GUIDED GENERATION OF B-CELL RECEPTORS WITH CONDITIONAL WALK-JUMP SAMPLING

Taylor Joren¹, Sarah A. Robinson¹, Homa MohammadiPeyhani¹, Sai Pooja Mahajan¹, Edith Lee¹, Stephen Lillington², Qixin Bei³, Saeed Saremi¹, Jae Hyeon Lee¹, Richard Bonneau¹, Isidro Hötzel², Vladimir Gligorijević¹, Simon Kelow¹, Nathan C. Frey¹

¹Prescient Design, Genentech ²Antibody Engineering, Genentech ³Proteomic & Genomic Technologies, Genentech joren.taylor@gene.com

Abstract

Antibody drug discovery campaigns often leverage immune repertoires from antigen-exposed animals, which can be divided into clonotypes, subclasses of sequences derived from the same progenitor B-cell. In this work, we adapt discrete walk-jump sampling (dWJS) to condition generation on categorical variables like clonotype, extending both energy-based and score-based dWJS with predictor-free guidance during Langevin dynamics ("walking") and denoising ("jumping"). Categorical and numerical variables are learned during training and specified during sampling, producing diverse and novel sequences from target clonotype classes. We train conditional WJS models on datasets of over 1.5M sequences obtained from antigen-exposed rats and human patients post-vaccination. Surprisingly, increasing guidance improves both sample quality and sequence diversity, enabling controllable sampling from thousands of distinct modes.

1 INTRODUCTION

The success of protein engineering depends on the quality of starting sequences, or leads. For antibodies, lead discovery traditionally relies on B-cell receptor (BCR) repertoires isolated from immunized animals using hybridoma or B cell cloning technologies (Zheng et al., 2024) (Pedrioli & Oxenius, 2021). These repertoires contain diverse antibody sequences that may bind to different epitopes of the target antigen, organized into canonical subclasses called clonotypes that likely share binding mechanisms (Mhanna et al., 2024). The identification and expansion of promising clonotypes is crucial, as subsequent optimization efforts are constrained by the quality of these initial leads.

Machine learning approaches have transformed antibody discovery by replacing traditional mutational scanning and selection with generative sequence models coupled with discriminative models that predict protein fitness (Hie & Yang, 2022) (Wu et al., 2021). However, controlling sequence generation to stay within relevant clonotype neighborhoods remains underexplored. Training on individual clonotypes (Erlach et al., 2025) restricts sample diversity, while training on multiple clonotypes without conditioning risks generating chimeric sequences that combine segments from non-complementary clonotypes. Discrete walk-jump sampling (dWJS) (Frey et al., 2024) has shown promise for generating viable antibody sequences, but its unconditional sampling approach can make it difficult to target generation toward specific clonotypes or other subclasses of interest.

We present an improved dWJS that enables controlled generation by conditioning on categorical and numerical variables. Unlike approaches that require separate classifier or regressor models to guide sampling (Ikram et al., 2024), our method incorporates conditioning variables directly into the walk-jump sampling process. We adapt both energy-based and score-based dWJS with predictor-free guidance towards a protein family, in this case antibody clonotypes. Building on predictor-free guidance approaches in other domains (Ho & Salimans, 2022) (Liu et al., 2024), our method allows for generating samples with specific properties without relying on gradients from external property prediction models during sampling. Here, we improve the fast, efficient, and high-quality sampling of

dWJS with predictor-free guidance that enables controllable mode exploration. Each model is trained on a large corpus of BCR repertoire data spanning many clonotypes, with conditioning variables learned during training and specified explicitly during sampling. As demonstrated in other predictorfree diffusion methods, this approach allows control of sample diversity and quality, however we show evidence that one does not come at the expense of the other. To provide mechanistic insights into the sampling process, the sampling trajectories are inspected by denoising intermediate states during Langevin dynamics, or the "walking" step of dWJS. We show that these trajectories converge to the target context with high precision over the course of sampling.

In summary, the main contributions of this work include:

- 1. A novel formulation of dWJS with predictor-free guidance for steered generation towards desired sequence subclasses, simultaneously preserving the strengths of dWJS while improving sample quality and diversity.
- 2. Antibody sequence generation from specific clonotype modes using both categorical and numerical conditioning variables.
- 3. Large-scale pretraining of generative protein sequence models trained on 1.5M+ samples, enabling sampling from \sim 3,000 unique clonotype modes.

2 ADAPTING WALK-JUMP SAMPLING WITH DIRECT CONDITIONING

2.1 DIRECT CONTEXT GUIDANCE WITHOUT PREDICTOR GRADIENTS

Given a dataset of BCR (antibody) sequences x of many pre-determined clonotype subclasses \mathbf{c}_k , we aim to generate a diverse set of sequences that belong to a target clonotype $\mathbf{c}_t = (c_v, c_j, c_l)$, where c_v, c_j and c_l respectively correspond to categorical variables V gene, J gene and numerical variable CDR3 length. These three canonical variables together define a clonotype subclass (Mhanna et al., 2024), a group of sequences that stem from the same starting BCR.

We introduce **conditional discrete walk-jump samping** (cWJS) as a predictor-free method for conditional sequence generation. As in discrete WJS, cWJS uses an optional energy-based model (EBM) and required score-based or denoising model.

Our conditioning approach follows the predictor-free guidance paradigm (Ho & Salimans, 2022) (Liu et al., 2024) which differs from predictor-based guidance methods that rely on the gradients of external prediction models to steer the sampling process. In predictor-based approaches (Dhariwal & Nichol, 2021) Ikram et al. (2024), an unconditional generative model produces initial samples, and the gradients from separate classifier or regressor models are then used to adjust these samples toward desired attributes. In contrast, our method directly incorporates the conditioning information into the walk-jump sampling processes, eliminating the need for external predictor gradients during sampling. This simplifies the generation pipeline and enables more efficient conditioning on multiple clonotype properties simultaneously.

CONDITIONAL ENERGY-BASED MODEL

The energy-based model (EBM) is implemented in a similar way to unconditional dWJS but adapted to additionally take in a context $\mathbf{c} = (c_1, c_2...c_n)$. Each context c_i is embedded separately, followed by concatenation and projection to size $d_{context}$. Categorical variables including V & J gene are one-hot encoded prior to an embedding layer. Numerical variables, such as CDR3 length and number of somatic hypermutations (SHMs), are encoded using a clustering encoding method following (Liu et al., 2024). The full context embedding is split into gain and bias terms to modify the embedding of the noisy, sequence embedding, following the implementation of (Du et al., 2020). Figure 4 shows the architecture with context included.

The EBM is trained using contrastive divergence with data augmentation. The contrastive divergence loss encourages the model to assign lower energy to positive samples and higher energy to negative samples. A positive sample (y_{pos}, c_{pos}) includes a noised sequence $y_{pos} = x + \mathcal{N}(0, \sigma^2 I_d)$ from the training data and its corresponding context **c** without noise added. More details on augmenting the set of negatives samples and training with contrastive divergence may be found in Section A.3.



Figure 1: Sequence design labels closely match the target clonotypes (TCs) used in conditioning, showing controllable and diverse mode sampling. Left: DeepNGS Navigator embeddings of training set and designed sequences. Middle: embedding of designs, colored by TCs. Right: the labeled clonotypes of designs

CONDITIONAL DENOISER AND SCORE-BASED MODEL

Rather than training an EBM, a score-based model $g_{\phi}(y, c)$ may instead be trained to directly approximate the score function. $g_{\phi}(y, c)$ is trained with a denoising objective as in dWJS, but the context is embedded and added to intermediate representations of the sequence to focus denoising on the intended clonotype. We modify ByteNet for conditioning, as shown in Figure 4.

During training, the context is randomly replaced with a null \emptyset token with probability $p_{uncond} = 0.2$, following (Ho & Salimans, 2022). This enables the model to be sampled unconditionally when \emptyset is provided. Having easy access to an unconditional model $g_{\phi}(y)$ also allows for tuning the level of guidance during sampling. That is, sampling uses a linear combination of the conditional and unconditional score estimates with w as the level of guidance:

$$s_{\phi}(y,c,w) = (1+w)g_{\phi}(y,c) - wg_{\phi}(y,\varnothing) \tag{1}$$

2.2 SAMPLING WITH CONTEXT GUIDANCE

To generate sequences during inference, we follow (Frey et al., 2024) and use Langevin dynamics, though we explicitly include the target context. When a conditional EBM is used in energy-based cWJS, the update rule for each step k is given by:

$$y_{(k+1)} = y_{(k)} - \delta \nabla_y f_\theta(y_{(k)}, \mathbf{c}_t) + \sqrt{2\delta} \epsilon_k$$

where δ is the step size, ∇_y is the gradient of the energy function f_{θ} with respect to noisy data y, and ϵ is additional noise. In scored-based dWJS, $\nabla_y f_{\theta}$ may be simply replaced with $s_{\phi}(y, c, w)$, the learned approximation of the score function. After T steps, clean discrete sequences are recovered by "jumping" using the least-squares estimator from Neural Empirical Bayes (Saremi & Hyvärinen, 2019):

$$\hat{x}_{\phi}(y_T, \mathbf{c}_t, w) = y_T + \sigma^2 s_{\phi}(y_T, \mathbf{c}_t, w)$$

Note that regardless of the method used, energy-based or scored-based cWJS, s_{ϕ} is always used to denoise. Context **c**_t is specified explicitly during sampling to simplify inference.

3 EXPERIMENTS

We evaluate cWJS using two BCR repertoire datasets: (1) 1.5 million heavy-chain sequences from rats immunized with a transmembrane (TM) antigen, annotated with V genes, J genes, and CDR3 lengths. The TM dataset comprises 122 single-context modes (V genes) and 2,994 multi-context modes (V/J genes & CDR3 length). (2) consists of 167,537 paired-chain sequences from patients receiving SARS-CoV-2 mRNA vaccines targeting the Spike (S) protein, annotated with heavy-chain and light-chain V genes. The S dataset contains 2,797 multi-context modes (heavy- & light- V genes).



Figure 2: Increasing guidance both improves fidelity and diversity metrics. Multi-context conditioning is done on heavy-chain (HC) and light chain (LC) V genes from human anti-SARS-CoV-2 repertoires in (Wang et al., 2024)

For each dataset, we train cWJS with relevant context included. For sampling, we select context conditions with a range of relative abundances in the training repertoire. Given a target context, 100 samples are generated from each starting seed where each seed is a representative sequence from each V gene cluster in the training set. This ensures no starting seeds are out-of-distribution to the trained model. Both the training sequences and final samples are labeled by Absolve (Genentech, 2018), a bioinformatics tool for antibody sequence annotation, providing consistent labels for context conditioning during training and sample labeling.

We evaluate cWJS across three metrics. Fidelity is measured through both accuracy of individual and joint context conditions and position-wise KL-divergence, calculated as: AA-Position-KL = $\sqrt{\frac{1}{n}\sum_{i=1}^{n} D_{KL}(S_i||T_i)^2}$, where S_i and T_i are the amino acid distributions of the samples and conditioned-matched training subset at position *i*. Diversity is quantified as the ratio of unique to total samples, while novelty is measured by edit distance to the nearest training sequence.

3.1 CONDITIONING ON V GENE

We first evaluate cWJS on single variable conditioning using V gene as the target context. Samples are generated by fixing a target V gene and undergoing 10 steps of Langevin MCMC prior to denoising. Nine V genes from the TM dataset were chosen as conditioning tokens, representing various levels of abundance and proximity to other V genes.

The generated sequences were embedded and visualized using Deep NGS Navigator (MohammadiPeyhani et al., 2025), as shown in Figure 1. The assigned labels of samples closely match the target V gene context and always match the target gene family, a set of related V genes. Mismatches are likely due to the presence of training sequences from neighboring V genes within the same family and the blurry delineation between V gene definitions, explained further in Section A.2 (Bentley & Rabbitts, 1983). Interestingly, mismatches still occur when conditioning on V genes highly represented in the training set (e.g. IGHV2-12), suggesting that subclass *definitions* may matter more for conditioning than subclass *representation* in the training repertoire.

3.2 CONDITIONING ON MULTIPLE CONTEXTS

We examine cWJS's ability to consider multiple categorical and numerical contexts simultaneously. For the paired-chain SARS-Cov2 dataset, we condition on heavy-chain V genes and consider ten different context combinations found in the training repertoire. With the TM dataset, the context is VH, JH, and CDR3 length, the last of which is numerical ranging from 5 to 24. We choose fifteen (VH, JH, CDR3) clonotypes, three from the top 1% most represented clonotypes, the top 25%, the top 50%, the bottom 25% and three that are not found in the training repertoire to assess the ability to generalize to unseen context pairs.

Figures 2 and 6 show results for the S and TM datasets, respectively. As the guidance level w increases, the accuracy of samples being in the target context generally increases and AA-Pos-KL decreases, implying guidance improves the fidelity of generated sequences and their distributions compared to a conditioned-matched training subset. Interestingly, the uniqueness and edit distance to the nearest training neighbor also *increase* with guidance, which indicates that cWJS avoids mode

collapse and recapitulating training sequences at high guidance levels. We posit this occurs because increasing w decreases the contributed score of an unconditional model during sampling, as shown in Equation 1. The unconditional model was tasked with denoising samples from *all* clonotypes without context guidance, a challenging task that may have required it to rely on a limited set of representative samples for each clonotype. This lowers the diversity of sampled sequences from the unconditional model. Thus, decreasing its contribution would explain why both accuracy and diversity increase with w.

cWJS is evaluated against two baselines: unconditional dWJS trained only on sequences from a target context and one trained on sequences from all context modes. Table 1 shows cWJS produces higher quality and more diverse samples, even when all training sequences are from the target subclass. The limited sample uniqueness of dWJS trained on all clonotypes further supports why increasing w improves sample diversity, as alluded to above.

3.3 LEVERAGING AFFINITY DATA FOR GUIDANCE

To demonstrate practical utility, we use the human S dataset containing 15,538 sequences with binding labels for the S protein. We identify the five VH/VL pairs most enriched and five least enriched among S binders, using these for conditioning. For each context, 7,200 sequences are sampled and compared by edit distance to known S binders from CoVAbDab (Raybould et al., 2021).

Table 2 shows the number of samples within 5 mutations of a known S binders in CoVAbDab. When conditioned on context enriched in S binders, more samples are within 5 edits of CoVAbDab than conditioned on non-enriched context and thus more likely to be S-binders. This suggests only knowing the VH and VL genes of binders and using this alone as context for conditioning can dramatically increase the efficiency of discovering novel binders, showing the practicality of using cWJS to accelerate protein discovery.

3.4 EVALUATING TRAJECTORIES TAKEN DURING SAMPLING

We analyze sampling trajectories by examining intermediate states during Langevin MCMC. For each state, a sample is denoised to clean sequence space and annotated with its V gene. For a given target V gene, we sample 1000 sequences per starting seed and track their V gene annotations across ten Langevin steps (Figure 3), also showing V gene frequencies at the final state.

Most samples converge to the target V gene within 1-3 steps, with some exploration of nearby genes within the same gene family (e.g. IGHV2-X when conditioned on IGHV2-63), but never outside a family. Convergence is slower when related V genes exist in the training repertoire (Figure 8). In some cases, sequences converge to a closely related but incorrect V gene (Figure 8), likely due to imperfect separation between V gene definitions. Across the nine contexts assessed using the TM dataset, 79.7% (± 33.9) of samples converged to the target context, six of which showed 99% convergence.

4 CONCLUSIONS

We present a conditional adaptation of discrete walk-jump sampling that enables controlled generation of antibody sequences from specific categorical and numerical variables, including clonotype. We demonstrate that increasing the level of guidance improves sample fidelity to target contexts, novelty,

Method	Accuracy ↑	Uniqueness ↑	$KL(tr_match, gen) \downarrow$
cWJS w=4.0 (ours)	0.99 ± 0.02	$\textbf{1.0} \pm \textbf{0.01}$	$\textbf{0.49} \pm \textbf{0.17}$
cWJS w=0.1 (ours)	0.89 ± 0.07	0.97 ± 0.06	0.80 ± 0.04
Unc dWJS (single clonotype)	0.40 ± 0.55	$\textbf{1.0} \pm \textbf{0.00}$	2.27 ± 0.63
Unc dWJS (all clonotypes)	0.08 ± 0.25	0.77 ± 0.32	4.48 ± 6.67

Table 1: cWJS outperforms unconditional dWJS trained in a variety of scenarios on accuracy, uniqueness, and KL-div. Values shown are mean and std. across 10 contexts from the SARS-CoV-2 dataset.

Min edit distance to CovAbDab binders	Binding-enriched context (HC only) ↑	Binding-depleted context (HC only) \downarrow	Binding-enriched context (HC+LC) ↑	Binding-depleted context (HC+LC) \downarrow
<= 5	2444	238	147	0
<= 3	255	2	6	0
<= 1	2	0	0	0

Table 2: When using binding-enriched context, cWJS generates more sequences (out of 36,000) near known SARS-CoV-2 binders in CoVAbDab compared to using binding-depleted context. (HC only), e.d. computed with heavy chain only; (HC+LC), e.d. computed with both chains



Figure 3: Sampled sequences (N=1000) converge on the target V-gene IGHV5-10 over 10 steps of Langevin MCMC "walking" with w = 1. The target V-gene is outlined and the fraction of samples converging to each class is shown (right).

and diversity when trained on BCR repertoire data. Our analysis of sampling trajectories shows quick convergence to target contexts and local exploration within the target and nearby v-genes. When applied to SARS-CoV-2 repertoire data, conditioning only on enriched heavy and light v-gene pairs produced sequences proximal to known binders while remaining novel, highlighting the method's potential for accelerating lead discovery.

This work opens several promising directions for future research. While this work focuses on clonotype-level conditioning, this framework can be extended to other categorical and continuous properties relevant for antibody discovery, such as developability metrics or binding affinity. Overall, our cWJS approach provides a practical tool for targeted antibody sequence generation.

5 ACKNOWLEDGEMENTS

We'd like to thank Donald Lee and members of the Prescient Design team for helpful discussions that motivated the need for this work.

REFERENCES

- David L. Bentley and Terence H. Rabbitts. Evolution of immunoglobulin V genes: Evidence indicating that recently duplicated human V sequences have diverged by gene conversion. *Cell*, 32(1):181–189, January 1983. ISSN 0092-8674. doi: 10.1016/0092-8674(83)90508-1. URL https://www.sciencedirect. com/science/article/pii/0092867483905081.
- Prafulla Dhariwal and Alexander Nichol. Diffusion Models Beat GANs on Image Synthesis. In Advances in Neural Information Processing Systems, volume 34, pp. 8780–8794. Curran Associates, Inc., 2021. URL https://papers.nips.cc/paper/2021/hash/ 49ad23d1ec9fa4bd8d77d02681df5cfa-Abstract.html.
- Yilun Du, Shuang Li, and Igor Mordatch. Compositional Visual Generation with Energy Based Models. In Advances in Neural Information Processing Systems, volume 33, pp. 6637–6647. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/hash/ 49856ed476ad01fcff881d57e161d73f-Abstract.html.

- Lena Erlach, Simon Friedensohn, Daniel Neumeier, Derek M. Mason, and Sai T. Reddy. Antibody affinity engineering using antibody repertoire data and machine learning, January 2025. URL https://www.biorxiv.org/content/10.1101/2025.01.10.632313v1. Pages: 2025.01.10.632313 Section: New Results.
- Nathan C. Frey, Daniel Berenberg, Karina Zadorozhny, Joseph Kleinhenz, Julien Lafrance-Vanasse, Isidro Hotzel, Yan Wu, Stephen Ra, Richard Bonneau, Kyunghyun Cho, Andreas Loukas, Vladimir Gligorijevic, and Saeed Saremi. Protein Discovery with Discrete Walk-Jump Sampling, March 2024. URL http: //arxiv.org/abs/2306.12360. arXiv:2306.12360 [q-bio].
- Genentech. Genentech/Absolve, August 2018. URL https://github.com/Genentech/Absolve. original-date: 2018-09-12T22:18:02Z.
- Brian L. Hie and Kevin K. Yang. Adaptive machine learning for protein engineering. *Current Opinion in Structural Biology*, 72:145–152, February 2022. ISSN 1879-033X. doi: 10.1016/j.sbi.2021.11.002.
- Jonathan Ho and Tim Salimans. Classifier-Free Diffusion Guidance, July 2022. URL http://arxiv.org/ abs/2207.12598. arXiv:2207.12598 [cs].
- Zarif Ikram, Dianbo Liu, and M. Saifur Rahman. Antibody sequence optimization with gradient-guided discrete walk-jump sampling. April 2024. URL https://openreview.net/forum?id=3aqm3bi0jh.
- Gang Liu, Jiaxin Xu, Tengfei Luo, and Meng Jiang. Graph Diffusion Transformers for Multi-Conditional Molecular Generation, October 2024. URL http://arxiv.org/abs/2401.13858. arXiv:2401.13858 [cs].
- Vanessa Mhanna, Habib Bashour, Khang Lê Quý, Pierre Barennes, Puneet Rawat, Victor Greiff, and Encarnita Mariotti-Ferrandiz. Adaptive immune receptor repertoire analysis. *Nature Reviews Methods Primers*, 4(1): 1–25, January 2024. ISSN 2662-8449. doi: 10.1038/s43586-023-00284-1. URL https://www.nature. com/articles/s43586-023-00284-1. Publisher: Nature Publishing Group.
- Homa MohammadiPeyhani, Edith Lee, Richard Bonneau, Vladimir Gligorijevic, and Jae Hyeon Lee. deepNGS Navigator: Exploring antibody NGS datasets using deep contrastive learning, January 2025. URL https: //www.biorxiv.org/content/10.1101/2025.01.27.634805v1. Pages: 2025.01.27.634805 Section: New Results.
- Alessandro Pedrioli and Annette Oxenius. Single B cell technologies for monoclonal antibody discovery. *Trends in Immunology*, 42(12):1143–1158, December 2021. ISSN 1471-4906. doi: 10.1016/j.it.2021.10.008. URL https://www.sciencedirect.com/science/article/pii/S1471490621002131.
- Matthew I. J. Raybould, Aleksandr Kovaltsuk, Claire Marks, and Charlotte M. Deane. CoV-AbDab: the coronavirus antibody database. *Bioinformatics (Oxford, England)*, 37(5):734–735, May 2021. ISSN 1367-4811. doi: 10.1093/bioinformatics/btaa739.
- Matthias Sachs, Benedict Leimkuhler, and Vincent Danos. Langevin Dynamics with Variable Coefficients and Nonconservative Forces: From Stationary States to Numerical Methods. URL https://www.mdpi.com/1099-4300/19/12/647.
- Saeed Saremi and Aapo Hyvärinen. Neural Empirical Bayes. *Journal of Machine Learning Research*, 20(181): 1–23, 2019. ISSN 1533-7928. URL http://jmlr.org/papers/v20/19-216.html.
- Meng Wang, Jonathan Patsenker, Henry Li, Yuval Kluger, and Steven H. Kleinstein. Supervised fine-tuning of pre-trained antibody language models improves antigen specificity prediction, May 2024. URL https://www.biorxiv.org/content/10.1101/2024.05.13.593807v1. Pages: 2024.05.13.593807 Section: New Results.
- Zachary Wu, Kadina E. Johnston, Frances H. Arnold, and Kevin K. Yang. Protein sequence design with deep generative models. *Current Opinion in Chemical Biology*, 65:18–27, December 2021. ISSN 13675931. doi: 10.1016/j.cbpa.2021.04.004. URL http://arxiv.org/abs/2104.04457. arXiv:2104.04457 [q-bio].
- Jiayao Zheng, Yu Wang, Qianying Liang, Lun Cui, and Liqun Wang. The Application of Machine Learning on Antibody Discovery and Optimization. *Molecules*, 29(24):5923, January 2024. ISSN 1420-3049. doi: 10.3390/molecules29245923. URL https://www.mdpi.com/1420-3049/29/24/5923. Number: 24 Publisher: Multidisciplinary Digital Publishing Institute.

A APPENDIX



Figure 4: An overview of context conditioning. a) B-cell receptor (BCR) repertoire data is taken from an antigen-exposed animal's immune system. The subclasses of chosen leads provide context for coditional sampling from a cWJS trained on the repertoire. b) V gene, J gene and CDR3 length define the clonotype subclass of a sequence and the context for conditioning. c) The EBM and score-based/denoiser are conditioned by adding a context embedding to a noisy sequence y prior to outputting energy (cEBM) or the predicted score.

A.1 TRAINING AND SAMPLING PARAMETERS

All models, including cWJS and baseline dWJS, are trained with a learning rate of 2e-4 that decays by 1% after each epoch. Antibody sequences are one-hot encoded and noised with $\sigma = 0.5$. Sampling takes 10 Langevin MCMC steps (Sachs et al.) with $\sigma = 0.5$ and $\delta = 0.5$

A.2 ABSOLVE FOR ANNOTATION

Absolve is a command-line tool for antibody variable domain sequence annotation (Genentech, 2018). The tool annotates single or paired-chain antibody sequences with open reading frames (ORFs), framework and CDR regions, Kabat numbering, germline assignments, and somatic hypermutations based on a set of reference germline sequences. Absolve is used for annotating the training BCR sequences and generated sequences with germline assignments, though its annotations are not perfect, especially when sequences are similar edit distances to two distinct germline sequences.

A.3 AUGMENTING NEGATIVE DATA FOR EBM TRAINING

Contrastive divergence loss is formulated as:

$$\mathcal{L} = \mathbb{E}_{(y_{\text{pos}}, c_{\text{pos}}) \sim \mathcal{D}_{\text{pos}}} [E(y_{\text{pos}}, c_{\text{pos}})] \\ - \mathbb{E}_{(y_{\text{neg}}, c_{\text{neg}}) \sim \mathcal{D}_{\text{neg}}} [E(y_{\text{neg}}, c_{\text{neg}})]$$

where \mathcal{D}_{pos} and \mathcal{D}_{neg} are the distributions of positive and negative samples, respectively.

To create a smoother energy lanscape, three negatives are sampled for each positive, each in a different way. The first negative comes from a traditional implementation of contrastive divergence, in which a random seed undergoes langevin MCMC, guided by the score of an underfit EBM (Frey et al., 2024). Additional negatives are sampled by pairing (y_{pos}) with an incorrect subclass. (y_{pos}, c_{rand}) is a negative with a randomly-drawn subclass and (y_{pos}, c_{family}) is in a related, but still incorrect subclass. In clonotype conditioning, c_{family} is from the same gene family but from a different gene (e.g. IGVH2-5 instead of IGHV2-11).



Figure 5: Sequence logo plots of the full training set (top), the subset of training data labeled as IGHV2-63 (middle), and the generated samples from cWJS with IGHV2-63 conditioning (bottom).



Figure 6: Multi-context conditioning on clonotypes (V-gene, J-gene, CDR3 length) from the TM dataset with each context condition shown separately, including context permutations not seen in the training set. All three contexts are sampled with high accuracy and diversity. Sample novelty also increases with guidance strength, as measured by edit distance to the nearest sequence in the training repertoire.



A.4 CONDITIONING ON NUMBER OF SOMATIC HYPERMUTATIONS

Figure 7: Accuracy and RMSE of generated samples with different number of somatic hypermutations (SHMs) used as conditioning. Accuracy is defined as within 3 SHMs of target; RMSE, root-mean squared error. This acts as a more challenging setting for conditioning on a numerical variable. RMSE drops as guidance increases for high SHM contexts, the most challenging.



(a) Trajectories when conditioned on IGHV6-5 and the ending fraction of samples in each V-gene. This is the most common convergence pattern seen using the TM dataset.



(b) Target V gene is IGHV2-63, with 98.8% of samples from IGHV2-63 after 10 steps while 1.2% are from other genes in the IGVH2 family.



(c) Only 18.6% of samples are from the target V gene (IGHV5-29) while 61.5% are from a nearby V gene, IGHV5-7.

Figure 8: Trajectories with varying convergence patterns. (a) Shows quick, targeted convergence, the most common pattern seen (b) Demonstrates slower, near-perfect convergence with exploration to nearby V-genes (c) Illustrates incomplete convergence with 61.5% of samples converging to a nearby V gene rather than the target.