scPerturb: Information Resource for Harmonized Single-Cell Perturbation Data

Anonymous Author(s)
Affiliation
Address
email

Recent biotechnological advances led to growing numbers of single-cell studies, which reveal molecular and phenotypic responses to large numbers of perturbations. However, analysis across diverse datasets is typically hampered by differences in format, naming conventions, data filtering and normalization. To facilitate development and benchmarking of computational methods in systems biology, we collect a set of 44 publicly available single-cell perturbation-response datasets with molecular readouts, including RNA, proteins and chromatin accessibility (Figure Panel A). We apply uniform pre-processing and quality control pipelines and harmonize feature annotations. The resulting information resource enables efficient development and testing of computational analysis methods, and facilitates direct comparison and integration across datasets. 32 RNA datasets in this resource were perturbed using CRISPR and 9 were perturbed with drugs (Figure Panel B). We also include three scATAC datasets, as well as three CITE-seq datasets with protein and RNA counts separately downloadable. For each scRNA-seq dataset we supply count matrices, where each cell has a perturbation annotation, quality control metrics including gene counts and mitochondrial read percentage. Quality control plots for each dataset are also available on the scPerturb website. Notably, more than 8000 CRISPR perturbations are shared across multiple datasets. We anticipate this data resource being useful for developing machine learning models for perturbation responses across datasets and other tasks.

To compare and evaluate perturbations within each dataset, we demonstrate the application of a distance measure between high-dimensional point clouds to quantify perturbation similarity and strength. The E-distance (short for energy distance, [1,2]) contextualizes the notion that two such point clouds are distinguishable if they are far apart compared to the width of both clouds, providing intuition about the signal-to-noise ratio in a dataset.

Let $x_1, ..., x_N \in \mathbb{R}^d$ and $y_1, ..., y_N \in \mathbb{R}^d$ be samples from two distributions X, Y, corresponding to two sets of N and M cells, respectively.

We define

$$\delta_{XY} = \frac{1}{NM} \sum_{i=1}^{M} \sum_{j=1}^{N} ||x_i - y_j||$$

$$\sigma_X = \frac{1}{N(N-1)} \sum_{i=1}^{N} \sum_{j=1}^{N} ||x_i - x_j||$$

and σ_Y defined accordingly. Intuitively, δ_{XY} is the mean distance between cells from the two distributions, while σ_X describes the mean distance between a cell from X to another cell from X.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

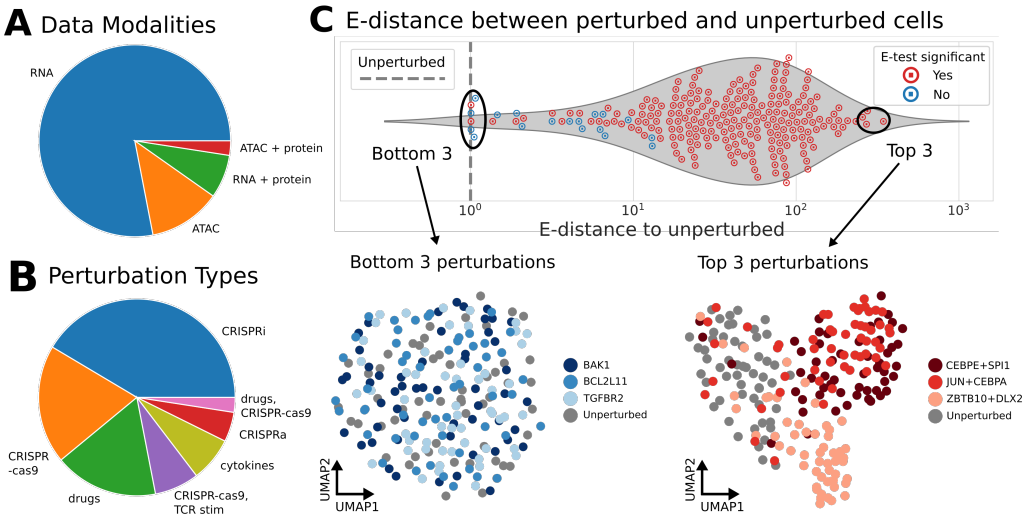

Figure 1: **Diverse single-cell perturbation-response datasets profiled with E-distances**. (A,B) The majority of included datasets are scRNA-seq data with CRISPR (DNA cut, inhibition or activation) perturbations using cell lines derived from various cancers. The studies performed on cells from primary tissues generally use drug perturbations. (C) E-distances between perturbed and unperturbed cells in NormanWeissman2019, log+1 scaled. Color indicates significant difference to unperturbed cells according to E-test. UMAP embeddings of cells from top and bottom perturbations w.r.t. E-distance to unperturbed exemplify the connection between E-distance and similarity of point clouds

The E-distance between X and Y is defined as:

$$E(X,Y) := 2\delta_{XY} = \sigma_X - \sigma_Y$$

We computed the E-distance in PCA space between each group of perturbed cells and the unperturbed cells for our collected datasets. This results in a distribution of E-distances for each perturbation in a dataset, exemplarily shown for one selected dataset (Figure Panel C, upper portion). The more a perturbation is on the right side of this distribution, the more distinctive its effect on the cells is captured in the data. We also used the E-distance as a test statistic for a Monte Carlo permutation test ("E-test"), robustly testing whether the distinctiveness is statistically significant. To visualize what a high or low E-distance to unperturbed cells means, we select the three closest / furthest perturbations in the example dataset and calculate a UMAP embedding for the cells of these perturbations jointly with the unperturbed cells (Figure Panel C, lower portion). The top three furthest perturbations are easily distinguishable from the gray unperturbed cells, while the bottom three closest perturbations form a single, uniform cloud

This work provides an information resource and guide for researchers working with single-cell perturbation data and highlights conceptual considerations for new experiments. The scPerturb data collections is publically available and is described in further detail in an associated preprint.

References

[1] Replogle JM, Saunders RA, Pogson AN, Hussmann JA, Lenail A, Guna A, et al. Mapping information-rich genotype-phenotype landscapes with genome-scale Perturb-seq. Cell. 2022;185:2559-2575.e28.

[2] Rizzo ML, Székely GJ. Energy distance. WIREs Comput Stat. 2016;8:27–38.