Supplemental Figures

Figure S1

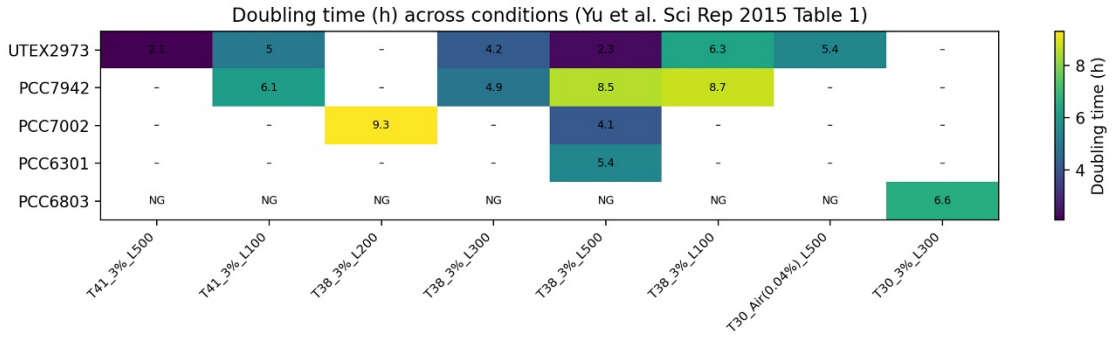


Figure S1: **Doubling-time benchmarks across cyanobacterial chassis and growth conditions (compiled from Yu et al., also in Supplemental Table 2).** Heatmap of reported doubling times (h) for multiple cyanobacterial strains across temperature (30–41 °C), CO_2 level (air vs 3% CO_2), and light intensity (100–500 µmol photons m^{-2} s^{-1}). Cells are annotated with the reported doubling time; "NG" indicates no growth under the specified condition; "–" indicates condition not reported/measured. This figure provides condition-matched context for comparing achievable growth throughput across commonly used chassis.

Figure S2

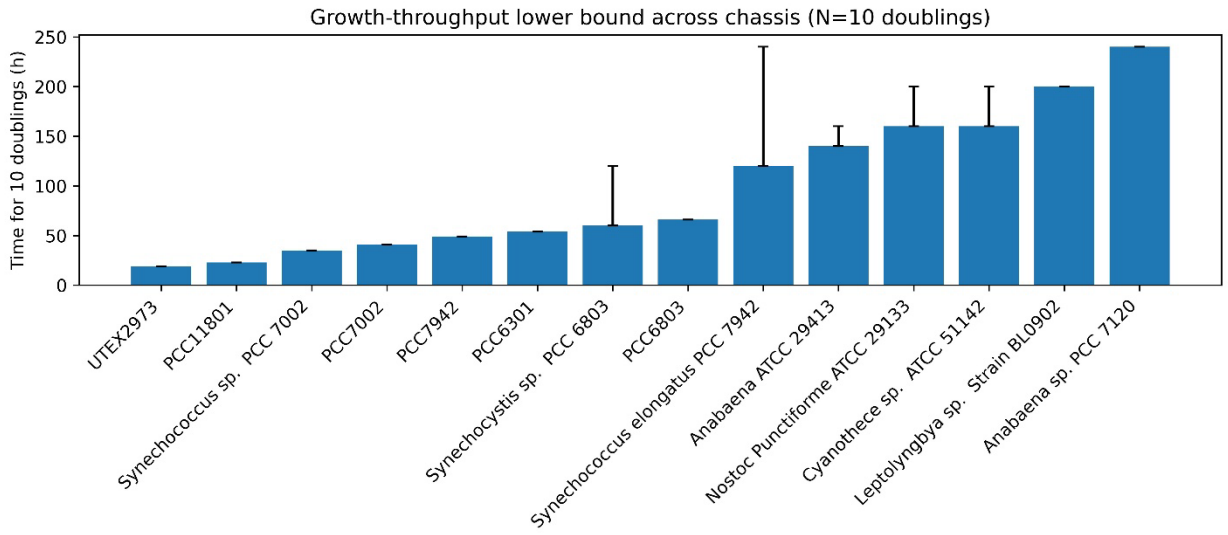


Figure S2: Growth-throughput lower bound across chassis from best-case doubling times, with literature-range whiskers (N = 10 doublings). Bars show the minimum reported doubling time for each chassis (Table S1) converted to a lower bound on growth-only throughput as the time required for N = 10 doublings (time_min = N × doubling_time_min). One-sided whiskers extend upward where available to the corresponding maximum reported doubling time under alternate conditions/sources (time_max = N × doubling_time_max), visualizing cross-condition/literature variability rather than within-experiment uncertainty. Accordingly, large whiskers indicate chassis for which published doubling times span a wide range (e.g., due to differences in CO_2, temperature, light, medium, or reporting source), while missing/short whiskers indicate cases where only a single best-case value (or a narrow range) was available.

Figure S3

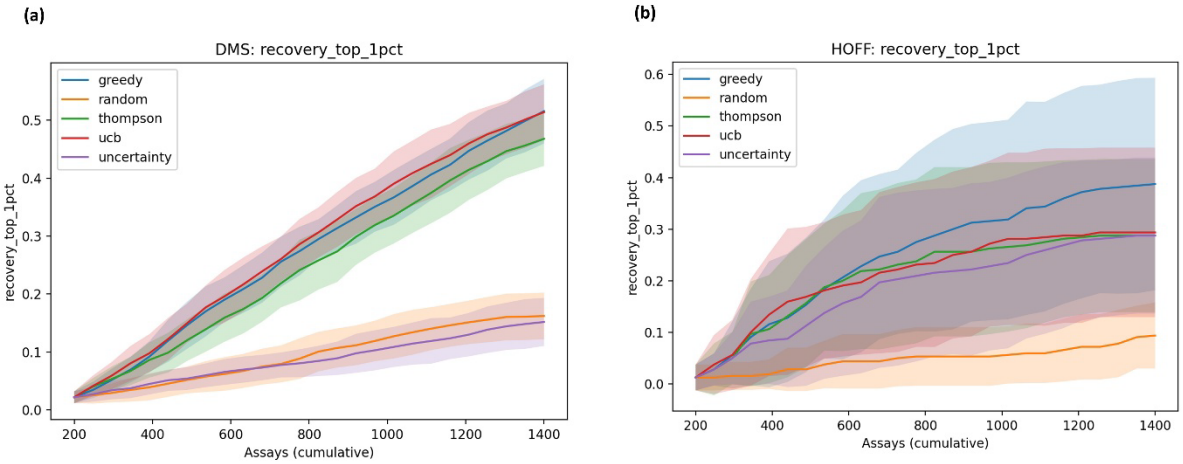


Figure S3: Recovery of the global top 1% of variants during retrospective active-learning simulations. (a) DMS landscape: fraction of the globally top 1% variants (by the specified DMS objective) recovered among assayed variants as a function of cumulative experimental budget. (b) Oxygen-tolerance landscape: analogous recovery on the multi-mutant screen objective under mutational-depth holdout. Curves summarize repeated retrospective simulations (initial labeled set and fixed batch additions per round) for each acquisition strategy; lines denote the mean across replicates and shaded bands indicate ±1 standard deviation across replicates. Higher values at a fixed budget indicate more effective prioritization of elite variants.

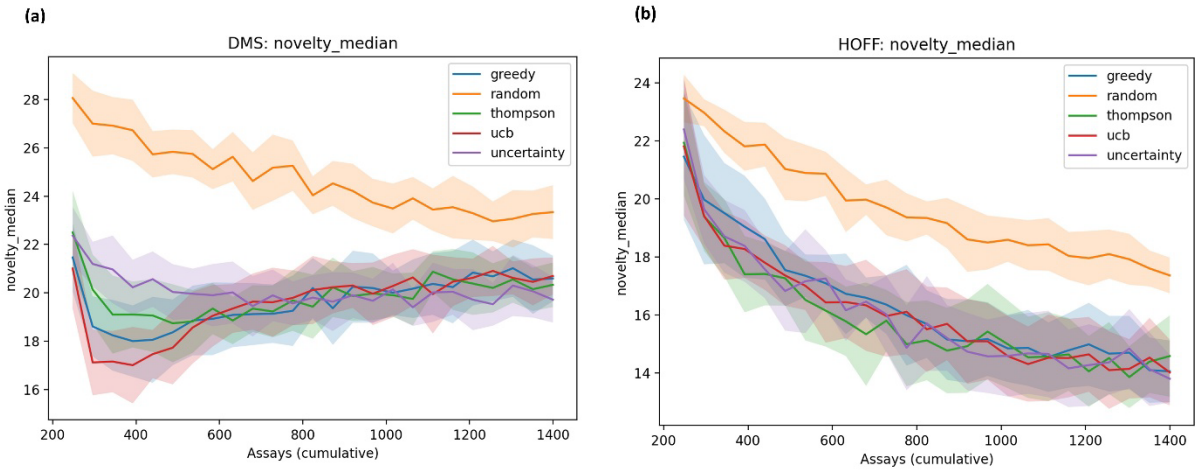


Figure S4: Exploration behavior quantified by median novelty of selected variants during active learning. (a) DMS landscape: median novelty of each selected batch over the campaign. (b) Oxygen-tolerance landscape: median novelty under mutational-depth holdout. Novelty is computed as the median nearest-neighbor distance (Euclidean) in the model representation space between newly selected variants and the set of previously assayed variants (i.e., median of the minimum distance to the labeled set), providing an operational measure of exploration versus exploitation. Lines denote the mean across replicate simulations and shaded bands indicate ±1 standard deviation; sustained higher novelty suggests broader exploration, whereas declining novelty indicates convergence toward local regions of the landscape.