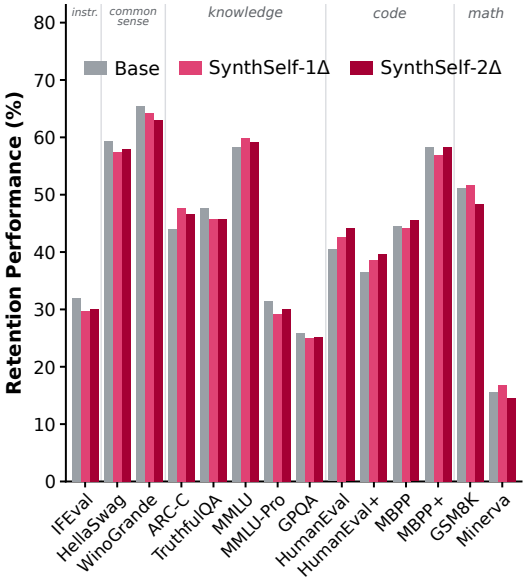
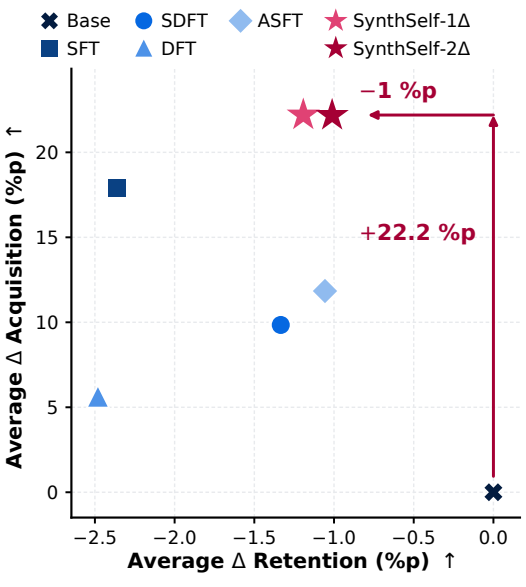


(a) Averaged across 3 base models, continual post-training via SYNTHSELF on the medical domain results in near-zero forgetting across all 14 benchmarks.



(b) Across 3 base models, 4 acquisition benchmarks, and 14 retention tasks, SYNTHSELF maximizes performance gains with minimal forgetting.