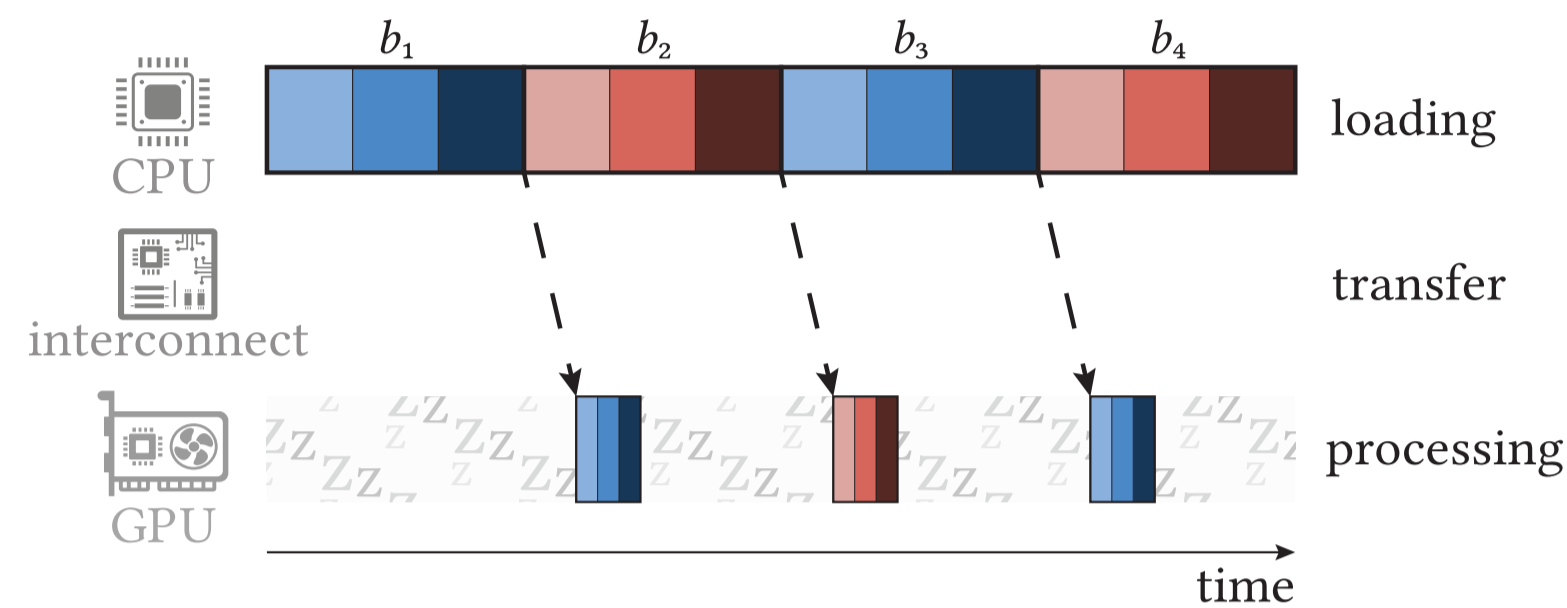
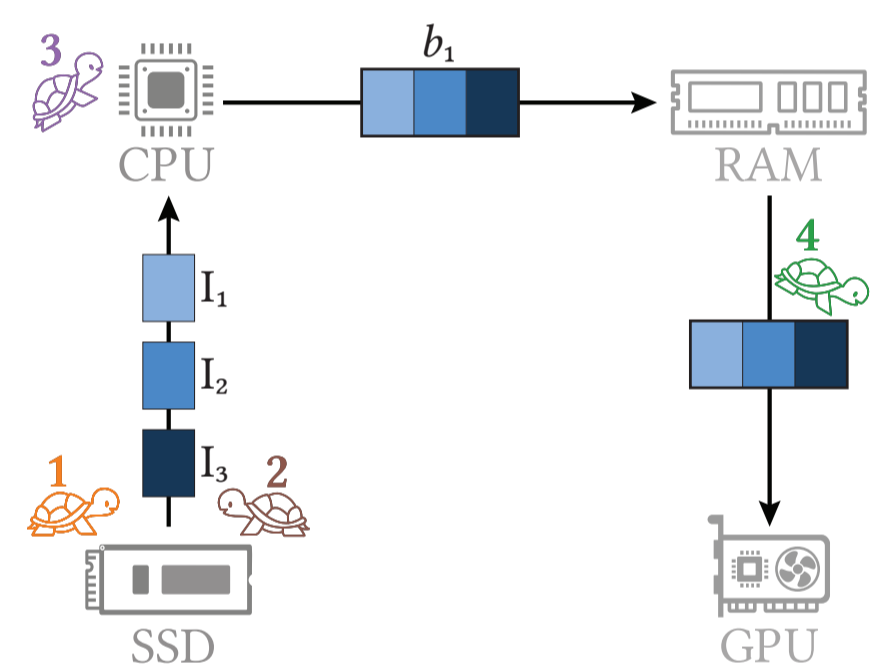


## Data loading bottlenecks

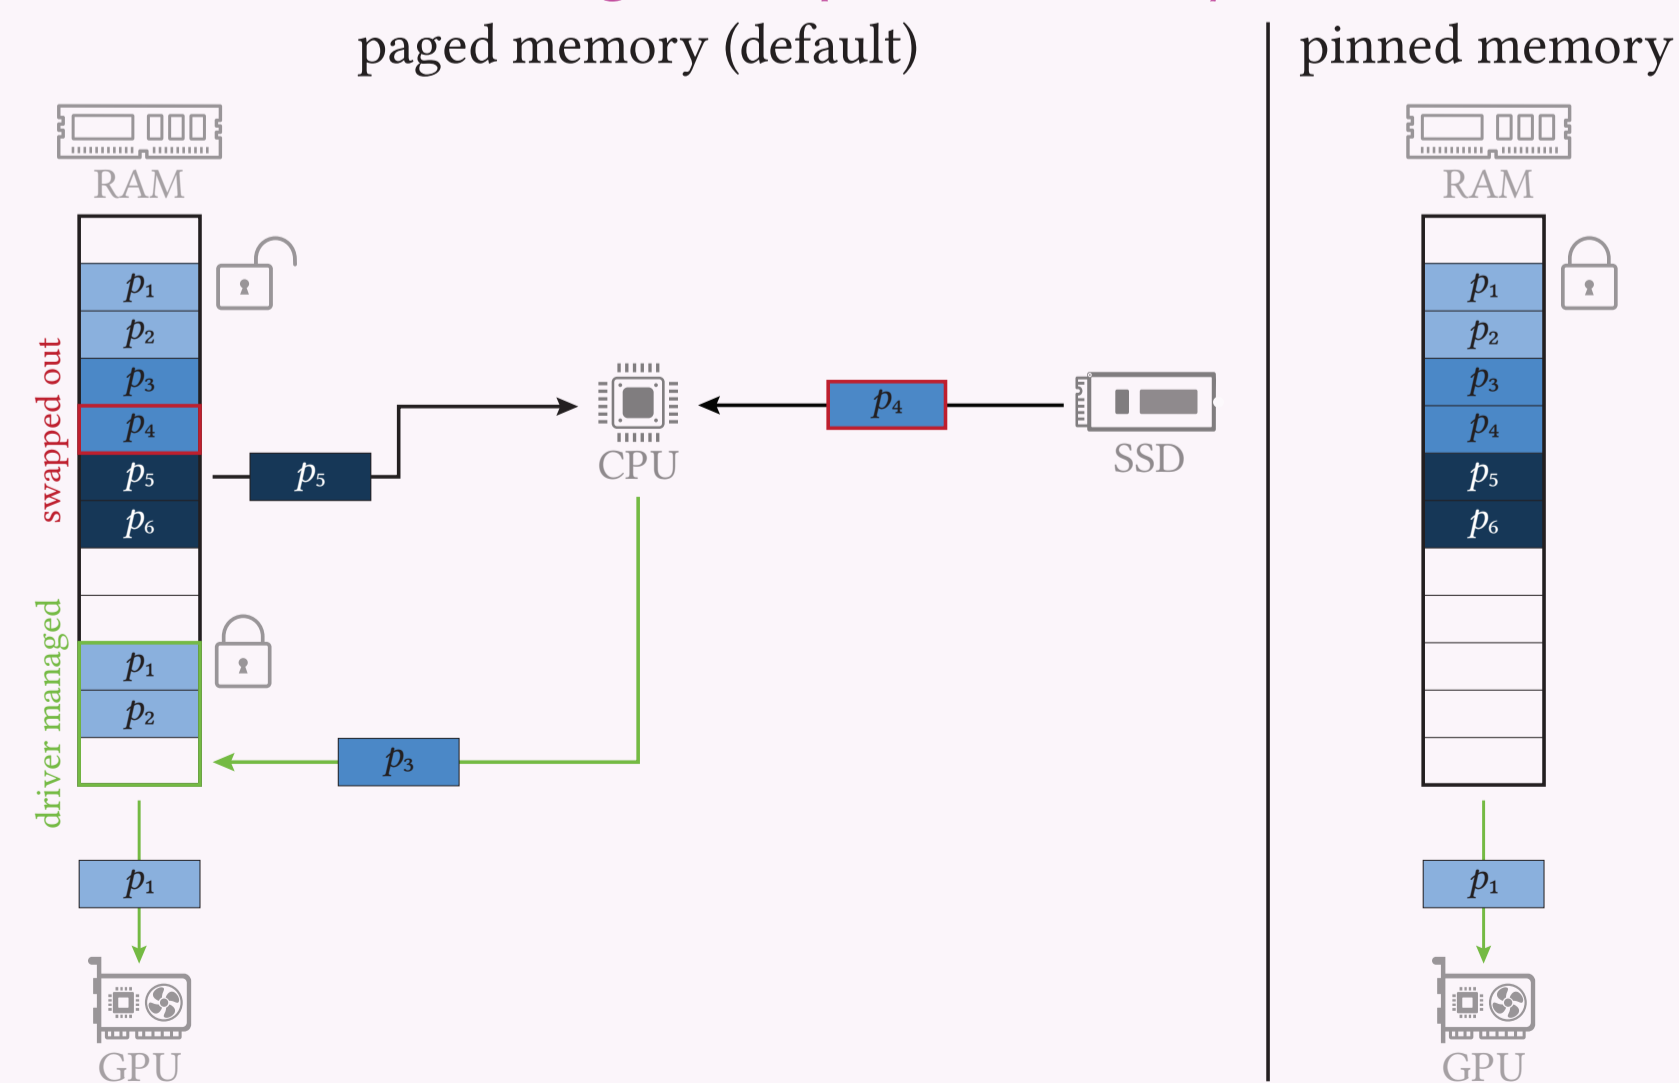


In high-throughput model training, the GPU is subject to idle times if the CPU doesn't deliver the data in time (CPU- and I/O-bound training).



Several bottlenecks retard the data loading from the SSD to the GPU.

### Background: pinned memory




With pinned memory, data can be moved directly from the RAM to the GPU since it is ensured that the memory pages aren't swapped out by the OS. This avoids unnecessary data transfers.

## Independent counter-strategies

- 0: original**  
 Data samples of shape (480, 640, 100) = (height, width, channels) with float32 values.
- 1: +blosc**  
 Storing data compressed on disk reduces disk load. An efficient decompressor (e.g. BLOSC<sup>1</sup>) should be used to avoid long decompression times.
- 2: +fp16**  
`dtype=torch.float16`  
 Minimizing input data precision whilst maintaining its information content reduces loading times from disk and transfer times to the GPU.
- 3: +gpu-aug**  
`device='cuda'`  
 If the CPU is the bottleneck, moving computations to the GPU (e.g. via Kornia<sup>2</sup>) frees resources on the CPU.
- 4: +ring-buffer**  
 Using a fixed (initialized on training start), shared (accessible by all workers) and pinned memory buffer avoids unnecessary memory transfers.

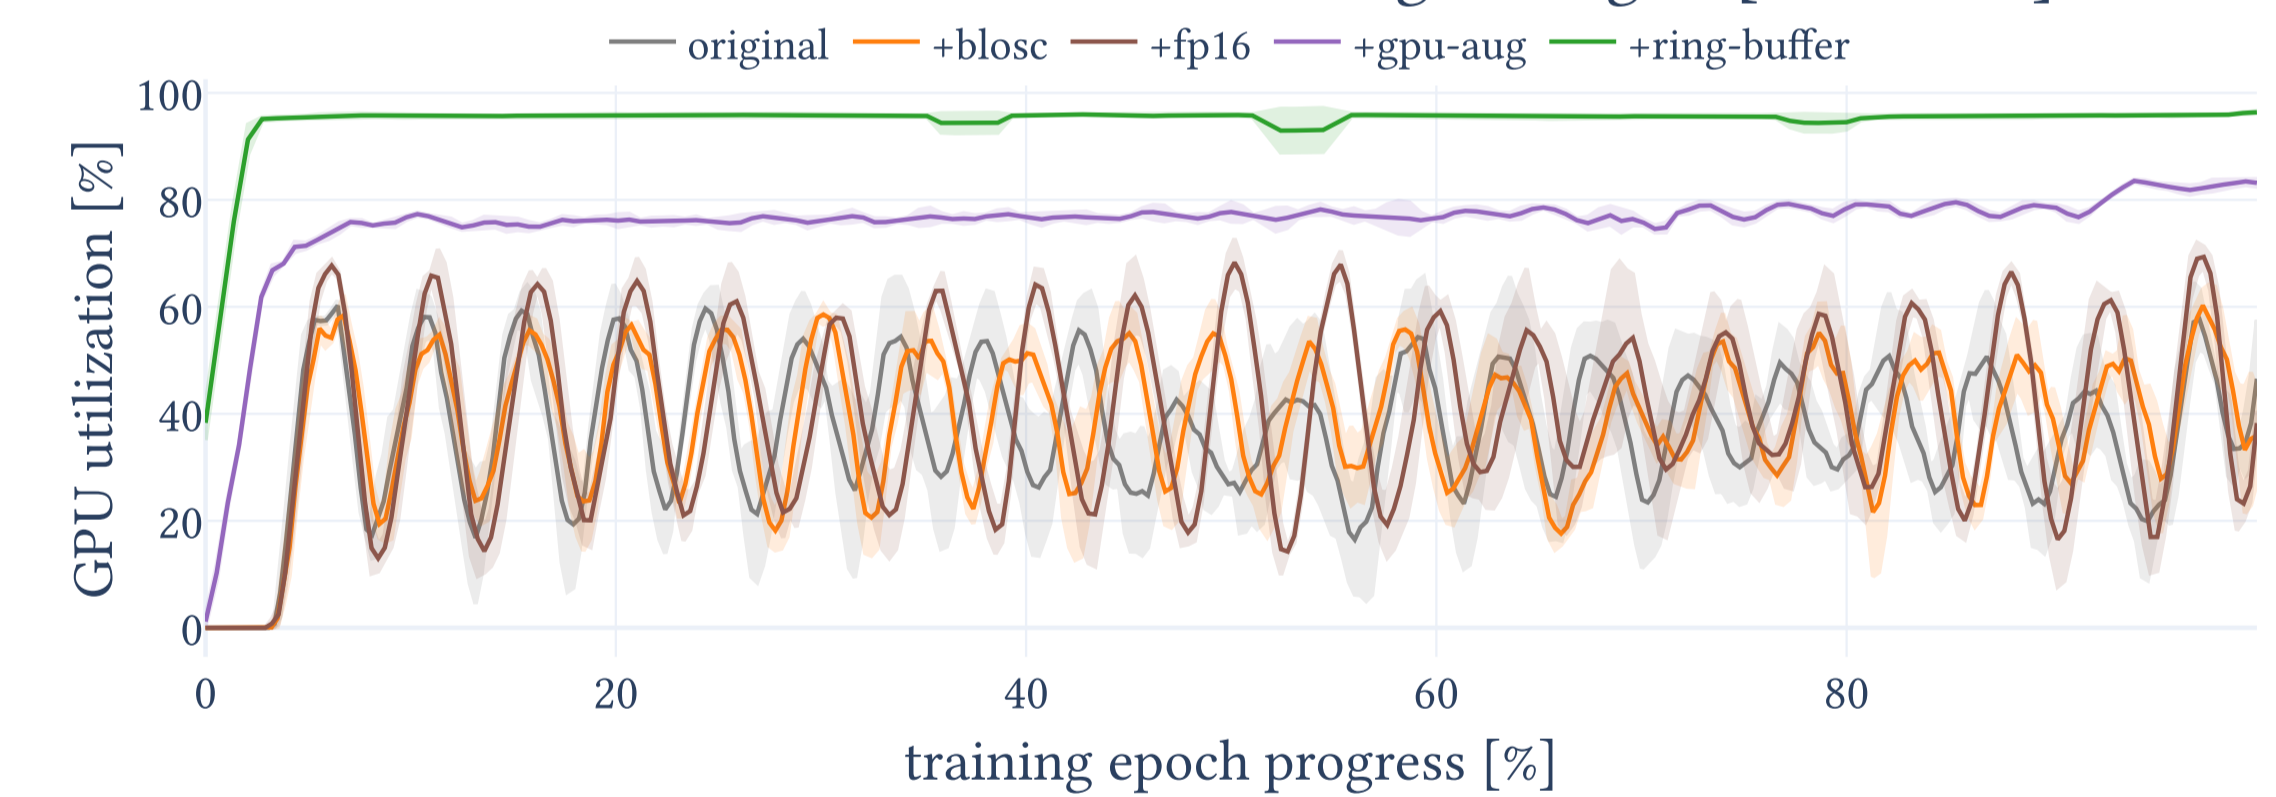
## Results

### Benchmarking of different data loading strategies



Through the presented counter-strategies, I/O and CPU load is reduced and training is up to 3.6 times faster.

### GPU utilization of different data loading strategies [unlimited]



If data is not delivered in time, GPU utilization is suboptimal. In these cases, alternating loading and waiting times can be observed.

### Application: Hyperspectral Imaging (HSI) in surgery

Postoperative death is the third leading cause of death as it accounts for 7.7% of all deaths globally.

Hyperspectral Tissue Classification

Hyperspectral Imaging extends the human vision by removing the arbitrary restriction of capturing only red, green and blue colors. Instead, a fine-grained reflectance spectrum is recorded, providing additional information about tissues and their functional properties (e.g. oxygenation).

<sup>1</sup> <https://www.blosc.org>  
<sup>2</sup> <https://kornia.readthedocs.io>