Briefings in Bioinformatics, 2025, **26(2)**, bbaf126
https://doi.org/10.1093/bib/bbaf126
Problem Solving Protocol

CoupleVAE: coupled variational autoencoders for predicting perturbational single-cell RNA sequencing data

Yahao Wu[1], Jing Liu[1], Yanni Xiao[1], Shuqin Zhang[2], Limin Li[1,*]

[1]School of Mathematics and Statistics, Xi'an Jiaotong University, No. 28 Xianning West Road, Xi'an, Shaanxi 710049, China
[2]School of Mathematical Sciences, Center for Applied Mathematics, Research Institute of Intelligent Complex Systems, and Shanghai Key Laboratory for Contemporary Applied Mathematics, Fudan University, 220 Handan Road, 200433 Shanghai, China
*Corresponding author. School of Mathematics and Statistics, Xi'an Jiaotong University, No. 28 Xianning West Road, Xi'an, Shaanxi 710049, China.
E-mail: liminli@mail.xjtu.edu.cn

Abstract

With the rapid advances in single-cell sequencing technology, it is now feasible to conduct in-depth genetic analysis in individual cells. Study on the dynamics of single cells in response to perturbations is of great significance for understanding the functions and behaviors of living organisms. However, the acquisition of post-perturbation cellular states via biological experiments is frequently cost-prohibitive. Predicting the single-cell perturbation responses poses a critical challenge in the field of computational biology. In this work, we propose a novel deep learning method called coupled variational autoencoders (CoupleVAE), devised to predict the postperturbation single-cell RNA-Seq data. CoupleVAE is composed of two coupled VAEs connected by a coupler, initially extracting latent features for controlled and perturbed cells via two encoders, subsequently engaging in mutual translation within the latent space through two nonlinear mappings via a coupler, and ultimately generating controlled and perturbed data by two separate decoders to process the encoded and translated features. CoupleVAE facilitates a more intricate state transformation of single cells within the latent space. Experiments in three real datasets on infection, stimulation and cross-species prediction show that CoupleVAE surpasses the existing comparative models in effectively predicting single-cell RNA-seq data for perturbed cells, achieving superior accuracy.

Keywords: single-cell perturbation; cross-species; VAE; deep learning

Introduction

Single-cell RNA sequencing (scRNA-Seq) has revolutionized biological and medical research by revealing the heterogeneity of individual cells. Advances in high-throughput data analysis have enabled biologists to understand changes in cell type heterogeneity and transcriptomes over a long period of life [1], explore divergent developmental lineages [2], and quantify the pluripotency landscape of cell differentiation through cellular decomposition and dynamics [3, 4]. Studying single-cell dynamics in response to perturbations such as virus infections [5], drug treatments [6, 7], gene knockouts [8–10], and the cellular differences for different species [11] is essential for understanding organism functions and aiding drug discovery. Identifying marker genes linked to dynamic differences can enhance cellular phenotype understanding [12]. Due to high experimental costs [13, 14], computational modeling has become essential for predicting responses, enabling the generation of in silico scRNA-Seq data under specific conditions.

While mechanistic models have been introduced in various studies for predicting single-cell perturbation responses [15, 16], such models usually have high computational costs and low precision in parameter estimation, and cannot explain the nonlinearity of the system. For example, a regularized linear model called MIMOSCA [8] is used to analyze the transcriptional effects of single-cell perturbations by modeling the expression level of each gene as a linear combination of the perturbation effects, which is not possible to model the complex nonlinear effects. In recent years, several deep generative methods for predicting single-cell perturbation responses have been proposed, especially using adversarial generative networks (GANs) and variational autoencoders (VAEs) [17, 18]. Particularly, the work in [19] designed a GAN-based generative model for simulating gene expression and predicting single-cell perturbations. Since GANs do not allow for mapping features into a latent space, another GAN-based method named scPreGAN [20] predicts single-cell by integrating autoencoder and GANs. However, it is often hard to train GANs with structured high-dimensional data, which limits the wide use of GANs for the single-cell perturbation prediction [21]. VAE-based generative models have been more popularly used [21, 22]. The method scGen [21] first learns the VAE latent feature differences between the perturbed and unperturbed cells and then maps the prediction in latent space to high-dimensional gene expression space. The method trVAE [22] uses CVAE [23] and maximum mean deviation to handle multiple perturbations. Both models consider the response of different cell types to the perturbation as the same and thus may ignore the different responses of different cell types to perturbations, which may lead to prediction bias. The MichiGAN model [24] combines VAEs and GANs to generate

Received: October 24, 2024. **Revised:** January 21, 2025. **Accepted:** March 03, 2025

better samples while sampling from the untangled representation, where predicting drug response is one of its applications. However, this combined model usually has a complex structure, which may slow down the convergence rate and take a lot of time to train.

Recent studies have explored various applications of VAEs for predicting perturbational single-cell responses. Approaches such as scPRAM [25], scVIDR [26], and scPerb [27] enhance the VAE framework by adding auxiliary loss terms to capture perturbation effects, enabling more accurate predictions of cell state changes under perturbations. These approaches are particularly effective at learning latent representations that capture broad shifts between control and perturbed conditions. However, they primarily focus on computing the latent space shift δ between control and perturbed states, which may oversimplify the underlying biological processes. This latent space shift typically represents a linear displacement, neglecting the complex, non-linear dynamics involved in cellular responses. As a result, these methods may fail to fully capture the intricate biological transitions that occur during perturbations, limiting their ability to accurately model the true cellular transformations.

To overcome the limitations of the existing methods, in this work, we propose a deep generative model called coupling VAE (CoupleVAE) for predicting perturbation responses of scRNA-Seq data. CoupleVAE is a probabilistic graphical model that contains an inference process, a coupling process and a generative process. The key advantage of CoupleVAE is that the coupling process could learn the nonlinear mutual transformations between the controlled and the perturbed cells in latent space, which explains how the controlled cells respond to the perturbations.

Our approach not only enables predictions of SARS-COV-2 infection response of cells, IFN-β stimulation across different cell types, and it could also predict lipopolysaccharide (LPS) perturbation of single cells from other species, with better performance than existing methods.

Materials and methods

We formulate the perturbational single-cell RNA sequencing data prediction problem as follows. Let $\mathbf{x}_c$ and $\mathbf{x}_p$ be the vectors representing the RNA sequencing data for a single cell before and after some perturbation. We aim to learn a predictive model $P(\mathbf{x}_p|\mathbf{x}_c)$ based on the training pairs $\{\mathbf{x}_c^i, \mathbf{x}_p^i\}_{i=1}^N$.

Variational autoencoder

VAE [28] is a network structure to learn a generative process for multi-dimensional variable $\mathbf{x}$ through a latent variable $\mathbf{z}$ based on N *i.i.d.* samples $\mathbf{X} = \{\mathbf{x}^i\}_{i=1}^N$ of $\mathbf{x}$. It assumes a generative process that consists of generating $\mathbf{z}^i$ with a prior distribution $P_\theta(\mathbf{z})$ and generating $\mathbf{x}^i$ from a conditional distribution $P_\theta(\mathbf{x}|\mathbf{z})$, where the parameters θ for the latent variables $\mathbf{z}^i$ are unknown. By introducing an encoder $Q_\phi(\mathbf{z}|\mathbf{x})$ to approximate the true intractable posterior $P_\theta(\mathbf{z}|\mathbf{x})$, the evidence lower bound (ELBO) of the marginal likelihood $\log P_\theta(\mathbf{x}^i)$ [29] can be written as:

$$\mathcal{L}(\theta, \phi; \mathbf{x}^i) = \mathbb{E}_{Q_\phi(\mathbf{z}|\mathbf{x}^i)}[\log P_\theta(\mathbf{x}|\mathbf{z})] - KL(Q_\phi(\mathbf{z}|\mathbf{x}^i) \parallel P_\theta(\mathbf{z})), \quad (1)$$

where $KL(p||q) = -\int_{\mathcal{X}} p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})} d\mathbf{x}$ is the KL Divergence for two distributions p and q. The VAE jointly learns parameters ϕ in the encoder $Q_\phi(\mathbf{z}|\mathbf{x})$ and parameters θ in the decoder $P_\theta(\mathbf{x}|\mathbf{z})$ by maximizing the ELBO $\mathcal{L}(\theta, \phi; \mathbf{x}^i)$, where the first term represents the reconstruction likelihood and the second term ensures that the learned distribution Q_ϕ is similar to the true prior distribution P_θ.

In order to estimate the ELBO based on the given data points, the VAE takes prior of $\mathbf{z}$ as $P_\theta(\mathbf{z}) = \mathcal{N}(\mathbf{z}; \mathbf{0}, \mathbf{I})$, and assumes the variational distribution $Q_\phi(\mathbf{z}|\mathbf{x})$ is a multivariate Gaussian distribution

$$\log Q_\phi(\mathbf{z}|\mathbf{x}^i) = \log \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}^i, {\boldsymbol{\sigma}^i}^2 \mathbf{I}),$$

where $\boldsymbol{\mu}^i$ and $\boldsymbol{\sigma}^i$ are outputs of encoder for $\mathbf{x}^i$. The VAE further uses the reparameterization trick $\mathbf{z}^{i,l} = \boldsymbol{\mu}^i + \boldsymbol{\sigma}^i \odot \boldsymbol{\epsilon}^l, \boldsymbol{\epsilon}^l \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, where $\odot$ means element-wise product, to sample from the posterior $Q_\phi(\mathbf{z}|\mathbf{x}^i)$. By sampling $\mathbf{z}^i$ L times, the ELBO of the model can then be estimated as:

$$\mathcal{L}(\theta, \phi; \mathbf{x}^i) \approx \frac{1}{L} \sum_{l=1}^{L} (\log P_\theta(\mathbf{x}^i|\mathbf{z}^{i,l})) - KL(Q_\phi(\mathbf{z}|\mathbf{x}^i)||P_\theta(\mathbf{z})). \quad (2)$$

The VAE makes use of two neural networks as encoder and decoder, and applies Adam [30] to optimize the objective. The encoder network with parameters ϕ takes $\mathbf{x}^i$ as input and outputs the parameters ($\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$) in the approximate posterior $Q_\phi(\mathbf{z}|\mathbf{x})$. By generating $\mathbf{z}^i$ from this posterior distribution as an input, the decoder network with parameters θ outputs the parameters in the conditional distribution $P_\theta(\mathbf{x}|\mathbf{z})$.

CoupleVAE for perturbational single-cell RNA-sequencing data prediction

CoupleVAE is developed under the framework of VAE. The workflow of CoupleVAE is shown in Figure 1. It aims to predict the responses of single cells measured by scRNA-Seq data after a specific type of perturbation. It takes the gene expression of controlled and perturbed cells as training data, and consists of an inference process (encoder), a coupling process (coupler), and a generative process (decoder). The inference process includes two independent encoders to obtain the lower-dimensional representation of cells under the control and the perturbation condition, respectively, in two separated latent spaces. Then the coupling process translates the two latent vectors using nonlinear transformations into the latent space of different conditions as a coupler. Finally, with the generative process by decoders, controlled cells and perturbed cells are reconstructed from the encoded vectors and the translated vectors in the same latent space, respectively. The details for the inference process, generative process, and variational lower bound are described as follows.

Inference process

CoupleVAE contains two recognition models in the inference process, which independently produce the latent variables $\mathbf{z}_c$ and $\mathbf{z}_p$ from the input $\mathbf{x}_c$ and $\mathbf{x}_p$, respectively. The true posteriors of $\mathbf{z}_c$ and $\mathbf{z}_p$ are denoted as $P_\theta(\mathbf{z}_c|\mathbf{x}_c)$ and $P_\theta(\mathbf{z}_p|\mathbf{x}_p)$, respectively. Since it is intractable to calculate the integral of posterior as in VAE, we introduce a distribution Q parameterized by ϕ to approximate the true posterior. Due to the independence of the two recognition models, the approximate posterior can be written as

$$Q_\phi(\mathbf{z}_c, \mathbf{z}_p|\mathbf{x}_c, \mathbf{x}_p) = Q_\phi(\mathbf{z}_c|\mathbf{x}_c) \cdot Q_\phi(\mathbf{z}_p|\mathbf{x}_p). \quad (3)$$

In practice, the CoupleVAE takes the prior over the latent variables $\mathbf{z}_c$ and $\mathbf{z}_p$ as $P_\theta(\mathbf{z}_c) = \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $P_\theta(\mathbf{z}_p) = \mathcal{N}(\mathbf{0}, \mathbf{I})$, and assumes approximate posterior distribution to be multivariate

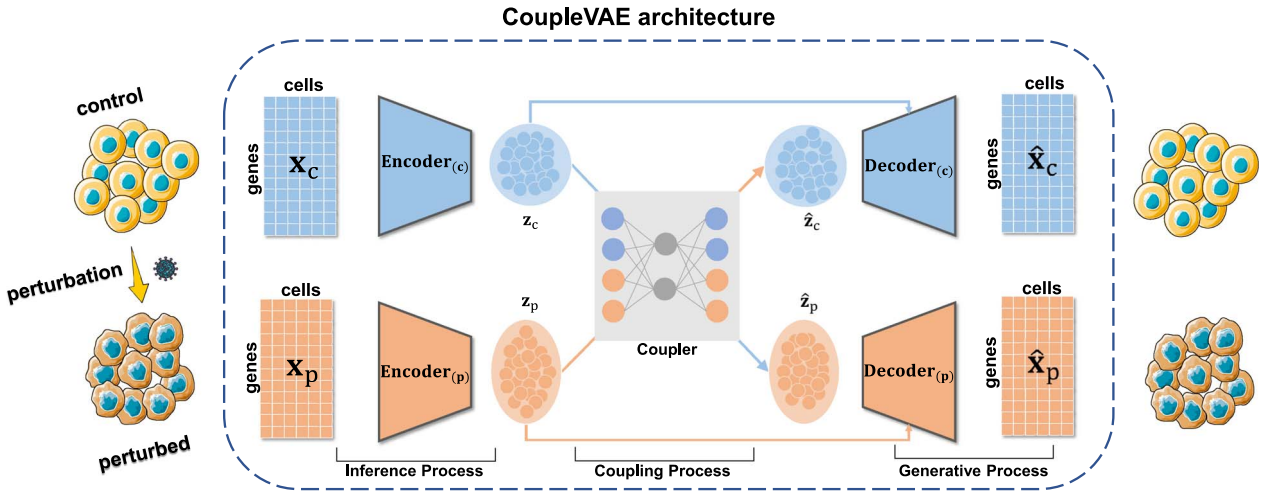

Figure 1. Workflow of CoupleVAE. CoupleVAE takes the gene expression of controlled cells ($\mathbf{x}_c$) and perturbed ($\mathbf{x}_p$) cells as input. The whole procedure includes three processes. Inference process: the latent representations $\mathbf{z}_c$ and $\mathbf{z}_p$ of controlled cells ($\mathbf{x}_c$) and perturbed cells ($\mathbf{x}_p$) are obtained by two encoders. Coupling process: $\mathbf{z}_c$ is transformed to $\hat{\mathbf{z}}_p$ through a nonlinear mapping, and $\mathbf{z}_p$ is transformed to $\hat{\mathbf{z}}_c$ through inverse mapping by the coupler. Generative process: gene expression of the controlled cell ($\hat{\mathbf{x}}_c$) and the perturbed cell ($\hat{\mathbf{x}}_p$) are reconstructed from $\mathbf{z}_c$, $\hat{\mathbf{z}}_c$ and $\mathbf{z}_p$, $\hat{\mathbf{z}}_p$ through the corresponding decoder. Components of this figure are created using Servier Medical Art templates, which are licensed under a Creative Commons Attribution 3.0 Unported License from https://smart.servier.com.

Gaussians

$$Q_\phi(\mathbf{z}_c|\mathbf{x}_c) = \mathcal{N}(\mathbf{z}_c|\boldsymbol{\mu}_{\phi_c}, \boldsymbol{\sigma}^2_{\phi_c}), Q_\phi(\mathbf{z}_p|\mathbf{x}_p) = \mathcal{N}(\mathbf{z}_c|\boldsymbol{\mu}_{\phi_p}, \boldsymbol{\sigma}^2_{\phi_p}). \quad (4)$$

Coupling process

CoupleVAE assumes a coupling process of two hidden variables in the latent space, that is, latent variable $\mathbf{z}_c$ and $\mathbf{z}_p$ are transformed to $\mathbf{z}_p$ and $\mathbf{z}_c$ via conditional distributions $P_\theta(\hat{\mathbf{z}}_p|\mathbf{z}_c)$ and $P_\theta(\hat{\mathbf{z}}_c|\mathbf{z}_p)$, respectively.

It realizes the mutual translation between the controlled cell and the perturbed cell in the latent space by forcing $\hat{\mathbf{z}}_p$ ($\hat{\mathbf{z}}_c$) and $\mathbf{z}_p$ ($\mathbf{z}_c$) to be close. This connects the inference process and generative process.

By assuming the independence between the two coupling distributions, as shown in Figure 1, we could have

$$P_\theta(\hat{\mathbf{z}}_c, \hat{\mathbf{z}}_p|\mathbf{z}_c, \mathbf{z}_p) = P_\theta(\hat{\mathbf{z}}_c|\mathbf{z}_p)P_\theta(\hat{\mathbf{z}}_p|\mathbf{z}_c). \quad (5)$$

In practice, CoupleVAE takes the coupling distributions as multivariate Gaussians:

$$P_\theta(\hat{\mathbf{z}}_p|\mathbf{z}_c) = \mathcal{N}(\hat{\mathbf{z}}_p|\boldsymbol{\mu}_p, \boldsymbol{\sigma}^2_p), P_\theta(\hat{\mathbf{z}}_c|\mathbf{z}_p) = \mathcal{N}(\hat{\mathbf{z}}_c|\boldsymbol{\mu}_c, \boldsymbol{\sigma}^2_c). \quad (6)$$

Generative process

CoupleVAE includes two generative processes. One is to reconstruct $\mathbf{x}_c$ and $\mathbf{x}_p$ from $\mathbf{z}_c$ and $\mathbf{z}_p$ respectively, which is the same as VAE, and the other is to reconstruct $\mathbf{x}_c$ and $\mathbf{x}_p$ from hidden variables $\hat{\mathbf{z}}_p$ and $\hat{\mathbf{z}}_c$, respectively. As shown in Figure 1, CoupleVAE generates $\mathbf{x}_p$ for the perturbed cell from the translated latent vector $\hat{\mathbf{z}}_c$ through $P_\theta(\mathbf{x}_p|\hat{\mathbf{z}}_c)$, while generates $\mathbf{x}_c$ for the controlled cell from the translated latent vector $\hat{\mathbf{z}}_p$ through $P_\theta(\mathbf{x}_c|\hat{\mathbf{z}}_p)$.

By assuming the independence of the generative process for $\mathbf{x}_p$ and $\mathbf{x}_c$, which is also shown in the graph model in Figure 1, we could have generative model $P_\theta(\mathbf{x}_c|\mathbf{z}_p)$, $P_\theta(\mathbf{x}_p|\mathbf{z}_c)$ (reconstruction) and $P_\theta(\mathbf{x}_c|\hat{\mathbf{z}}_p)$, $P_\theta(\mathbf{x}_p|\hat{\mathbf{z}}_c)$ (transform-reconstruction). In practice, the CoupleVAE takes the generative probability as

$$P_\theta(\mathbf{x}_p|\mathbf{z}_p) = \mathcal{N}(\mathbf{x}_p|\boldsymbol{\mu}_{\mathbf{z}_p}, \boldsymbol{\sigma}^2_{\mathbf{z}_p}), P_\theta(\mathbf{x}_c|\mathbf{z}_c) = \mathcal{N}(\mathbf{x}_c|\boldsymbol{\mu}_{\mathbf{z}_c}, \boldsymbol{\sigma}^2_{\mathbf{z}_c}), \quad (7)$$

$$P_\theta(\mathbf{x}_p|\hat{\mathbf{z}}_p) = \mathcal{N}(\mathbf{x}_p|\boldsymbol{\mu}_{\hat{\mathbf{z}}_p}, \boldsymbol{\sigma}^2_{\hat{\mathbf{z}}_p}), P_\theta(\mathbf{x}_c|\hat{\mathbf{z}}_c) = \mathcal{N}(\mathbf{x}_c|\boldsymbol{\mu}_{\hat{\mathbf{z}}_c}, \boldsymbol{\sigma}^2_{\hat{\mathbf{z}}_c}). \quad (8)$$

Variational lower bound

To learn the inference, coupling and generative processes based on training pairs, the objective of CoupleVAE is to maximize a sum over the log likelihood of the pairs of cells $(\mathbf{x}_c, \mathbf{x}_p)$ in the training set, i.e. $\sum \log P_\theta(\mathbf{x}_c, \mathbf{x}_p)$. We theoretically prove that the log likelihood of joint distribution for $(\mathbf{x}_c, \mathbf{x}_p)$ has the following lower bound:

$$\begin{aligned}
&\log P_\theta(\mathbf{x}_c, \mathbf{x}_p) \geqslant \mathcal{L}(\mathbf{x}_c, \mathbf{x}_p),\\
&\mathcal{L} = \frac{1}{2}(\underbrace{\mathbb{E}_{Q_\phi(\mathbf{z}_c|\mathbf{x}_c)}[\log P_\theta(\mathbf{x}_p|\hat{\mathbf{z}}_p)] + \mathbb{E}_{Q_\phi(\mathbf{z}_p|\mathbf{x}_p)}[\log P_\theta(\mathbf{x}_c|\hat{\mathbf{z}}_c)]}_{\mathcal{L}_{\text{trans-recon}}}\\
&\quad + \underbrace{\mathbb{E}_{Q_\phi(\mathbf{z}_p|\mathbf{x}_p)}[\log P_\theta(\mathbf{x}_p|\mathbf{z}_p)] + \mathbb{E}_{Q_\phi(\mathbf{z}_c|\mathbf{x}_c)}[\log P_\theta(\mathbf{x}_c|\mathbf{z}_c)]}_{\mathcal{L}_{\text{recon}}}\\
&\quad \underbrace{- \mathbb{KL}(Q_\phi(\mathbf{z}_c|\mathbf{x}_c)||P_\theta(\mathbf{z}_c)) - \mathbb{KL}(Q_\phi(\mathbf{z}_p|\mathbf{x}_p)||P_\theta(\mathbf{z}_p))}_{\mathcal{L}_{\text{regu}}}\\
&\quad + \underbrace{\mathbb{E}_{Q_\phi(\mathbf{z}_c|\mathbf{x}_c)}[\log P_\theta(\hat{\mathbf{z}}_p|\mathbf{z}_c)] + \mathbb{E}_{Q_\phi(\mathbf{z}_p|\mathbf{x}_p)}[\log P_\theta(\hat{\mathbf{z}}_c|\mathbf{z}_p)]}_{\mathcal{L}_{\text{couple}}}). \quad (9)
\end{aligned}$$

From the equation (9), the ELBO consisting of four items is formulated as:

$$\mathcal{L} = \frac{1}{2}(\mathcal{L}_{\text{trans-recon}} + \mathcal{L}_{\text{recon}} + \mathcal{L}_{\text{regu}} + \mathcal{L}_{\text{couple}}), \quad (10)$$

where reconstruction errors $\mathcal{L}_{\text{trans-recon}}$ and $\mathcal{L}_{\text{recon}}$ are to minimize the similarities between the generated cells and the true cells, regularization term $\mathcal{L}_{\text{regu}}$ is to ensure the closeness between posterior distributions of $\mathbf{z}_p$ and $\mathbf{z}_c$ and the prior distributions,

and the coupling loss $\mathcal{L}_{\text{couple}}$ enables the mutual translation in the latent space, i.e. the transformed $\mathbf{z}_c$ (or transformed $\mathbf{z}_p$) is similar with $\mathbf{z}_p$ (or $\mathbf{z}_c$) in the latent space. Furthermore, the independence assumption is mainly used to simplify the derivation of the ELBO. Nevertheless, the two processes are fundamentally connected through the joint likelihood, which is constrained by the ELBO. The CoupleVAE model learns the nonlinear intertransformation between control and perturbed cells in latent space through a coupling process. This approach captures the complex dynamics of cells when perturbed, providing new insights into the mechanisms of cell state transition. More details of the theoretical and empirical loss function are shown in Supplementary Note 9.

The workflow of CoupleVAE is shown in Figure 1. Let $\mathbf{x}_c$ and $\mathbf{x}_p$ represent the gene expression for a controlled and perturbed cell, respectively. CoupleVAE first learns the posterior distributions $Q_\phi(\mathbf{z}_c|\mathbf{x}_c)$ and $Q_\phi(\mathbf{z}_p|\mathbf{x}_p)$ corresponding to the latent variables $\mathbf{z}_c$ and $\mathbf{z}_p$ for the controlled and perturbed cell by two encoders, respectively. The two latent variables are then mutually translated with conditional distributions $P(\hat{\mathbf{z}}_p|\mathbf{z}_c)$ and $P(\hat{\mathbf{z}}_c|\mathbf{z}_p)$ into the latent space of different conditions via coupling process. Finally, CoupleVAE generates gene expressions for controlled cells and perturbed cells by two decoders from translations $\hat{\mathbf{z}}_p$ and $\hat{\mathbf{z}}_c$, respectively. In CoupleVAE, the coupler contains a perturbation process '$\mathbf{z}_c \to \hat{\mathbf{z}}_p$' and a restoration process '$\mathbf{z}_p \to \hat{\mathbf{z}}_c$' in the latent space. The perturbation process '$\mathbf{z}_c \to \hat{\mathbf{z}}_p$' of the model in the latent space corresponds to a perturbation applied to the control cells in real biological experiments, while the restoration process '$\mathbf{z}_p \to \hat{\mathbf{z}}_c$' in the latent space refers to the restoration of the cells from the perturbed state to the control state. Our model aims to infer the two processes by maximizing the joint likelihood of the given scRNA-seq for both states. Both the perturbation and restoration processes are assumed to occur in a latent space. Biologically, a specific perturbation primarily affects latent unknown factors of cells, influencing not only gene expression but also gene regulatory relationships, pathways, and cell-cell communications. Utilizing latent factors rather than gene expressions in the input space offers a more intuitive and systematic approach for generating profiles of perturbed cells.

Experimental design

Our experiments are divided into two setups: independent-and-identically distributed (i.i.d.) and out-of-distribution (o.o.d.).

- **i.i.d. Setup.** For the i.i.d. setup, we used the COVID-19 and peripheral blood mononuclear cell (PBMC) datasets. In this setup, the data are randomly split into training, validation, and test sets in an 8:1:1 ratio. After splitting the training and testing sets using cross-validation, we balanced the samples in the training and test sets so that each cell type contains an equal number of control and perturbational cells. Specifically, for each cell type, we first determined the minimum number of cells M across the two conditions, and randomly selected M cells from the condition with more cells (undersampling). This ensures that each cell type has an equal number of cells under both conditions. During the training process, CoupleVAE is fed cells in controlled and perturbed states from the same specific cell type. In the latent space, the model constructs the transformation from cells in the control state to cells in the perturbed state. Finally, the cells in the latent space are passed through the decoder to reconstruct cells. For prediction, we input the controlled cells from the testing set into the encoder, which transforms them into a latent vector that is then perturbed. After passing through the decoder, the output is the predicted data.
- **o.o.d. Setup.** In the o.o.d. setup, the model does not have access to cells in the holdout condition. This setup is used with the LPS6 dataset. When predicting a perturbation for a particular cell type, the training data consist of all cell types except the one being predicted. For the cross-species prediction task on the LPS6 dataset, we used two training strategies corresponding to two models: condition-CoupleVAE and species-CoupleVAE. For condition-CoupleVAE, we use the cells from all species in the control state of the training set (red box in Fig. 4a) as $\mathbf{x}_c$, and the cells from all the species in the LPS state as $\mathbf{x}_p$. For species-CoupleVAE, we use mouse cells in the training set (red box in Fig. 4e) as $\mathbf{x}_c$ and rabbit cells as $\mathbf{x}_p$. We input the mouse cells in the LPS state as the test data into species-CoupleVAE, and finally get the rabbit cells in the LPS state.

Results

The experiments are done for three datasets including: the COVID-19 dataset with SARS-COV-2 infected cells, the PBMC dataset with IFN-β stimulation cells, and the LPS6 dataset with LPS6 stimulation for different species, with all dataset having their corresponding controlled cells. The following experiments show that CoupleVAE is able to predict the cells' response to perturbation more accurately compared to the existing popular methods, and it enables cross-species perturbation prediction for an unseen species.

CoupleVAE accurately predicts response to SARS-COV-2 infection

This dataset contains cells from the COVID-19 dataset summarized by Lotfollahi *et al.* [31](Fig. 2a), where the samples are collected from patients of different ages. We select six cell types in two conditions (control and severe COVID-19) and use the independently-and-identically distributed (i.i.d.) setting proposed by Bunne *et al.* [32](see Section 2) to predict the single-cell responses to SARS-COV-2 infection. For example, in predicting macrophages' response to perturbation, the training data contain only the cells of type macrophages. To evaluate the performance of CoupleVAE, the scRNA-Seq data for all six cell types after infection with SARS-CoV-2 are predicted and compared with the real perturbational data. We compare CoupleVAE with five single-cell perturbation prediction methods: scPreGAN [20], trVAE [22], scGen [21], scPRAM [25] and scVIDR [26], and one machine learning method CVAE [23], which has recently been adapted to preprocessing, batch-correcting and differential testing of single-cell data [17].

We first calculate the correlations between the predicted and the real expression of all genes for each specific cell type, which are measured by the square of Pearson correlation coefficient (PCC) R^2. Figure 2b and Supplementary Figure S1a report the boxplot of R^2 for all genes and the top 100 deferentially expressed genes(DEGs), respectively. Results for the compared methods are also presented. Here, the DEGs are obtained by Scanpy [33] based on the controlled and real perturbational scRNA-Seq data. CoupleVAE can obtain the highest or competitive R^2 values for all cell types, which implies that it could simulate reliable gene expression levels after SARS-COV-2 infection. CoupleVAE also obtains significantly higher correlation value $R^2_{\text{top 100 DEGs}}$ for top 100 DEGs (shown in Supplementary Figure S1a and Table S21).

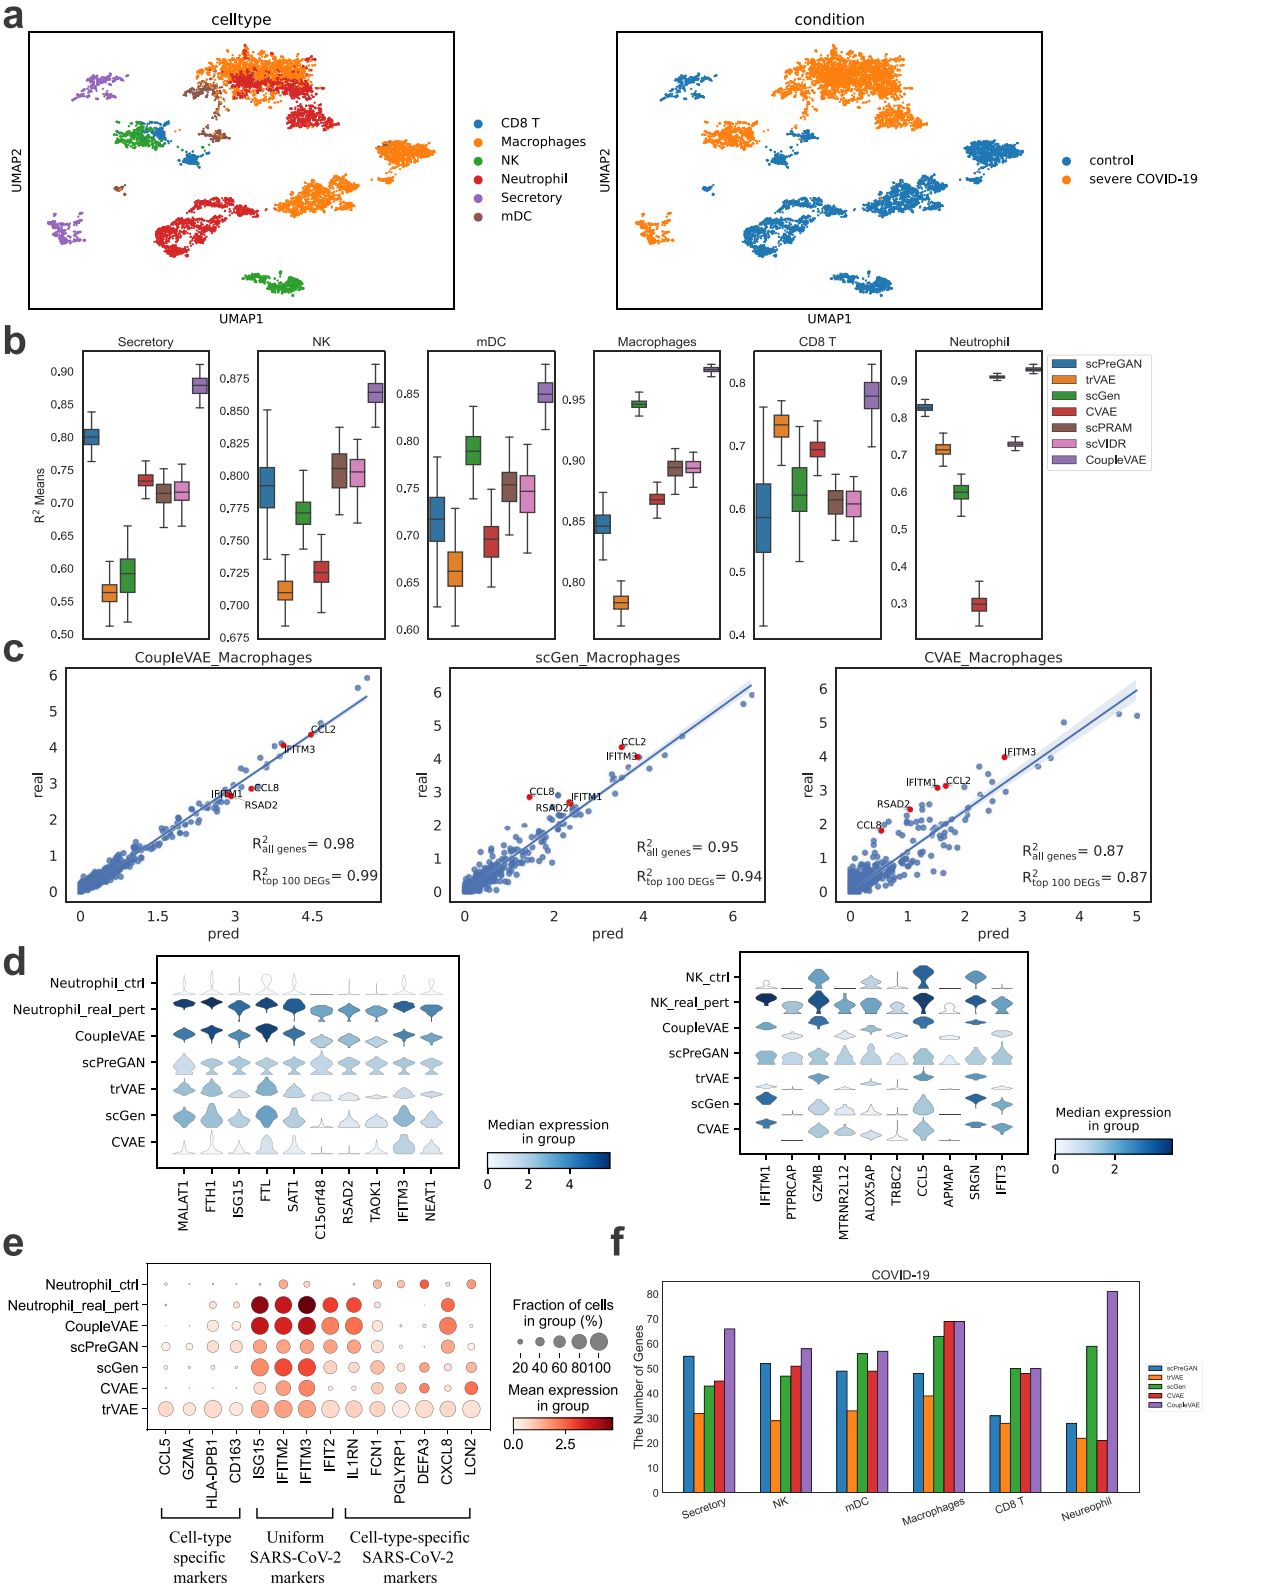

Figure 2. Single-cell response prediction after SARS-CoV-2 infection. (a) UMAP visualization of the distribution for cells of different types with and without SARS-COV-2 infection. (b) Comparison of correlations between the real expression of all genes with SARS-COV-2 infection and the predicted using seven methods. R^2 denotes the square of the PCC between the means. (c) Correlation comparison between the real mean values of 6000 genes'expression in 248 macrophage cells with SARS-CoV-2 infection and those predicted by CoupleVAE, scGen, and CVAE. The genes marked as significant are top five DEGs under real perturbed conditions. (d) Predicted gene distribution for 10 differentially expressed genes in neutrophil and NK cells after infected by SARS-CoV-2. The vertical axis shows the gene expression and distribution of neutrophil & NK cells in control, real perturbed and predicted by five different methods. (e) Predicted mean gene expression comparison for the neutrophil cells with SARS-COV-2 infection. (f) DEG identification results comparison. The number of genes represents the number of common genes between ground truth and predicted value among the top 100 DEGs.

To examine the details of the predicted gene expression, we take a further look at the average of the predicted and real expression level for each gene in the cell type macrophages in Figure 2c and other cell types in Supplementary Figures S2 and S3. Each point in the figure represents the average expression value of one gene across all cells in a given cell type. Compared to the two competitive methods scGen and CVAE, the expression level obtained from CoupleVAE tends to gather more closely to the real expression, and the $R^2_{\text{all genes}}$ reaches 0.98, which is higher than that of scGen (0.95) and CVAE (0.87). In particular, CoupleVAE predicts much more accurate expression values for the top 100 DEGs with $R^2_{\text{top 100 DEGs}}$ being 0.99, compared to that of scGen (0.94) and CVAE (0.87). The top five DEGs predicted by CoupleVAE also show a higher correlation with the real expression as marked in Figure 2c. We further examine whether the predicted gene expression distribution for each specific gene is close to the real distribution by looking at each of the top 10 DEGs in different cell types. Figure 2d shows the violin plots of the top 10 DEGs in neutrophil and NK cells, respectively, using the five different methods. CoupleVAE achieves more similar expression distributions for most of these genes in these two cell types, which implies that CoupleVAE could predict perturbational scRNA-Seq data with better accuracy. The violin plots for other cell types are also presented in Supplementary Figure S1b.

In addition, CoupleVAE can successfully predict gene regulation patterns such as no regulation, shared regulation across all cell types, and specific regulation for one cell type. We use previous description of marker genes for three different gene regulation patterns: cell-type-specific markers corresponding to no regulation [34–38], uniform SARS-Cov-2 markers corresponding to shared regulation [37], and cell-type-specific SARS-Cov-2 markers corresponding to specific regulation [37]. Figure 2e and Supplementary Figure S4 show that the gene regulation pattern after infection could be captured accurately by CoupleVAE. Cell-type-specific markers show no expression changes before and after perturbation such as gene CCL5 and HLA-DPB1. Uniform SARS-Cov-2 markers change similarly for all cell types, for example, ISG15, IFIM2, IFIM3, and IFIT2 are upregulated for all cell types after perturbation. Cell-type-specific SARS-Cov-2 markers change differently for different cell types, for example, SARS-Cov-2 gene DEFA3 is downregulated only for neutrophils cells, and LCN2 is downregulated for neutrophils cells and upregulated in secretory cells. These results imply the ability of CoupleVAE in accurately predicting the gene regulation patterns after infection. Different cell types may respond differently to the same perturbation, and CoupleVAE is able to distinguish between these cell type-specific responses to avoid prediction bias.

Finally, CoupleVAE is further evaluated by comparing the identified DEGs. Once the perturbational gene expressions are predicted, one of key downstream applications can be estimating DEGs by comparing the expression levels of control cells with those predicted for the perturbed cells, and thus detecting perturbation effects. The DEGs are identified using the methods in Scanpy [33] following the setting in scPreGAN [20]. That is, we take the top 100 DEGs between the controlled data and the real perturbational data as the ground truth, and identify the top 100 DEGs from the controlled data and the predicted data as the predicted DEGs. Then the number of the overlapping DEGs is computed. A higher percentage implies a more reliable model for predicting perturbational data. As shown in Figure 2f, among the five methods on the six cell types, CoupleVAE shows its better ability in identifying the ground truth DEGs. Particularly, CoupleVAE could identify more than 50 ground truth DEGs for all cell types, and even more than 80 DEGs for neutrophil cells among the top 100 DGEs.

CoupleVAE accurately predicts cell response for $IFN-\beta$ treatment

We benchmark against a single-cell gene expression dataset containing 7217 $IFN-\beta$ stimulated and 6359 controlled PBMCs from eight different human lupus patients. Stimulation with $IFN-\beta$ induced significant changes in the transcriptional profile of immune cells, which resulted in visible changes between control and stimulated cells. CoupleVAE and the four comparison models are applied to predict the single-cell responses after IFN-β treatment. Similar to the strategy used for the COVID-19 dataset, we follow the i.i.d. settings for training and testing.

The performance of CoupleVAE is first evaluated by computing the square of PCCs R^2 between the predicted and the real gene expression levels of either all genes or the top 100 DEGs, as shown in Figure 3a and Supplementary Figure S5a, respectively. The results show that CoupleVAE more reliably predicts perturbation responses in most cell types. For NK cells and CD8 T cells, the predictions from CoupleVAE and scGen are comparable. For CD14+Mono cells, the predictions from CoupleVAE closely resemble those from scPRAM and scVIDR.

The details of the predicted expression level for each gene are further explored by looking at the distribution compared to the real expression in each specific cell type. Figure S5b plots the mean of the predicted expression by CoupleVAE and the real expression for each gene in cell type FCGR3A+Mono, which shows that the dots are concentrated around the straight line. Particularly, among all the genes, the predicted of the top five DEGs lie closely to the real ones. This implies that CoupleVAE could produce gene expression close to the real expression. Results for other cell types are shown in Supplementary Figures S6 and S7. Furthermore, Figure 3c compares distribution of the predicted and real gene expression for gene IFITM3 (the top response gene to $IFN-\beta$) in cell type FCGR3A+Mono(F-Mono). Supplementary Figure S5b shows the violin plots for the other five top DEGs. The violin plots show that CoupleVAE predicts the expression distribution most similar to the real one.

CoupleVAE could predict the gene regulations responding to the treatment similar to those in the real treatment data. The visualization of CD4T cells by UMAP [40] before and after treatment in Figure 3d shows that the predicted perturbational scRNA-Seq data by CoupleVAE is similar to the real one, and Figure 3e further gives some examples where the predicted expression values of three genes (ISG15, IFI6, ISG20) show the same tendency of upregulation as the real ones in CD4 cells. Supplementary Figures S8 and S9 show the same results for the other six cell types. Supplementary Figure S10 shows the results of the four comparison methods on seven cell types. Figure 3f and Supplementary Figure S11 show how the different methods predict regulation styles for cell-type-specific markers, uniform IFN-β markers and cell-type-specific IFN-β markers for F-mono cells [41]. CoupleVAE could predict the regulation pattern better than the other methods. For example, in F-Mono cells, scGen fails to identify the upregulation for some uniform cell-specific IFN-β markers DEFB1 and CCL8 after treatment, CVAE misses the upregulation of other cell-specific IFN-β markers, including IFIT1, CXCL20, and CXCL11. Meanwhile, trVAE predicts all genes as upregulated. Figure 3g finally gives the number of common DEGs between the top 100 predicted DEGs obtained by predicted perturbational scRNA-seq and top 100 ground truth DEGs obtained by real scRNA-seq data. For six of the seven cell types, CoupleVAE could identify the

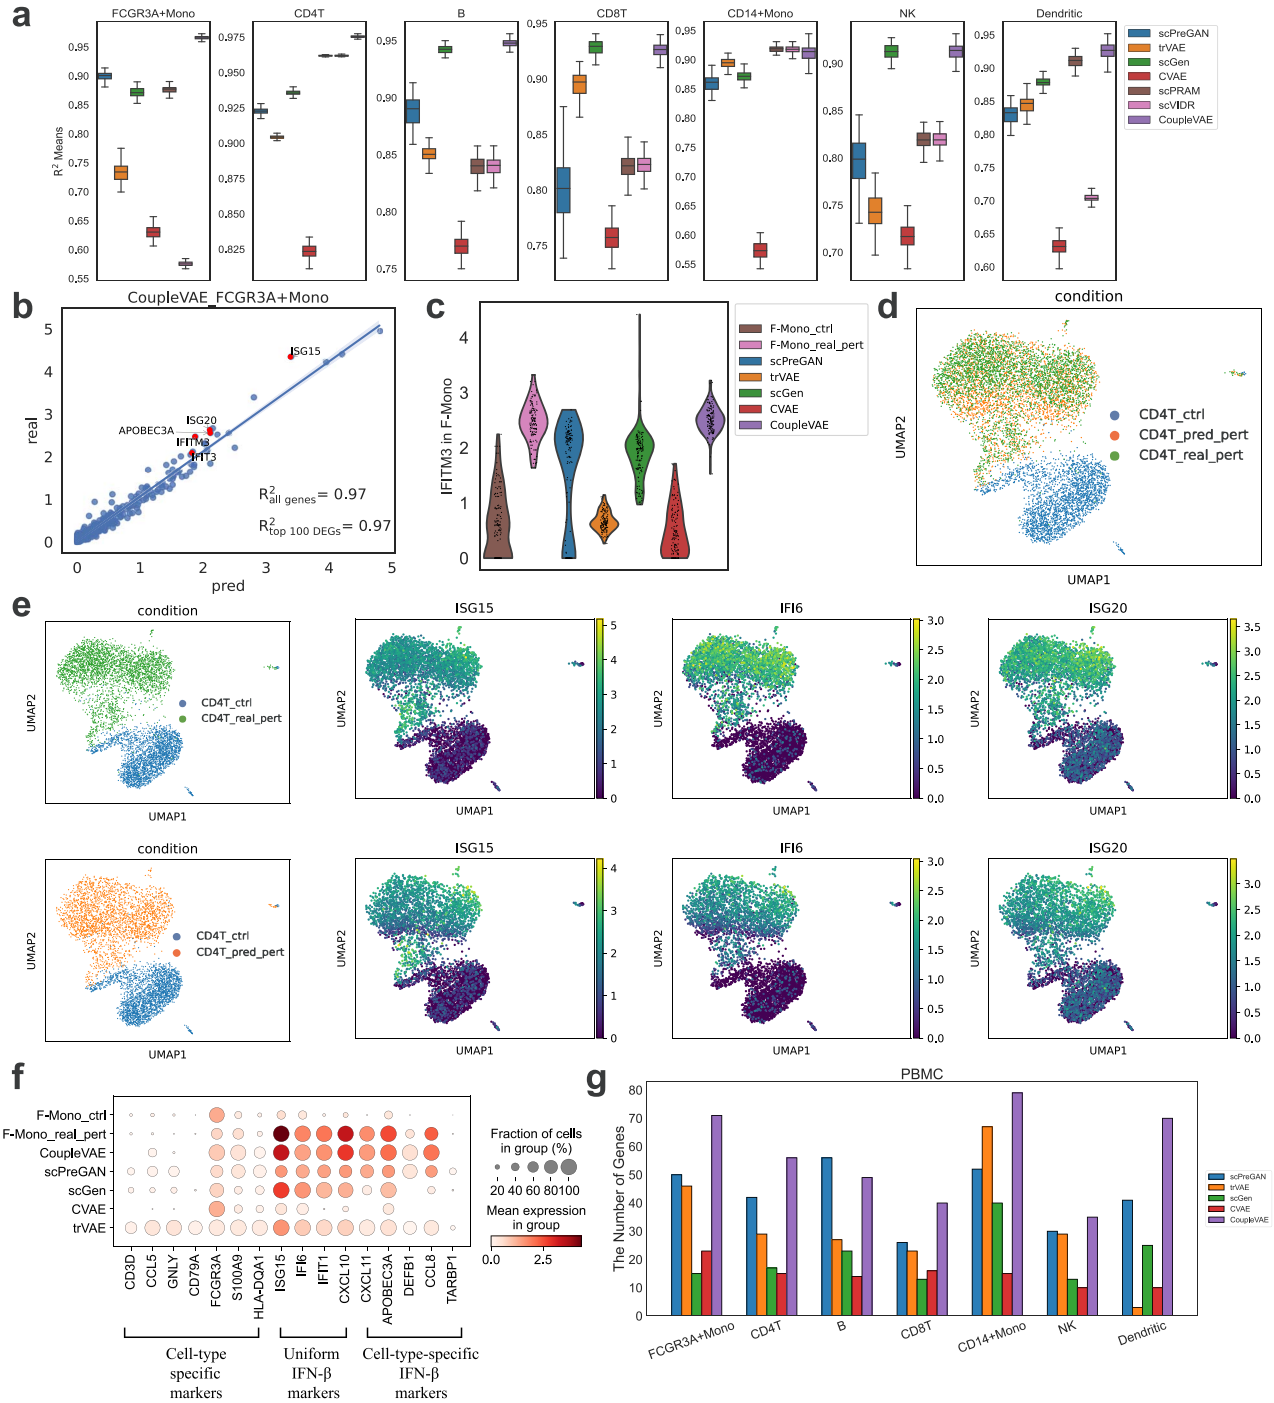

Figure 3. Single-cell response prediction for $IFN-\beta$ treatment. (a) Boxplot of correlations between the predicted gene expressions and real ones. R^2 denotes the square of the PCC. (b) The mean value of 6998 gene expressions in 137 FCGR3A+Mono predicted using CoupleVAE compared to the real expressions. The genes marked as significant are top 5 DEGs under real perturbed conditions. (c) Distribution of the most strongly changed gene IFITM3 in FCGR3A+Mono with $IFN-\beta$ perturbation in control, real and predicted stimulated FCGR3A+Mono cells of CoupleVAE compared to other prediction models. (d) UMAP visualization of controlled data, real perturbed data and predicted data using CoupleVAE for CD4T cells. (e) UMAP visualization of the expression of three genes, ISG15, IFI6 and ISG20 in controlled data, real perturbed data and predicted data using CoupleVAE for CD4T cells. The left figure shows the pattern of all genes under different conditions. The upper line shows real perturbational data and the lower line shows predicted. (f) Comparison of gene expression for control, real and predicted perturbation on the FCGR3A+Mono(F-Mono) from PBMC for five methods. (g) Comparison of DEG identification with five models in PBMC. The number of genes represents the number of common genes between ground truth and predicted value in the top 100 DEGs.

highest number of true DEGs, which implies that the predicted perturbational scRNA-Seq captures the regulation information the best.

By predicting changes in gene expression, CoupleVAE helps us to understand changes in cellular function in response to different perturbations. For example, up- or down-regulation of specific genes can reveal the mechanisms of cells in response to stimuli or infections. This predictive capability is crucial for deciphering the complex regulatory networks that govern cellular behavior. It allows researchers to identify key regulatory genes and pathways that are modulated under various conditions, providing insights into the underlying biological processes.

CoupleVAE accurately predicts cross-species perturbation responses

CoupleVAE could also be applied to predict cross-species perturbation responses, either under the same or different conditions. We collect the scRNA-Seq data [39] that study the evolution of innate immunity programs of mononuclear phagocytes within different species, including pigs, rabbits, mice and rats. The primary bone marrow-derived mononuclear phagocytes are stimulated using lipopolysaccharide (LPS) for 6 hours. The data include two groups: control phagocytes and phagocytes perturbed by LPS for 6h.

We refer to the task of predicting perturbations across species as out-of-distribution (o.o.d.), where unlike the previous two datasets, we expect different species respond to the stimulation differently. The cross-species perturbation prediction can be formulated as two scenarios. One is to train the prediction model using data from all species to predict the response of one of these species (Fig. 4a), and the other is to train the model using one specific species to predict the perturbation response of another species (Fig. 4e). Corresponding to these two scenarios, we set up two different training strategies: condition-CoupleVAE and species-CoupleVAE, where condition-CoupleVAE explains the differences of stimulation responses in all the considered species and species-CoupleVAE explains the transformations across different species.

Condition-CoupleVAE trains the model using controlled and perturbed data for single cells from different species as shown in Figure 4a, and aims to predict the responses to perturbation for one specific species based on its own controlled data. For example, mouse controlled cells can be used as test data to predict the responses to LPS of mouse with the trained model. Figure 4b shows the results for different models in predicting the response of the four species to the perturbation. CoupleVAE achieves the highest R^2 in all the four species. To further explore whether the predicted data from each method match the distribution of the real data, we combine the predicted data from the four species with the real data, and visualize the distribution using UMAP. As shown in Figure 4c, the predicted data do not get close to the real perturbational data using the four comparison methods, and the predicted data of trVAE and CVAE even overlap with the controlled data. In contrast, the predicted data of CoupleVAE is closer to the real data. This is due to the fact that CoupleVAE interconverts the distributions of the data under the two conditions in the latent space, which makes the generated data closer to the real data. More detailed explanations of the data distribution in the latent space are provided in the Supplementary Figure S13 and Supplementary Note 5. We then select six DEGs that are differentially expressed before and after the mouse is perturbed by LPS, and study their predicted distribution using five methods (Fig. 4d). For gene *Ccl3, Sod2, Cflar*, CoupleVAE predicts the entire distribution well. For some other genes (*Cd44, Srgn*), their distributions are closer to the true distribution compared to other methods. The predicted gene distributions associated with LPS perturbations for the other three species are shown in the Supplementary Figure S14.

Species-CoupleVAE is to predict the perturbation responses of one species based on the perturbational data from another species and the controlled data of both species. As shown in Figure 4e, the training data are rabbit and mouse cells in the control state, and the model predicts mouse cells perturbed by LPS from rabbit cells perturbed by LPS. Figure 4f demonstrates the R^2 for the perturbational mouse cells predicted by different models from the other three species. CoupleVAE gives much higher R^2 than all the other four methods. Figure 4g shows UMAP plots of the predicted perturbational data of mouse cells from rabbit cells. Each plot contains rabbit and mouse cells in the control and LPS states, as well as the predicted expression for mouse cells perturbed by LPS for the five methods. Among them, the predictions of scGen and trVAE are close to the mouse cells in the control state. For CVAE, its predicted data partially overlap with rabbit cells in the control state. As for scPreGAN, the predicted data still partially overlap with the mouse cells in the control state. The predicted data of CoupleVAE do not get close to rabbit and mouse cells in control state, and they are closer to the real data distribution. Figure 4h is similar to Figure 4d, where the same genes are picked to study the prediction of gene distribution by species-CoupleVAE (training data are in Fig. 4e). CoupleVAE predicts the distribution of the above mentioned genes better.

CoupleVAE is able to make predictions of perturbation responses across species, which is important for investigating mechanisms of differential responses between species. For example, the response of different species to the same drug can help determine the drug's generalizability or specificity. Understanding these differential responses is crucial for translational research, as it can reveal species-specific pathways and genetic factors that influence drug efficacy and safety. This cross-species predictive capability enables researchers to identify conserved and divergent biological mechanisms, which can inform the development of more effective and targeted therapies. By leveraging CoupleVAE's cross-species prediction power, scientists can gain a more comprehensive understanding of how different organisms respond to perturbations, leading to better models of human diseases and more accurate predictions of treatment outcomes.

Discussion

Perturbations in biology reveal how cells respond to external stimuli, such as drug therapy, gene knockout, and viral infection. Studying these responses enhances our understanding of cell function and biological dynamics. This work introduces CoupleVAE, a novel deep generative model for predicting single-cell responses to perturbations using scRNA-seq data. CoupleVAE consists of two encoders, a coupler and two decoders, which do not share parameters between them. It transforms latent variables to ensure predicted cells closely match real perturbational cells. By constructing the transformations between the latent vectors for different conditions, it effectively learns the effects of perturbations in a low-dimensional gene expression representation. Experiments on COVID-19, PBMC and LPS6 datasets show CoupleVAE's effectiveness in mean gene expression prediction, DEG detection, and response marker expression prediction.

Although CoupleVAE works well for most cases, it is far from perfect, for example, when predicting NK cells in the PBMC dataset and the COVID-19 dataset. This may be due to the small

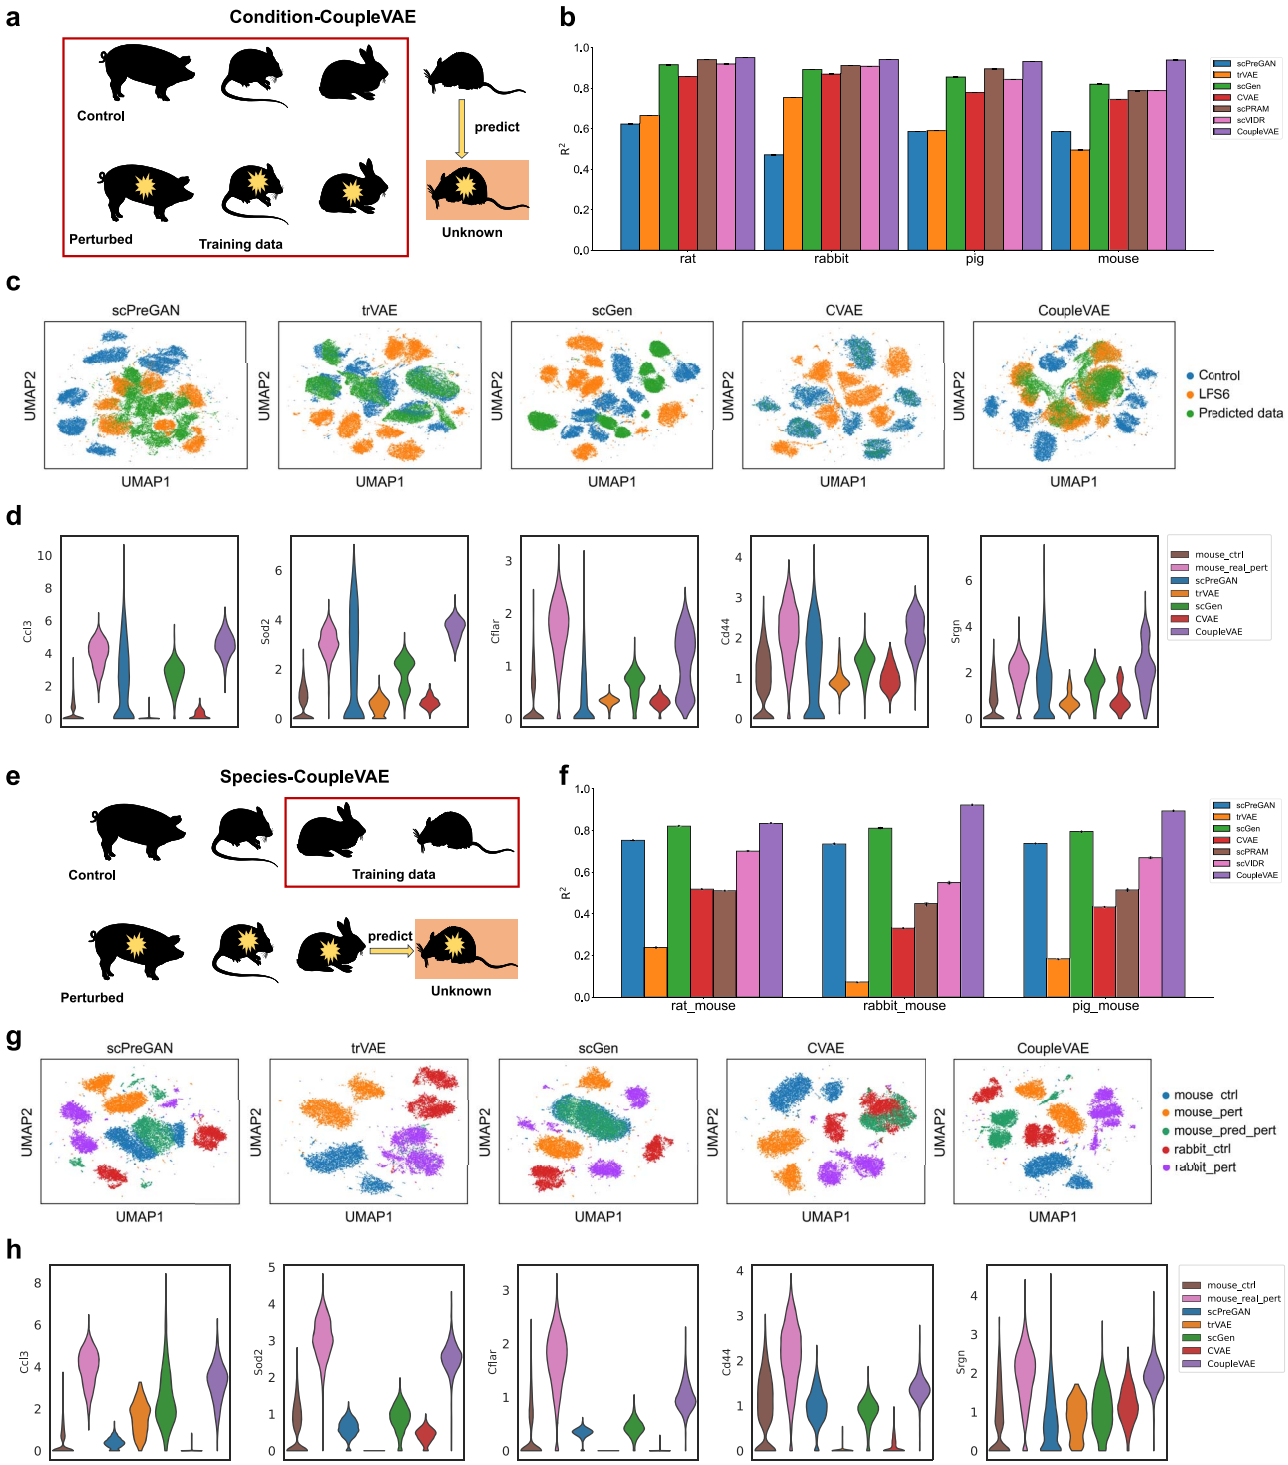

Figure 4. Cross-species single-cell response prediction. (a) The LPS dataset contains four species: rat, rabbit, pig, and mouse. The first row represents the controlled species, and the second row is the species perturbed with LPS for 6 hours. Experimental setup of condition-CoupleVAE: data in box are used for training to predict unseen LPS perturbed mouse cells from mouse controlled cells. (b) Prediction comparison for rat, rabbit, pig, and mouse cells perturbed by LPS from the corresponding cells in a healthy state, respectively. R^2 for the average expression of all genes are compared between real and predicted cells from different models. (c) UMAP visualization of controlled data, real perturbational data and predicted perturbed data using five methods for all species. (d) Distribution of LPS-responsive genes (*Ccl3*, *Sod2*, *Cflar*, *Cd44*, *Srgn* [39]). (e) Experimental setup of species-CoupleVAE: data in red box are used for training to predict unseen LPS perturbed mouse cells from LPS perturbed rat cells. (f) Prediction of mouse cells' responses to LPS from rat, rabbit, and pig cells perturbed by LPS, respectively. (g) UMAP visualization of controlled data, real perturbational data for mouse and rabbit, and predicted mouse perturbational data. (h) Distribution of LPS-responsive genes(*Ccl3*, *Sod2*, *Cflar*, *Cd44*, *Srgn* [39]). Vertical axis: expression distribution for genes. Horizontal axis: control, real and predicted distribution by different models. Part figures (a) and (e) were created using Servier Medical Art templates, which are licensed under a Creative Commons Attribution 3.0 Unported License; https://smart.servier.com.

sample size of NK cells, or the complex role of NK cells in the immune system. For example, it has been shown that the severity of COVID-19 disease is associated with reduced NK cell numbers, NK cell exhaustion and lack of certain mature effective NK cell phenotypes [42, 43]. The mapping of interconversions between conditions in latent space will be designed to be more complex in order to capture complex perturbation response information. In addition, cross-species perturbation prediction is an interesting direction of research, since there are still many species with differential response mechanisms among them that are not yet known. This brings a great challenge to the prediction problem. Recent advancements in perturbation modeling based on gene regulatory networks (GRNs), such as GenKI [44], offer new avenues for exploration. GenKI utilizes a variational graph autoencoder to learn latent relationships between genes through GRNs. In future work, we can use advanced network-based algorithms such as SLNMF [45], MCGL [46], scGDC [47], and sLMIC [48] to process GRNs. By leveraging these algorithms, we can understand deeper insights into the effects of perturbations on biological mechanisms.

Conclusion

In summary, the application of CoupleVAE in biological research demonstrates the great potential of deep learning models in predicting and explaining complex biological phenomena, helping us to better understand the mechanisms of cellular responses to different perturbations, and thus advancing disease research and drug discovery.

Key Points

- Coupled variational autoencoders (CoupleVAE) is a deep generative model based on coupled variational autoencoders designed to predict single-cell RNA sequencing responses to perturbations. A theoretical variational lower bound on the log-likelihood of the joint distribution was proved in the work.
- CoupleVAE encodes gene expressions of controlled and perturbed cells into latent variables, which are then translated through a coupling process with conditional distributions. It generates gene expressions from these translated variables via two decoders.
- The experiments were conducted on three datasets: COVID-19, IFN-β stimulation, and cross-species lipopolysaccharide perturbation. The results demonstrated that the CoupleVAE model surpasses existing comparative models in effectively predicting single-cell RNA-seq data for perturbed cells, achieving superior accuracies.

Supplementary data

Supplementary data is available at *Briefings in Bioinformatics* online.

Conflict of interest: No competing interest is declared.

Funding

This work was supported by National Natural Science Foundation of China projects under Grant Nos 12222115 and 92470106 to L.L., Grant No. 12471350 to S.Z and Science and Technology Commission of Shanghai Municipality under Grant No. 23JC1401000 to S.Z.

Data availability

All datasets used in this study are publicly available. These datasets include COVID-19 (https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE145926), PBMC (https://www.ncbi.nlm.nih.gov/geo/query/acc.cgi?acc=GSE96583), and LPS6 (https://www.ebi.ac.uk/arrayexpress/experiments/E-MTAB-6754/?query=tzachi+hagai). They are downloadable at https://github.com/LiminLi-xjtu/CoupleVAE.

Code availability

The software is available at https://github.com/LiminLi-xjtu/CoupleVAE.

References

1. Lu J, Ahmad R, Nguyen T. *et al.* Heterogeneity and transcriptome changes of human cd8+ t cells across nine decades of life. *Nat Commun* 2022;**13**:1–13. https://doi.org/10.1038/s41467-022-32869-x
2. Zhang S, Cui Y, Ma X. *et al.* Single-cell transcriptomics identifies divergent developmental lineage trajectories during human pituitary development. *Nat Commun* 2020;**11**:5275. https://doi.org/10.1038/s41467-020-19012-4
3. Shi J, Li T, Chen L. *et al.* Quantifying pluripotency landscape of cell differentiation from scrna-seq data by continuous birth-death process. *PLoS Comput Biol* 2019;**15**:e1007488. https://doi.org/10.1371/journal.pcbi.1007488
4. Sha Y, Qiu Y, Zhou P. *et al.* Reconstructing growth and dynamic trajectories from single-cell transcriptomics data. *Nat Mach Intell* 2023;**6**:25–39. https://doi.org/10.1038/s42256-023-00763-w
5. Jin K, Schnell D, Li G. *et al.* Celldrift: Inferring perturbation responses in temporally sampled single-cell data. *Brief Bioinform* 2022;**23**:bbac324. https://doi.org/10.1093/bib/bbac324
6. Xin Z, Lin M, Hao Z. *et al.* The immune landscape of human thymic epithelial tumors. *Nat Commun* 2022;**13**:1–16. https://doi.org/10.1038/s41467-022-33170-7
7. Hetzel L, Böhm S, Kilbertus N. *et al.* Predicting cellular responses to novel drug perturbations at a single-cell resolution. *Adv Neural Inform Process Syst* 2022;**35**:26711–22.
8. Dixit A, Parnas O, Li B. *et al.* Perturb-seq: Dissecting molecular circuits with scalable single-cell rna profiling of pooled genetic screens. *Cell* 2016;**167**:1853–1866.e17. https://doi.org/10.1016/j.cell.2016.11.038
9. Adamson B, Norman TM, Jost M. *et al.* A multiplexed single-cell crispr screening platform enables systematic dissection of the unfolded protein response. *Cell* 2016;**167**:1867–82. https://doi.org/10.1016/j.cell.2016.11.048
10. Datlinger P, Rendeiro AF, Schmidl C. *et al.* Pooled crispr screening with single-cell transcriptome readout. *Nat Methods* 2017;**14**:297–301. https://doi.org/10.1038/nmeth.4177
11. Ubah OC, Lake EW, Gunaratne GS. *et al.* Mechanisms of sars-cov-2 neutralization by shark variable new antigen receptors elucidated through X-ray crystallography. *Nat Commun* 2021;**12**:1–12. https://doi.org/10.1038/s41467-021-27611-y
12. Ji Y, Lotfollahi M, Wolf FA. *et al.* Machine learning for perturbational single-cell omics. *Cell Syst* 2021;**12**:522–37. https://doi.org/10.1016/j.cels.2021.05.016

13. Srivatsan SR, McFaline-Figueroa JL, Ramani V. *et al.* Massively multiplex chemical transcriptomics at single-cell resolution. *Science* 2020;**367**:45–51. https://doi.org/10.1126/science.aax6234
14. Russkikh N, Antonets D, Shtokalo D. *et al.* Style transfer with variational autoencoders is a promising approach to rna-seq data harmonization and analysis. *Bioinformatics* 2020;**36**:5076–85. https://doi.org/10.1093/bioinformatics/btaa624
15. Fröhlich F, Kessler T, Weindl D. *et al.* Efficient parameter estimation enables the prediction of drug response using a mechanistic pan-cancer pathway model. *Cell Syst* 2018;**7**:567–579.e6. https://doi.org/10.1016/j.cels.2018.10.013
16. Choi K, Hellerstein J, Wiley HS. *et al.* Inferring reaction networks using perturbation data. *BioRxiv* 2018;351767.
17. Lopez R, Regier J, Cole MB. *et al.* Deep generative modeling for single-cell transcriptomics. *Nat Methods* 2018;**15**:1053–8. https://doi.org/10.1038/s41592-018-0229-2
18. Eraslan G, Simon LM, Mircea M. *et al.* Single-cell rna-seq denoising using a deep count autoencoder. *Nat Commun* 2019;**10**:390. https://doi.org/10.1038/s41467-018-07931-2
19. Ghahramani A, Watt FM, Luscombe NM. Generative adversarial networks simulate gene expression and predict perturbations in single cells. *BioRxiv* 2018;262501.
20. Wei X, Dong J, Wang F. Scpregan, a deep generative model for predicting the response of single-cell expression to perturbation. *Bioinformatics* 2022;**38**:3377–84. https://doi.org/10.1093/bioinformatics/btac357
21. Lotfollahi M, Wolf FA, Theis FJ. Scgen predicts single-cell perturbation responses. *Nat Methods* 2019;**16**:715–21. https://doi.org/10.1038/s41592-019-0494-8
22. Lotfollahi M, Naghipourfar M, Theis FJ. *et al.* Conditional out-of-distribution generation for unpaired data using transfer vae. *Bioinformatics* 2020;**36**:i610–7. https://doi.org/10.1093/bioinformatics/btaa800
23. Sohn K, Lee H, Yan X. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems* 2015;**28**.
24. Yu H, Welch JD. Michigan: Sampling from disentangled representations of single-cell data using generative adversarial networks. *Genome Biol* 2021;**22**:1–26. https://doi.org/10.1186/s13059-021-02373-4
25. Jiang Q, Chen S, Chen X. *et al.* Scpram accurately predicts single-cell gene expression perturbation response based on attention mechanism. *Bioinformatics* 2024;**40**:btae265. https://doi.org/10.1093/bioinformatics/btae265
26. Kana O, Nault R, Filipovic D. *et al.* Generative modeling of single-cell gene expression for dose-dependent chemical perturbations. *Patterns* 2023;**4**:100817. https://doi.org/10.1016/j.patter.2023.100817
27. Tang Z, Zhou M, Zhang K. *et al.* Scperb: Predict single-cell perturbation via style transfer-based variational autoencoder. *J Adv Res* 2024. https://doi.org/10.1016/j.jare.2024.10.035
28. Kingma DP, Welling M. Auto-encoding variational bayes arXiv preprint arXiv:1312.6114. 2013.
29. Doersch C. Tutorial on variational autoencoders arXiv preprint arXiv:1606.05908. 2016.
30. Kingma DP, Ba J. Adam: A method for stochastic optimization arXiv preprint arXiv:1412.6980. 2014.
31. Lotfollahi M, Naghipourfar M, Luecken MD. *et al.* Mapping single-cell data to reference atlases by transfer learning. *Nat Biotechnol* 2022;**40**:121–30. https://doi.org/10.1038/s41587-021-01001-7
32. Bunne C, Stark SG, Gut G. *et al.* Learning single-cell perturbation responses using neural optimal transport. *Nat Methods* 2023;**20**:1759–68.
33. Wolf FA, Angerer P, Theis FJ. Scanpy: Large-scale single-cell gene expression data analysis. *Genome Biol* 2018;**19**:1–5. https://doi.org/10.1186/s13059-017-1382-0
34. Wilk AJ, Rustagi A, Zhao NQ. *et al.* A single-cell atlas of the peripheral immune response in patients with severe covid-19. *Nat Med* 2020;**26**:1070–6. https://doi.org/10.1038/s41591-020-0944-y
35. Liao M, Liu Y, Yuan J. *et al.* Single-cell landscape of bronchoalveolar immune cells in patients with covid-19. *Nat Med* 2020;**26**:842–4. https://doi.org/10.1038/s41591-020-0901-9
36. Silvin A, Chapuis N, Dunsmore G. *et al.* Elevated calprotectin and abnormal myeloid cell subsets discriminate severe from mild Covid-19. *Cell* 2020;**182**:1401–1418.e18. https://doi.org/10.1016/j.cell.2020.08.002
37. Schulte-Schrepping J, Reusch N, Paclik D. *et al.* Severe covid-19 is marked by a dysregulated myeloid cell compartment. *Cell* 2020;**182**:1419–1440.e23. https://doi.org/10.1016/j.cell.2020.08.001
38. Wen W, Su W, Tang H. *et al.* Immune cell profiling of covid-19 patients in the recovery stage by single-cell sequencing. *Cell Discovery* 2020;**6**:31. https://doi.org/10.1038/s41421-020-0168-9
39. Hagai T, Chen X, Miragaia RJ. *et al.* Gene expression variability across cells and species shapes innate immunity. *Nature* 2018;**563**:197–202. https://doi.org/10.1038/s41586-018-0657-2
40. McInnes L, Healy J, Melville J. Umap: Uniform manifold approximation and projection for dimension reduction arXiv preprint arXiv:1802.03426. 2018.
41. Butler A, Hoffman P, Smibert P. *et al.* Integrating single-cell transcriptomic data across different conditions, technologies, and species. *Nat Biotechnol* 2018;**36**:411–20. https://doi.org/10.1038/nbt.4096
42. Ankerhold J, Giese S, Kolb P. *et al.* Circulating multimeric immune complexes contribute to immunopathology in Covid-19. *Nat Commun* 2022;**13**:1–15. https://doi.org/10.1038/s41467-022-32867-z
43. Zhang W, Chua BY, Selva KJ. *et al.* Sars-cov-2 infection results in immune responses in the respiratory tract and peripheral blood that suggest mechanisms of disease severity. *Nat Commun* 2022;**13**:1–18.
44. Yang Y, Li G, Zhong Y. *et al.* Gene knockout inference with variational graph autoencoder learning single-cell gene regulatory networks. *Nucleic Acids Res* 2023;**51**:6578–92. https://doi.org/10.1093/nar/gkad450
45. Wu W, Ma X. Network-based structural learning nonnegative matrix factorization algorithm for clustering of scrna-seq data. *IEEE/ACM Trans Comput Biol Bioinform* 2022;**20**:566–75. https://doi.org/10.1109/TCBB.2022.3161131
46. Wu W, Zhang W, Hou W. *et al.* Multi-view clustering with graph learning for scrna-seq data. *IEEE/ACM Trans Comput Biol Bioinform* 2023;**20**:3535–46. https://doi.org/10.1109/TCBB.2023.3298334
47. Wang H, Ma X. Learning deep features and topological structure of cells for clustering of scrna-sequencing data. *Brief Bioinform* 2022;**23**:bbac068. https://doi.org/10.1093/bib/bbac068
48. Wang H, Liu Z, Ma X. Learning consistency and specificity of cells from single-cell multi-omic data. *IEEE J Biomed Health Inform* 2024;**28**:3134–45. https://doi.org/10.1109/JBHI.2024.3370868

Briefings in Bioinformatics, 2025, **26(2)**, bbaf126
https://doi.org/10.1093/bib/bbaf126
Problem Solving Protocol