Steering Generative Models with Experimental Data for Protein Fitness Optimization

Jason Yang*[1], **Wenda Chu***[1], **Daniel P. Khalil**[1], **Raul Astudillo**[1], Bruce J. Wittmann[2], Frances H. Arnold[1], and Yisong Yue[1]
[1]Caltech, [2]Microsoft

Motivation: AI-assisted protein optimization

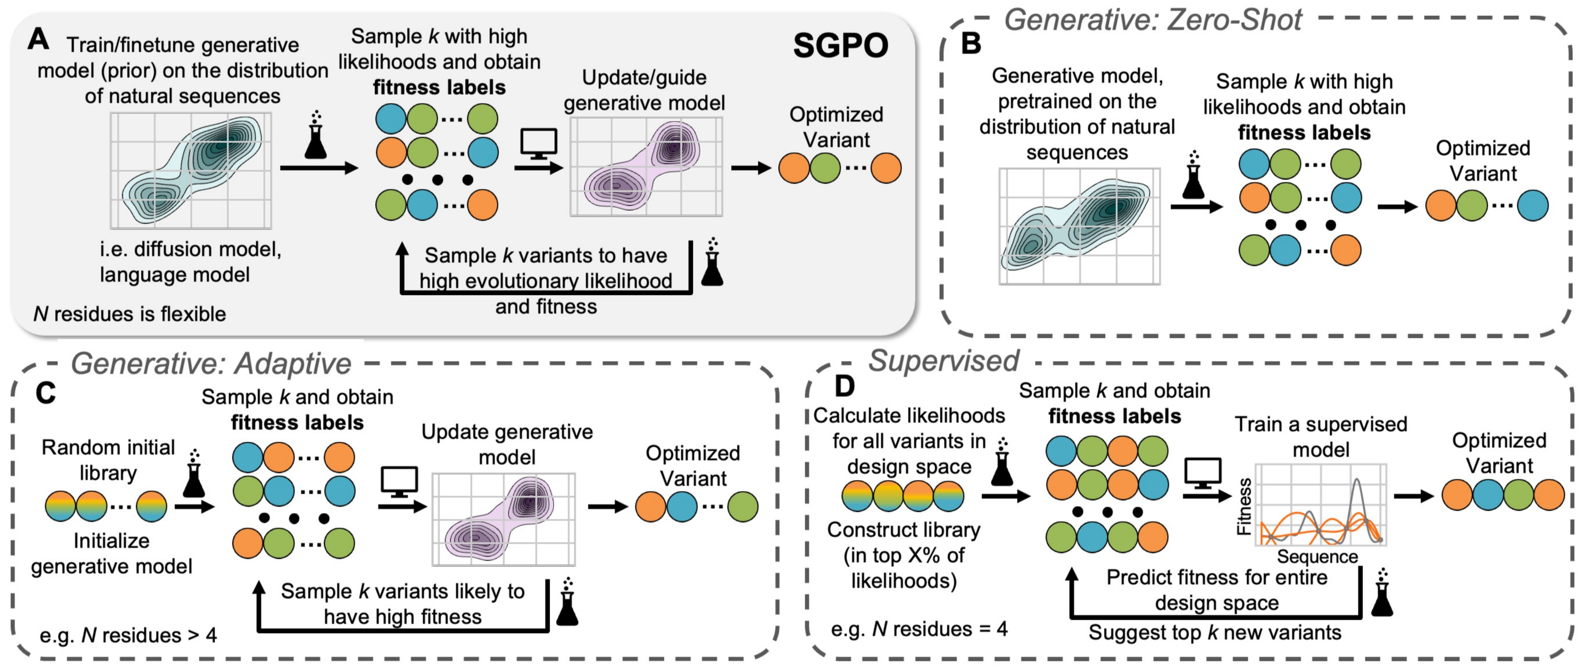

Approach	Prior Info Used?	Assay Fitness Used?	Scales to Large *N*?
SGPO	✓	✓	✓
Generative: Zero-Shot	✓	×	✓
Generative: Adaptive	×	✓	✓
Supervised	✓	✓	×

Steered **g**eneration for **p**rotein **o**ptimization (**SGPO**) approaches are advantageous, as the utilize both *unlabeled* (natural sequences) and *labeled* (assay fitness) data and scale to large design spaces In the real world, only ~10^2 - 10^3 true fitness labels are available, as they must be collected through expensive wet-lab assays (screens).

Background: generative models for protein sequences and steering strategies

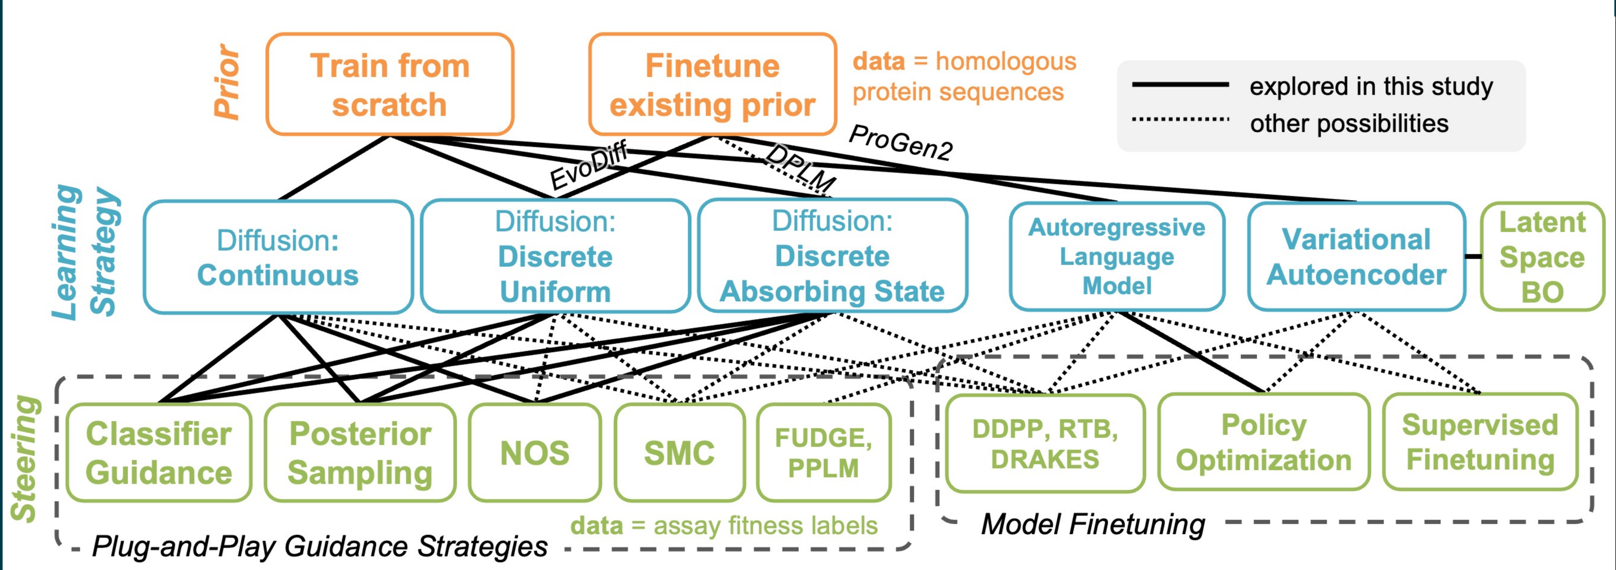

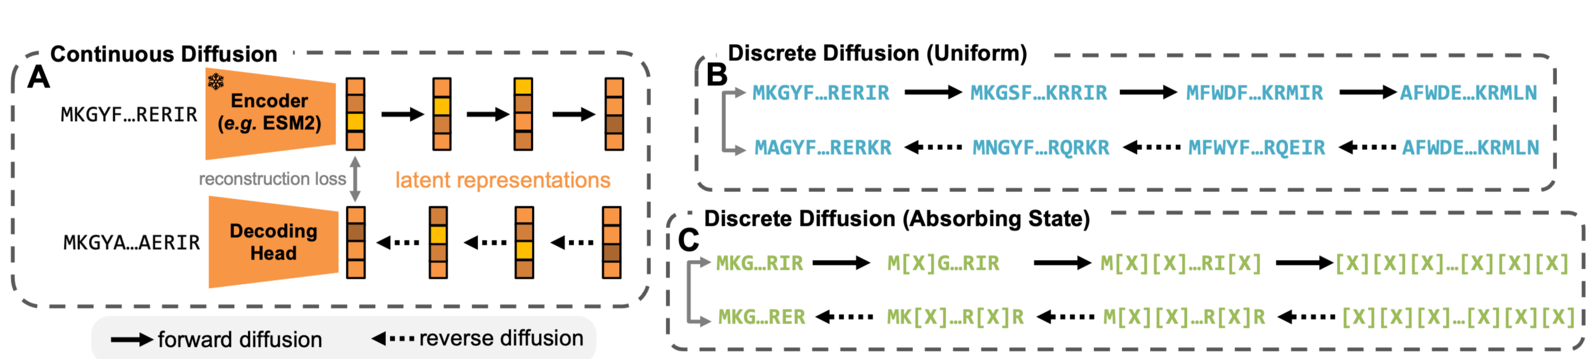

Our study focuses on different guidance strategies (namely *classifier guidance* and *posterior sampling*) for diffusion models of discrete protein sequences, with comparison to language models steered with direct preference optimization (DPO). Our models are trained or finetuned with naturally occurring sequences belonging to the protein family being optimized. We are also testing masked diffusion language models (MDLMs) and other state-of-the-art guidance strategies.

Summary of Contributions

- We motivate SGPO as a useful, general framework and contextualize existing methods for protein optimization under this umbrella.
- We comprehensively evaluate design decisions for SGPO, including different generative models for sequences and steering strategies, offering best practices for protein optimization with few fitness labels.
- We introduce ideas from adaptive optimization into SGPO by proposing a method that ensembles multiple plug-and-play fitness predictors and leverages their predictive uncertainty to enable more efficient exploration.
- We are the first to adapt *decoupled annealing posterior sampling* for SGPO, and this type of plug-and-play guidance has the strongest performance overall.

Task: maximize protein fitness

Protein Dataset	Fitness Metric	Length	Design Space	MSA Size (for prior pretraining)
TrpB	Enzyme Activity	389	*N* = 15	5.7e4
CreiLOV	Fluorescence	119	*N* = 119	3.7e5
GB1	Binding	56	*N* = 56	126

During evaluation, due to the massive size of the design space, a supervised oracle trained on available fitness data was used as an approximation for true fitness.

Generative priors capture the distribution of natural sequences

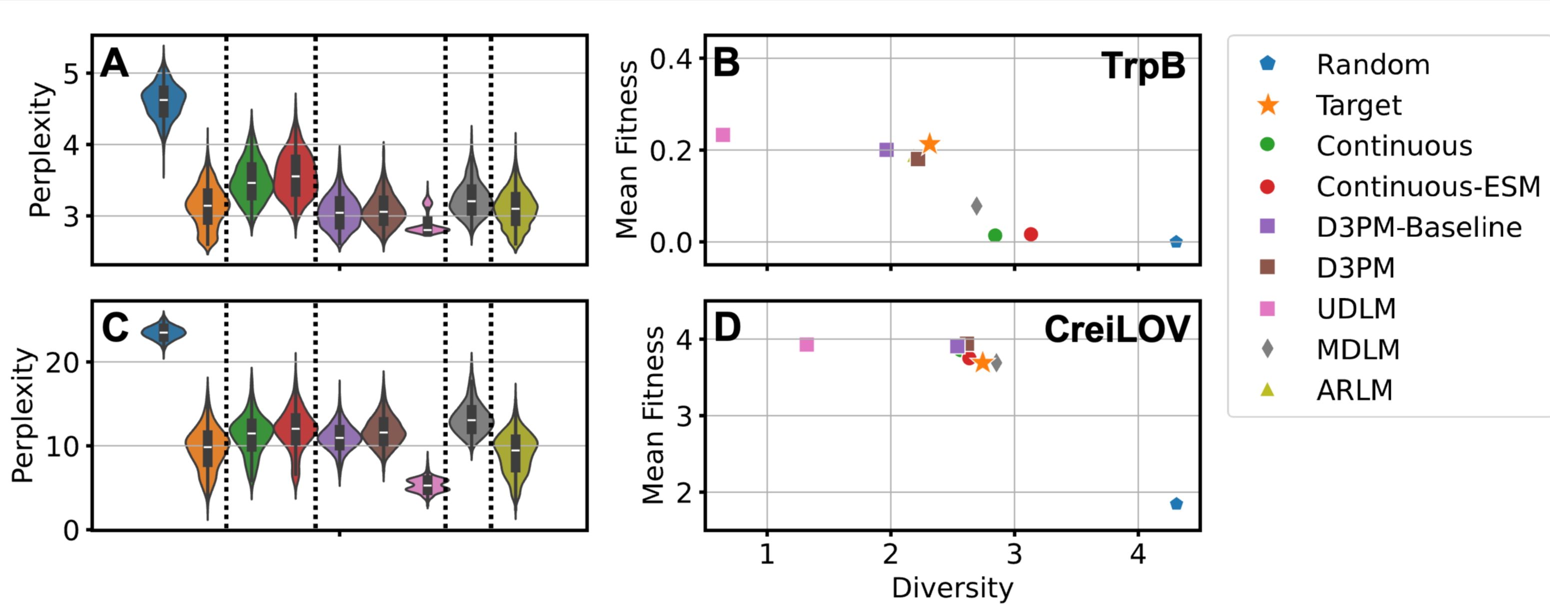

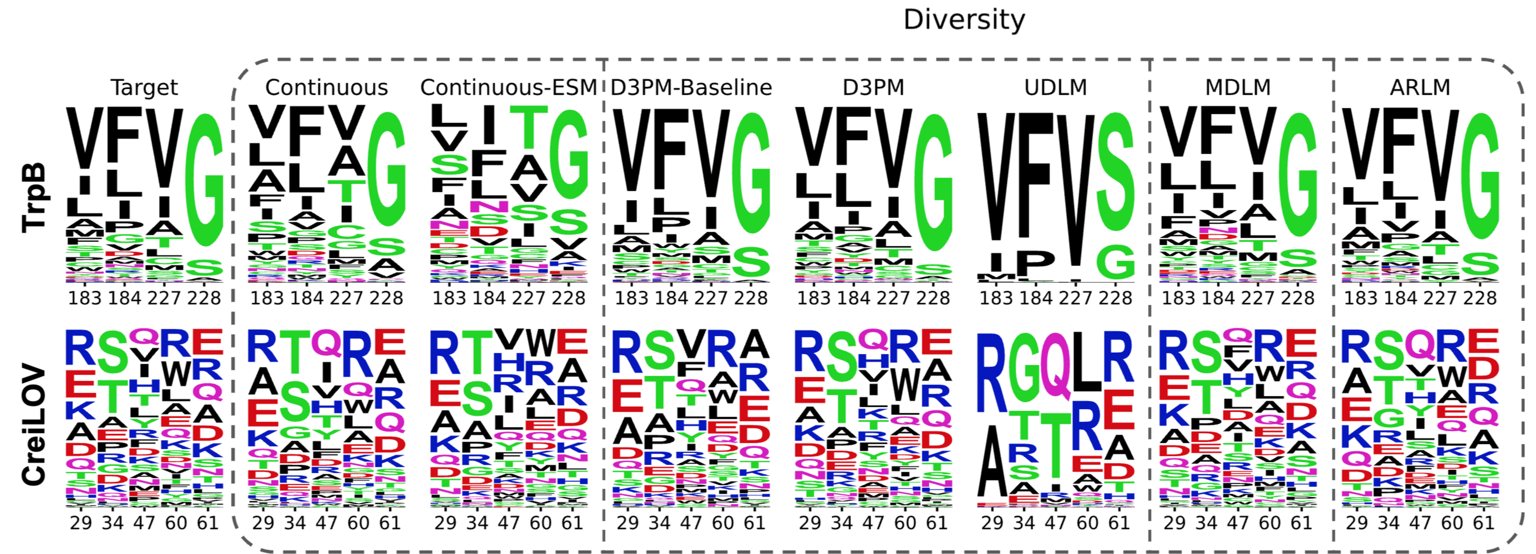

The distribution learned by pretrained generative priors, approximated by 1000 generations, largely matches that of the target (naturally occurring) distribution. The perplexities of generated sequences also match the target training distributions (MSAs), as evaluated by passing generations through ProGen2.

Evaluating SGPO design choices

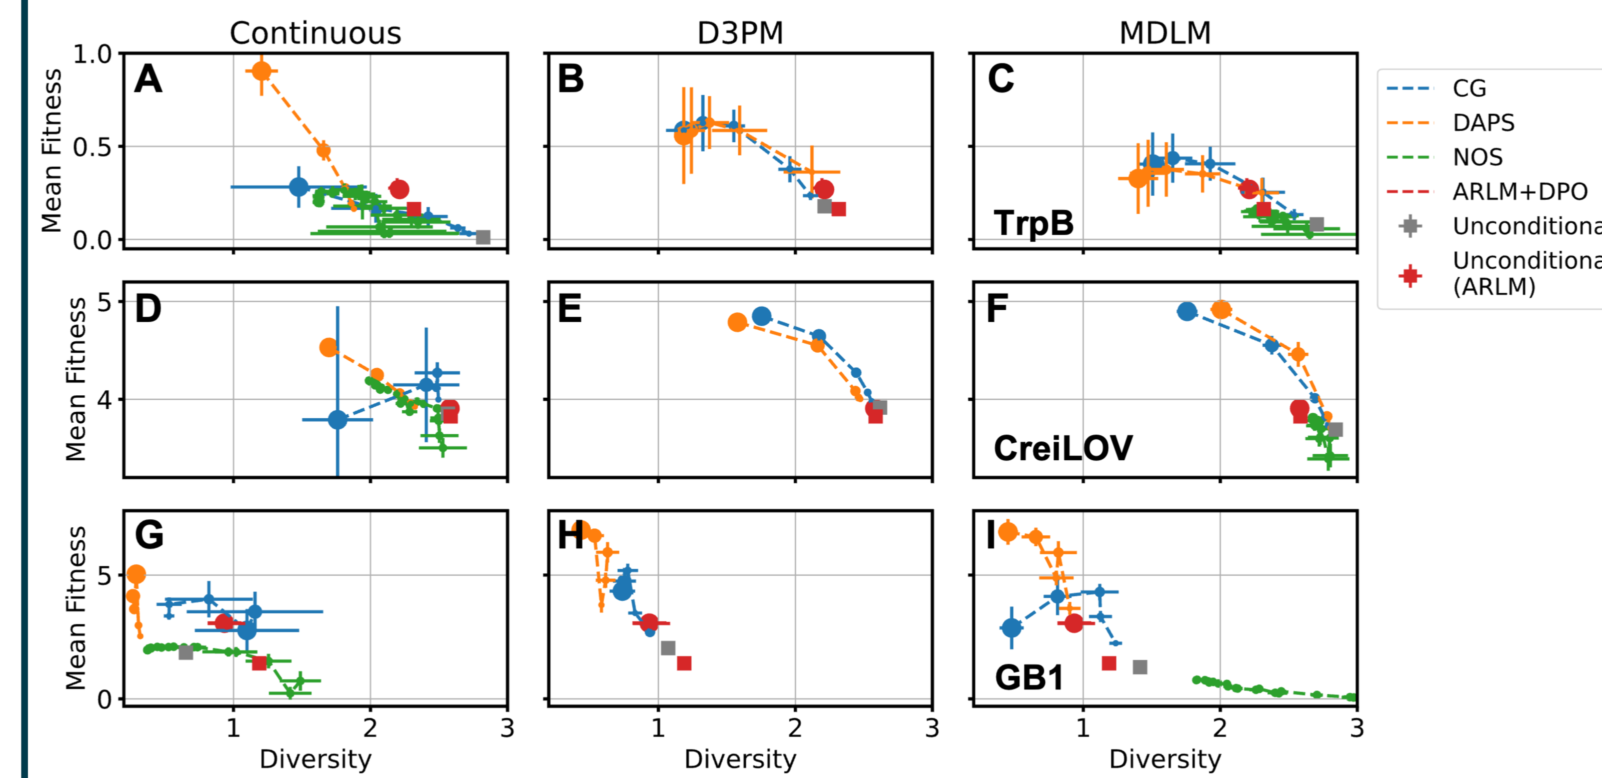

Pareto boundaries demonstrate the tradeoff between generating sequences with high fitness and high diversity for TrpB, CreiLOV, and GB1. Sequences sampled from the generative models (Continuous, D3PM, and MDLM), after guidance with labeled fitness data, are enriched in high-fitness protein variants, and most methods show higher performance than the ARLM+DPO baseline. Larger circle indicates a stronger guidance strength hyperparameter (excluding NOS). 200 unique sequences were used for steering, drawn from a prior distribution. Unconditional refers to sequences sampled from the prior with no guidance.

Adaptive optimization with SGPO

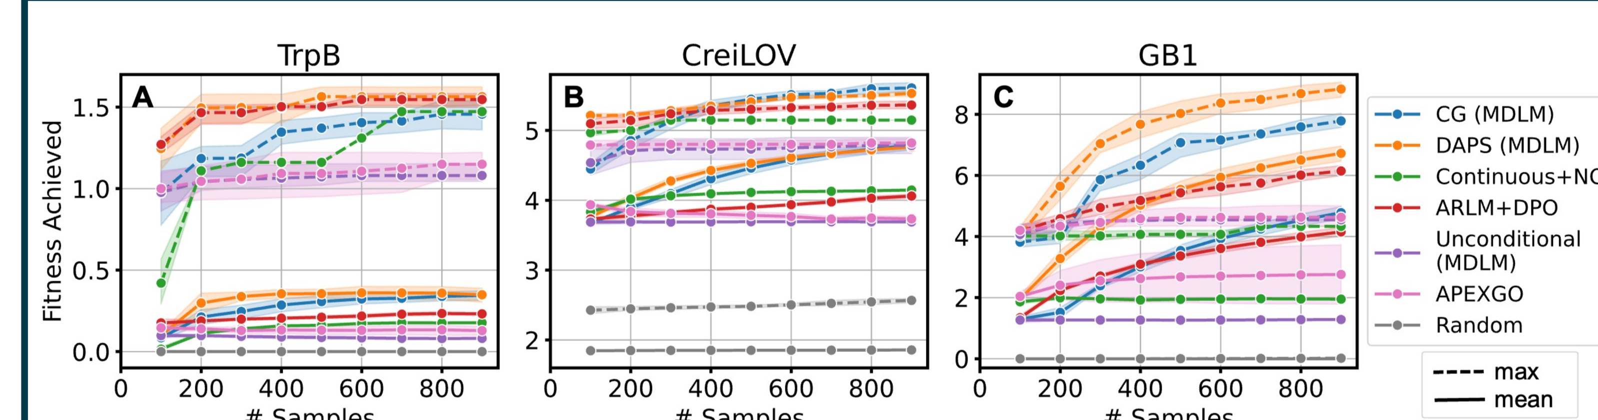

Maximum/mean fitness achieved improves over multiple iterations of steering in an adaptive setting similar to batch Bayesian optimization. 100 sequences were sampled in each round. Within each round, an ensemble of 10 value functions (classifiers) was trained on fitness data from all previously queried samples, and each new sample was generated by the MDLM model guided with a value function sampled from the ensemble (akin to Thompson sampling).

Takeaways

- Plug-and-play guidance-based strategies are generally more effective than finetuning language models; the latter can be difficult when only few fitness labels are available.
- Furthermore, SGPO approaches outperform latent space Bayesian optimization (namely APEXGo), which we attribute to over-reliance on the structure of the pretrained latent space.
- An advantage of guidance strategies is that only one hyperparameter (guidance strength) needs to be tuned, which is practical for real-world engineering scenarios.
- Interesting directions for future research include enabling sampling from other acquisition functions (formulated under a Bayesian perspective, multi-objective, etc.) and steering outside the distribution of the generative prior