

SVDAG Compression for Segmentation Volume Path Tracing

Mirco Werner[†], Max Piochowiak[†], and Carsten Dachsbacher

(† joint first authors)



Hi, I'm Mirco and I'm presenting our paper SVDAG Compression for Segmentation Volume Path Tracing. I want to start with a quick introduction about segmentation volumes.



Segmentation Volumes



Segmentation Volume: $V(x) = l \in \mathbb{N}$

Segmentation volumes are voxel data sets that are commonly used in a variety of domains like medicine, connectomics, or computational biology. These volumes store an integer label for each voxel. This segments the space into separate object regions, for example neurons in a mouse brain which you can see on the right.

Segmentation Volumes



Segmentation Volume: $V(x) = l \in \mathbb{N}$

Segmentation volumes can become quite large, even for small label counts that allow to store volumes with few bits per voxel. These sizes easily exceed GPU memory, hindering interactive visual exploration and rendering.





Google Neuroglancer Therefore, rendering workflows are typically a two stage process: First, interactive exploration is done in a low-fidelity visualization framework. Pictured here is Google neuroglancer which can display 2D slices of volumes as well as a small number of precomputed meshes of object region.

Brick-Wise Compression

- stores voxels in bricks
- label palette per brick
- $\lceil log_2(\|palette\|) \rceil$ bits per voxel







Google neuroglancer uses a form a brick-wise compression: the neuroglancer precomputed format. This compression technique operates on a voxel basis. The volume is split into small bricks. For each brick, a palette of occurring voxel labels is stored. Voxels inside a brick then index into the palette using only a small number of bits. This voxel format is fast to compute and is suitable for 2D visualization wethods. However, for 3D visualization, it does not directly map to raytracing hardware and typically relies on slower ray marching. More evolved brick-wise formats that achieve stronger compression additionally require larger caches for on-the-fly decompression during rendering.



Neuroglancer Precomputed Format



Blender manual CC-BY-SA 4.0

Therefore, in a second step, often other external software is used to generate high quality renderings for science communication or publication purposes, for example with path tracing in Blender Cycles. To that end, the volume data has to be meshed or converted into another format first.

[Velicky*2023] Dense 4D nanoscale reconstruction of living brain tissue, Nature methods









Speaker notes

This existing process to generate high quality images is therefore typically slow: It requires switching between different data formats and software. The often CPU-bound rendering process takes minutes or hours to render one frame.





Compression for Segmentation Volumes (our data structure)

Speaker notes

With our method, we solve these problems using a compressed data format that allows hardware accelerated GPU path tracing directly on the compressed large scale segmentation volumes. Our method offers precision and compression as in voxel-based representations as well as the rendering performance and hardware suitability of mesh based representations.



Speaker notes

Here I'm showing a video of our final framework. As you can see, we can interactively explore the segmentation volume. Once the camera is static, path tracing noise quickly vanishes and the image converges to the path-traced reference.

• key observation: large uniform areas with the same label



How do we compress the segmentation volumes? For this, let's first take a look at a segmentation volume, on the right side in 3D and on the left side in 2D. You can clearly observe that there are large uniform areas where all voxels are assigned to the same label

- key observation: large uniform areas with the same label
- compression idea: decouple label and voxel occupancy



Based on this observation, our compression idea is to decouple the label information from the voxel occupancy.

- key observation: large uniform areas with the same label
- compression idea: decouple label and voxel occupancy



Therefore, we can define an AABB per label region. All solid voxels that are assigned to this AABB now all implicitly have the same label and the label does not have to be stored for each individual voxel.

- key observation: large uniform areas with the same label
- compression idea: decouple label and voxel occupancy

 - voxel occupancy becomes binary attribute inside region



This allows for much easier and stronger compression of voxel occupancy, since voxel occupancy, i.e. solid/empty, becomes a binary attribute inside the region.

Data Structure - Level 1 (blocks)

• split AABBs into blocks of fixed size (16^3)



To make our compression work, we introduce a three level data structure. The goal of the first level is to decouple the label information from the voxel occupancy. Therefore, we split the AABBs into blocks of fixed size, 16 by 16 by 16.

decouple label from voxel occupancy



Data Structure - Level 1 (BVH w/ blocks)

- split AABBs into blocks of fixed size (16^3)
- BVH for hardware-accelerated raytracing
 - blocks = BVH leaves, store label information



For efficient rendering, we can then build a BVH for hardware-accelerated raytracing. Our blocks are the BVH leaves and store the label information. All solid voxels that are assigned to an block now all implicitly have the same label.





Data Structure - Level 2 (SVOs)

• sparse voxel octree (SVO) [Laine*2010] in each block



Speaker notes

Our second level of our data structure now encodes the actual voxel occupancy inside the blocks. Our first idea was to construct a SVO in each block. You can see a visualization of this on the left, where each block that contains the label information, here colored, stores this SVO.

encode voxel occupancy



Data Structure - Level 2 (SVOs)

 sparse voxel octree (SVO) [Laine*2010] in each block identical subtrees



Now when we take a look at the voxel occupancy inside each block, we can see that there are often identical subregions. This means, the SVO contains identical subtrees and therefore stores redundant information.



Data Structure - Level 2 (SVDAGs)

 sparse voxel directed acyclic graph (SVDAG) [Kämpe*2013] in each block



Luckily, Kämpe and colleques found a solution for this by using a SVDAG instead. The idea here is that equal subtrees from the SVO are merged to transform this tree structure into a directed acyclic graph. This process is again shown here on the left with some of the now unused original SVO nodes grayed out and re-directed pointers. After this step we now have one SVDAG in each block.



Data Structure - Level 2 (SVDAGs)

- sparse voxel directed acyclic graph (SVDAG) [Kämpe*2013] in each block
 - identical subtrees across blocks with different labels
 - (no label information stored)



Of course, there are also identical subtrees across blocks with different labels. Since we just encode the binary voxel occupancy in the SVDAGs, we can merge all these individual SVDAGs...



Data Structure - Level 2 (SVDAG)

single large SVDAG across all blocks



into a single large SVDAG that is shared across all blocks. Again visualized here on the left.



store occupancy directly in 4^3 SVDAG leaf in bitmask $4^3 = 64$ bits

less VRAM accesses

The goal of our third and last level is to improve traversal performance. As you can imagine, loading large SVDAG nodes during traversal introduces many unaligned and random VRAM accesses. We can re-define an area of 4 by 4 by 4 voxels as an SVDAG leaf node. The voxel occupancy inside a leaf can be stored directly in the leaf node using a bitmask using just 64 bits. We call this bitmask or bitfield occupancy field. With this optimization less VRAM accesses are required during traversal since we reach leaf nodes quicker with less steps.



Data Structure



Compressed Segmentation Volume Path Tracing (traversal of our data structure)

Speaker notes

Now let's continue with how we perform path tracing on our compressed segmentation volume. So how we traverse our data structure.

GPU Traversal



You might have already noticed that this data structure maps really well to GPU hardware. This means for our first level, we can just input our blocks into a BVH builder and then use hardware-accelerated ray tracing to find intersections with the blocks. Once an block is intersected with, a custom intersection shader is invoked.



hardware-accelerated ray tracing invokes custom intersection shader

GPU Traversal



Speaker notes

In this custom intersection shader, we query the label information to load material properties and query the pointer to the SVDAG root node. Using multi-level 3D-DDA we can then manually traverse the SVDAG and the occupancy fields.



Level of Detail





block (16^3)

SVDAG leaf (4^3)

Speaker notes

Our method naturally allows for a level of detail scheme. Here on the slide you can see three possible levels of detail. On the left, we stop traversal directly after hitting a block. In the center, we traverse the SVDAG all the way down to a leaf node. And on the right, the individual voxels can be rendered by also traversing the occupancy fields. It's of course possible to dynamically choose the level of the detail based on the distance from the intersected block to the camera.



occupancy field (1^3)

Transfer Function Editing

Settings		I a kin	Y	1 1× 1
oxel (X,Y,Z) = (7298,7834,-3095) osition (X,Y,Z) = (7298.374,7834.136,- otation (r,theta,phi) = (0.000,2.562,2	-3094, 255) 2. 005)		19.	
rames: 104/1024		Sea 1		
.024	- + Max Frames	7 . 7	79 104.	
	Trace Manalan			
SMMMA_CORRECTION ONLY	Tone-Mapping			
SEGMENTATION_VOLUMES	Renderer	N A No	And the second sec	IN Fish
	- + Samples per Pixel per Fr			
32	- + Max Bounces	Pail In a		AND IT
🖊 Level of Detail (LOD)				1 Martin Part
LOD: Disable on First Frame				ere the state
LOD: Smooth Level Decision		2 de	S. S	
Spheres			Part A	
► Environment		-1-1	SI NORP	AL AL
🗸 Scene Hierarchy				
nable All			Nor The The	Au
Disable All				
<pre>v mouse/svdag_occupancy_field_test</pre>			N N Y	1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1
neuroni				1 . 1
reuron2		20		
			Nº A WAY	
Peuron14		Contraction of the second seco	The 2/2	
Peurop17			2 A A A	and the second of the
v neuron19			1	A State A State
		Jac 20-18 martin	The of the second	
▼ Materials		200	A STATE AND A STATE OF	
n_folder/m_file/0/1 (1)	Material			Contraction of the second second
R:145 G: 0 B	Base Color			F. Start
0.000	Specular Transmission	81	111	
0.000	Metallic	7 3 43	A TACK I	
	Subsurface		States of the	ALL MRXL
0.000	and the second sec		V 1 K 12A	TLEK VE
0.906 0.906	Specular			
0,909 0,909 0, <u>5</u> 09	Roughness		S. 14 1.193	A ALA ALA
0.000 0.000 0. <u>50</u> 0 0.000	Specular Roughness Specular Tint	1 1.50	- Contract	
0.909 0.909 0.559 0.999 0.999	Specular Roughness Specular Tint Anisotropic	1 ACT		
0.909 0.909 0.509 0.909 0.909 0.909	Specular Roughness Specular Tint Anisotropic Sheen	1 AC		
0,909 0,909 0,509 0,909 0,909 0,909 0,909 0,909	Specular Roughness Specular Tint Anisotropic Sheen Sheen Tint	H		
0.000 0.000 0.500 0.000 0.000 0.000 0.000	Specular Roughness Specular Tint Anisotropic Sheen Sheen Tint Clearcoat	H	and a	A

Speaker notes

Another interseting property of our method is that it allows for interactive editing of the transfer function. We can quickly adjust the material attributes corresponding to a specific label by updating the material buffer.



Transfer Function Editing

1024 ▼ Rendering GAMMA_CORRECTION ONLY SEGMENTATION_VOLUMES 1 32 ✓ Level of Detail (LOD) ✓ LOD: Disable on First Frame LOD: Smooth Level Decision Spheres ► Environment ▼ Scene Hierarchy Enable All	 + Max Frames Tone-Mapping Renderer + Samples per Pixel per Frame + Max Bounces
 Kendering GAMMA_CORRECTION ONLY SEGMENTATION_VOLUMES 1 32 ✓ Level of Detail (LOD) ✓ LOD: Disable on First Frame LOD: Smooth Level Decision Spheres > Environment ▼ Scene Hierarchy Enable All 	 ▼ Tone-Mapping ▼ Renderer - + Samples per Pixel per Frage - + Max Bounces
SEGMENTATION_VOLUMES SEGMENTATION_VOLUMES Level of Detail (LOD) Level of Detail (LOD) LoD: Disable on First Frame LOD: Smooth Level Decision Spheres Environment Scene Hierarchy Enable All	 Renderer + Samples per Pixel per Frances + Max Bounces
1 32 ✓ Level of Detail (LOD) ✓ LOD: Disable on First Frame LOD: Smooth Level Decision Spheres ► Environment ▼ Scene Hierarchy Enable All	- + Samples per Pixel per Fra - + Max Bounces
32 ✓ Level of Detail (LOD) ✓ LOD: Disable on First Frame LOD: Smooth Level Decision Spheres ► Environment ✓ Scene Hierarchy Enable All	- + Max Bounces
 Level of Detail (LOD) LOD: Disable on First Frame LOD: Smooth Level Decision Spheres Environment Scene Hierarchy Enable All 	
 LOD: Disable on First Frame LOD: Smooth Level Decision Spheres Environment Scene Hierarchy Enable All 	
LOD: Smooth Level Decision Spheres ► Environment ▼ Scene Hierarchy Enable All	
Spheres ► Environment ▼ Scene Hierarchy Enable All	
Environment Scene Hierarchy Enable All	
Enable All	
Disable All	
▼ mouse/svdag_occupancy_field_test	
v neuron1	
neuron2	
<pre>v neuron3 v neuron18 v neuron18</pre>	
v neuron14	
v neuron17	
v neuron19	
v neuron20	
▶ Materials	R. (1931)
Camera	
► Utility ▼ Timings	
SegmentationVolumesPassPathtrace	7.938
PassTonemapper	0.197
0.00 / 0.00	

Speaker notes

When we would like to disable entire label regions, we can disable the corresponding blocks and quickly re-build the BVH to achieve an optimized traversal performance.



Results

Speaker notes

Now let's take a look at some results.

29

Datasets



Speaker notes

We evaluated our method on these three datasets. The rather small CELLS dataset with only a raw size of 4 GB and the large CELEGANS and MOUSE datasets with up to 150 GB raw size.

112.8 GB) MOUSE (raw 150.0 GB)

	blocks + BVH + SVOS		
	mem. [MB]	CR	
Cells	253.97	6.35%	
C.ELEGANS	14322.37	12.7%	
Mouse	62911.29	41.95%	



Let's start with the memory and the compression rate. Remember that the reported memory is exactly the memory required for the rendering on the GPU.









The total memory consists of a comparitively small share required for the blocks and for the constructed BVH. Most memory is required for the SVO or SVDAG to encode the voxel occupancy.





	blocks + BVH + SVOS		
	mem. [MB]	CR	
Cells	253.97	6.35%	
C.ELEGANS	14322.37	12.7%	

62911.29 41.95% MOUSE



With our initial idea of using SVOs in each block you can see...









...that the MOUSE cortex requires a total of 63 GB to encode the voxel occupancy which results in a poor compression rate of only 41.95%.







CELLS	253.97	6.35%
C.ELEGANS	14322.37	12.7%
Mouse	62911.29	41.95%



Speaker notes

When we use our large shared SVDAG instead...

blocks + BVH + SVDAG
mem. [MB]CR86.342.16%2534.282.25%12095.488.06%





CELLS	253.97	6.35%
C.ELEGANS	14322.37	12.7%
Mouse	62911.29	41.95%



Speaker notes

...the MOUSE cortex requires only 12 GB of memory which is an overall compression rate of 8%.

blocks + BVH + SVDAG
mem. [MB]CR86.342.16%2534.282.25%12095.488.06%





CELLS	253.97	6.35%
C.ELEGANS	14322.37	12.7%
Mouse	62911.29	41.95%



And for the other two datasets we achieve even compression rates up to 2%. These volume sizes now fit into the VRAM of a modern high-end consumer GPU.



blocks + BVH + SVDAG mem. [MB] CR

86.34 2534.28 12095.48





		SVOs [ms]		
		finest LOD	coarsest LOD	
lce	CELLS	8.19	0.4	
1 bour	C.ELEGANS	—	_	
	Mouse	—	_	
IC.	Cells	11.42	0.81	
pn	C.ELEGANS	_	_	
32	Mouse	_	_	



Speaker notes

Moving on to rendering performance, we show here path tracing with 1 bounce in the first row and up to 32 bounces using throughput-based Russian roulette in the second row for the three datasets. Also we report numbers either using the finest LOD, i.e. terminate traversal directly when hitting a block. Again, using our initial SVOs we can only render the CELLS dataset since the other two exceed our GPU memory.



		SVOs [ms]		SVD.	AG [ms]
		finest LOD	coarsest LOD	finest LOD	coarsest L
JCe	CELLS	8.19	0.4	2.95	0
Ino	C.ELEGANS	_	—	5.61	0
1 p	Mouse	—		126.73	0
IC.	CELLS	11.42	0.81	5.31	1
bn	C.ELEGANS	_	_	9.25	0
32	Mouse	_	_	189.54	



Speaker notes

When using the SVDAG...

OD 0.51 0.55 0.98 1.02 0.98 2.0



		SVOs [ms]		SVD.	AG [ms]
		finest LOD	coarsest LOD	finest LOD	coarsest LOD
Jce	Cells	8.19	0.4	2.95	0.51
Ino	C.ELEGANS	_	_	5.61	0.55
1 þ	Mouse	_	_	126.73	0.98
IC.	CELLS	11.42	0.81	5.31	1.02
bn	C.ELEGANS	_	_	9.25	0.98
32	Mouse	_	_	189.54	2.0



Speaker notes

...you can see that the frametime is reduced significantly.



		SVOs [ms]		SVD.	AG [ms]
		finest LOD	coarsest LOD	finest LOD	coarsest LOD
lce	Cells	8.19	0.4	2.95	0.51
Ino	C.ELEGANS	—	—	5.61	0.55
1 b	Mouse	_	_	126.73	0.98
c.	CELLS	11.42	0.81	5.31	1.02
hn	C.ELEGANS	_	_	9.25	0.98
32	Mouse	_	_	189.54	2.0



Speaker notes

The CELLS and CELEGANS datasets can be fully path traced even with up to 32 bounces with over 60 FPS on this lossless representation.



		SVOs [ms]		SVD.	AG [ms]
		finest LOD	coarsest LOD	finest LOD	coarsest LOD
Jce	CELLS	8.19	0.4	2.95	0.51
Ino	C.ELEGANS	—	_	5.61	0.55
1 b	Mouse	_	_	126.73	0.98
c.	CELLS	11.42	0.81	5.31	1.02
hn	C.ELEGANS	_	_	9.25	0.98
32	Mouse	_	_	189.54	2.0



Speaker notes

And still, the much more sparse and harder to ray trace MOUSE dataset can still be interactively path-traced to some degree.



		SVC	DS [ms]	SVDAG [ms]		
		finest LOD	coarsest LOD	finest LOD	coarsest LOD	
lce	CELLS	8.19	0.4	2.95	0.51	
out	C.ELEGANS	_	-	5.61	0.55	
1 p	MOUSE	_	_	126.73	0.98	
c.	Cells	11.42	0.81	5.31	1.02	
hn	C.ELEGANS	_	_	9.25	0.98	
32	Mouse	_		189.54	2.0	



Speaker notes

By using dynamic LOD or by always using the coarsest LOD while the camera is moving, we can ensure exploration of the MOUSE dataset with more than 60 FPS Once the camera is static, the finest LOD can be used to sender the converged image. On the right side you can see the MOUSE cortex with coarsest LOD and with the finest lossless LOD. For these large datasets, from further away, you can barely see a differences. This means that using this coarser LOD during exploration is indeed a viable option.

Speaker notes

In conclusion, we are able to unify this process of exploring and rendering large segmentation volumes into a single framework...

44

- introduced a
 - Iossless SVDAG compression method

Speaker notes

... by introducing a lossless SVDAG compression method...



C.ELEGANS (raw 112.8 GB)

- introduced a
 - Iossless SVDAG compression method
 - and a hardware-accelerated path tracing framework

Speaker notes

... and a hardware-accelerated path tracing framework...



46

- introduced a
 - Iossless SVDAG compression method
 - and a hardware-accelerated path tracing framework
- for interactive rendering of segmentation volumes

Speaker notes

that allows for interactive rendering of segmentation volumes.



C.ELEGANS (raw 112.8 GB)

- introduced a
 - Iossless SVDAG compression method
 - and a hardware-accelerated path tracing framework
- for interactive rendering of segmentation volumes

strong compression and rendering performance

We thereby achieve strong compression rates and high rendering performance.



C.ELEGANS (raw 112.8 GB)

Thank you!



Speaker notes

Our code is publicly available on GitHub. With that said, thank you for your attention!



Code available on GitHub:



https://github.com/MircoWerner/ SegmentationVolumeCompression

References

- [Rosenbauer*2020]
 - ROSENBAUER, J., BERGHOFF, M., and SCHUG, A. "Emerging Tumor Development by Simulating Single-cell Events". bioRxiv (2020). DOI: 10.1101/2020.08.24.264150
- [Witvliet*2021]
 - WITVLIET, D., MULCAHY, B., MITCHELL, J. K., et al. "Connectomes across development reveal principles of brain maturation". Nature 596.7871 (2021), 257-261. DOI: 10.1038/s41586-021-03778-8
- [Motta*2019]
 - MOTTA, A., BERNING, M., BOERGENS, K. M., et al. "Dense connectomic reconstruction in layer 4 of the somatosensory cortex". Science 366.6469 (2019), eaay3134. DOI: 10.1126/science.aay3134
- [Velicky*2023]
 - VELICKY, P., MIGUEL, E., MICHALSKA, J.M. et al. Dense 4D nanoscale reconstruction of living brain tissue. Nat Methods 20, 1256-1265 (2023). https://doi.org/10.1038/s41592-023-01936-6
- [Laine*2010]
 - LAINE, S. and KARRAS, T. "Efficient sparse voxel octrees". Proc. ACM SIGGRAPH Symposium on Interact. 3D Graph. and Games. New York, NY, USA: ACM, 2010, 55-63. DOI: 10.1145/1730804.1730814
- [Kämpe*2013]
 - KÄMPE, V., SINTORN, E., and ASSARSSON, U. "High resolution sparse voxel DAGs". ACM Transactions on Graphics 32.4 (July 2013). DOI: 10.1145/2461912.2462024

Appendix

Speaker notes

51

SVDAG Node Memory Layout

child pointer 0	
child pointer 1	
child pointer 2	
child pointer 3	
child pointer 7	
8× 32 bit	

leaf flag (0xFFFFFFF)	
occupancy field	
occupancy field	
empty	
•••	
empty	
8×32 bit	

Datasets



We evaluated our method on these three datasets. The rather small CELLS dataset with only a raw size of 4 GB and the large CELEGANS and MOUSE datasets with up to 150 GB raw size.

Datasets



Speaker notes

As you can see in the bottom row, the CELLS dataset contains more than 100k labels while the other two datasets contain only a few hundreds of labels. However, keep in mind, that since we decouple labels and voxel occupancy, our compression of the voxels works and scales independently from the actual number of labels.

	AABBs + BVH +	SVOs		SVDAG	
	mem. [MB]	mem. [MB]	CR	mem. [MB]	
CELLS	5.63 + 4.42	+ 243.92	6.35%	+ 76.29	2
C.ELEGANS	248.13 + 186.24	+13888.0	12.7%	+ 2099.91	2
Mouse	1032.11 + 786.44	+ 61092.74	41.95%	+ 10276.93	8





Rendering Performance [ms] using finest (coarsest) LOD

		SVOs	w/ occupancy field SVDAG	w/o occupa
1 bounce	Cells C.Elegans Mouse	8.19 (0.4) 	2.95 (0.51) 5.61 (0.55) 126.73 (0.98)	5.8 12.7 134.7
32 bnc.	Cells C.Elegans Mouse	11.42 (0.81) _ _	5.31 (1.02) 9.25 (0.98) 189.54 (2.0)	8.0 18.7 203.0



- ancy field **SVDAG**
- 87 (0.42) 75 (0.47) 72 (0.78)
- 66 (0.85) 72 (0.83) 65 (1.67)

Rendering Performance [ms] using finest (coarsest) LOD

		SVOs	w/ occupancy field SVDAG	w/o occupa
1 bounce	Cells C.Elegans Mouse	8.19 (0.4) 	2.95 (0.51) 5.61 (0.55) 126.73 (0.98)	5.8 12.7 134.7
32 bnc.	Cells C.Elegans Mouse	11.42 (0.81) _ _	5.31 (1.02) 9.25 (0.98) 189.54 (2.0)	8.0 18.7 203.0

- ancy field SVDAG
- 87 (0.42) 75 (0.47) 72 (0.78)
- 66 (0.85) 72 (0.83) 65 (1.67)
- MOUSE (raw 150.0 GB)
 - 15.57% solid 96 labels

Compression Timings

	CELLS	C.ELEGANS	Mouse
SVOS (AABB) CPU [s]	5.97	905.26	4127.16
SVDAGS (TYPE) CPU [s]	2.24	117.92	683.45
SVDAG (SHARED) CPU [ms] GPU [ms]	408.49 101.96	33089.85 2865.33	147721.0

Table 5: Timings for each step to compress segmentation volumes: Computing AABBs and their SVOs from lists of solid voxel positions per label, merging them to SVDAGS (TYPE), and subsequently merging these partial SVDAGs to one SVDAG (SHARED). We provide a GPU implementation of the last step for fast merging of visualized types during rendering start-up or runtime.

CSGV [Piochowiak*2023] (memory)

	blocks + BVH	I + SVOS	blocks + BVH	+ SVDAG
	mem. [MB]	CR	mem. [MB]	CR
CELLS	253.97	6.35%	86.34	2.16%
C.ELEGANS	14322.37	12.7%	2534.28	2.25%
Mouse	62911.29	41.95%	12095.48	8.06%

CSGV

15.4 736.98 1910.1

Table 2: Memory and compression rate of CSGV [PD24]. Theoretical compression rate when including a GPU cache size to decompress visible regions to fitting levels of detail during rendering in ().

mem. [MB]		CSC	GV CR
6	(+544)	0.39%	(13.99%)
1	(+768)	0.65%	(1.33%)
7	(+768)	1.27%	(1.79%)

CSGV [Piochowiak*2023] (rendering performance)

		SVOs [ms]		SVD.	AG [ms]	
		finest LOD	coarsest LOD	finest LOD	coarsest LOD	CSGV
1 bounce	CELLS	8.19	0.4	2.95	0.51	23.21 (2.28)
	C.ELEGANS Mouse	_	_	5.61 126.73	0.55 0.98	87.18 (16.56) 124.89 (8.47)
bnc.	CELLS	11.42	0.81	5.31	1.02	25.64 (2.29)
	C.ELEGANS	—	—	9.25	0.98	97.63 (17.06)
32	Mouse	_	_	189.54	2.0	136.16 (10.55)