

AIRIVA: A Deep Generative Model of Adaptive Immune Repertoires

Melanie F. Pradier

Microsoft Research, Cambridge, UK

MELANIEF@MICROSOFT.COM

Niranjani Prasad*

Microsoft Research, Cambridge, UK

NIRANJANI.PRASAD@MICROSOFT.COM

Paidamoyo Chapfuwa*

Microsoft Research, Redmond, USA

PCHAPFUWA@MICROSOFT.COM

Sahra Ghalebikesabi

Oxford University, Oxford, UK

SAHRA.GHALEBIKESABI@UNIV.OX.AC.UK

Maximilian Ilse

Microsoft Research, Amsterdam, The Netherlands

MAX.ILSE@MICROSOFT.COM

Steven Woodhouse

Adaptive Biotechnologies, Seattle, USA

SWOODHOUSE@ADAPTIVEBIOTECH.COM

Rebecca Elyanow

Adaptive Biotechnologies, Seattle, USA

RELYANOW@ADAPTIVEBIOTECH.COM

Javier Zazo

Microsoft Research, Cambridge, UK

JAVIERZAZO@MICROSOFT.COM

Javier Gonzalez Hernandez

Microsoft Research, Cambridge, UK

GONZALEZ.JAVIER@MICROSOFT.COM

Julia Greissl

Microsoft Research, Redmond, USA

JULIA.GREISSL@MICROSOFT.COM

Edward Meeds

Microsoft Research, Cambridge, UK

TED.MEEDS@MICROSOFT.COM

Abstract

Recent advances in immunomics have shown that T-cell receptor (TCR) signatures can accurately predict active or recent infection by leveraging the high specificity of TCR binding to disease antigens. However, the extreme diversity of the adaptive immune repertoire presents challenges in reliably identifying disease-specific TCRs. Population genetics and sequencing depth can also have strong systematic effects on repertoires, which requires careful consideration when developing diagnostic models. We present an Adaptive Immune Repertoire-Invariant Variational Autoencoder (AIRIVA), a generative model that learns a low-dimensional, interpretable, and compositional representation of TCR repertoires to disentangle such systematic effects in repertoires. We apply AIRIVA to two infectious disease case-studies: COVID-19 (natural infection and vaccination) and the Herpes Simplex Virus (HSV-1 and HSV-2), and empirically show that we can disentangle the individual disease signals. We further demonstrate AIRIVA’s capability to: learn from unlabelled samples; generate in-silico TCR repertoires by intervening on the latent factors; and identify disease-associated TCRs validated using TCR annotations from external assay data.

*. These authors contributed equally.




Figure 1: **In-vivo generative process of immune repertoires.** *Antigens* are presented by *human leukocyte antigens* (HLAs) to T cells. When a T cell binds to a specific antigen-HLA complex, it undergoes clonal expansion (A), which is captured in a blood sample (B) and quantified via high-throughput immuno-sequencing resulting in a TCR *repertoire* (C).

1. Introduction

Precision medicine, that is diagnosis and treatment targeted to an individual, is one of the most promising applications of machine learning (ML) in healthcare. Because of an explosion in biomarkers and corresponding large datasets, there has been increased interest in applying deep learning artificial intelligence (AI) models in a clinical setting (Miotto et al., 2018; Sandeep Kumar and Satya Jayadev, 2020). In radiology and pathology settings, AI-based diagnostics and human-in-the-loop systems are now common (Rajpurkar et al., 2022; Oktay et al., 2020). Data complexity and clinical risk have spurred the development of ML models that are robust, interpretable and fair. This includes works that explicitly incorporate interpretability and resilience to distribution shifts in real-world domains such as longitudinal biomarker modeling (Hussain et al., 2021), biomarker discovery (Pradier et al., 2019) and medical imaging (Chartsias et al., 2019; Ilse et al., 2020).

Our focus is **immunomics**, where recent advances in high-throughput T cell sequencing (Robins et al., 2012) have opened the door to a new form of precision diagnostic based on adaptive immune repertoires (see Figure 1). Immunomics-based diagnostic models have shown promise in detecting new or past infection in Cytomegalovirus (Emerson et al., 2017a), Lyme (Greissl et al., 2021), and COVID-19 (Snyder et al., 2020). TCR sequencing can also be used to identify neoantigen-specific T cells that are present in most cancers and aid in the development of cancer therapies, such as mRNA vaccines (Cafri et al., 2020). Moving forward, we expect improved knowledge of the TCR repertoires to be useful for the development of novel diagnostics and therapeutics, especially for immune-mediated disease contexts where standard of care is lacking. But with such promise comes the challenge of validating signal and building clinical trust within an extremely complex, diverse, high-dimensional data domain that lacks human-understandable structure.

To address these challenges, a *generative* model framework that explicitly incorporates *interpretable* latent factors linked to repertoire targets (e.g. disease labels, genetic variants, sequencing depth, batch id) is highly desirable. We propose the **Adaptive Immune Repertoire-Invariant Autoencoder** (AIRIVA) as a solution: a deep generative model

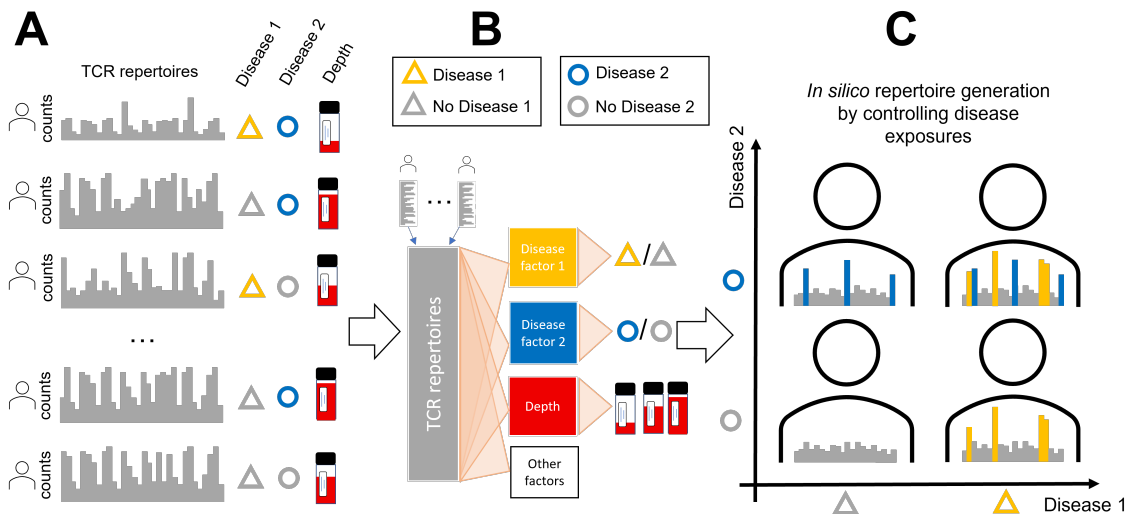


Figure 2: **Predictive and generative data flow in AIRIVA.** (A) Training data consists of TCR counts and repertoire labels. (B) TCRs are mapped to independent latent factors (factorized representation) which are subsequently used for prediction. (C) By intervening on targets and depth (toggles of a simulator), in-silico repertoires are generated.

for TCR repertoires, an adaptation of recent work (Ilse et al., 2020; Joy et al., 2021) to immunomics. AIRIVA is a semi-supervised generative model trained to explicitly learn disentangled and interpretable latent representations of TCR repertoires and can be used to both predict multiple disease labels and to generate in-silico repertoires, which can in turn be used for assigning TCR-disease associations. Figure 2 illustrates these key ideas.

Generalizable Insights about Machine Learning in the Context of Healthcare

- Immunomics is a relatively new field that has the potential to revolutionize precision medicine. We propose a generative model for handling complex immunomics data which is high dimensional, sparse, and heterogeneous.
- By enforcing a disentangled latent representation, we are able to learn interpretable *label-specific* factors of variation.
- We propose an approach for generating in-silico (synthetic) *counterfactual* repertoires with AIRIVA by *intervening* on specific latent factors of interest. We leverage the generated counterfactuals for model selection and identification of label-specific TCRs, validated with an external dataset of TCR annotations.
- We empirically demonstrate AIRIVA’s capacity to disentangle disease signal in the context of two infectious diseases, COVID-19 and Herpes Simplex Virus, improving model robustness across subgroups of interest.
- We demonstrate improved AIRIVA label prediction by learning with additional TCR repertoires with missing labels. Leveraging samples with missing labels is crucial for the immunomics setting where cohorts are heterogeneous with small sample sizes.

2. Immunomics

T cells are a central component of the adaptive immune response, complementary to antibodies secreted by B cells. Their role is to identify and destroy infected or cancerous cells that are expressing non-native antigens—small fragments of proteins presented by *human leukocyte antigens* (HLAs)—and to recognize antigens collected from other immune cells as self or non-self, to regulate the overall response (Murphy and Weaver, 2016). The adaptive immune system “learns” to recognize antigens by first generating random, non-self binding naive cells, which then undergo clonal expansion upon encountering their cognate antigen (Mazza and Malissen, 2007). Thus an immune repertoire, sampled from $O(10^7)$ unique TCRs (Emerson et al., 2017a) in an individual, is a mixture of naive (random) and expanded (antigen-specific) T cells. Individuals with shared HLAs who are exposed to similar antigens (e.g. viral infections, vaccinations), tend to share similar and even identical TCRs (Dendrou et al., 2018); this mechanism underpins statistical immunomics modeling (Greiff et al., 2020) (see Section 2.2).

2.1. Modeling Challenges

High-throughput sequencing given a standard blood sample (Robins et al., 2012) has opened the door to a new form of precision diagnostic. The complexity and diversity of repertoires presents several important modeling challenges, see (Pavlovic et al., 2022) for a detailed overview. Repertoire composition is influenced by many factors, primarily antigen **exposure history**, which is in turn influenced by risk factors, geography, personal health choices, genetics, age, sex, *etc.* Further variation stems from **sequencing depth** (here, the total number of unique TCRs), as well as potential batch effects introduced by differences in collection methods and sequencing protocols. Based on past studies, T cell signatures are estimated to account for $< 1\%$ of the T cells in a repertoire (Grifoni et al., 2020) in the majority of viral infections. This means that only a small fraction of the $O(10^9)$ total unique TCR sequence set in a population of repertoires is relevant to a disease of interest.

Another challenge is that viral genomes often have high homology resulting in a significant number of **shared antigens** (Kieff et al., 1972). The same is true for vaccines derived from specific sub-units of viral proteins. In both cases, we would expect predictive models to disentangle the shared TCRs from label-specific signals. Finally, due to their fundamental role in antigen presentation (Reynisson et al., 2020) and TCR-antigen binding (La Gruta et al., 2018) **HLAs** directly influence an individual’s TCR repertoire. In some auto-immune diseases such as celiac disease and multiple sclerosis, HLAs can act as confounders because they directly influence both disease risk and the TCR repertoire (Romanos et al., 2009; Baranzini, 2011). Our work focuses on demonstrating disease disentanglement and we leave HLA associations to future work.

2.2. Enhanced Sequence Models

Since only a small fraction of TCR sequences are statistically significantly associated with a single disease, it is standard in immunomics modeling to perform Fisher’s Exact Test (FET) per TCR to identify these sequences (Emerson et al., 2017a). Specifically, J TCRs, called **enhanced sequences** (ES), are selected based on some p -value threshold, often reducing

the feature set to numbers in the 100s or 1000s down from the *billions* in the full training set of repertoires. Given an ES list, a simple but robust logistic regression model can be trained using two features: the sum of ES counts in each repertoire and sequencing depth. Prior work has established *Enhanced Sequence Logistic Growth* (ESLG) as a competitive benchmark for TCR-based diagnostics (Greissl et al., 2021; Snyder et al., 2020). ESLG calibrates classification scores as a function of sequencing depth on the control population, which makes it generally more robust across depth. This simple model is suitable when we are learning from a single binary label but is ill-equipped to handle more complex immunomics data characterized by confounding or covariate shifts (such as age or HLA status) between cases and controls that may influence TCR selection. Moreover, ESLG assumes that the ES counts is only a function of the sequencing depth from the control population, which is restrictive and rarely satisfied in practice.

2.3. Deep Learning Applied to TCR Sequences

Our work focuses on statistical models of the discrete count representation of TCR repertoires and we do not directly use the amino-acid sequence information. However, there are several recent developments applying deep learning to TCR sequences, see (Pertseva et al., 2021) for an overview. DeepTCR (Sidhom et al., 2021) uses deep VAEs of TCR sequences to learn features useful for antigen binding prediction and repertoire classification using a multiple instance learning framework. DeepRC (Widrich et al., 2020) incorporates Hopfield networks into BERT transformers to classify repertoires in the context of existing meta-data. Although promising, deep sequence models struggle to outperform simple regularized logistic regression models (Kanduri et al., 2022; Emerson et al., 2017b).

3. The Adaptive Immune Repertoire Invariant Variational Autoencoder

A variational autoencoder (VAE) is a probabilistic model that learns a mapping from some input data \mathbf{X} , here corresponding to TCR counts within a repertoire, to a latent representation \mathbf{Z} (Kingma and Welling, 2013; Rezende et al., 2014). This latent representation summarizes the information of the input data such that we can reconstruct the original data \mathbf{X} from the latent representation \mathbf{Z} with high-fidelity. AIRIVA is based upon the Domain Invariant Variational Autoencoder (DIVA) (Ilse et al., 2020) and the Capturing Characteristic VAE (CC-VAE) (Joy et al., 2021), but applied and extended to immunomics. Both CC-VAE and DIVA leverage label information \mathbf{Y} to disentangle the latent representation, where some of the dimensions of \mathbf{Z} are constrained to also predict repertoire labels \mathbf{Y} accurately. We can thus decompose \mathbf{Z} as the concatenation of two kinds of latent variables. We call *predictive latents* $\mathbf{Z}_{\mathbf{Y}}$ those latent dimensions that are label-specific, and *residual latents* \mathbf{Z}_{ϵ} those that are label-agnostic. Further, as illustrated in DIVA and CC-VAE, these models are capable of learning from samples with missing labels by marginalizing the unobserved labels. In the immunomics setting this is very advantageous because it allows the model to learn from large cohorts of unlabelled repertoires. Figure 3 shows the probabilistic graphical model of AIRIVA, which matches that of Ilse et al. (2020) and Joy et al. (2021).




Figure 3: **Probabilistic graphical model for AIRIVA.** The generative and inference models are defined by solid and dashed arrows, respectively. We assume J TCRs \mathbf{X}_j in each repertoire along K associated labels (observed or missing) \mathbf{Y}^k , *e.g.*, disease, sequencing depth, *etc.* We generate and infer the latent factors $\{\mathbf{Z}_y^k, \mathbf{Z}_\epsilon\}$ independently from other factors which leads to an interpretable disentangled representation, enabling counterfactual generation, as described in Section 3.2. All distributions over latents \mathbf{Z} are parameterized Normal distributions; TCR counts \mathbf{X} are generated by Poisson distributions; labels \mathbf{Y} are distributed by Bernoulli, categorical, or log-normal depending on the setting.

3.1. Model Description and Objective Function

For simplicity and brevity, we restrict the AIRIVA objective function formulation to the **fully-supervised** setting (see Appendices B.2 and B.3 for the **semi-supervised** setting). Our goal is to maximize the joint log-likelihood defined by the *intractable* integral over latent variables \mathbf{Z} :

$$\log p_\theta(\mathbf{X}, \mathbf{Y}) = \log \int_{\mathbf{Z}} p_\theta(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) d\mathbf{Z}, \quad (1)$$

where $p_\theta(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) = p_\theta(\mathbf{X}|\mathbf{Z})p_\theta(\mathbf{Z}|\mathbf{Y})$ and θ represent the neural network parameters of the generative model. By introducing a variational distribution $q_\phi(\mathbf{Z}|\mathbf{X})$, where ϕ are the neural network parameters of the inference model, we can tractably maximize the variational evidence lower bound (ELBO):

$$\log p_{\theta, \phi}(\mathbf{X}, \mathbf{Y}) \geq \ell_{\text{ELBO}}(\theta, \phi; \mathbf{X}, \mathbf{Y}) \quad (2)$$

By factorizing both the generative and inference model over K labels, we have:

$$p_\theta(\mathbf{Z}|\mathbf{Y}) = p(\mathbf{Z}_\epsilon) \prod_k p_\theta^k(\mathbf{Z}_Y^k | Y^k) \quad (3)$$

$$q_\phi(\mathbf{Z}|\mathbf{X}) = q_\phi^\epsilon(\mathbf{Z}_\epsilon|\mathbf{X}) \prod_k q_\phi^k(\mathbf{Z}_Y^k|\mathbf{X}). \quad (4)$$

Following the formulation from Ilse et al. (2020), the ELBO simplifies as follows:

$$\begin{aligned} \ell_{\text{ELBO}}(\theta, \phi; \mathbf{X}, \mathbf{Y}) &= \mathbb{E}_{q_\phi(\mathbf{Z}|\mathbf{X})} [\log p_\theta(\mathbf{X}|\mathbf{Z})] \\ &\quad - \beta_\epsilon \cdot KL(q_\phi^\epsilon(\mathbf{Z}_\epsilon|\mathbf{X}) \| p(\mathbf{Z}_\epsilon)) \\ &\quad - \sum_k \beta_k \cdot KL(q_\phi^k(\mathbf{Z}_Y^k|\mathbf{X}) \| p_\theta^k(\mathbf{Z}_Y^k|Y^k)), \end{aligned} \quad (5)$$

where $\{\beta_\epsilon, \beta_k\}$ are additional hyperparameters that can be scheduled during training to prevent posterior collapse by KL divergence annealing (Higgins et al., 2017; Bowman et al., 2015) or otherwise set to trade-off various parts of the objective function.

This objective function is missing a key driver required to enforce disentanglement, a label classification term (Locatello et al., 2019; Kim and Mnih, 2018). Following Ilse et al. (2020) we introduce an auxiliary objective:

$$\ell_{\text{aux}}(\boldsymbol{\omega}, \boldsymbol{\phi}; \mathbf{X}, \mathbf{Y}) = \sum_k \alpha_k \mathbb{E}_{q_\phi^k(\mathbf{Z}_Y^k | \mathbf{X})} \left[\log q_\omega^k(Y^k | \mathbf{Z}_Y^k) \right], \quad (6)$$

with independent label-specific prediction neural networks with parameters $\boldsymbol{\omega}$ to encourage learning of label-specific independent factors¹ and α_k hyperparameters for weighting the auxiliary predictor. For the **fully supervised** setting, the complete AIRIVA regularized variational objective for jointly learning all model parameters $\{\boldsymbol{\theta}, \boldsymbol{\psi}, \boldsymbol{\omega}\}$ via stochastic gradient ascent is:

$$\ell(\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\omega}; \mathbf{X}, \mathbf{Y}) = \ell_{\text{ELBO}}(\boldsymbol{\theta}, \boldsymbol{\phi}; \mathbf{X}, \mathbf{Y}) + \ell_{\text{aux}}(\boldsymbol{\omega}, \boldsymbol{\phi}; \mathbf{X}, \mathbf{Y}). \quad (7)$$

3.2. In-silico Generation of Repertoires

Following a standard VAE setup, we generate synthetic (in-silico) repertoires $\mathbf{X} \sim p_\theta(\mathbf{X} | \mathbf{Z})$ from AIRIVA, where \mathbf{Z} is drawn according to the prior distribution from Eq. (3) also known as **conditional generation**. With conditional generation, we can synthesize samples that match the *factual* (empirical) TCR repertoire distribution. This can be particularly useful to balance training datasets, by increasing the number of repertoires for minority subgroups via targeted data augmentation. However, we may also want to explore *what if* questions to interpret the change in the outcome variable that we would observe, had we intervened on one of the inputs. In the immunomics context, this could help answer questions such as: “*what would the immune response look like if the patient had not been exposed to certain antigens*”, or “*which set of TCRs would we expect to observe if the healthy control had been vaccinated?*”, holding other factors constant. Reliably estimating counterfactuals allows us to better characterize the individual-level TCR response to the disease as well as identify disease-specific TCRs at the population level, given only observational data with potential biases across case and control populations. Below we detail the procedure to generate counterfactuals and infer label-specific TCRs.

Counterfactual Generation Following the potential outcomes framework (Rubin, 2005), we assume a binary *treatment* Y^k where each TCR X_j has two potential outcomes denoted by $X_j(Y^k = 0)$ and $X_j(Y^k = 1)$. In practice, for each sample, we only observe the *factual* outcome $X_j(Y^k)$, while the *counterfactual* $X_j(1 - Y^k)$ is unobserved. Hence, we generate counterfactuals with AIRIVA as follows:

1. sample from posterior $\mathbf{Z} \sim q_\phi(\mathbf{Z} | \mathbf{X})$ in Eq. (4) (equiv. to computing exogenous noise).

1. Using a similar graphical model, Joy et al. (2021) derives the objective assuming $q_\phi(\mathbf{Z} | \mathbf{X}, \mathbf{Y})$, rather than $q_\phi(\mathbf{Z} | \mathbf{X})$ in Ilse et al. (2020), resulting in the classifier naturally appearing in the objective. However, the final objective in both cases is similar, modulo a weighting in the expectation.

2. assign posterior *label-specific* latent factor \mathbf{Z}_Y^k to the sample from the conditional prior in Eq. (3) $\mathbf{Z}_Y^k \sim p_\theta(\mathbf{Z}_Y^k | 1 - Y^k)$ (equivalent to *intervening* on Y^k since \mathbf{Z}_Y^k is independent to all the other latent factors by construction).
3. sample *counterfactual* immune repertoire $\mathbf{X}(1 - Y^k) \sim p_\theta(\mathbf{X} | \mathbf{Z})$.

Validity of Counterfactuals The counterfactuals $\mathbf{X}(1 - Y^k)$ are valid under the following conditions: (i) the $KL(q_\phi^k(\mathbf{Z}_Y^k | \mathbf{X}) || p_\theta^k(\mathbf{Z}_Y^k | Y^k))$ divergence between the label-specific posterior and prior is close to zero; (ii) the underlying causal model follows (Suter et al., 2019), *i.e.*, there are no direct arrows between the labels; (iii) the overlap assumption holds, *i.e.*, $0 < P(Y^k | \mathbf{X}) < 1$, almost surely if $P(\mathbf{X}) > 0$. See Appendix A for more details.

Inferring Label-Specific TCRs Finally, we leverage generated counterfactuals from AIRIVA to infer label-specific TCRs. To quantify the expected effect of a given label Y^k on each individual TCR X_j for a given subpopulation, we can compute the *conditional average treatment effect* (CATE), defined as:

$$\text{CATE}_j(\mathbf{y}^{\setminus k}) := \mathbb{E}_{X_j} \left[X_j(Y^k = 1) - X_j(Y^k = 0) \mid \mathbf{Y}^{\setminus k} = \mathbf{y}^{\setminus k} \right], \quad (8)$$

where a subpopulation is defined by a fixed value assignment for the other labels $\mathbf{Y}^{\setminus k} = \mathbf{y}^{\setminus k}$. In our experiments, we rank TCRs according to CATE_j such that the inferred label-specific TCRs are those with the highest CATE_j .

4. Experiments

In this section, we demonstrate the capacity of AIRIVA to disentangle disease-specific TCRs in the context of two infectious diseases: COVID-19 and the Herpes Simplex Virus. We also show that AIRIVA can learn from unlabelled samples, generate realistic in-silico TCR repertoires, and learn a latent space that aligns with biological intuition. We first outline general procedures for training, model selection and model evaluation used in our experiments. This is followed by detailed descriptions of experimental procedures and results for specific case studies. See Appendix for further experimental details, results and analyses.

4.1. Experimental Details

Before training, we rank TCRs using one-sided FET p -values on occurrence in cases vs. controls and select the top TCRs. This ensures selection of relatively public TCRs (Emerson et al., 2017b). **Training** consists of stochastic gradient ascent using the objective function defined by Eq. (7). Because the objective function is composed of several (hyperparameter-weighted) sub-optimization problems, the typical **model selection** procedure of choosing the model with the best maximum objective on the validation set sometimes results in poor models, *e.g.*, that have good reconstruction but have ignored the latent factor constraints. We have empirically found that a uniform mixture of classification metrics² on the factual (validation) *and* on counterfactual (validation) data reliably selects good models. This

2. We use *concentrated* area under the ROC curve (AUCROC) (Swamidass et al., 2010) which modifies typical AUROC by emphasizing sensitivity.

aligns with previous works that link counterfactual performance to disentanglement and model robustness (Suter et al., 2019; Besserve et al., 2018).

We are interested in **evaluating** the following key aspects of AIRIVA: (1) predictive performance of AIRIVA when compared to ESLG (see Section 2.2) and a standard multi-label feed forward network (FFN) classifier; (2) qualitative inspection of latent factor disentanglement, both from its conditional prior and posterior distributions; (3) similar evaluation of counterfactuals in the latent space for correctness; (4) when applicable, validation of AIRIVA estimated TCR-associations based on CATE (see Section 3.2) with an independently verified set of annotated TCRs from an external assay dataset (see Section 4.2).

4.2. COVID-19 Case Study

The viral envelope of SARS-CoV-2 is made up of four proteins: the membrane protein, the envelope protein, the nucleocapsid protein and the spike protein (Jackson et al., 2022). The spike protein is believed to be responsible for binding to host cells and initiating viral infection, and has therefore been the target of vaccine development (McCarthy et al., 2021). During natural infection, we expect to see the expansion of TCRs associated with any one of these proteins, while in the case of vaccination³ we should only observe the expansion of TCRs binding the spike protein. We would like to build a model that can disentangle natural infection from vaccination. Because TCRs clonally expanded during natural infection are a superset of the spike-only TCRs expanded following vaccination, disentangling both sets of TCRs is a challenging task.

Data Cohorts To train our models we used 1,954 samples from donors with natural SARS-CoV-2 infection, 477 healthy donors post SARS-CoV-2 vaccination, 5,198 healthy controls sampled prior to March 2020, and no donors who were both naturally infected and later vaccinated. The classifiers are tested on a holdout set of 525 samples from naturally infected donors, 400 vaccinated donors, 100 donors who were naturally infected and later vaccinated, and 4,606 healthy controls.

We posit $\mathbf{Y} = \{Y_{\text{ns}}, Y_{\text{s}}, Y_{\text{depth}}\}$, where Y_{ns} and Y_{s} refer to exposure to the non-spike and spike proteins respectively, and Y_{depth} refers to sequencing depth—a standardized measure of log total template count—as there is a strong depth imbalance between the vaccinated subpopulation and others. For Y_{ns} , naturally infected samples are positives, and vaccinated and healthy controls are negatives. For Y_{s} , both naturally infected and vaccinated samples are positives and healthy controls are negatives.

TCR Selection We ran a sweep of 100 ESLG models for each prediction task (spike and non-spike) to determine the best threshold number η of input TCRs, where $\eta \in [200, 500, 1000, 5000, 10000]$. Further, we select η that maximizes the predictive performance of ESLG in the validation set. The final input TCRs used to train AIRIVA are based on: *i*) 1,254 TCR sequences from the union of 1,000 most significant η TCRs selected by FET per binary label (non-spike and spike); and *ii*) 342 TCRs associated with cytomegalovirus (CMV) to explicitly test AIRIVA’s robustness to noise induced by an unrelated signal from another infectious disease, yielding a total of $J = 1,596$ input TCRs.

3. We consider viral vector vaccines AZD1222 (Oxford-AstraZeneca) and Ad26.COV2.S (Johnson & Johnson), as well as mRNA vaccines BNT162b2 (Pfizer) and mRNA-1273 (Moderna).

External Spike and Non-Spike Annotations In addition to repertoire labels, in the case of COVID-19, we have access to in-vitro labels for spike and nonspike TCR associations from external MIRA (**M**ultiplexed **I**dentification of T cell **R**eceptor **A**ntigen specificity) assay data (Klinger et al., 2015). This data comprises over 400k experimentally-derived TCR-antigen binding pairs (*hits*) for the spike protein, and 1.2M hits for non-spike proteins (Nolan et al., 2020). A total of 291 spike and 462 non-spike MIRA hits appear in our input set of 1,596 TCRs. By intersecting public COVID-19-associated TCRs with this database, it is possible to associate a subset of the TCRs to an antigen (spike or non-spike) with high confidence (Li et al., 2022). We use these spike and non-spike protein-specific TCRs from MIRA to assess the quality of the inferred TCR-label associations from AIRIVA.

COVID-19 RESULTS

Disease Classification FET sequence selection for the non-spike label uses the natural infection subgroup as cases, however, these contain both the spike and non-spike signal. Hence, we expect non-spike AIRIVA predictions to outperform predictions from the non-spike ESLG model given natural infection samples as cases and vaccinated samples as controls. This is due to ESLG’s lack of robustness to noisy input TCRs (*i.e.*, shared TCRs across disease labels). To demonstrate this, we report holdout performance on the subgroups listed in Table 1.

Table 1: **Subgroups used for evaluation of COVID-19 models.**

	Cases	Controls
Overall	all natural infection	vaccinated, healthy controls
Unvaccinated	natural infection	healthy controls
Vaccinated	natural infection + vaccinated	vaccinated

Table 2 compares the discriminative value of non-spike AIRIVA against FFN and ESLG for these different case/control groups.⁴ We can see that for the entire population as well as when comparing natural infection to healthy controls all three models perform comparably. However, as expected, AIRIVA significantly outperforms ESLG and FFN when comparing natural infection samples with subsequent vaccination against vaccinated samples. Figure 4 shows the corresponding Receiver-Operating Characteristic (ROC) curves for the three groups considered, highlighting the degradation in model performance of ESLG for the vaccinated subgroup. These results indicate that AIRIVA is robust to shared antigens in disease labels and can disentangle spike-associated disease signal from non-spike signal.

Interpreting Learnt Latent Space Figure 5(a) plots posterior samples from $q_\phi(\mathbf{Z}_Y|\mathbf{X})$ for the holdout data along with the exact learned prior predictive distributions $p_\theta(\mathbf{Z}_Y|\mathbf{Y})$ in the latent subspace ($\mathbf{Z}_Y = \{Z_{\text{non-spike}}, Z_{\text{spike}}\}$), stratified by subgroups. Note that one of the label groups, corresponding to “*non-spike only*” exposure, is never observed in our data since both spike and non-spike co-occur in natural infection. As expected, the vaccinated

4. In the Appendix, we report comparative performance for the spike label; AIRIVA achieves similar performance against ESLG and FFN, which is expected and further validates the modeling framework.

Table 2: **Comparison of non-spike disease models.** Sensitivity at 98% specificity and AUROC for predicting natural infection, overall and stratified by vaccination status (in both cases and controls), using 100 bootstraps samples. AIRIVA performs comparatively against Enhanced Sequence Logistic Growth (ESLG) overall, but is more robust at separating naturally infected samples with subsequent vaccination from vaccinated samples.

Model	Overall		Unvaccinated		Vaccinated	
	Sensitivity	AUROC	Sensitivity	AUROC	Sensitivity	AUROC
ESLG	0.74 ± 0.07	0.94 ± 0.02	0.85 ± 0.04	0.94 ± 0.02	0.27 ± 0.15	0.86 ± 0.06
FFN	0.76 ± 0.05	0.92 ± 0.02	0.77 ± 0.05	0.92 ± 0.02	0.63 ± 0.16	0.89 ± 0.06
AIRIVA	0.76 ± 0.06	0.93 ± 0.02	0.76 ± 0.05	0.93 ± 0.02	0.73 ± 0.12	0.94 ± 0.04

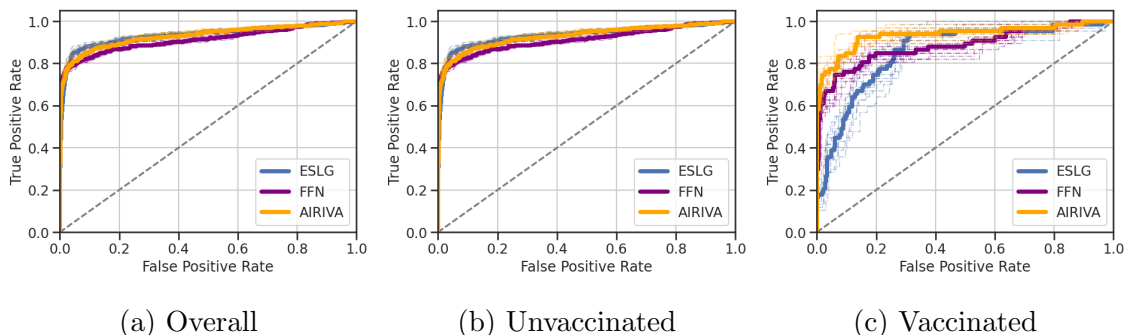


Figure 4: **ROC curves for non-spike disease models.** Bold lines correspond to the ROC curve on holdout; colored dashed lines correspond to 10 additional bootstrap samples. ESLG performs poorly for the vaccinated subgroup, while AIRIVA performance remains competitive, with similar predictive performance compared to the overall population.

subgroup exhibits only spike signal (top-left), whereas COVID-positive cases with or without vaccination show exposure to both spike and non-spike antigens (top-right).

Such a structured latent space provides a way to interpret both the composition and the strength of the observed immune response and could elicit new biological insights. For example, Figure 5(b) shows posterior samples of the vaccinated-only repertoires in holdout stratified by vaccine type; samples from other subgroups are plotted in grey. Our model suggests that repertoires that were administered the AstraZeneca viral vector AZD1222 vaccine exhibit lower spike T-cell response compared to mRNA vaccines from Johnson & Johnson, Pfizer, or Moderna which is consistent with findings in the literature (Schmidt et al., 2021; Predecki et al., 2021; Marking et al., 2022).

Conditional Generation To demonstrate the disentanglement of spike and non-spike TCR response in AIRIVA, we evaluate whether conditionally generated repertoires have the characteristics we expect. Specifically, repertoires of vaccinated and COVID-positive individuals should contain higher counts of spike-specific TCRs, while non-spike-specific TCRs should only activate for COVID-positive individuals. To test this, we generate repertoires conditionally by intervening on the spike and non-spike labels, at a fixed mean repertoire sequencing depth.

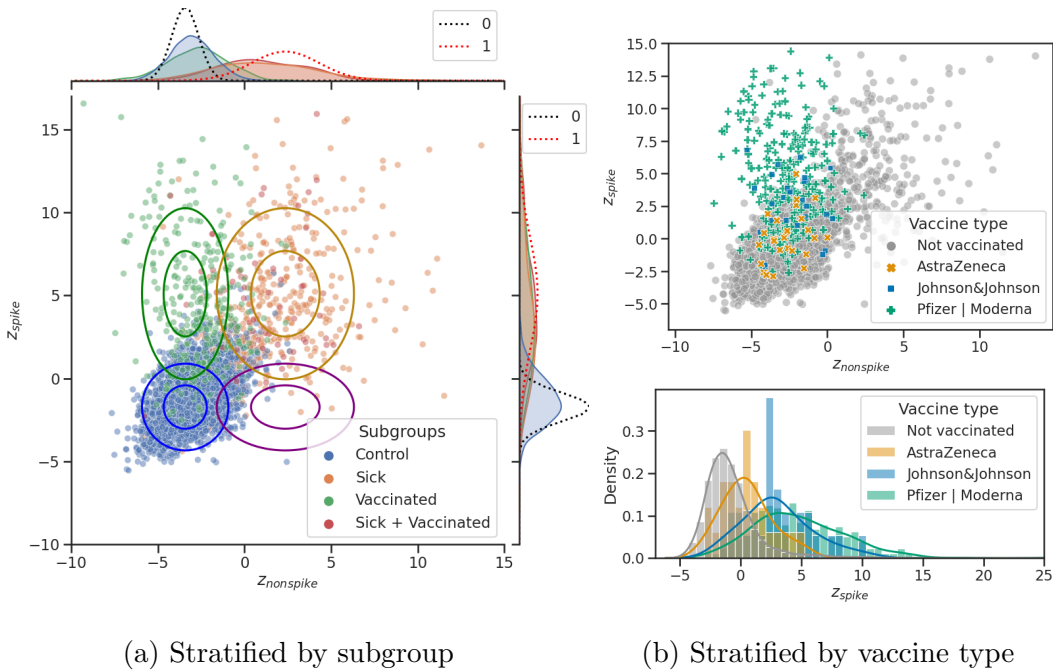


Figure 5: **Latent space $q_\phi(\mathbf{Z}_Y|\mathbf{X})$ for COVID-19 in holdout.** AIRIVA learns an interpretable latent space that matches biological intuition. (a) Each point corresponds to one posterior sample. Conditional priors $p_\theta(\mathbf{Z}_Y|\mathbf{Y})$ for each combination of labels are represented as ellipses, each circle denoting one standard deviation. On the sides, we show the marginal posteriors as continuous lines and conditional priors as dashed lines. (b) Top: posterior samples for vaccinated subgroups (colored), samples for other subgroups are shown in grey; Bottom: histogram of posterior samples on the Z_{spike} factor dimension.

Figure 6 shows the average count per TCR estimated by AIRIVA (*i.e.*, per-TCR Poisson rate for conditionally generated repertoires, averaged across repertoires) and empirical average count per TCR. We split TCRs based on external annotations of spike and non-spike MIRA hits (described in section 4.2), and CMV enhanced sequences. As expected, generating “vaccinated” repertoires with spike only yields increased average per-TCR count for spike MIRA hits. Moreover, generating repertoires with natural infection (conditioning on positive spike and non-spike labels) shows increased average per-TCR counts in both spike and non-spike MIRA TCRs. Average count of CMV-associated sequences remains negligible across all label groups. This provides empirical evidence that AIRIVA is robust to irrelevant background signal, which is captured by the *label-agnostic* latent factor \mathbf{Z}_ϵ . The conditionally generated TCR counts match factual counts well in the controls, however, AIRIVA overestimates counts in the natural infection samples. One hypothesis for this behavior is that sequence occurrence in controls is mostly due to non-expanded naive TCRs which occur randomly in repertoires in accordance with a well-defined generative process. This sampling should match our assumed Poisson distribution of TCR counts well. However, TCR counts in subpopulations that contain expanded TCRs may not follow a Poisson distribution because of extra variance introduced in the TCR counts by clonal expansion.

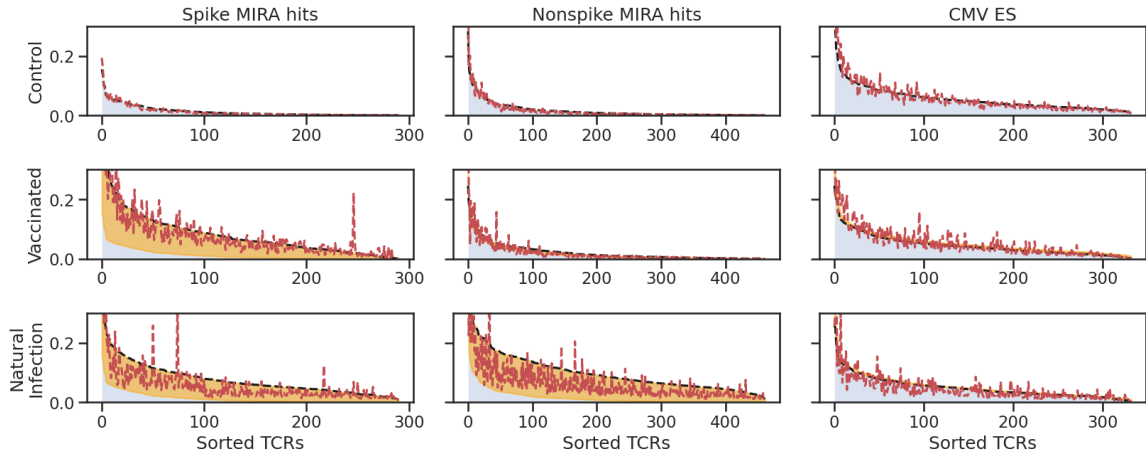


Figure 6: **Conditional generation with AIRIVA.** Per-TCR mean counts estimated by AIRIVA (average Poisson rate as dashed black line) or empirically (dashed red line), averaged across repertoires. We generate 5,000 in-silico repertoires for each subgroup; the pool of TCRs is the same as the initial set of 1596 input TCRs for all subgroups. TCRs are stratified according to TCR annotations and sorted in descending order per estimated CATE. The yellow highlights the difference w.r.t controls. In-silico repertoires generated by AIRIVA exhibit per-TCR mean counts similar to the observational data.

Consistent Counterfactual Generation for COVID-19 Figure 7 shows posterior samples (first row) and posterior predictive samples (second row) for one of the observed subgroups (first column), and counterfactuals (other columns) transformed from one subgroup to another. This figure demonstrates that we can generate consistent counterfactuals that exhibit expected properties by manipulating specific labels. Importantly, the latent representation of such counterfactuals maps to a similar region as their corresponding factual subgroups in Figure 5. Note that the last column corresponds to an *unobserved* repertoire subgroup: these are out-of-distribution counterfactuals. Interestingly, AIRIVA captures the correct direction of variation in the latent space for that novel subgroup without introducing additional inductive biases often necessary for zero-shot learning.

Inferring Label-Specific TCRs We use AIRIVA to explicitly infer label associations for input TCRs, by evaluating CATEs for spike and non-spike exposure respectively per TCR, as described in Section 3.2. Specifically, we define Spike CATE and Non-spike CATE for each TCR X_j as follows:

$$\begin{aligned} \text{Spike CATE} &:= \mathbb{E}_{X_j} [X_j(Y_s = 1) - X_j(Y_s = 0) | Y_{ns} = 0] \\ \text{Non-spike CATE} &:= \mathbb{E}_{X_j} [X_j(Y_{ns} = 1) - X_j(Y_{ns} = 0) | Y_s = 1], \end{aligned} \quad (9)$$

where each potential outcome is either observed (factual) or estimated via counterfactuals generated by AIRIVA⁵. We rank COVID-19 TCRs in descending order according to their

5. Note that we condition on Spike-positive for Non-spike CATE as both source and target subgroups in counterfactual generation need to be observed for counterfactuals to be well-defined, as discussed in Section 3.2

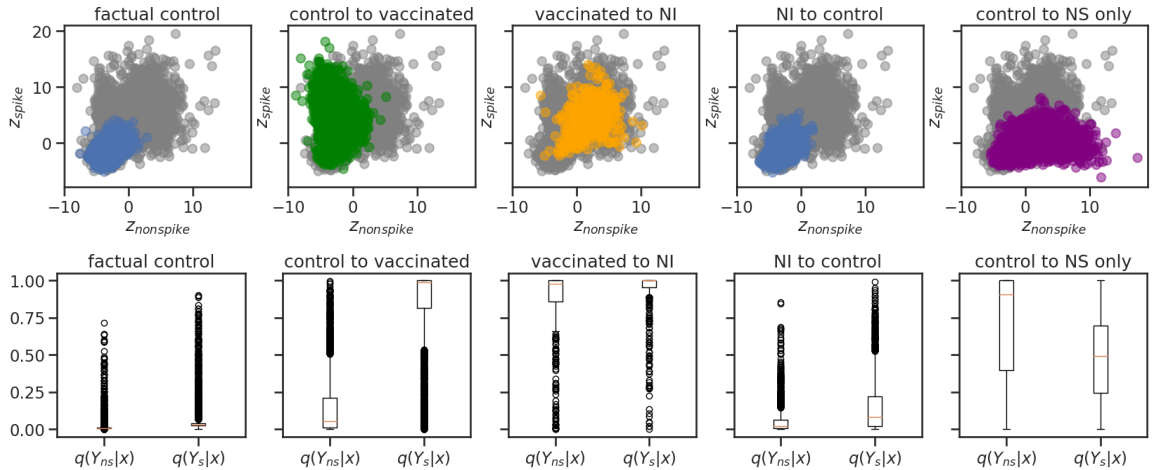


Figure 7: **Counterfactual generation for COVID-19.** (First row) Posterior samples $\mathbf{Z}_{\mathbf{Y}} \sim q_{\phi}(\mathbf{Z}_{\mathbf{Y}}|\mathbf{X})$ for factual data (first column) and counterfactual data (other columns). (Second row) Posterior predictive samples $\mathbf{Y} \sim q_{\omega}(\mathbf{Y}|\mathbf{Z}_{\mathbf{Y}})$ for each binary label. See the Appendix for all possible counterfactual combinations stratified by subgroups.

spike or non-spike CATE values. Then we leverage ground truth MIRA labels to compute a cumulative count of spike or non-spike MIRA TCRs (illustrated as blue or orange, respectively in Figure 8). We expect TCRs with the largest spike CATE to be mostly spike MIRA TCRs and vice-versa which is indeed what we observe. Note that we are only able to assign a little over half of our total TCRs from FET to spike or non-spike via MIRA and are not showing unassigned TCRs. We evaluate AIRIVA’s ability to recover ground truth spike on non-spike MIRA sequences using CATE (9) via average precision and false discovery rates (FDR), defined as:

$$\text{average precision} = \frac{1}{J} \sum_j \frac{\text{TP}_j}{\text{TP}_j + \text{FP}_j}$$

$$\text{average FDR} = 1 - \text{average precision} \quad (10)$$

To compute the *average spike precision rate*, we assign true positives (TP) to the computed cumulative spike MIRA hits and false positives (FP) to cumulative non-spike MIRA hits at TCR index J . Similarly, for the *average non-spike precision rate* we assign TP and FP to the cumulative non-spike and spike MIRA hits, respectively. Table 3 reports the precision and FDR for TCRs ranked according to spike and non-spike CATE, averaged over $J = 10, \dots, 100$. We can interpret the table as a confusion matrix for spike and non-spike hits: we want high values on the diagonal (precision), and low values in off-diagonal elements (FDR). Indeed, we report a high precision of 0.95 and 0.91 for spike and non-spike predictions, respectively. For instance, this implies that 95% of our spike sequence assignments are correct. Conversely, we report low FDRs for both spike and non-spike predictions. Finally, as expected we report low spike and non-spike FDRs for CMV TCRs.

Table 3: **Percentage of each type of TCR among top-TCRs ranked by AIRIVA.** We report fraction of Spike MIRA hits, Non-spike MIRA hits, and CMV Enhanced Sequences among all annotated TCRs present in the top J as ranked by Spike CATE and Non-spike CATE, averaging over J between 10 and 100.

	Spike hits	Non-spike hits	CMV ES
Spike CATE	0.95	0.05	0.0
Non-spike CATE	0.09	0.91	0.06




Figure 8: **Identification of Spike and Non-spike specific TCRs with CATE.** We show the proportion of spike and non-spike MIRA TCRs recovered when traversing spike and non-spike CATEs in descending order (only TCRs with MIRA annotations are considered). The x-axis for both plots corresponds to the index where all 291 spike and 462 non-spike MIRA hits are recovered (the maximum value of the y-axes).

4.3. HSV Case Study

HSV-1 and HSV-2 are two members of the Herpes Simplex Virus (HSV) family. The two viruses share approximately 80% of their genome (Greninger et al., 2018), and will therefore have partially overlapping antigen sets and corresponding TCRs, making it challenging to build a model which accurately disentangles their labels. Infections for both are frequent in the population, with $\approx 67\%$ prevalence for HSV-1 and $\approx 13\%$ prevalence for HSV-2 (James et al., 2020). Infection status for either virus is independent of the other and infections frequently co-occur.

Data Cohorts Our dataset consists of samples labeled for both HSV-1 and HSV-2 using multiplexed immunoglobulin G (IgG) serology. The dataset contains 900 repertoires in the training set, 225 in the validation set, and 289 in the holdout set. We are unable to confidently assign a label for HSV-1 and/or HSV-2 in roughly 20% of our samples. We include these unlabeled samples for AIRIVA in a semi-supervised context (SS-AIRIVA). Here we build an AIRIVA model to disentangle the two labels. We posit the set of labels $\mathbf{Y} = \{Y_{\text{hsv1}}, Y_{\text{hsv2}}\}$.

TCR Selection We ran a sweep of 100 ESLG models for each prediction task (HSV-1 and HSV-2) with a p -value threshold for FET, where $p \in [0.0001, 0.005, 0.001, 0.05]$ and

selected p that maximizes ESLG predictive performance in the validation set. Based on this criteria, we selected a p -value threshold of 0.001 for the FET-based TCR selection step, yielding a total of 158 input TCRs.

HSV RESULTS

Disease Classification Table 4 compares ESLG models trained on HSV-1 and HSV-2 independently against AIRIVA models trained to predict HSV-1 and HSV-2 jointly, either on the labeled repertoires only (AIRIVA), or on all the samples including unlabelled repertoires (SS-AIRIVA). We also compare against FFN, which predicts both HSV-1 and HSV-2 labels jointly. Figure 9 reports the corresponding ROC curves. ESLG models show reasonable performance for both HSV-1 and HSV-2 prediction tasks in the overall population (0.62 and 0.75 AUROC respectively), but upon closer inspection, these models do not separate the two labels well. In particular, the HSV-1 model performance drops to AUROC of 0.5 for samples that are HSV-2 positive. In comparison, AIRIVA overall HSV-1 performance is comparable to performance on HSV-2 positives (0.67 AUROC), indicating that the model has learned to disentangle the two labels. In all cases, AIRIVA outperforms the ESLG baseline for HSV-1 both in terms of sensitivity at 98% specificity and AUROC.

Learning from Unlabelled Repertoires As AIRIVA is a generative model, we expect that including unlabeled data should improve the model’s ability to construct a factorized latent space. Roughly 20% of our HSV data is missing at least one of the HSV-1 or HSV-2 labels, see Appendix for details. As expected, including these samples in SS-AIRIVA improves model performance, particularly for the HSV-1 prediction task.

Table 4: **Comparison of HSV disease models:** Sensitivity at 98% specificity and AUROC, overall and stratified by the performance in the presence of the other subtype, using 100 bootstrap samples. AIRIVA trains with only labeled repertoires. SS-AIRIVA refers to the semi-supervised formulation of AIRIVA, training with additional unlabelled repertoires.

(a) HSV-1 Prediction Task

HSV-1 Model	Overall		HSV-2 negative		HSV-2 positive	
	Sensitivity	AUROC	Sensitivity	AUROC	Sensitivity	AUROC
ESLG	0.12 ± 0.10	0.62 ± 0.09	0.18 ± 0.15	0.63 ± 0.12	0.14 ± 0.17	0.50 ± 0.19
FFN	0.35 ± 0.13	0.80 ± 0.05	0.45 ± 0.17	0.82 ± 0.05	0.30 ± 0.20	0.72 ± 0.14
AIRIVA	0.30 ± 0.12	0.74 ± 0.09	0.35 ± 0.20	0.74 ± 0.10	0.32 ± 0.22	0.67 ± 0.16
SS-AIRIVA	0.45 ± 0.18	0.79 ± 0.07	0.49 ± 0.18	0.78 ± 0.09	0.53 ± 0.22	0.81 ± 0.13

(b) HSV-2 Prediction Task

HSV-2 Model	Overall		HSV-1 negative		HSV-1 positive	
	Sensitivity	AUROC	Sensitivity	AUROC	Sensitivity	AUROC
ESLG	0.11 ± 0.10	0.75 ± 0.07	0.16 ± 0.21	0.79 ± 0.16	0.12 ± 0.15	0.75 ± 0.10
FFN	0.38 ± 0.20	0.82 ± 0.08	0.66 ± 0.29	0.88 ± 0.13	0.30 ± 0.21	0.80 ± 0.09
AIRIVA	0.37 ± 0.18	0.78 ± 0.10	0.57 ± 0.26	0.86 ± 0.12	0.32 ± 0.20	0.77 ± 0.10
SS-AIRIVA	0.38 ± 0.19	0.84 ± 0.06	0.53 ± 0.32	0.88 ± 0.10	0.37 ± 0.25	0.84 ± 0.07

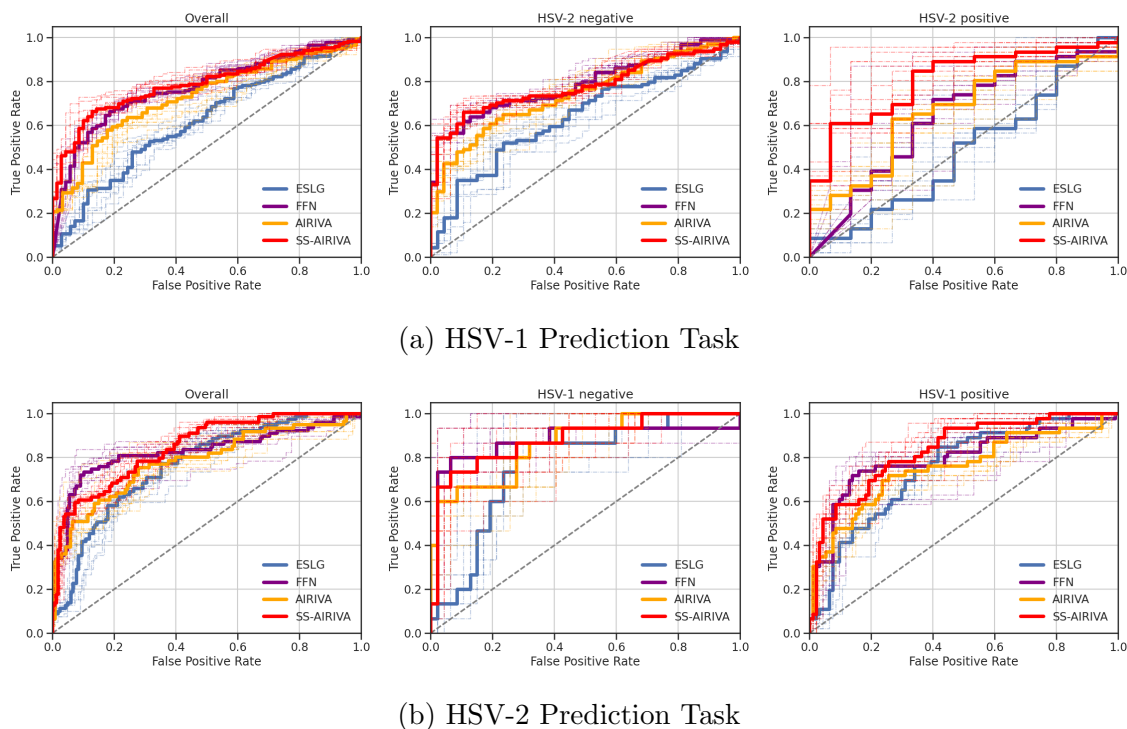


Figure 9: **ROC curves for HSV disease prediction models.** Bold lines correspond to the ROC curve on the test set; colored dashed lines correspond to 10 additional bootstrap samples. First row: HSV-1 Prediction Task; Second row: HSV-2 Prediction Task. SS-AIRIVA exhibits the best performance overall and across different subgroups. While ESLG performs poorly at the HSV-1 prediction task for the HSV2-positive subgroup, AIRIVA maintains the same predictive performance overall.

5. Discussion

A trained AIRIVA model is an effective **simulator of TCR repertoires**, providing manipulable “switches” and “knobs” that control factors of the data-generating process, such as disease and sequencing depth. Crucially, these controls act *independently* and yield counterfactual repertoires that are *consistent*. For instance, these properties are satisfied when intervening on the disease state only affects disease-associated TCRs and not TCR repertoire depth. Conversely, increasing TCR repertoire depth should affect the prevalence of all TCRs without affecting disease-associated TCRs. In our case studies, we have shown that we can effectively intervene and generate counterfactual repertoires that add or remove spike/non-spike TCRs, or add HSV-1 TCRs to HSV-2 positive repertoires, without significant changes to predictions of the non-intervened labels. Moreover, we leverage the generated counterfactuals to consistently recover label-specific TCRs when validated with an external dataset of TCR annotations.

The ability of AIRIVA to generate consistent counterfactuals is just one of several benefits of learning a **disentangled** latent representation, each with important potential impact in immunology. With **interpretable factors**, latent-space projections of repertoires

could explain diagnoses (predictions), monitor disease factors over time, or pre- and post-treatment. Our results demonstrate how disentanglement leads to robust predictive performance. AIRIVA’s ability to distinguish natural infection from vaccinated-only repertoires in COVID-19 is one example of the real-world impact that interpretable models can provide. The ability to distinguish between a large panel of, say, infectious diseases (having multiple and also overlapping sets of TCRs) could have significant clinical impact. Moving from repertoire-level to TCR-level effects, disentangled latent factors can be used to assign TCR-label associations using counterfactuals and CATE computations. Our experiments have validated spike and non-spike binding associations using external data. By generating CATE on reasonable interventions (see Limitations), AIRIVA can provide impact on both scientific discovery of disease-related TCRs and identification of candidates for further investigation for cellular therapy.

While the results of our case studies demonstrate the value of learning disentangled representations of immunomics data, their limited scope (two labels) restricts their immediate impact. However, since we have validated disentanglement on two few-label scenarios with overlapping label settings (requiring disentanglement), and due to AIRIVA’s improvement with unlabelled repertoires, we believe that AIRIVA can scale to problems involving a large number of labels and diverse cohorts. A natural extension of this work would be to expand labels to include a) a comprehensive infectious disease panel, b) HLA typing, and c) sequencing batch IDs. An AIRIVA model configured with these parameters could help identify HLA-disease associated TCRs (thus increasing the precision of TCR annotations), as well as batch effects (by predicting TCRs useful for determining batch IDs).

Clinical Significance Although current standard of care already provides accurate diagnostics for many common infectious diseases such as COVID and HSV, we expect AIRIVA and its extensions to be applicable to more heterogeneous disease contexts, including autoimmune disorders and cancer. Demonstrating the use of AIRIVA in the well-understood space of infectious diseases enables us to better validate and build trust in the methodology, with the goal of extending to other more clinically meaningful use cases in future work.

Limitations While we believe AIRIVA is a useful framework for studying TCR repertoire data, we note the following limitations and leave suggested improvements for future work:

- CATE estimation might not be identifiable given more complex immunomics datasets, *i.e.*, HLA-mediated antigen presentation in autoimmune diseases, where HLA status causes exposure to disease-causing antigens.
- Understanding TCR response requires us to model a complex three-body problem: TCR-antigen-HLA. However, for simplicity we avoided modeling HLAs; this could be achieved with AIRIVA by extending the set of labels by HLA status.
- Disentanglement is weakly enforced via two inductive biases: the factorized conditional prior and classifier-guided loss, leading to challenges in training AIRIVA’s complex objective function. Supported by disentanglement literature (Shu et al., 2019) and our own experience using counterfactual consistency metrics for model selection, disentanglement could be improved by directly incorporating consistency of counterfactuals into the objective function.

- Generating counterfactuals for label combinations that have never been observed in training remains an open problem, which could potentially be addressed by imposing further inductive biases for the decoder to extrapolate to out-of-distribution regions.
- Other limitations and weaknesses are: our Poisson likelihood is not the best choice for zero-inflated TCR counts, AIRIVA works best with a limited number of TCRs; scaling AIRIVA to handle higher dimensional inputs as well as more labels, potentially of lower-quality/with high rates of missingness is an important avenue for future work.

Conclusion AIRIVA provides a powerful tool for analysing TCR repertoires, for both diagnosis and scientific discovery. The inferred latent factors are human-interpretable and can be manipulated to simulate in-silico repertoires and counterfactuals, enabling robust disease diagnosis and identification of disease-specific TCRs. To our knowledge, this is the first attempt leveraging deep generative models to build a simulator of T-cell Receptor Immune-repertoires and generating in-silico repertoires via conditional and counterfactual generation. In a similar way that fields like natural language processing and image analysis have been the focus of a large part of the community, we believe that the results of this paper constitute a first step towards the use of advanced machine learning methods in immunomics and will bring awareness about this domain to the ML community.

Acknowledgments

We would like to thank all our collaborators within the Antigen Map Project across Adaptive Biotechnologies and Microsoft Research.

References

- S. E. Baranzini. Revealing the genetic basis of multiple sclerosis: are we there yet? *Curr Opin Genet Dev*, 21(3):317–324, Jun 2011.
- Michel Besserve, Arash Mehrjou, Rémy Sun, and Bernhard Schölkopf. Counterfactuals uncover the modular structure of deep generative models. *arXiv preprint arXiv:1812.03253*, 2018.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew M Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. *arXiv*, 2015.
- Gal Cafri, Jared J Gartner, Tal Zaks, Kristen Hopson, Noam Levin, Biman C Paria, Maria R Parkhurst, Rami Yossef, Frank J Lowery, Mohammad S Jafferji, et al. mrna vaccine-induced neoantigen-specific t cell immunity in patients with gastrointestinal cancer. *The Journal of clinical investigation*, 130(11):5976–5988, 2020.
- Agisilaos Chartsias, Thomas Joyce, Giorgos Papanastasiou, Scott Semple, Michelle Williams, David E Newby, Rohan Dharmakumar, and Sotirios A Tsaftaris. Disentangled representation learning in cardiac image analysis. *Medical image analysis*, 58:101535, 2019.
- Calliope A. Dendrou, Jan Petersen, Jamie Rossjohn, and Lars Fugger. Hla variation and disease. *Nature Reviews Immunology*, 18:325–339, 2018. doi: 10.1038/nri.2017.143.

- Ryan O Emerson, William S DeWitt, Marissa Vignali, Jenna Gravley, Joyce K Hu, Edward J Osborne, Cindy Desmarais, Mark Klinger, Christopher S Carlson, John A Hansen, Mark Rieder, and Harlan S Robins. Immunosequencing identifies signatures of cytomegalovirus exposure history and HLA-mediated effects on the T cell repertoire. *Nature Genetics*, 49(5):659–665, May 2017a. ISSN 1061-4036, 1546-1718. doi: 10.1038/ng.3822. URL <http://www.nature.com/articles/ng.3822>.
- Ryan O Emerson, William S DeWitt, Marissa Vignali, Jenna Gravley, Joyce K Hu, Edward J Osborne, Cindy Desmarais, Mark Klinger, Christopher S Carlson, John A Hansen, et al. Immunosequencing identifies signatures of cytomegalovirus exposure history and hla-mediated effects on the t cell repertoire. *Nature genetics*, 49(5):659–665, 2017b.
- Madelyn Glymour, Judea Pearl, and Nicholas P Jewell. *Causal inference in statistics: A primer*. John Wiley & Sons, 2016.
- Victor Greiff, Gur Yaari, and Lindsay G. Cowell. Mining adaptive immune receptor repertoires for biological and clinical information using machine learning. *Current Opinion in Systems Biology*, 24:109–119, 2020. ISSN 2452-3100. doi: <https://doi.org/10.1016/j.coisb.2020.10.010>. URL <https://www.sciencedirect.com/science/article/pii/S2452310020300524>. Systems immunology & host-pathogen interaction (2020).
- Julia Greissl, Mitch Pesesky, Sudeb C Dalai, Alison W Rebman, Mark J Soloski, Elizabeth J Horn, Jennifer N Dines, Rachel M Gittelman, Thomas M Snyder, Ryan O Emerson, et al. Immunosequencing of the t-cell receptor repertoire reveals signatures specific for diagnosis and characterization of early lyme disease. *medRxiv*, 2021.
- Alexander L Greninger, Pavitra Roychoudhury, Hong Xie, Amanda Casto, Anne Cent, Gregory Pepper, David M Koelle, Meei-Li Huang, Anna Wald, Christine Johnston, et al. Ultrasensitive capture of human herpes simplex virus genomes directly from clinical samples reveals extraordinarily limited evolution in cell culture. *MSphere*, 3(3):e00283–18, 2018.
- A Grifoni, D Weiskopf, SI Ramirez, J Mateus, JM Dan, CR Moderbacher, SA Rawlings, A Sutherland, L Premkumar, RS Jadi, D Marrama, AM de Silva, A Frazier, AF Carlin, JA Greenbaum, B Peters, F Krammer, DM Smith, S Crotty, and A. Sette. Targets of t cell responses to sars-cov-2 coronavirus in humans with covid-19 disease and unexposed individuals. *Cell*, 181(7):1489–1501, 2020.
- Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. In *ICLR*, 2017.
- Zeshan Hussain, Rahul G Krishnan, and David Sontag. Neural pharmacodynamic state space modeling. In *International Conference on Machine Learning*, 2021.
- Maximilian Ilse, Jakub M. Tomczak, Christos Louizos, and Max Welling. Diva: Domain invariant variational autoencoders. In *Proceedings of the Third Conference on Medical Imaging with Deep Learning*, volume 121 of *Proceedings of Machine Learning Research*, pages 322–348, 2020.

- Cody B. Jackson, Michael Farzan, Bing Chen, and Hyeryun Choe. Mechanisms of sars-cov-2 entry into cells. *Nature Reviews Molecular Cell Biology*, 23:3–20, 2022. doi: 10.1038/s41580-021-00418-x.
- Charlotte James, Manale Harfouche, Nicky J Welton, Katherine ME Turner, Laith J Abu-Raddad, Sami L Gottlieb, and Katharine J Looker. Herpes simplex virus: global infection prevalence and incidence estimates, 2016. *Bulletin of the World Health Organization*, 98(5):315, 2020.
- Zhuxi Jiang, Yin Zheng, Huachun Tan, Bangsheng Tang, and Hanning Zhou. Variational deep embedding: An unsupervised and generative approach to clustering. *arXiv preprint arXiv:1611.05148*, 2016.
- Tom Joy, Sebastian M Schmon, Philip HS Torr, N Siddharth, and Tom Rainforth. Capturing label characteristics in vaes. In *ICLR*, 2021.
- Chakravarthi Kanduri, Milena Pavlović, Lonneke Scheffer, Keshav Motwani, Maria Chernigovskaya, Victor Greiff, and Geir K Sandve. Profiling the baseline performance and limits of machine learning models for adaptive immune receptor repertoire classification. *GigaScience*, 11, 2022.
- Elliott Kieff, Bill Hoyer, Steven Bachenheimer, and Bernard Roizman. Genetic relatedness of type 1 and type 2 herpes simplex viruses. *Journal of Virology*, 9(5):738–745, May 1972. doi: 10.1128/jvi.9.5.738-745.1972. URL <https://doi.org/10.1128/jvi.9.5.738-745.1972>.
- Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In *ICML*, 2018.
- Diederik P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *2th International Conference on Learning Representations, ICLR 2013*, 2013.
- Mark Klinger, Francois Pepin, Jen Wilkins, Thomas Asbury, Tobias Wittkop, Jianbiao Zheng, Martin Moorhead, and Malek Faham. Multiplex identification of antigen-specific t cell receptors using a combination of immune assays and immune receptor sequencing. *PLoS One*, 10(10):e0141561, 2015.
- N. L. La Gruta, S. Gras, S. R. Daley, P. G. Thomas, and J. Rossjohn. Understanding the drivers of MHC restriction of T cell receptors. *Nat Rev Immunol*, 18(7):467–478, Jul 2018.
- Dalin Li, Alexander Xu, Emebet Mengesha, Rebecca Elyanow, Rachel M Gittelman, Heidi Chapman, John C Prostko, Edwin C Frias, James L Stewart, Valeriya Pozdnyakova, Philip Debbas, Angela Mujukian, Arash A Horizon, Noah Merin, Sandy Joung, Gregory J Botwin, Kimia Sobhani, Jane C Figueiredo, Susan Cheng, Ian M Kaplan, Dermot P B McGovern, Akil Merchant, Gil Y Melmed, and Jonathan Braun. The t-cell response to SARS-CoV-2 vaccination in inflammatory bowel disease is augmented with anti-TNF therapy. *Inflammatory Bowel Diseases*, 28(7):1130–1133, April 2022. doi: 10.1093/ibd/izac071. URL <https://doi.org/10.1093/ibd/izac071>.

- Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Raetsch, Sylvain Gelly, Bernhard Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learning of disentangled representations. In *international conference on machine learning*, pages 4114–4124. PMLR, 2019.
- Ulrika Marking, Sebastian Havervall, Nina Greilert-Norin, Henry Ng, Kim Blom, Peter Nilsson, Mia Phillipson, Sophia Hober, Charlotta Nilsson, Sara Mangsbo, Wanda Christ, Jonas Klingström, Max Gordon, Mikael Åberg, and Charlotte Thålin. Duration of SARS-CoV-2 immune responses up to six months following homologous or heterologous primary immunization with ChAdOx1 nCoV-19 and BNT162b2 mRNA vaccines. *Vaccines*, 10(3): 359, February 2022. doi: 10.3390/vaccines10030359. URL <https://doi.org/10.3390/vaccines10030359>.
- Catherine Mazza and Bernard Malissen. What guides mhc-restricted tcr recognition? *Seminars in Immunology*, 19(4):225–35, 2007. doi: 10.1016/j.smim.2007.03.003.
- Kevin R. McCarthy, Linda J. Rennick, Sham Nambulli, Lindsey R. Robinson-McCarthy, William G. Bain, Ghady Haidar, and W. Paul Duprex. Recurrent deletions in the sars-cov-2 spike glycoprotein drive antibody escape. *Science*, 371(6534):1139–1142, 2021. doi: 10.1126/science.abf6950. URL <https://www.science.org/doi/abs/10.1126/science.abf6950>.
- Riccardo Miotto, Fei Wang, Shuang Wang, Xiaoqian Jiang, and Joel T Dudley. Deep learning for healthcare: review, opportunities and challenges. *Briefings in bioinformatics*, 19(6):1236–1246, 2018.
- Kenneth Murphy and Casey Weaver. *Janeway’s immunobiology*. Garland science, 2016.
- Sean Nolan, Marissa Vignali, Mark Klinger, Jennifer N Dines, Ian M Kaplan, Emily Svejnoha, Tracy Craft, Katie Boland, Mitch Pesesky, Rachel M Gittelman, et al. A large-scale database of t-cell receptor beta (tcr β) sequences and binding associations from natural and synthetic exposure to sars-cov-2. *Research square*, 2020.
- Ozan Oktay, Jay Nanavati, Anton Schwaighofer, David Carter, Melissa Bristow, Ryutaro Tanno, Rajesh Jena, Gill Barnett, David Noble, Yvonne Rimmer, et al. Evaluation of deep learning to augment image-guided radiotherapy for head and neck and prostate cancers. *JAMA network open*, 3(11):e2027426–e2027426, 2020.
- Milena Pavlovic, Ghadi S Al Hajj, Victor Greiff, Johan Pensar, and Geir Kjetil Sandve. Using causal modeling to analyze generalization of biomarkers in high-dimensional domains: a case study of adaptive immune repertoires. In *ICML 2022: Workshop on Spurious Correlations, Invariance and Stability*, 2022.
- Margarita Pertseva, Beichen Gao, Daniel Neumeier, Alexander Yermanos, and Sai T. Reddy. Applications of machine and deep learning in adaptive immunity. *Annual Review of Chemical and Biomolecular Engineering*, 12(1):39–62, 2021.
- Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. *Elements of causal inference: foundations and learning algorithms*. The MIT Press, 2017.

- Melanie F Pradier, Bernhard Reis, Lori Jukofsky, Francesca Milletti, Toshihiko Ohtomo, Fernando Perez-Cruz, and Oscar Puig. Case-control indian buffet process identifies biomarkers of response to codrituzumab. *BMC cancer*, 19:1–7, 2019.
- Maria Predecki, Tina Thomson, Candice L. Clarke, Paul Martin, Sarah Gleeson, Rute Cardoso De Aguiar, Helena Edwards, Paige Mortimer, Stacey McIntyre, Shanice Lewis, Jaid Deborah, Alison Cox, Graham Pickard, Liz Lightstone, David Thomas, Stephen P. McAdoo, Peter Kelleher, and Michelle Willicombe and. Comparison of humoral and cellular responses in kidney transplant recipients receiving BNT162b2 and ChAdOx1 SARS-CoV-2 vaccines. July 2021. doi: 10.1101/2021.07.09.21260192. URL <https://doi.org/10.1101/2021.07.09.21260192>.
- Pranav Rajpurkar, Emma Chen, Oishi Banerjee, and Eric J. Topol. Ai in health and medicine. *Nature medicine*, 28:31–38, 2022.
- Birkir Reynisson, Bruno Alvarez, Sinu Paul, Bjoern Peters, and Morten Nielsen. NetMHCpan-4.1 and NetMHCIIpan-4.0: improved predictions of MHC antigen presentation by concurrent motif deconvolution and integration of MS MHC eluted ligand data. *Nucleic Acids Research*, 48(W1):W449–W454, 05 2020. ISSN 0305-1048. doi: 10.1093/nar/gkaa379. URL <https://doi.org/10.1093/nar/gkaa379>.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *ICML*, 2014.
- Harlan Robins, Cindy Desmarais, Jessica Matthis, Robert Livingston, Jessica Andriesen, Helena Reijonen, Christopher Carlson, Gerold Nepon, Cassian Yee, and Karen Cerosaletti. Ultra-sensitive detection of rare t cell clones. *Journal of immunological methods*, 375(0):14–19, 2012.
- Jihane Romanos, Cleo C Van Diemen, Ilja M Nolte, Gosia Trynka, Alexandra Zhernakova, Jingyuan Fu, Maria Teresa Bardella, Donatella Barisani, Ross McManus, David A Van Heel, et al. Analysis of hla and non-hla alleles can identify individuals at high risk for celiac disease. *Gastroenterology*, 137(3):834–840, 2009.
- Donald B Rubin. Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, 100(469):322–331, 2005.
- E Sandeep Kumar and Pappu Satya Jayadev. Deep learning for clinical decision support systems: a review from the panorama of smart healthcare. *Deep learning techniques for biomedical and health informatics*, pages 79–99, 2020.
- Tina Schmidt, Verena Klemis, David Schub, Janine Mihm, Franziska Hielscher, Stefanie Marx, Amina Abu-Omar, Laura Ziegler, Candida Guckelmus, Rebecca Urschel, Sophie Schneitler, Sören L. Becker, Barbara C. Gärtner, Urban Sester, and Martina Sester. Immunogenicity and reactogenicity of heterologous ChAdOx1 nCoV-19/mRNA vaccination. *Nature Medicine*, 27(9):1530–1535, July 2021. doi: 10.1038/s41591-021-01464-w. URL <https://doi.org/10.1038/s41591-021-01464-w>.


- Rui Shu, Yining Chen, Abhishek Kumar, Stefano Ermon, and Ben Poole. Weakly supervised disentanglement with guarantees. *arXiv preprint arXiv:1910.09772*, 2019.
- John-William Sidhom, H Benjamin Larman, Drew M Pardoll, and Alexander S Baras. DeepTCR is a deep learning framework for revealing sequence concepts within t-cell repertoires. *Nature communications*, 12(1):1–12, 2021.
- Thomas M Snyder, Rachel M Gittelman, Mark Klinger, Damon H May, Edward J Osborne, Ruth Taniguchi, H Jabran Zahid, Ian M Kaplan, Jennifer N Dines, Matthew N Noakes, et al. Magnitude and dynamics of the t-cell response to sars-cov-2 infection at both individual and population levels. *MedRxiv*, 2020.
- Raphael Suter, Djordje Miladinovic, Bernhard Schölkopf, and Stefan Bauer. Robustly disentangled causal mechanisms: Validating deep representations for interventional robustness. In *International Conference on Machine Learning*, pages 6056–6065. PMLR, 2019.
- S. J. Swamidass, C.-A. Azencott, K. Daily, and P. Baldi. A CROC stronger than ROC: measuring, visualizing and optimizing early retrieval. *Bioinformatics*, 26(10):1348–1356, April 2010. doi: 10.1093/bioinformatics/btq140. URL <https://doi.org/10.1093/bioinformatics/btq140>.
- Michael Widrich, Bernhard Schäfl, Milena Pavlovic, Hubert Ramsauer, Lukas Gruber, Markus Holzleitner, Johannes Brandstetter, Geir Kjetil Sandve, Victor Greiff, Sepp Hochreiter, et al. Modern hopfield networks and attention for immune repertoire classification. In *NeurIPS*, 2020.

Appendix A. Links to Causality

A.1. Disentangled Causal Process.

AIRIVA assumes that data comes from a disentangled causal process, a specific family of causal graphs where the true generative factors of variation do not cause each other, as shown in Figure 10. Let \mathbf{X} be the outcome (TCRs) and $\mathbf{Y} = \{Y^1, \dots, Y^K\}$ denote the K generative factors of variation underlying the true data generation process. We assume there exists some confounding factors C^1, \dots, C^R on top of \mathbf{Y} , which results in statistical dependencies between factors Y^i and Y^j . The true factors of variation \mathbf{Y} are *elementary* in the sense that they can be changed without affecting others (there is no direct causal effect between them).

Figure 10: **Assumed causal graph: Disentangled Causal Process:** The observed TCR counts repertoire \mathbf{X} results from the combination of factors \mathbf{Y} , i.e., exposures to certain antigens, age, genetic factors, etc. which may be correlated by unobserved confounders \mathbf{C} , but do not cause one another.



Formally, a causal model for \mathbf{X} is called *disentangled* (Suter et al., 2019) if and only if it can be described by the causal graph of Figure 10 and a structural causal model (SCM) of the form:

$$\begin{aligned}
 C &\leftarrow N_c \\
 Y^k &\leftarrow f_k(\mathbf{PA}_k^C, N_k), \mathbf{PA}_k^C \in \{C^1, \dots, C^R\}, k = 1, \dots, K \\
 \mathbf{X} &\leftarrow g(\mathbf{Y}, N_x),
 \end{aligned} \tag{11}$$

where $N_x, N_c, N_1, \dots, N_K$ are independent noise variables.

A.2. Counterfactual Generation

A *counterfactual* is a hypothetical observation $\mathbf{X} = \mathbf{x}^*$ given a hypothetical treatment $Y^j = y_\Delta^j$ for a specific datapoint $(\mathbf{x}, y^j, \mathbf{y}^{\setminus j})$, denoted as:

$$\mathbf{x}^* \sim P(\mathbf{X}(y_\Delta^j) \mid \mathbf{X} = \mathbf{x}, Y^j = y^j, \mathbf{Y}^{\setminus j} = \mathbf{y}^{\setminus j}). \tag{12}$$

Let \mathbf{U} refer to the set of observed and unobserved random variables in the SCM (endogenous and exogenous variables) excluding the treatment Y^j and outcome \mathbf{X} . If we had access to the true underlying SCM in Eq. 11, we could compute valid probabilistic counterfactuals following three steps (Glymour et al., 2016; Peters et al., 2017):

1. **Abduction:** Use an observation (\mathbf{x}, \mathbf{y}) to update the distribution of the observation-specific noise variables \mathbf{U} .
2. **Action:** Make an intervention in the SCM, i.e., replacing the structural equation for the treatment Y^j with $Y^j := y_{\Delta}^j$.
3. **Prediction:** Use the updated distribution of \mathbf{U} from step 1 and the modified SCM from step 2 to compute the distribution of $p(\mathbf{X}|do(Y^j), \mathbf{U})$. We can then sample a counterfactual $\mathbf{x}^* \sim p(\mathbf{X}|do(Y^j), \mathbf{U})$.

In practice, we do not know the true underlying SCM, which is why we turn to a model like AIRIVA to estimate the true generative process for \mathbf{X} . Shu et al. (2019) shows that performing distribution matching in the augmented space of observations \mathbf{X} and labels \mathbf{Y} provides disentanglement guarantees on the learned representations. Thus, in the limits of infinite data, maximizing the regularized ELBO in AIRIVA yields the recovery of a disentangled causal process.

We can generate counterfactuals using AIRIVA following 3 steps: i) infer the latent factors from an observation; ii) perform an intervention where we turn on/off one of the generative factors (labels); iii) generate a counterfactual. Each of these steps directly corresponds to the steps described above, namely:

- **Abduction:** Given an observation \mathbf{x} , sample its latent representation $\mathbf{z} \sim p(\mathbf{Z}|\mathbf{X})$ which captures observation-specific characteristics and noise.
- **Action:** Replace the repertoire label value y^j by an alternative label y_{Δ}^j and sample its corresponding latent factor from the conditional prior, i.e., $\mathbf{z}_{\Delta}^j \sim p(\mathbf{Z}|do(Y^j \leftarrow y_{\Delta}^j))$.
- **Prediction:** Combine the latent representation from step 1 and the modified latent factor from step 2 to compute the value of $\mathbf{x}^* \sim p(\mathbf{X}|\mathbf{z}_{\Delta})$, where $\mathbf{z}^{\Delta} = [\mathbf{z}^{\setminus j}, \mathbf{z}_{\Delta}^j]$ is the intervened latent representation.⁶

A.3. Validity of Counterfactuals

As a generative model, AIRIVA is trained to infer the joint data generation process of observed factual data $p(\mathbf{X}, \mathbf{Y})$. Hence, conditional factual data generation of $\mathbf{X}(Y^k)$ is enforced by maximizing the ELBO in Eq. (5). By including the label-specific auxiliary objective in Eq. (6), we encourage a disentangled latent space \mathbf{Z} , where the label-specific posterior $q(\mathbf{Z}_{\mathbf{Y}^k}|\mathbf{X})$ is independent to all other latent factors by construction. Therefore, we can conclude that the counterfactuals $\mathbf{X}(1 - Y^k)$ are valid under the following conditions: (i) the $KL(q(\mathbf{Z}_{\mathbf{Y}^k}|\mathbf{X})||p(\mathbf{Z}_{\mathbf{Y}^k}|\mathbf{X}))$ divergence between the label-specific posterior and prior is close to zero; (ii) the underlying causal model follows Figure 10 in (Suter et al., 2019), i.e., there are no direct arrows between the labels; (iii) the overlap assumption holds, i.e., $0 < P(\mathbf{Y}|\mathbf{X}) < 1$, almost surely if $P(\mathbf{X}) > 0$.

6. Without loss of generality, here we assume that we are intervening on the last factor to simplify notation.

Appendix B. Derivations for supervised and semi-supervised AIRIVA

B.1. Supervised AIRIVA

The objective function for supervised AIRIVA⁷ looks as follows:

$$\begin{aligned} \text{ELBO}_l(x, y) &= \mathbb{E}_{q_\phi(z|x)}[\log p_\theta(x|z_y, z_x)] - \beta D_{KL}(q_\phi(z_x|x)||p_\theta(z_x)) - \beta D_{KL}(q_\phi(z_y|x)||p_\theta(z_y|y)) \\ \mathcal{L}_{\text{AIRIVA}} &= \sum_{n=1}^N \left(\text{ELBO}_l(x^{(n)}, y^{(n)}) + \alpha \mathbb{E}_{q_\phi(z^{(n)}|x^{(n)})}[\log q_\phi(y^{(n)}|z_y^{(n)})] \right), \end{aligned} \quad (13)$$

where x is a high-dimensional vector of TCR counts, y is a vector of repertoire labels, $z = [z_y; z_x]$ denote the latent representation of a repertoire, and is composed of predictive latents z_y and residual latents z_x , N is the number of available repertoires, and $\text{ELBO}_l(x, y)$ refers to the evidence lower bound for a single labelled repertoire. In the following, we will omit the distribution parameters ϕ and θ to simplify notation.

Derivation of the ELBO for labelled repertoires.

$$\log p(x, y) = \log \int p(x|z)p(z|y)p(y)dz \quad (14)$$

$$\begin{aligned} &= \log \int p(x|z)p(z|y)p(y) \frac{q(z|x)}{q(z|x)} dz \\ &\geq \text{ELBO}_l(x, y). \end{aligned} \quad (15)$$

$$\text{ELBO}_l(x, y) = \mathbb{E}_{q(z|x)}[\log p(x|z) + \log p(z|y) - \log q(z|x)] + \log p(y) \quad (16)$$

$$\begin{aligned} &= \mathbb{E}_{q(z_y|x)q(z_x|x)}[\log p(x|z_y, z_x) + \log p(z_y|y) + \log p(z_x) \\ &\quad - \log q(z_x|x) - \log q(z_y|x)] + \log p(y) \end{aligned} \quad (17)$$

$$\begin{aligned} &= \mathbb{E}_{q(z_y|x)q(z_x|x)}[\log p(x|z_y, z_x)] + \mathbb{E}_{q(z_x|x)}[\log p(z_x) - \log q(z_x|x)] \\ &\quad + \mathbb{E}_{q(z_y|x)}[\log p(z_y|y) - \log q(z_y|x)] + \log p(y) \end{aligned} \quad (18)$$

$$\begin{aligned} &= \mathbb{E}_{q(z_y|x)q(z_x|x)}[\log p(x|z_y, z_x)] - D_{KL}(q(z_x|x)||p(z_x)) \\ &\quad - D_{KL}(q(z_y|x)||p(z_y|y)) + \log p(y), \end{aligned} \quad (19)$$

where:

$$\begin{aligned} \mathbb{E}_{q(z|x)}[\log p(z_y|y) - \log q(z_y|x)] &= \mathbb{E}_{q(z_y|x)q(z_x|x)}[\log p(z_y|y) - \log q(z_y|x)] \\ &= \mathbb{E}_{q(z_y|x)} [\mathbb{E}_{q(z_x|x)} [\log p(z_y|y) - \log q(z_y|x)]] \\ &= \mathbb{E}_{q(z_y|x)} [\log p(z_y|y) - \log q(z_y|x)] \\ &= -D_{KL}(q(z_y|x)||p(z_y|y)). \end{aligned}$$

7. Note that in (Ilse et al., 2020) there was an additional domain variable d to account for observed labels against which we want to be invariant. Since their formulation treats y and d identically, here we have simplified the formulation such that d gets absorbed into a generic vector of labels y .

B.2. Semi-supervised AIRIVA

The ELBO for unlabelled data as in (Ilse et al., 2020) looks as follows:

$$\begin{aligned} \text{ELBO}_u(x) &= \mathbb{E}_{q(z|x)}[\log p(x|z_y, z_x)] - \beta D_{KL}(q(z_x|x)||p(z_x)) \\ &\quad + \beta \mathbb{E}_{q(z_y|x)q(y|z_y)}[\log p(z_y|y) - \log q(z_y|x)] \\ &\quad + \mathbb{E}_{q(z_y|x)q(y|z_y)}[\log p(y) - \log q(y|z_y)]. \end{aligned} \quad (20)$$

$$(21)$$

Derivation of the ELBO for unlabelled repertoires In the following, we derive this expression from scratch. The probability distribution of an unlabelled repertoire $p(x)$ in AIRIVA can be written as:

$$p(x) = \int \int \int p(x|z_y, z_x)p(z_y|y)p(z_x)p(y)dz_xdz_ydy. \quad (22)$$

We assume the following variational approximation:

$$q(z_y, z_d, z_x, y) = q(z_y|x)q(z_x|x)q(y|x). \quad (23)$$

$$\log p(x) = \log \int \int p(x|z)p(z|y)p(y)dzdy \quad (24)$$

$$= \log \int \int p(x|z)p(z|y)p(y) \frac{q(z|x)q(y|x)}{q(z|x)q(y|x)} dzdy \quad (25)$$

$$\geq \mathbb{E}_{q(z|x)q(y|x)}[\log p(x|z) + \log p(z|y) + \log p(y) - \log q(z|x) - \log q(y|x)] \quad (26)$$

$$\geq \text{ELBO}_u(x) \quad (27)$$

$$\begin{aligned} \text{ELBO}_u(x) &= \mathbb{E}_{q(z_y|x)q(z_x|x)q(y|x)}[\log p(x|z_y, z_x) + \log p(z_y|y) + \log p(z_x) + \log p(y) \\ &\quad - \log q(z_y|x) - \log q(z_x|x) - \log q(y|x)] \end{aligned} \quad (28)$$

$$\begin{aligned} &= \mathbb{E}_{q(z_y|x)q(z_x|x)}[\log p(x|z_y, z_x)] + \mathbb{E}_{q(z_x|x)}[\log p(z_x) - \log q(z_x|x)] \\ &\quad + \mathbb{E}_{q(z_y|x)q(y|x)}[\log p(z_y|y) - \log q(z_y|x) + \log p(y) - \log q(y|z_y)] \end{aligned} \quad (29)$$

$$\begin{aligned} &= \mathbb{E}_{q(z_y|x)q(z_x|x)}[\log p(x|z_y, z_x)] - D_{KL}(q(z_x|x)||p(z_x)) \\ &\quad - \mathbb{E}_{q(y|x)}[D_{KL}(q(z_y|x)||p(z_y|y))] - \mathbb{E}_{q(z_y|x)}[D_{KL}(q(y|z_y)||p(y))]. \end{aligned} \quad (30)$$

Objective function for semi-supervised AIRIVA The semi-supervised AIRIVA objective can be written as follows:

$$\begin{aligned} \mathcal{L}_{\text{SS-AIRIVA}} &= \underbrace{\sum_{n=1}^N \text{ELBO}_l(x^{(n)}, y^{(n)})}_{\text{ELBO for labelled data}} + \underbrace{\sum_{m=1}^M \text{ELBO}_u(x^{(m)}, y^{(m)})}_{\text{ELBO for unlabelled data}} + \underbrace{\sum_{i \in \mathcal{I}_y} \alpha_y \mathbb{E}_{q(z_y^{(i)}|x^{(i)})}[\log q_\phi(y^{(i)}|z_y^{(i)})]}_{\text{auxiliary predictive loss}} \end{aligned} \quad (31)$$

where N and M refer to the number of labelled and unlabelled repertoires respectively, and \mathcal{I}_y corresponds to the set of indexes for datapoints that have partially observed labels y .

$$\begin{aligned}
 \mathcal{L}_{\text{SS-AIRIVA}} = & \sum_{i=1}^{N+M} \left(\mathbb{E}_{q(z_y|x^{(i)})q(z_x|x^{(i)})} \left[\log p(x^{(i)}|z_y, z_x) \right] - \mathbb{E}_{q(z_x|x^{(i)})} \left[D_{\text{KL}}(q(z_x|x^{(i)})||p(z_x)) \right] \right) \\
 & - \sum_{n=1}^N D_{\text{KL}}(q(z_y|x^{(n)})||p(z_y|y^{(n)})) \\
 & - \sum_{m=1}^M \mathbb{E}_{q(y|x^{(m)})} \left[D_{\text{KL}}(q(z_y|x^{(m)})||p(z_y|y)) \right] + \sum_{m=1}^M \mathbb{E}_{q(z_y|x^{(m)})} \left[\mathcal{H}[q(y|x^{(m)})] \right] \\
 & + \sum_{i \in \mathcal{I}_y} \alpha_y \mathbb{E}_{q(z_y^{(i)}|x^{(i)})} \left[\log q(y^{(i)}|z_y^{(i)}) \right], \tag{32}
 \end{aligned}$$

where we choose $q(y|x^{(m)}) = q(y|\text{enc}(x^{(m)})) = q(y|z_y^{(m)})$ as in (Jiang et al., 2016), we parametrize $q(y|z_y)$ by a categorical distribution $\text{Cat}(y; \pi(z_y))$ and $\pi(\cdot)$ is a neural network.

B.3. Comparison with semi-supervised CC-VAE (Joy et al., 2021)

Supervised formulation. DIVA or AIRIVA and CC-VAE rely on the same generative process, but the assumed variational approximations are different. In CC-VAE, the authors work with the quantity $q(z|x, y)$ and apply Bayes rule to introduce the parametrized networks $q(y|z_y)$ and $q(z|x)$:

$$q(z|x, y) = \frac{q(y|z_y)q(z|x)}{q(y|x)}.$$

In contrast, the derivation in AIRIVA only considers: $q(y|z_y)q(z|x)$, the key difference is the denominator $q(y|x)$. A drawback from the DIVA derivation is that there is no emergence of a natural classifier in the supervised case, whereas this arises naturally in CC-VAE. That means that the DIVA objective function is not a true lower bound anymore, but rather, an approximation to the marginal likelihood of our data.

Semi-supervised formulation. CC-VAE and AIRIVA exhibit the same formula for the unlabelled case. The objective of both models only differ in the ELBO derivation for the labelled datapoints, in particular:

$$\text{ELBO}_{\text{AIRIVA}}(x, y) = \mathbb{E}_{q(z_y|x)q(z_x|x)} \left[\log \frac{p_\theta(x|z_y, z_x)p(z_y|y)p(z_x)}{q(z_y|x)q(z_x|x)} \right] \tag{33}$$

$$\begin{aligned}
 \text{ELBO}_{\text{CCVAE}}(x, y) &= \mathbb{E}_{q(z|x, y)} \left[\log \frac{p_\theta(x|z_y, z_x)p(z_y|y)p(z_x)}{q(z|x, y)} \right] \\
 &= \mathbb{E}_{q(z_y|x)q(z_x|x)} \left[\frac{q(y|z_y)}{q(y|x)} \log \frac{p_\theta(x|z_y, z_x)p(z_y|y)p(z_x)}{q(y|z_y)q(z_y|x)q(z_x|x)} \right] + \log q(y|x) + \log p(y). \tag{34}
 \end{aligned}$$

Appendix C. Further Experimental Details

- To train AIRIVA, we implemented softmax annealing in the variance of the reparameterization trick, which can be interpreted as sampling from a surrogate deterministic version of the posterior $\tilde{q}(z|x)$ with variance increasing with the number of iterations until it matches the variance of the original posterior $q(z|x)$.
- We restricted the predictive latents associated with each binary label of interest to be scalar variables.
- For interpretability, we restrict the regressors of binary labels to be linear, such that any complex feature transformation occurs in the encoder.
- We performed a hyperparameter sweeps of 1000 runs on a random grid search over a wide range of parameters, including encoder / decoder architectures, dimensionality of the residual latent space, and objective function weights. The detailed specification for the sweep can be found below:

```

alpha_kl_predictive: [1, 10, 100]
alpha_kl_residual: [1,10,100]
alpha_predictive: [1,10,100]
batch_norm: [False, True]
batch_size: [300, 3000]
discriminator_lr: [0.001, 0.01, 0.1]
dropout_prob: [0, 0.4]
hidden_layers_str: ["", "256", "256,64", "256,128,64"]
L1_activations: [0, 1, 10]
L1_embedding_layer: [0, 1, 10]
learning_rate: [0.0005, 0.001, 0.005]
num_z_residuals: [10, 50, 100, 300]
reg_ydecoder: [False, True]
initialization seed: [5, 6, 7, 8, 1]
yencoder_hidden_str: ["", "16,4"]

```

Appendix D. Further Details on the Datasets

D.1. COVID-19 Dataset

The COVID-19 diagnostic classifiers are trained on 1,954 samples from donors with natural SARS-CoV-2 infection, 477 healthy donors post SARS-CoV-2 vaccination, and 5,198 healthy controls sampled prior to December 2020. The classifiers are tested on a holdout set of 525 samples from naturally infected donors, 400 vaccinated donors, 100 donors who were naturally infected and later vaccinated, and 4606 healthy controls. The naturally infected donors all tested positive by RT-PCR, with an average of 13 days from diagnosis to blood draw (Figure 11g). The vaccinated donors were sampled following their final scheduled dose—dose 2 for BNT162b2 (Pfizer) and mRNA-1273 (Moderna) and dose 1 for Ad26.COV2.S (Johnson & Johnson), with an average of 56 days from vaccination to blood draw (Figure 11i).

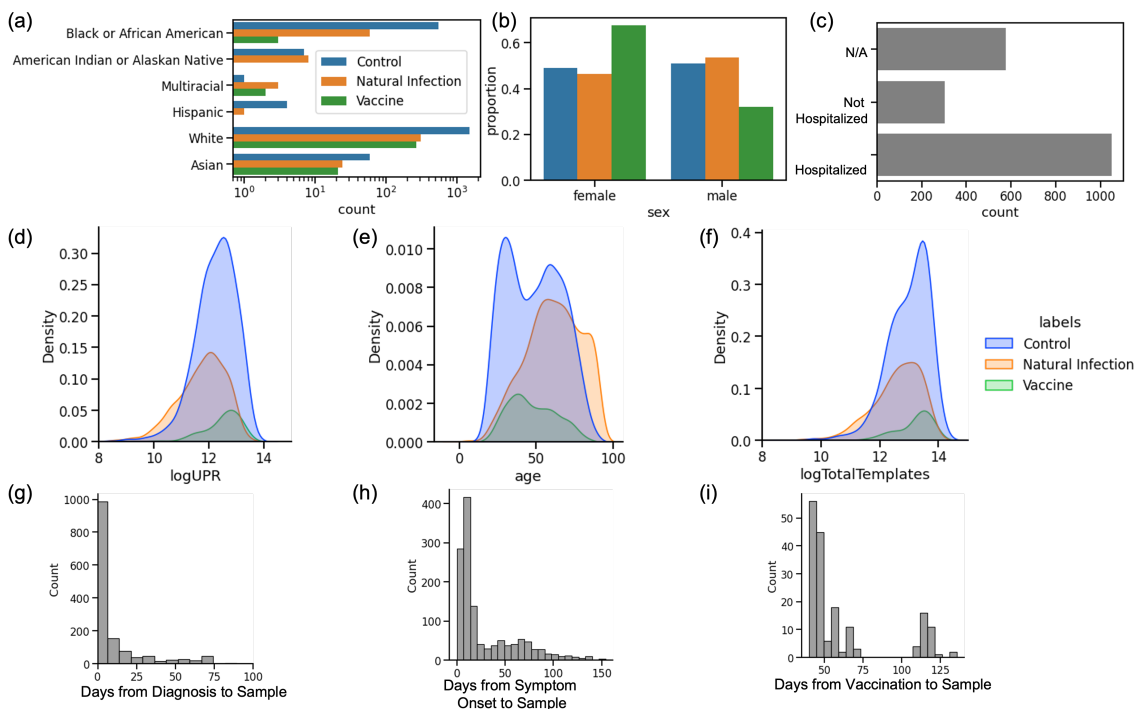


Figure 11: **COVID-19 Cohort Metadata** (a) Race distribution between control, natural infection, and vaccinated cohorts, (b) sex distribution, (c) hospitalization status of natural infection cohort, (d) Log Unique Productive Rearrangement (UPR) distribution, measuring the number of unique T cell clones in a repertoire, (e) age distribution, (f) log total template distribution, measuring the total number of T cell clones in a repertoire, (g) days from diagnosis to blood draw, (h) days from symptom onset to blood draw, (i) days from vaccination to blood draw.

The total sequencing depth measured by number of productive T-cell templates is lowest in the natural infection cohort (mean 380k), with the control and vaccinated cohorts having similar depth (mean 530k and 540) (Figure 11f). The natural infection and control cohorts are split roughly equally between female and male donors (mean 45% and 50% control), while the vaccinated cohort has more samples from female donors (61%) (Figure 11b). Vaccinated and control donors have similar age distributions (mean 47 years) while donors with natural infection skew slightly older (mean 60 years) (Figure 11e). Of the 43% samples with race metadata, the control cohort is the most diverse, with 42% non-white donors, followed by the natural infection cohort with 36% non-white donors, and the vaccine cohort with 22% non-white donors (Figure 11a).

D.2. HSV Dataset

HSV infection labels were created by detecting IgG antibodies against HSV-1 and HSV-2 antigen in the serum of associated TCR repertoires using a multiplexed immunoassay platform (Uplex from Meso Scale Discovery). The gG antigen in HSV-1 and gG in HSV-2 were

Table 5: Description of HSV dataset

	Repertoires	HSV1+	HSV2+	HSV1+ HSV2+	HSV1- HSV2-
Training	1125	645	295	177	208
Holdout	289	169	79	46	47

Table 6: Missing labels in HSV dataset

	Repertoires	HSV-1 missing	HSV-2 missing	Both missing
Training	1125	193	158	27
Holdout	289	50	42	5

chosen as antigen to detect the presence of HSV-1 and/or HSV-2 antibody responses since they share low sequence homology ($\ll 30\%$). The gG protein was linked to its respective spot on U-PLEX plates according to the manufacturer protocol. Antigen specific antibodies in the serum were detected with SULFO-TAG labeled anti-human IgG antibodies using MSD MESO QuickPlex SQ 120 instrument. Each sample was tested in triplicate.

We assigned binary labels to this data set by determining thresholds on the continuous experimental readout by comparing to the background signal level, leaving some samples where we were unable to confidently assign a label. The total number of HSV-1 positives, HSV-2 positives, double positives and unlabeled repertoires are shown in Table 5 and Table 6.

Are datasets and code be publicly available? The majority of the TCRB sequencing data is publicly available through the ImmuneACCESS database, maintained by Adaptive Biotechnologies. Relevant COVID repertoires can be accessed at: <https://clients.adaptivebiotech.com/pub/covid-2020>, <https://clients.adaptivebiotech.com/pub/elyanow-2022-jci>, <https://clients.adaptivebiotech.com/pub/gittelman-2022-jci>. While the code is proprietary, AIRIVA is an adaptation of the publicly available code for DIVA (Ilse et al., 2020) which hopefully facilitates replicability.

Appendix E. Additional Results for COVID

E.1. Spike Prediction Task

Table 7 and Figure 12 show the performance for spike AIRIVA, ESLG and FFN on the spike prediction task both matching the spike training setup (Overall) and for vaccinated samples versus healthy controls (Healthy). All three models are well matched in performance.

E.2. Concentrated ROC curves for COVID Disease Model


Figure 13 shows concentrated ROC curves for the non-spike disease models, matching the same ROC curves reported in the main text.

E.3. Latent space for train samples

Figure 14 shows the latent space for the COVID-19 training data.

Table 7: Comparison of Spike prediction models: Sensitivity at 98% specificity and AU-CROC) overall and for healthy repertoires.


Model	Overall		Healthy	
	Sensitivity	AUROC	Sensitivity	AUROC
ESLG	0.87 ± 0.03	0.95 ± 0.02	0.87 ± 0.05	0.97 ± 0.01
FFN	0.82 ± 0.04	0.93 ± 0.02	0.79 ± 0.06	0.93 ± 0.02
AIRIVA	0.79 ± 0.04	0.94 ± 0.01	0.81 ± 0.06	0.95 ± 0.02



(a) Overall

(b) Healthy

Figure 12: ROC Curves for Spike prediction model.



(a) Overall

(b) Unvaccinated

(c) Vaccinated

Figure 13: **Concentrated ROC Curves (Swamidass et al., 2010) for non-spike disease models.** Bold lines correspond to the CROC curve on the holdout data; colored dashed lines correspond to 10 additional bootstrap samples. ESLG performs poorly for the vaccinated subgroup, while AIRIVA performance remains competitive, with similar predictive performance compared to the overall population.

E.4. Depth Distribution of Conditionally Generated In-silico Repertoires

On the clinical significance of sequencing depth The main driver of differences in sequencing depth is the amount of input DNA when sequencing. We generally do not observe an increase in sequencing depth (total unique rearrangements) as these infections change. Only the composition of TCRs in a repertoire, that is, the clonality of disease associated T cells within the repertoire, changes in response to these exposures. In severely




Figure 14: **Latent space $q(\mathbf{Z}|\mathbf{X})$ for COVID-19 in train set.** AIRIVA learns an interpretable latent space that matches biological intuition. Each point corresponds to one posterior sample. Predictive priors $p(\mathbf{Z}_y|\mathbf{Y})$ for each combination of labels are represented as ellipses, each circle denoting one standard deviation. Purple circles correspond to the prior for the *non-spike only* subgroup, which is never observed in reality. On the sides, we show the marginal posteriors as continuous lines, and predictive priors as dashed lines.


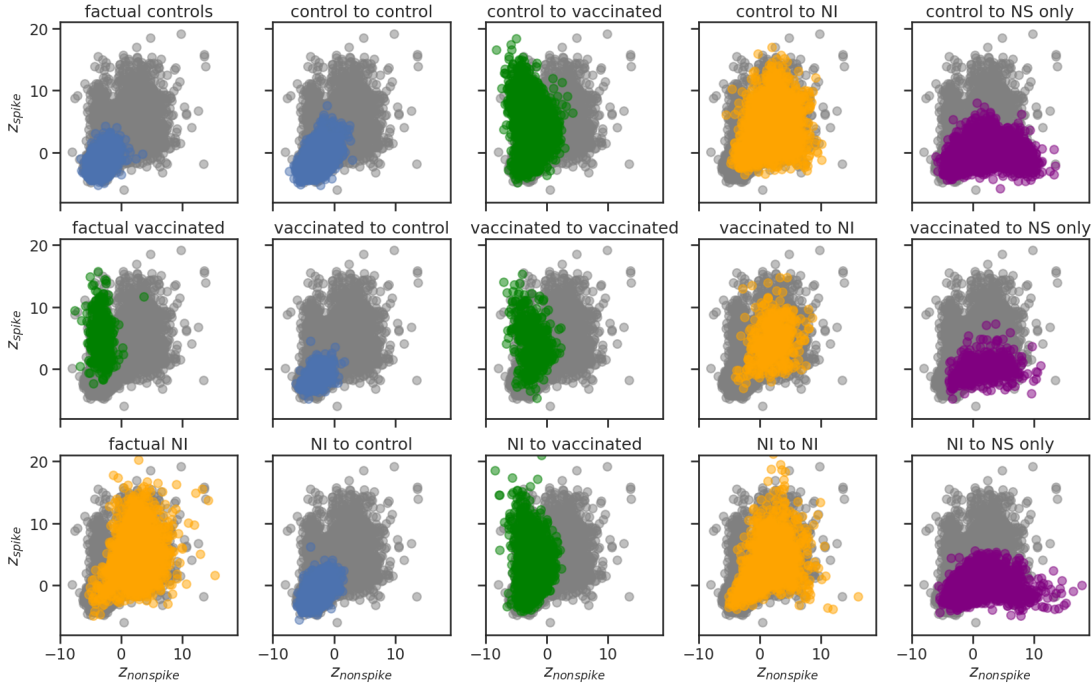


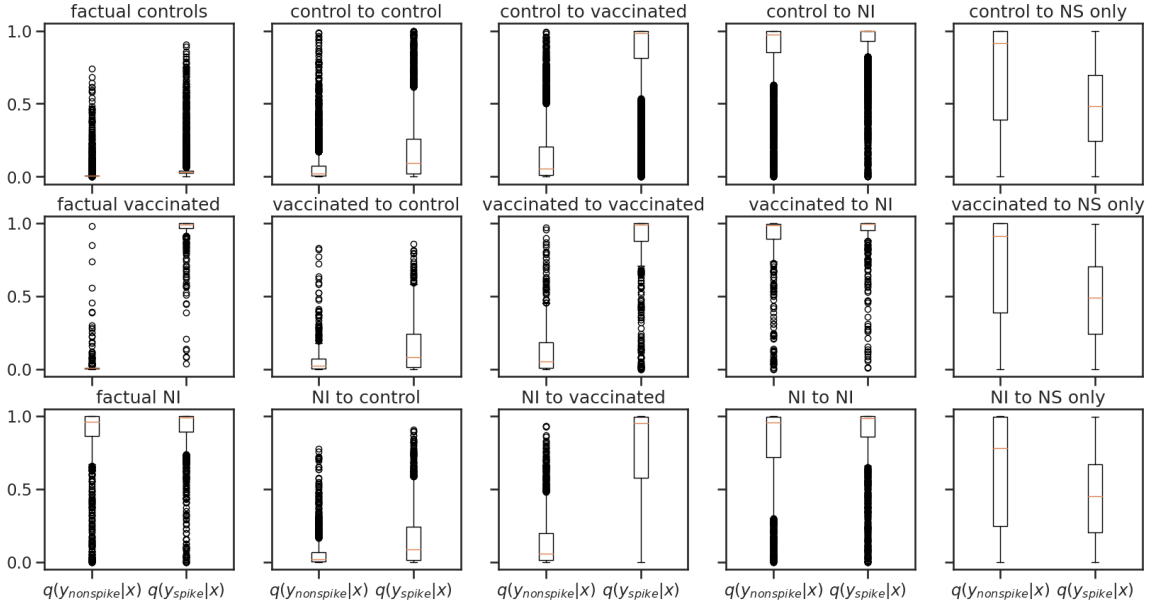
Figure 15: Distribution of TCR counts for observed and in-silico repertoires.

ill individuals however, such as those with acute COVID-19 infection, we often observe lower sequencing depth, associated with lymphopenia (Snyder et al., 2020).

Figure 15 compares the distribution of total sum of TCR counts for conditionally generated in-silico repertoires against the empirical distribution when stratifying by depth. Although the in-silico repertoires tend to exhibit a higher number of outliers due to the assumed Poisson distribution, the marginal distribution of in-silico repertoires follows the same trend than the empirical one, enabling us to control depth in a realistic manner.



(a) Posterior samples $z \sim q(\mathbf{Z}|\mathbf{X})$



(b) Posterior predictive $\mathbf{y} \sim q(\mathbf{Y}|\mathbf{X})$

Figure 16: **Counterfactual generation for COVID-19.** (a) Latent representation $z \sim q(\mathbf{Z}|\mathbf{X})$ for factual data (first column) and counterfactual data (other columns) stratified by subgroups. (b) Posterior predictive distribution $q(\mathbf{Y}|\mathbf{X})$ for factual (first column) and counterfactual data (columns 2-5) stratified by subgroups. The labels assignment is: control = [0 0], natural infection [1 0], and vaccinated repertoires [0 1]. The “nonspike only” subgroup can never be observed in reality.

E.5. Faithful Counterfactual Generation for COVID-19

Figure 16 shows posterior and posterior predictive samples for all possible label combinations to generate counterfactuals. This is the complete version of Figure 7 in the main text.

Appendix F. Additional Results for HSV

Figure 17 shows the latent space representation for AIRIVA trained on HSV. We see changes in HSV-1/HSV2 status of repertoires along the corresponding predictive latent axes. Figure 18 shows counterfactuals for each HSV subtype. Similarly to the COVID-19 case study, AIRIVA is able to generate counterfactuals whose latent representation is consistent with the latent representation of the observed data. Note that unlike COVID-19, all subgroups are observed.


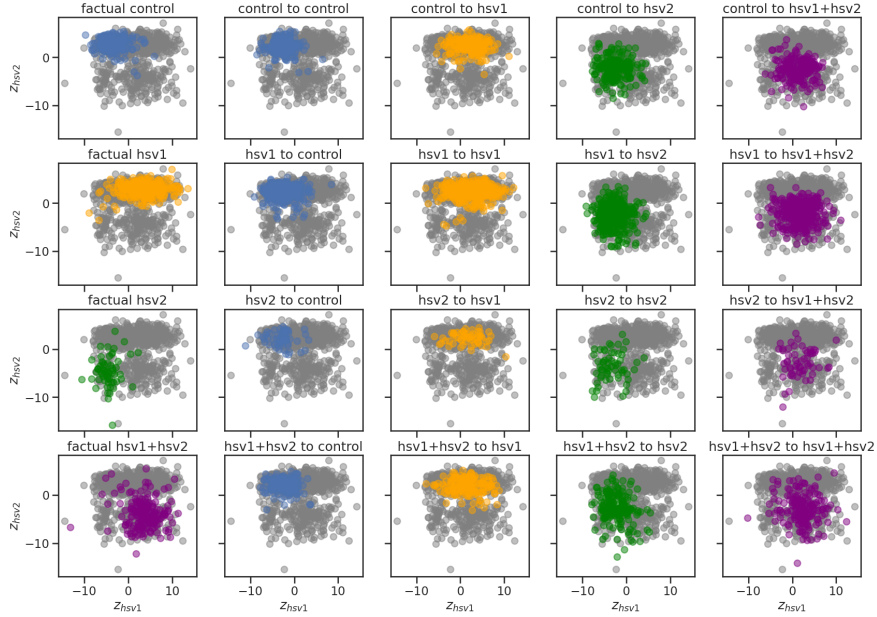
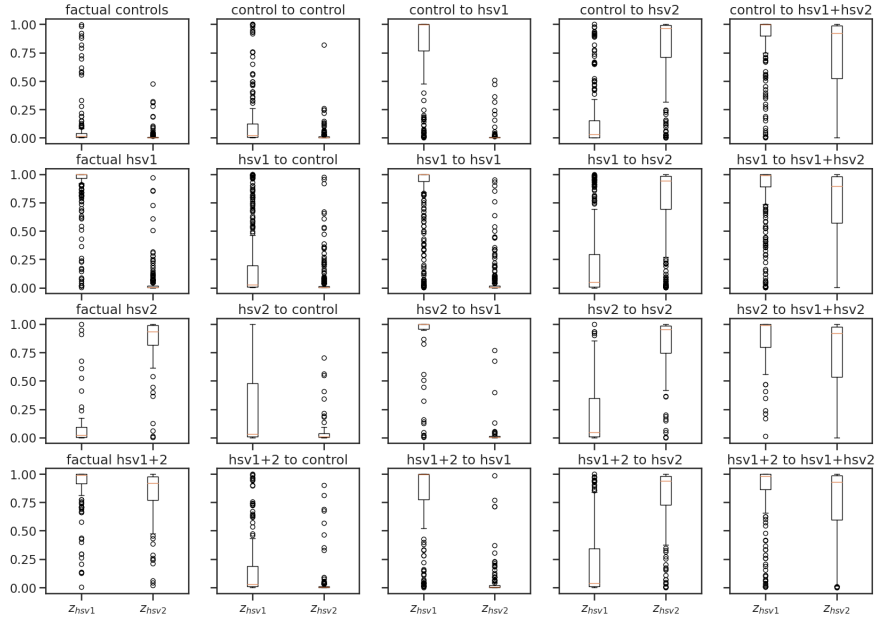


Figure 17: **Latent space for HSV.** Each point corresponds to one sample z from the posterior $q(z|x)$ in holdout. On the sides, we show marginal posterior as continuous lines, and predictive priors as dashed lines. Predictive priors $p(z_y|y)$ for each combination of labels y are represented as ellipses, each circle corresponding to one standard deviation.



(a) Posterior samples $z \sim q(\mathbf{Z}|\mathbf{X})$



(b) Posterior predictive $y \sim q(\mathbf{Y}|\mathbf{X})$

Figure 18: **Counterfactual generation for HSV.** (left) Latent representation $z \sim q(\mathbf{Z}|\mathbf{X})$ for factual data (first column) and counterfactual data (other columns) stratified by subgroups. (right) Posterior predictive distribution $q(\mathbf{Y}|\mathbf{X})$ for factual (first column) and counterfactual data (columns 2-5) stratified by subgroups. Label assignments are: control = [0 0], HSV1 [1 0], HSV2 [0 1] and HSV1+HSV2 [0 1].