Combinatorial Optimization of Antibody Libraries via Constrained Integer Programming

Extended Abstract

Conor F. Hayes*
Cognizant AI Lab
San Francisco, CA, USA
conor.hayes@cognizant.com

Andre R. Goncalves
Lawrence Livermore National Laboratory
Livermore, CA, USA
goncalves1@llnl.gov

Steven Magana-Zook
Lawrence Livermore National Laboratory
Livermore, CA, USA
maganazook1@llnl.gov

Jacob Pettit
Lawrence Livermore National Laboratory
Livermore, CA, USA
pettit8@llnl.gov

Ahmet Can Solak
Lawrence Livermore National Laboratory
Livermore, CA, USA
solak1@llnl.gov

Daniel Faissol
Lawrence Livermore National Laboratory
Livermore, CA, USA
faissol1@llnl.gov

Mikel Landajuela*†
Lawrence Livermore National Laboratory
Livermore, CA, USA
landajuelala1@llnl.gov

ABSTRACT

Designing effective antibody libraries is a challenging combinatorial search problem in computational biology. We propose ProtLib-Designer (PLD), a constrained integer linear programming (ILP) method that generates libraries while explicitly controlling library size and diversity. PLD uses in silico deep mutational scanning scores from protein language models and antibody inverse-folding models to guide optimization in a cold-start setting, without requiring target-specific experimental data. Across three antibody–antigen systems (Trastuzumab/HER2, D44.1/HEL, and Spesolimab/IL36R), PLD produces libraries that improve predicted binding-related objectives while enforcing user-specified diversity constraints. Code is available at github.com/LLNL/protlib-designer.

KEYWORDS

Combinatorial Optimization; Integer Programming; Antibody Library Design; Diversity Constraints; Protein Language Models; Inverse Folding

*Conor F. Hayes and Mikel Landajuela contributed equally to this work. Work performed while at Lawrence Livermore National Laboratory.
†Mikel Landajuela is the corresponding author.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/DUGN7746

ACM Reference Format:
Conor F. Hayes, Andre R. Goncalves, Steven Magana-Zook, Jacob Pettit, Ahmet Can Solak, Daniel Faissol, and Mikel Landajuela. 2026. Combinatorial Optimization of Antibody Libraries via Constrained Integer Programming: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/DUGN7746

1 PROBLEM SETTING

Antibody library design is a key step for bootstrapping directed evolution in early-stage drug discovery [1, 10]. Traditional approaches rely on random mutagenesis and/or experimental deep mutational scanning (DMS) [8, 12], but these are costly and often provide limited control over library diversity. Recent protein language and structure-conditioned inverse-folding models enable *in silico* mutational scanning at low cost [6, 7]. We target a *cold-start* setting (no target-specific fitness data), and aim to produce a *size-controlled* library that balances predicted quality with explicit diversity controls to hedge against model bias [3].

Goal. Given a wild-type antibody sequence $\mathbf{w}$ and a set of mutable interface positions $\mathbf{r}$, we construct a library of K mutants that (i) improves model-predicted objectives and (ii) satisfies user-specified diversity constraints that limit redundancy across the batch.

Contributions. We (i) formulate cold-start library design as constrained combinatorial optimization over additive mutation scores derived from deep learning models; (ii) introduce a diversity-constrained *solve-and-remove* procedure that guarantees exact library size; and (iii) demonstrate improved library quality on three antibody systems spanning 10–47 mutable positions.

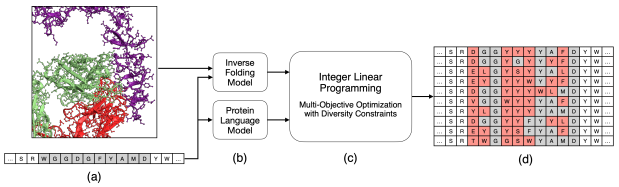

Figure 1: PLD combines in silico deep mutational scanning (sequence and structure models) with constrained ILP to generate a size-controlled, diverse antibody library.

2 APPROACH

In silico mutational scanning. For each interface position $i \in \mathbf{r}$ and amino acid a_j, we compute additive mutation scores from (i) a protein language model (intrinsic fitness) and (ii) an antibody inverse-folding model conditioned on antigen context (extrinsic fitness):

$$s_{ij}^{\text{PLM}} = -\log(p(x_i = a_j \mid \mathbf{w})) + \log(p(x_i = w_i \mid \mathbf{w})), \tag{1}$$
$$s_{ij}^{\text{IFOLD}} = -\log(p(x_i = a_j \mid \mathbf{w}, \mathbf{g})) + \log(p(x_i = w_i \mid \mathbf{w}, \mathbf{g})).$$

In our experiments, we use ProtBERT [4] and AntiFold [7], respectively. Because scores are additive across positions, they enable scalable optimization over large combinatorial mutation spaces while retaining interpretability: each selected substitution has an explicit contribution to each objective.

Multi-objective ILP. We encode a mutant with binary variables $z_{ij} \in \{0, 1\}$ indicating whether substitution (i, j) is chosen, and solve a scalarized ILP with weights $\boldsymbol{\lambda}$:

$$\min_{\mathbf{z}} \quad \sum_{q \in Q} \lambda_q \sum_{i=1}^{N} \sum_{j=1}^{M} s_{ij}^{q} z_{ij} \tag{2}$$
$$\text{s.t.} \quad \sum_{j=1}^{M} z_{ij} \le 1 \ \ \forall i, \quad n_{\min} \le \sum_{i=1}^{N} \sum_{j=1}^{M} z_{ij} \le n_{\max}, \quad z_{ij} \in \{0, 1\}.$$

We solve each instance with CBC [5].

Diversity-constrained library generation. PLD constructs a library of size K via a *solve-and-remove* strategy: at each iteration, it samples $\boldsymbol{\lambda}^{(k)}$ (Dirichlet) to vary objective trade-offs, solves (2), and then adds constraints to avoid near-duplicates and reduce overuse of positions and specific substitutions across the batch. Concretely, PLD (i) excludes a Hamming ball of radius ϵ around previously selected mutants, (ii) limits how often a position can be mutated (threshold δ_1), and (iii) limits how often a specific mutation can appear (threshold δ_2). The Hamming-ball exclusion can be written as a linear constraint for each previously selected mutant $\mathbf{z}^{(l)}$:

$$\sum_{i,j} z_{ij} - 2 \sum_{i,j} z_{ij}^{(l)} z_{ij} \ge 1 + \epsilon - \sum_{i,j} z_{ij}^{(l)}. \tag{3}$$

With $\epsilon = 0$, these constraints guarantee K *unique* mutants and explicit control of the diversity–quality trade-off; the number of constraints grows linearly with K (one ball constraint per previously chosen solution).

Table 1: Trastuzumab/HER2 results (1000 sequences).

Method	Ent.↑	BEU↓	HV↑	Oracle↑
PLD (div.)	3.22	4.62	2232	58.2%
PLD (no div.)	3.11	4.30	2232	61.8%
LMG	4.75	10.70	1938	17.1%
MODIFY	3.97	10.14	2012	18.6%
SPEA2	2.74	6.24	2232	38.0%

3 RESULTS

Setup. We evaluate three antibody systems: Trastuzumab/HER2 (10 positions) and D44.1/HEL (34 positions) [9], and Spesolimab/IL36R (47 positions) [11]. We generate $K = 1000$ mutants per system with 5–8 mutations each. Baselines include a linear mutant generator (LMG) [2], a learned diversity method (MODIFY) [3], and a multi-objective evolutionary algorithm (SPEA2) [13]. We report residue entropy and two multi-objective metrics: batch expected utility (BEU; lower is better) and hypervolume (HV; higher is better), plus an oracle fitness proxy for Trastuzumab. Unless noted otherwise, we set $\epsilon = 0$ to prevent duplicate mutants and use $(\delta_1, \delta_2) = (1000, 300)$ for Trastuzumab and $(1000, 500)$ for D44.1 and Spesolimab.

Findings. Table 1 shows that PLD substantially improves predicted oracle fitness over baselines while maintaining competitive (and controllable) diversity. In Trastuzumab, adding diversity constraints increases entropy from 3.11 to 3.22 while only modestly reducing oracle fitness (61.8% → 58.2%), illustrating a tunable trade-off. Compared to learned and heuristic baselines, PLD raises the oracle proxy from ∼17–19% (LMG/MODIFY) to 58% while matching the best observed HV (2232), indicating broad coverage of high-quality trade-offs. For D44.1, PLD achieves BEU 1.84 (vs. 10.81 for LMG) and HV 1262.6 (vs. 705.8), and for Spesolimab, PLD achieves BEU −7.28 (vs. 4.50) and HV 552.3 (vs. 118.7), indicating stronger objective trade-offs in larger search spaces. Notably, PLD can match or recover the Pareto front in objective space while still producing diverse *batch* solutions, avoiding the common failure mode where Pareto-optimal sets are too small or too redundant for practical wet-lab screening. We observe an explicit diversity–fitness trade-off: tightening the mutation-frequency cap δ_2 increases entropy but can reduce oracle fitness, allowing practitioners to tune diversity to match screening capacity and risk tolerance.

4 CONCLUSION

PLD casts cold-start antibody library design as diversity-constrained combinatorial optimization. By combining in silico mutational landscapes with ILP constraints, it produces interpretable, size-controlled libraries that can be tuned to experimental requirements. Limitations include dependence on structural context and the fidelity of learned scores; future work will extend PLD to breadth optimization against sets of targets and richer (e.g., quadratic) interaction models.

ACKNOWLEDGMENTS

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344 (LLNL-CONF-868784).

REFERENCES

[1] P. Carter and G. A. Lazar. 2017. Next generation antibody drugs: pursuit of the 'high-hanging fruit'. *Nature Reviews Drug Discovery* 17 (2017), 197–223.

[2] Thomas A Desautels, Kathryn T Arrildt, Adam T Zemla, Edmond Y Lau, Fangqiang Zhu, Dante Ricci, Stephanie Cronin, Seth J Zost, Elad Binshtein, Suzanne M Scheaffer, et al. 2024. Computationally restoring the potency of a clinical antibody against Omicron. *Nature* (2024), 1–8.

[3] Kerr Ding, Michael Chin, Yunlong Zhao, Wei Huang, Binh Khanh Mai, Huanan Wang, Peng Liu, Yang Yang, and Yunan Luo. 2024. Machine learning-guided co-optimization of fitness and diversity facilitates combinatorial library design in enzyme engineering. *Nature Communications* 15 (2024).

[4] Ahmed Elnaggar, Michael Heinzinger, Christian Dallago, Ghalia Rehawi, Yu Wang, Llion Jones, Tom Gibbs, Tamas Feher, Christoph Angerer, Martin Steinegger, et al. 2021. Prottrans: Toward understanding the language of life through self-supervised learning. *IEEE transactions on pattern analysis and machine intelligence* 44, 10 (2021), 7112–7127.

[5] John Forrest and Robin Lougee-Heimer. 2005. CBC user guide. In *Emerging theory, methods, and applications*. INFORMS, 257–277.

[6] Brian L Hie, Varun R Shanker, Duo Xu, Theodora UJ Bruun, Payton A Weidenbacher, Shaogeng Tang, Wesley Wu, John E Pak, and Peter S Kim. 2023. Efficient evolution of human antibodies from general protein language models. *Nature Biotechnology* (2023).

[7] M. H. Høie, Alissa M. Hummer, T. H. Olsen, Broncio Aguilar-Sanjuan, Morten Nielsen, and Charlotte M. Deane. 2024. AntiFold: Improved antibody structure-based design using inverse folding.

[8] P. Koenig, Chingwei V. Lee, Benjamin T. Walters, V. Janakiraman, Jeremy Stinson, T. Patapoff, and G. Fuh. 2017. Mutational landscape of antibody variable domains reveals a switch modulating the interdomain conformational dynamics and antigen binding. *Proceedings of the National Academy of Sciences* 114 (2017), E486 – E495.

[9] Derek M. Mason, Simon Friedensohn, Cédric R. Weber, Christian Jordi, Bastian Wagner, Simon M Meng, R. Ehling, Lucia Bonati, J. Dahinden, P. Gainza, B. Correia, and S. Reddy. 2021. Optimization of therapeutic antibodies by predicting antigen specificity from antibody sequence via deep learning. *Nature Biomedical Engineering* 5 (2021), 600 – 612.

[10] Mary S Morrison, Christopher J Podracky, and David R Liu. 2020. The developing toolkit of continuous directed evolution. *Nature chemical biology* 16, 6 (2020), 610–619.

[11] Amir Shanehsazzadeh, Julian Alverio, George Kasun, Simon Levine, Ido Calman, Jibran A. Khan, Chelsea Chung, Nicolas Diaz, Breanna K. Luton, Ysis Tarter, Cailen M. McCloskey, Katherine B. Bateman, Hayley Carter, Dalton Chapman, Rebecca Consbruck, Alec Jaeger, Christa Kohnert, Gaelin Kopec-Belliveau, John M. Sutton, Zheyuan Guo, Gustavo Canales, Kai Ejan, Emily Marsh, Alyssa Ruelos, Rylee Ripley, Brooke Stoddard, Rodante Caguiat, Kyra Chapman, Matthew Saunders, Jared Sharp, Douglas Ganini da Silva, Audree Feltner, Jake Ripley, Megan E. Bryant, Danni Castillo, Joshua Meier, Christian M. Stegmann, Katherine M. Moran, Christine Lemke, S. Abdulhaqq, Lillian R. Klug, and S. Bachas. 2024. IgDesign: In vitro validated antibody design against multiple therapeutic antigens using inverse folding. *bioRxiv* (2024).

[12] Timothy A. Whitehead, A. Chevalier, Yifan Song, C. Dreyfus, S. Fleishman, C. D. Mattos, C. A. Myers, Hetunandan Kamisetty, P. Blair, I. Wilson, and David Baker. 2012. Optimization of affinity, specificity and function of designed influenza inhibitors using deep sequencing. *Nature Biotechnology* 30 (2012), 543–548.

[13] Eckart Zitzler, Marco Laumanns, and Lothar Thiele. 2001-05. *SPEA2: Improving the strength pareto evolutionary algorithm*. Report. ETH Zurich, Computer Engineering and Networks Laboratory, Zurich. https://doi.org/10.3929/ethz-a-004284029