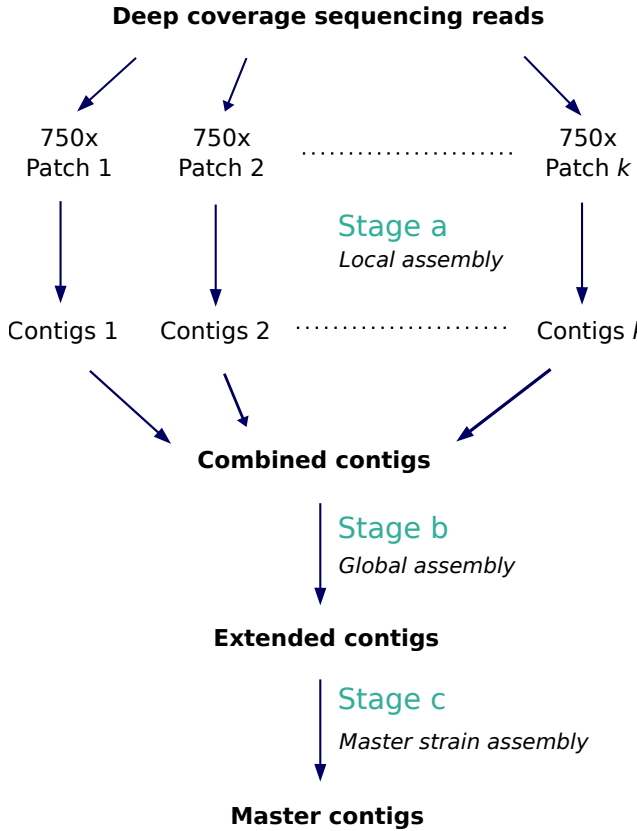

Supplemental Fig S1: Deep coverage datasets are handled by splitting them into k patches of 750x each, because SAVAGE *Stage a* has been tailored to perform optimally on such coverage levels. For the simulated benchmarks and the lab mix $k = 30$, for the ZIKV data $k = 50$, and for the HCV data $k = 75$. The contigs resulting from *Stage a* are combined into one collection of contigs on which we apply SAVAGE *Stage b*. This gives us a set of maximally extended contigs, which can then (optionally) be merged into master strains using *Stage c*.