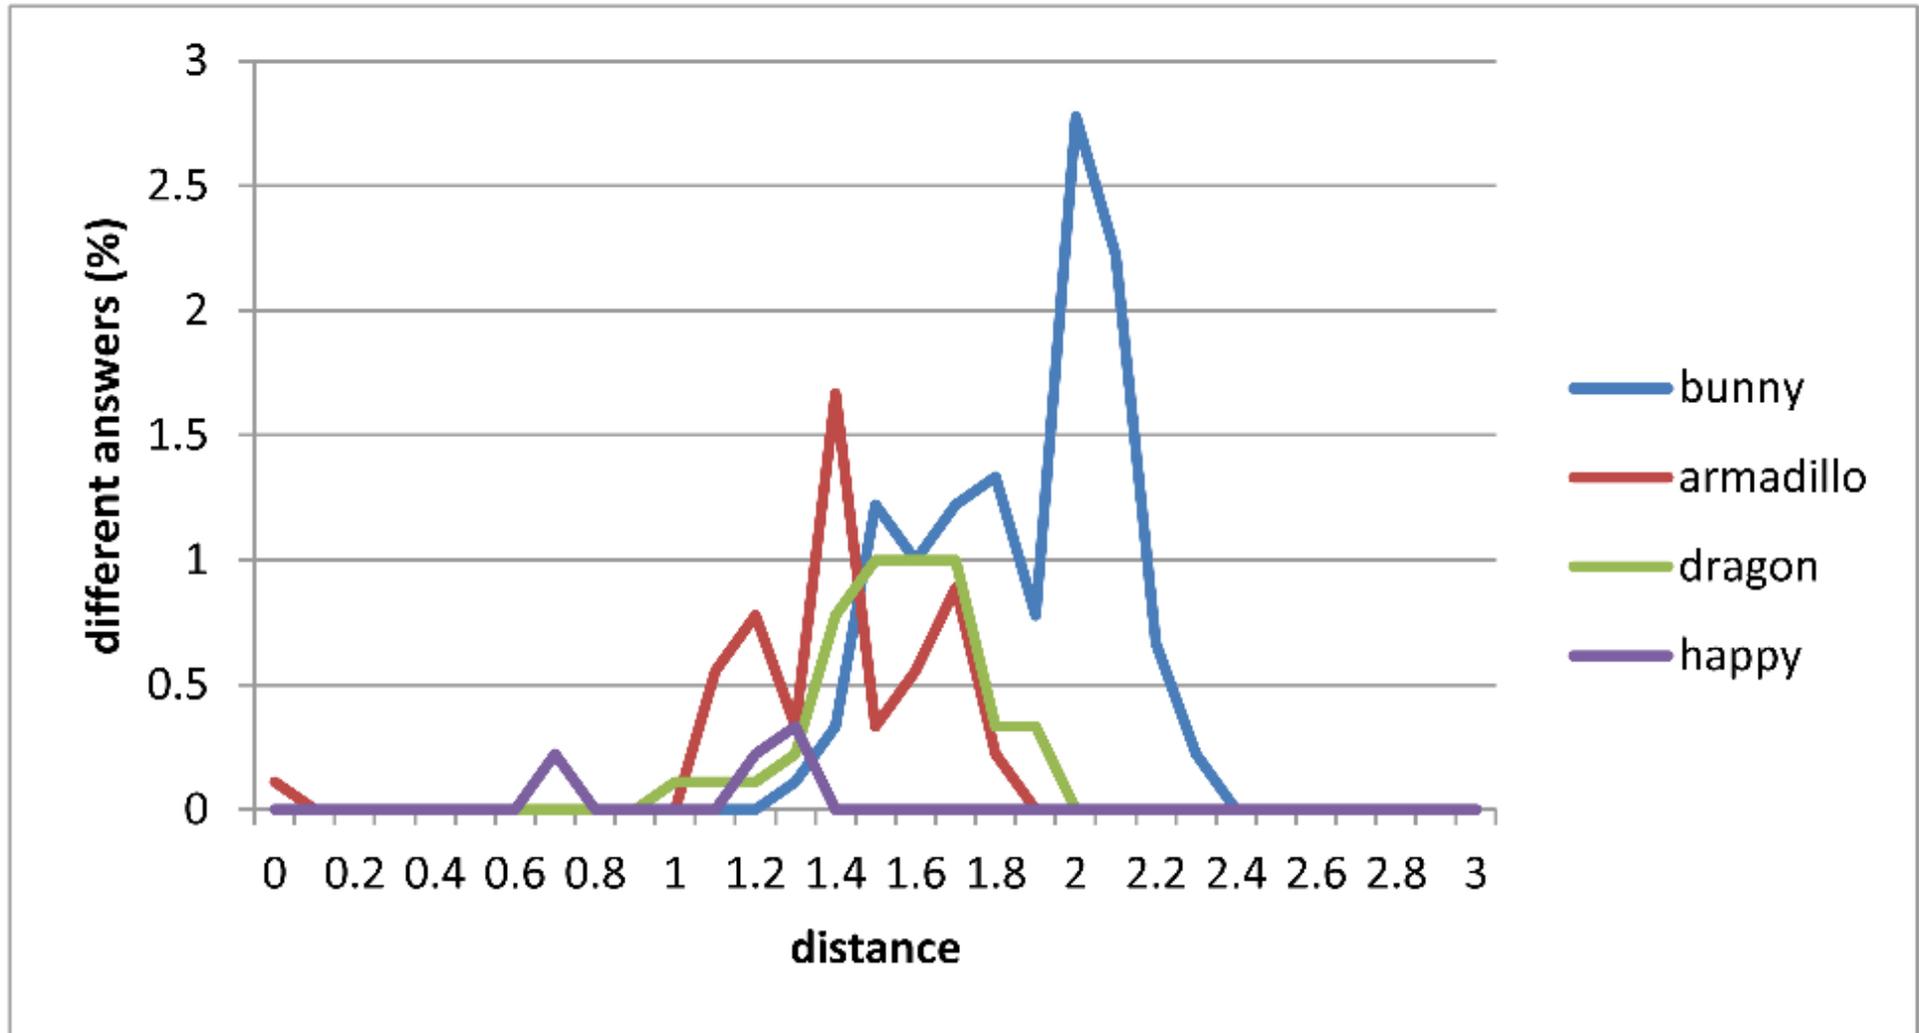


Result 1: Enhanced accuracy



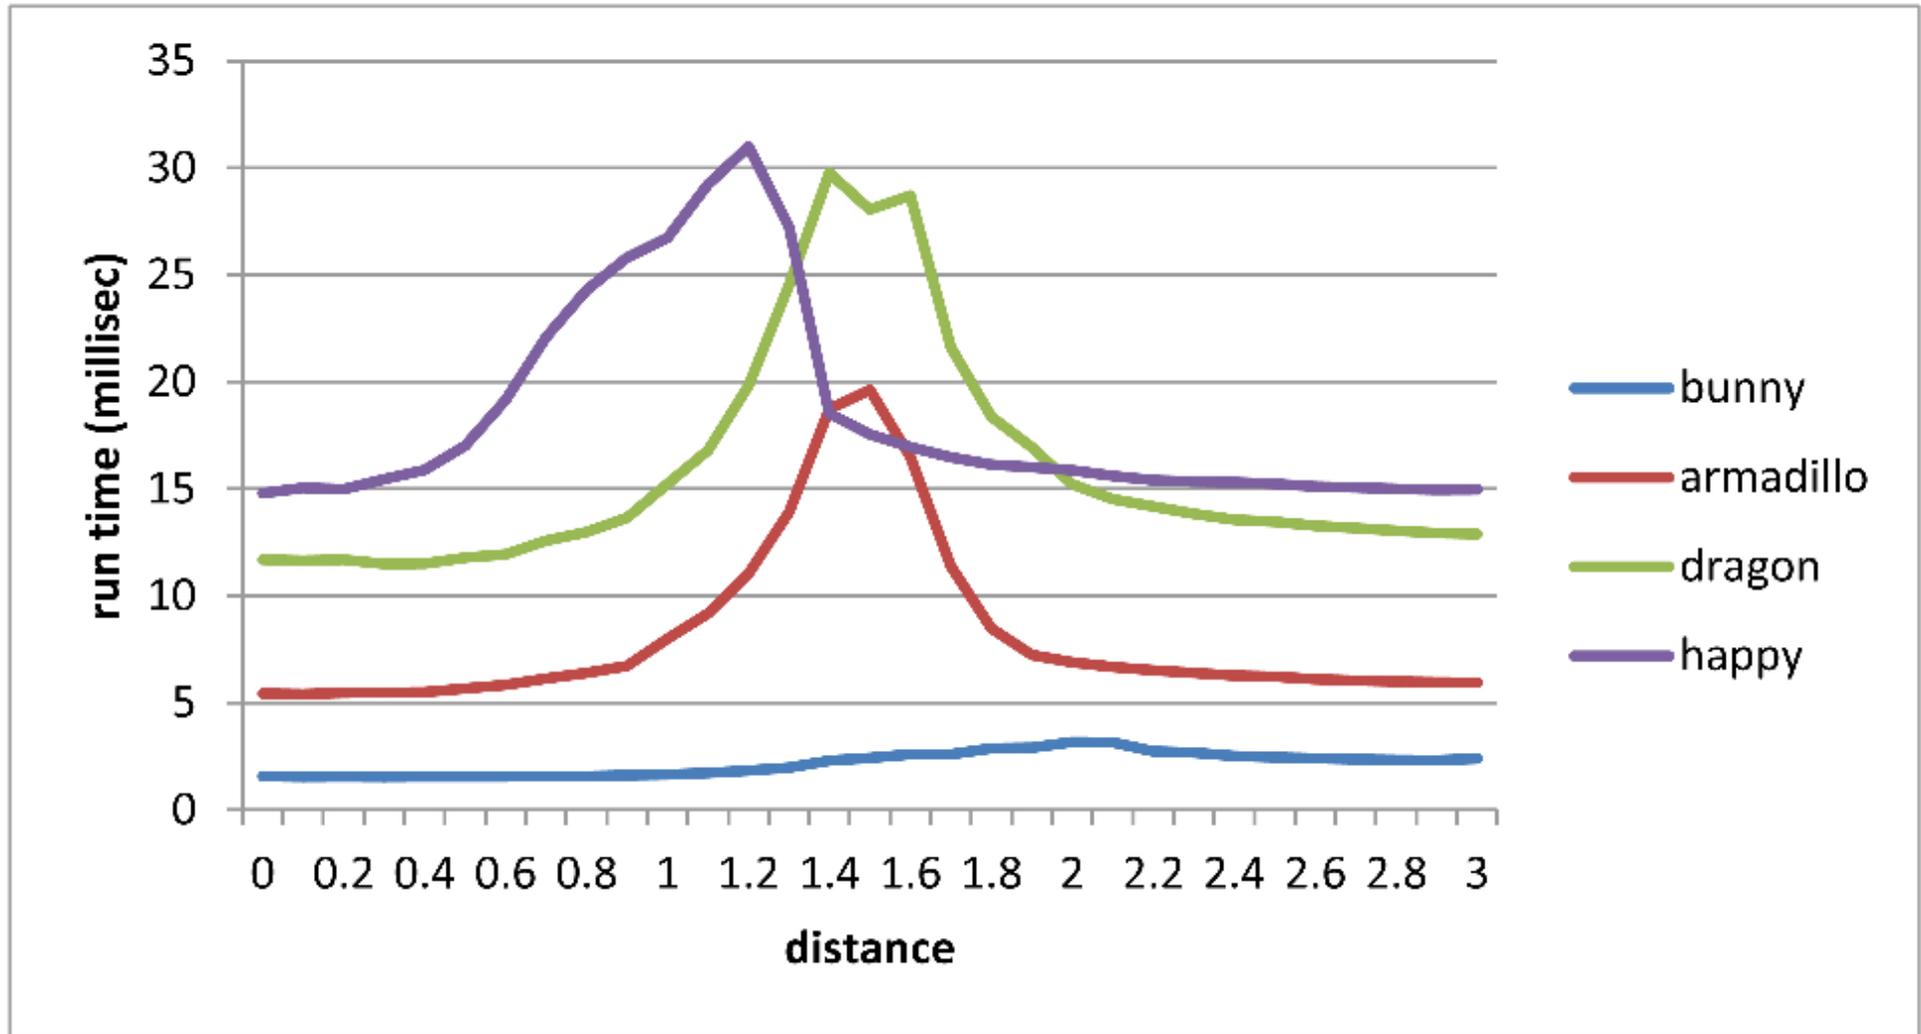
[Zachmann 2002] ~7% accuracy, ours: from 0.3%





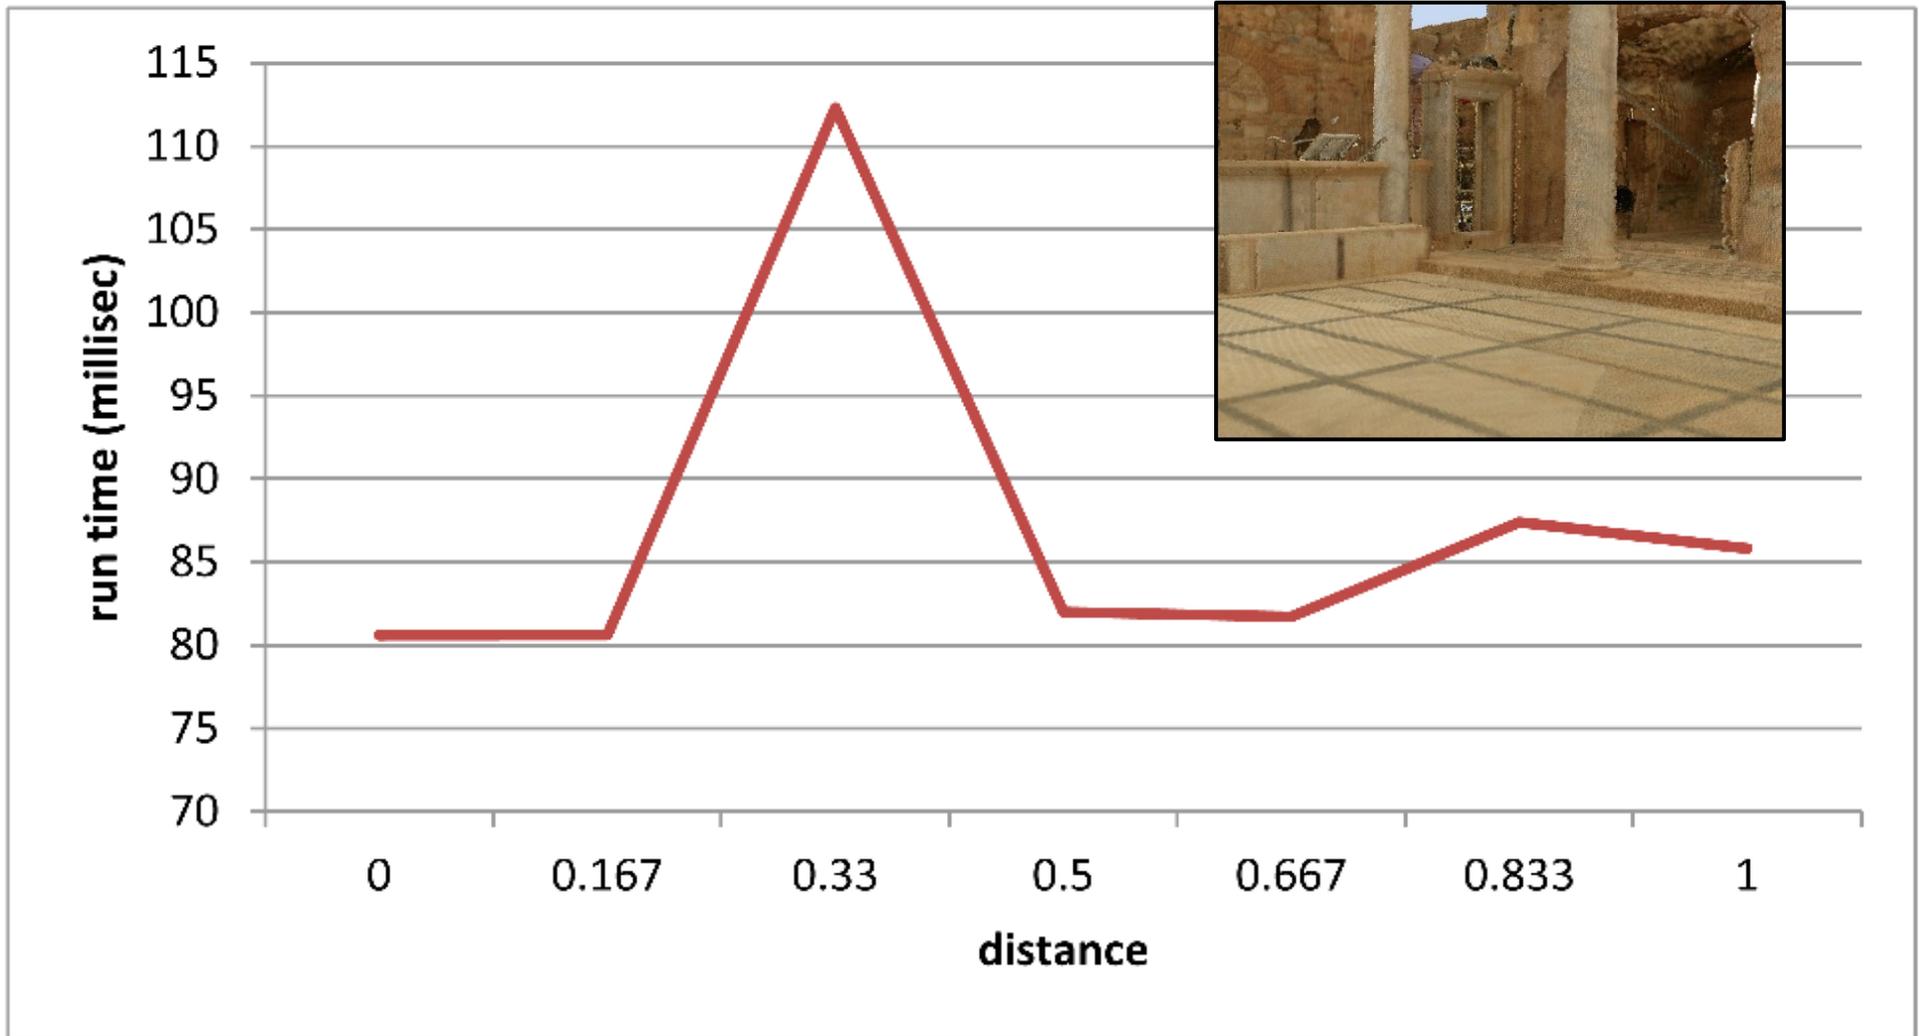
within 0.3-3% of reference

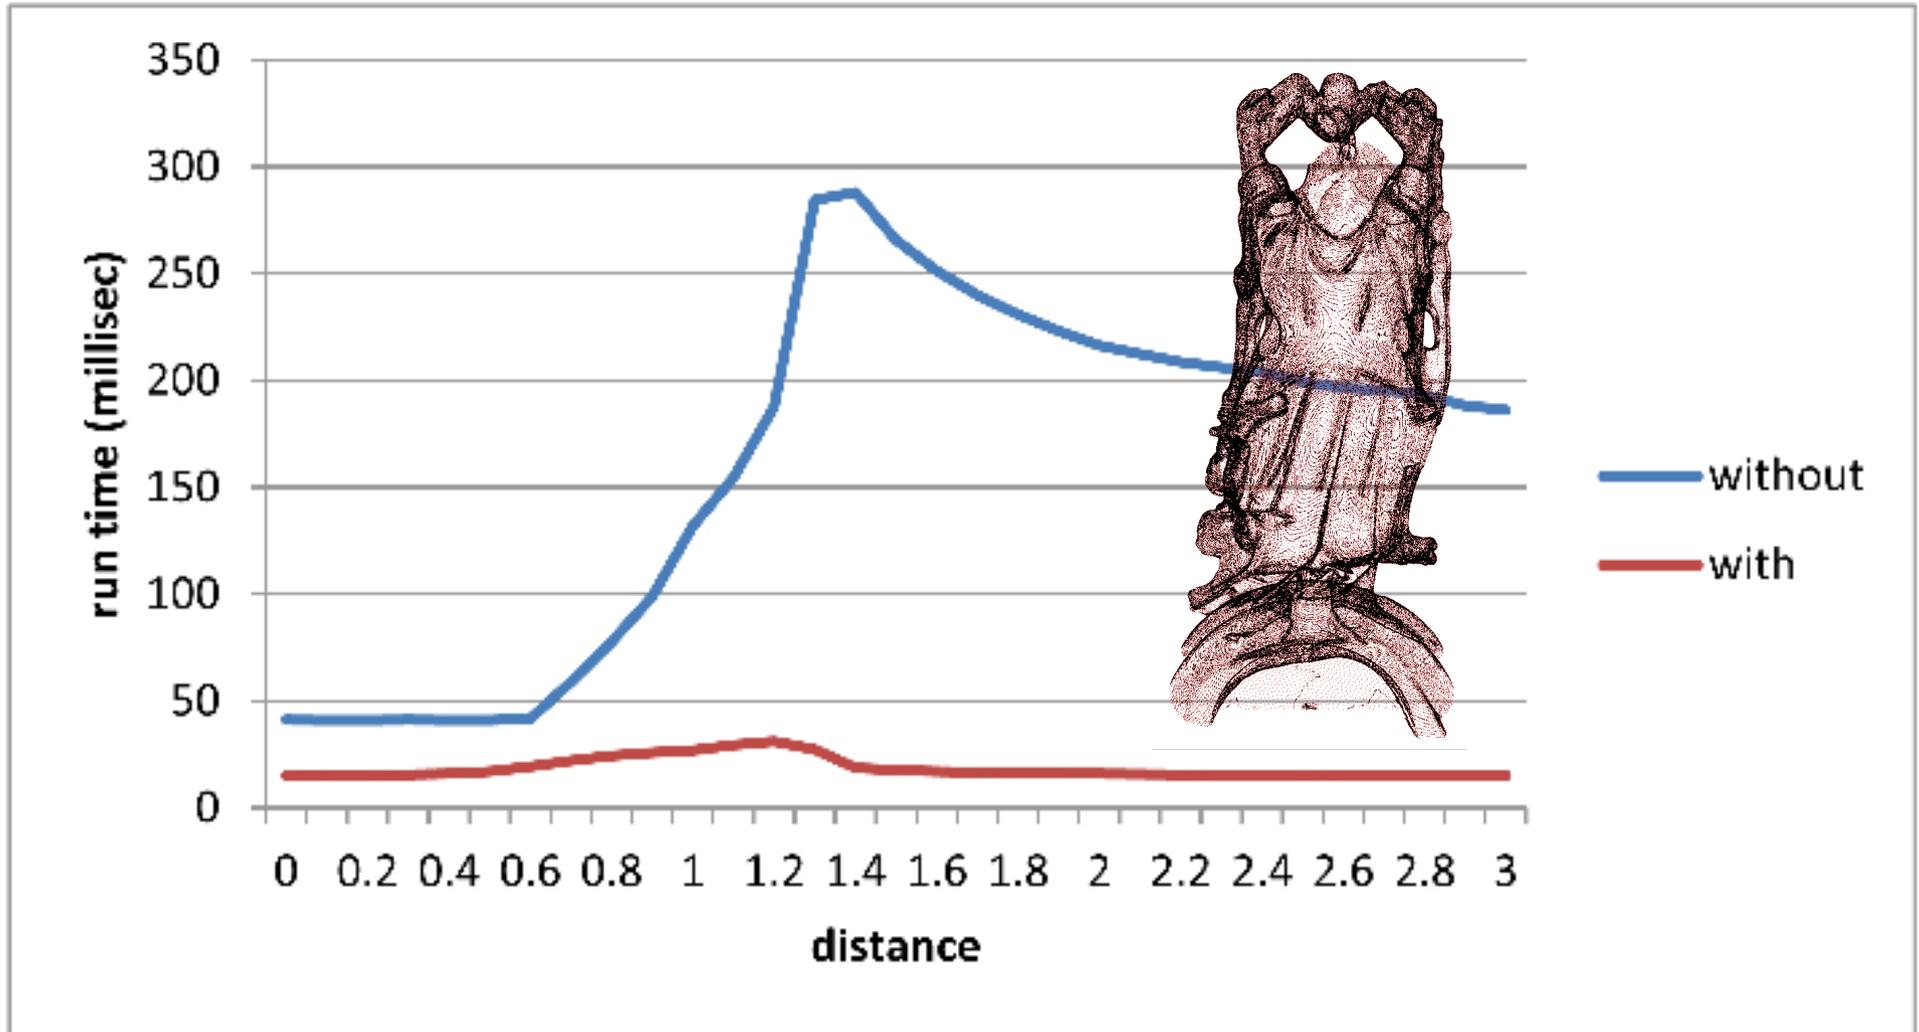




Previously only feasible using probability







Test first layer against last layer of other point cloud



For collision detection:

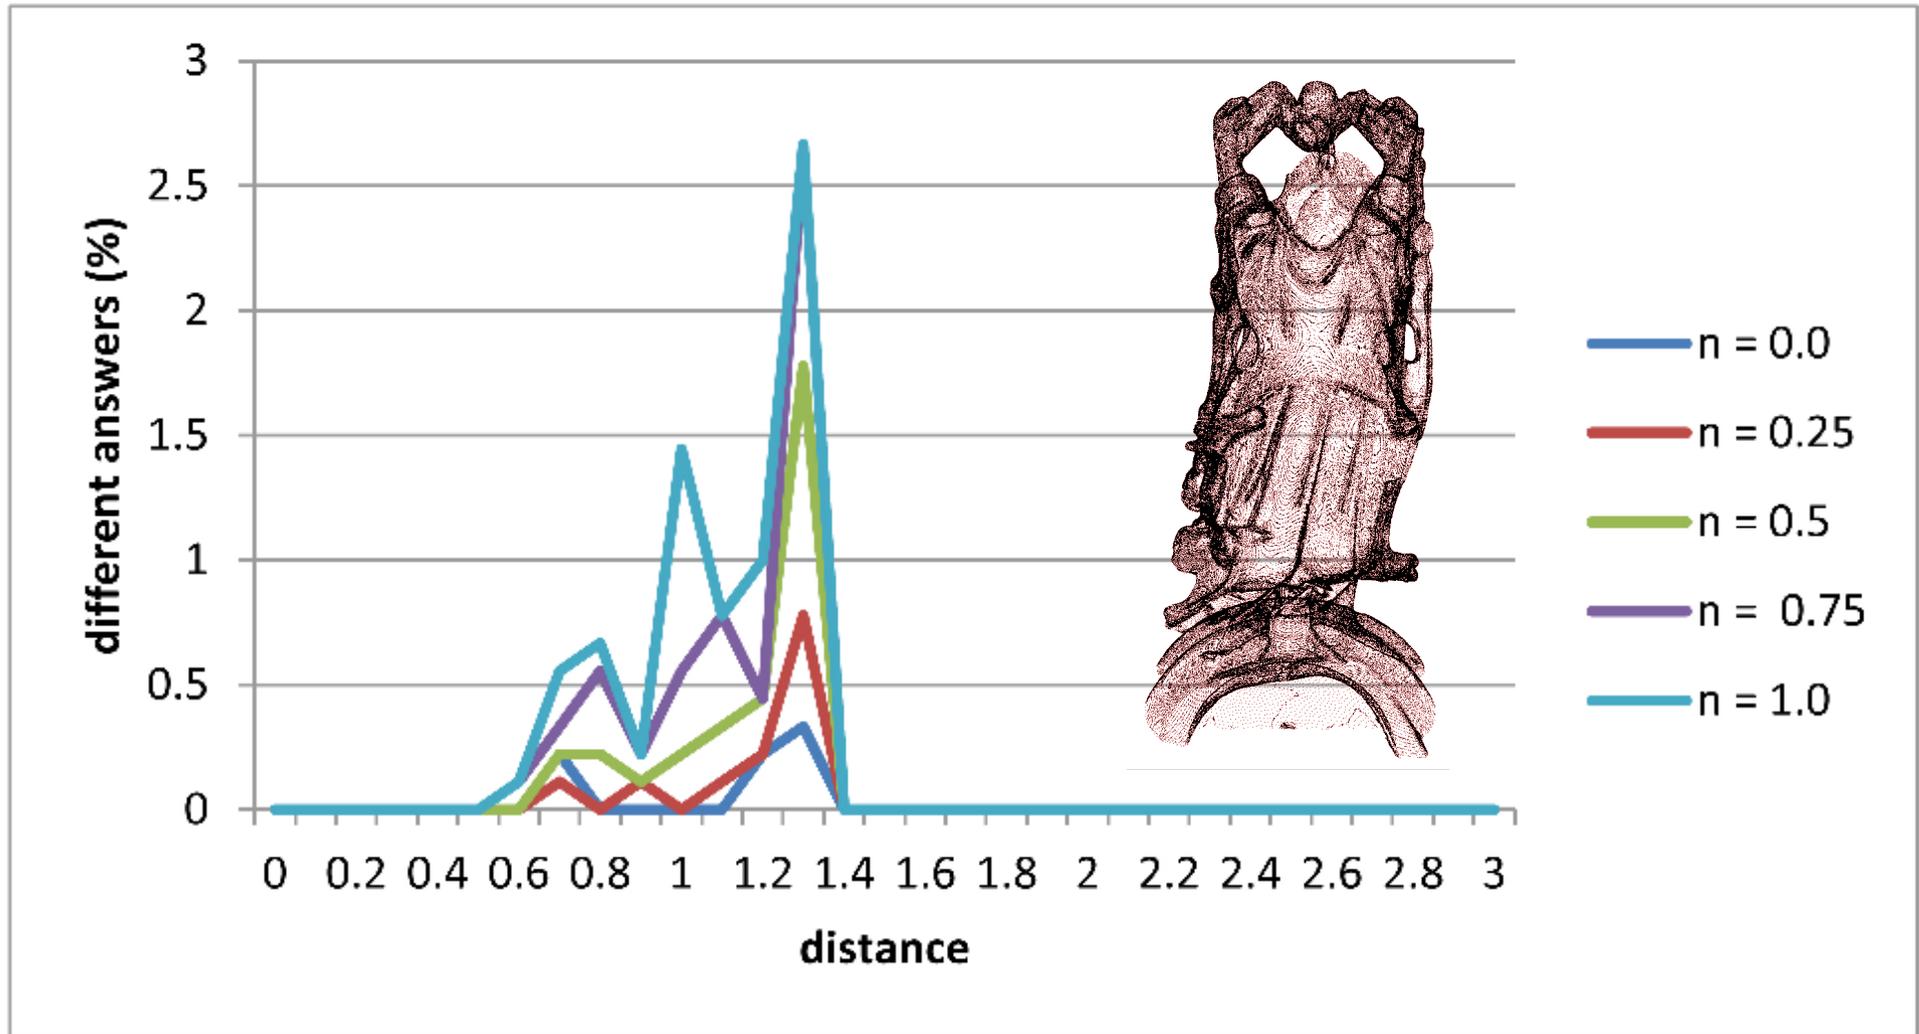
- For each view ray (pixel), test if intervals overlap

In case of non-collision:

- Keep shortest distance between view ray intervals
→ separation distance in view direction



Result 3: Robust to noise



Test with added uniform Gaussian noise $\sigma = nr_{avg}$



- $O(LN)$, L = number of layers (depth complexity)
- Very little output-sensitivity measured

Colliding $m > 2$ objects, adds factor m^2

- But need only construct TLDIs once
- Can also combine small point clouds to reduce m



Novel structure bounds surface of dynamic points

- Real-time, accurate, robust to noise
- Potential other applications: everything which benefits from fast surface queries: GI, ray-tracing

Work in progress

- Compact TLDI data structure
- Speed up construction+queries

