

Learning Single-Cell Perturbation Responses using Neural Optimal Transport

Charlotte Bunne

Department of Computer Science, ETH Zurich <https://orcid.org/0000-0003-1431-103X>

Stefan G. Stark

Department of Computer Science, ETH Zurich <https://orcid.org/0000-0003-2478-9512>

Gabriele Gut

Department of Molecular Life Sciences, University of Zurich, Switzerland <https://orcid.org/0000-0001-8991-0040>

Jacobo Sarabia del Castillo

Department of Molecular Life Sciences, University of Zurich, Switzerland <https://orcid.org/0000-0001-9656-1416>

Mitchell Levesque

Department of Dermatology, University of Zurich Hospital, University of Zurich, Switzerland
<https://orcid.org/0000-0001-9656-1416>

Kjong-Van Lehmann

Department of Computer Science, ETH Zurich <https://orcid.org/0000-0002-1936-298X>

Lucas Pelkmans

Department of Molecular Life Sciences, University of Zurich, Switzerland <https://orcid.org/0000-0002-6754-9730>

Andreas Krause

Department of Computer Science, ETH Zurich <https://orcid.org/0000-0001-7260-9673>

Gunnar Rätsch (✉ Gunnar.Ratsch@ratschlab.org)

Department of Computer Science, ETH Zurich <https://orcid.org/0000-0001-5486-8532>

Article

Keywords:

Posted Date: July 7th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1805107/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Learning Single-Cell Perturbation Responses using Neural Optimal Transport

Charlotte Bunne,^{1,2,*} Stefan G. Stark,^{1,2,3,4,*} Gabriele Gut,^{5,*}
Jacobo Sarabia del Castillo,⁵ Mitch Levesque,⁶ Kjong-Van Lehmann,^{1,2,3,4,†}
Lucas Pelkmans,^{5,†} Andreas Krause,^{1,2,†} Gunnar Rätsch^{1,2,3,4,7,†}

¹ Department of Computer Science, ETH Zurich, Switzerland;

² AI Center, ETH Zurich, Switzerland;

³ Medical Informatics Unit, University of Zurich Hospital, Switzerland;

⁴ Swiss Institute of Bioinformatics, Switzerland;

⁵ Department of Molecular Life Sciences, University of Zurich, Switzerland;

⁶ Department of Dermatology, University of Zurich Hospital, University of Zurich, Switzerland;

⁷ Department of Biology, ETH Zurich, Switzerland.

Abstract

Understanding and predicting molecular responses in single cells upon chemical, genetic, or mechanical perturbations is a core question in biology. Obtaining single-cell measurements typically requires the cells to be destroyed. This makes learning heterogeneous perturbation responses challenging as we only observe *unpaired* distributions of perturbed or non-perturbed cells. Here we leverage the theory of optimal transport and the recent advent of convex neural architectures to present CELLOT, a framework for learning the response of individual cells to a given perturbation by coupling these unpaired distributions. We achieve this alignment with a learned transport map that allows us to infer the treatment responses of unseen untreated cells. CELLOT outperforms current methods at predicting single-cell drug responses, as profiled by scRNA-seq and a multiplexed protein imaging technology. Further, we illustrate that CELLOT generalizes well on unseen settings by (a) predicting the scRNA-seq responses of holdout lupus patients exposed to IFN- β , and (b) modeling the hematopoietic developmental trajectories of different subpopulations. We expect CELLOT to lay the grounds for delineating the causes of heterogeneous single-cell responses to perturbations and predicting patient-specific drug response landscapes instead of population averages.

Characterizing and modeling perturbation responses at the single-cell level from non-time-resolved data remains one of biology’s grand challenges. It finds applications in predicting cellular reactions to environmental stress or a patient’s response to drug treatments. Accurate inference of perturbation responses at the single cell level allows us, for instance, to understand how and why individual tumor cells evade cancer therapies (19). More generally, it deepens the mechanistic understanding of the molecular machinery that determines the respective responses to perturbations. Single-cell responses to genetic or chemical perturbations are highly heterogeneous (38) due to multiple factors, including pre-existing variability in the abundance and subcellular organization of mRNA and proteins (4, 5, 24, 52), cellular states (35), and the cellular microenvironment (54). To effectively predict the drug response of each cell in a population, whether derived from tissue culture or as primary cells from a patient biopsy, it is thus crucial to incorporate this heterogeneous multivariate subpopulation structure into the analysis.

A fundamental difficulty in learning perturbation responses is that cells are usually fixed and stained or chemically destroyed to obtain these measurements. Hence, it is only possible to measure the same cells before or after a perturbation is applied. Therefore, while we do

*These authors contributed equally.

†To whom correspondence should be addressed: kjlehmann@ukaachen.de, lucas.pelkmans@mls.uzh.ch, krausea@ethz.ch, Gunnar.Raetsch@inf.ethz.ch.

not have access to a set of *paired* control/perturbed single-cell observations, we do have access to *sets* of single-cell observations separately from control and perturbed cells, respectively. To subsequently match single cells between conditions and, at the same time, account for cellular heterogeneity is a highly complex pairing problem that is unsolved.

Here, we seek to learn a perturbation model that robustly describes the cellular dynamics upon intervention while still accounting for underlying variability across samples. Learning the responses on an existing patient cohort enables inference of treatment responses for new, i.e., previously unseen patients, assuming that we captured the heterogeneous drug reactions of patients during training. It is crucial, however, to not simply model average perturbation responses of a patient cohort, but to capture the specificities of a single patient through personalized treatment effect predictions.

Previous methods to approximate single-cell perturbation responses fall short of solving this highly complex *pairing* problem while, at the same time, accounting for cellular heterogeneity and the strong subpopulation structure of cell samples. Current state-of-the-art methods (39, 40, 68) predict perturbation responses via *linear shifts* in a learned latent space. While this can capture nonlinear cell-type-specific responses, the use of linear interpolations reduces the alignment problem to the possibly more challenging task of learning representations that are invariant to their perturbation status.

In this work, we introduce CELLOT, a novel approach that predicts perturbation responses of single cells by *directly* learning and uncovering couplings between control and perturbed cell states that explicitly account for heterogeneous subpopulation structures in multiplexed molecular readouts. Assuming perturbations incrementally alter molecular profiles of cells, such as gene expression or signaling activities, we learn these changes and alignments using optimal transportation theory (OT) (63). Optimal transport provides natural geometric and mathematical tools to manipulate probability distributions. It has found recent successes modeling cellular development processes (37, 51), albeit in a *non-parameterized* setting, which cannot make predictions on unseen cells, such as those from unseen samples, e.g., new patients.

Based on recent developments in neural optimal transport (41), CELLOT learns an optimal transport coupling for each perturbation in a fully parameterized and highly scalable manner. Instead of directly learning a transport function (28, 67, 46), CELLOT parameterizes a pair of dual potentials with convex neural networks. This choice induces an important theory-motivated inductive bias essential to model stability (41).

We demonstrate CELLOT’s effectiveness by (i) learning single-cell marker responses to different cancer drugs in melanoma cell lines, (ii) predicting single-cell transcriptome responses in biopsies of patients with systemic lupus erythematosus, and (iii) modeling the transcriptome evolution of cell fates in hematopoiesis. Moreover, we benchmark CELLOT against current state-of-the-art methods on multiple tasks (39, 40).

Results

Modeling single-cell responses by directly learning couplings between cell states before and after perturbation

Small molecule drugs can have profound effects on the cellular phenotype by, for instance, altering signaling cascades. Most of these effects depend on the context in which the perturbation occurs. Given the heterogeneity of these contexts among single cells in cell populations and tissues, predicting cellular responses requires understanding the rules by which context shapes genome activity and its response to drugs. High-dimensional single-cell data measured via single-cell genomics or multiplexed imaging technologies can provide this contextual information but only return unpaired or unaligned observations of cell populations. Here, CELLOT allows us to utilize such unpaired data and enables learning cell state transitions upon perturbation.

In formal terms, we denote the unperturbed control population by ρ_c consisting of n cells x_i for $i = 1, \dots, n$. Upon perturbation k , the multivariate state of each cell x_i of the unperturbed population changes, which we observe as the perturbed population ρ_k (**Fig. 1a**). In order to

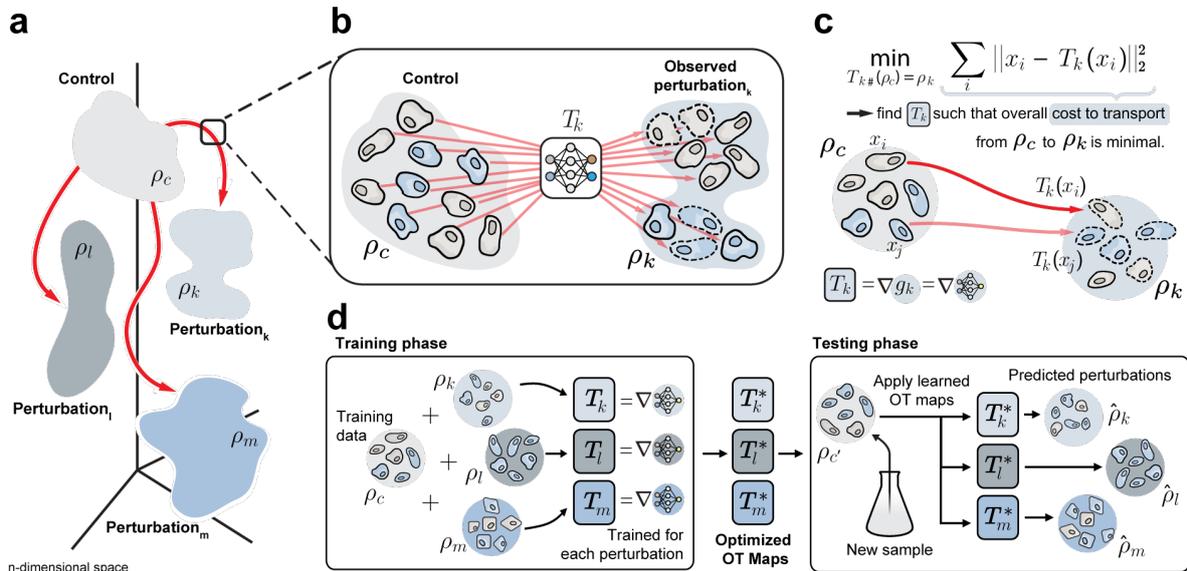


Figure 1: Overview of the CELLOT Model. **a)** Distributions of single cells were measured in either an untreated control state (ρ_c) or in one of several perturbed states ($\rho_k, \rho_l, \rho_m, \dots$). These distributions lie in a high-dimensional space of profiled features. **b)** For a perturbation k , we aim to model it with a function T_k that maps untreated cells in ρ_c to their treated counterparts in ρ_k . **c)** Lacking paired measurements, we assume that the perturbation transforms ρ_c into ρ_k under a principle of minimal effort. In particular, we learn T_k using optimal transport theory to directly estimate this distributional mapping as the gradient of the optimal transport dual potential ∇g_θ . **d)** OT maps are learned for all perturbations independently. Because these maps are fully parameterized, CELLOT can be trained, for example, on a set of initially provided samples to then make predictions on untreated cells originating from new, previously unseen samples.

understand the mode of action and effect of perturbations, we seek to learn the transition and alignment between populations ρ_c and ρ_k via parameterizing a map T_k (see **Fig. 1a-b**), which explains the transition of each cell from the unperturbed cell population ρ_c into their perturbed state ρ_k upon treatment k . Despite originating from different observations, map T_k determines for each cell x_i the most likely corresponding cell $T_k(x_i)$ in the perturbed population (**Fig. 1c**). Finding this map then not only allows us to model single-cell trajectories upon perturbation but also to predict the perturbed state of previously unseen control cells. As a result, we can forecast the outcome of a perturbation k by applying the learned map T_k to a new unperturbed population ρ_c' (**Fig. 1d**).

The optimal map T_k aligning the control and perturbed population, which we seek to find, should best describe the incremental changes in the multivariate profile of each cell after applying a perturbation k . Using optimal transportation theory (62, 50) to recover these maps and unveil single-cell reprogramming trajectories has been proposed as a strong modeling hypothesis in the domain of single-cell biology (51, 11, 16, 27, 37, 70). Optimal transport problems return the alignment between distributions ρ_c and ρ_k corresponding to the minimal overall cost between aligned molecular profiles, thus determining the most likely state of each cell upon perturbation (Fig. 1c). T_k is learned such that its image corresponds to ρ_k and mass is moved from ρ_c into ρ_k according to a principle of minimal effort. As directly parameterizing the optimal transport map T_k (28, 67, 46) is unstable (41, Table 1), we parameterize the convex potentials of the dual optimal transport problem f and g by convex neural networks (2) and recover the optimal map T_k using the gradient of a convex function g_k , i.e., ∇g_k (41). Supplementary Material Section S1 provides a more detailed review on optimal transport methods proposed for single-cell biology problems and how our approach deviates from previous methods.

To put CELLOT's performance in perspective, we benchmark it against current state-of-the-art methods based on autoencoders (40, 39), which attempt to add perturbation effects through the manipulation of a learned latent representation (review also in Supplementary Section S1). To further test the hypothesis of the optimal transport modeling prior, we compare the learned OT map ∇g_k for each perturbation k with naïve non-OT-based alignments.

CELLOT outperforms state-of-the-art methods in predicting cancer treatment effects

We apply CELLOT to predict the responses of cell populations to cancer treatments using a proteomic dataset consisting of two melanoma cell lines (M130219 and M130429) (47), profiled by 4i (24), and a scRNAseq dataset (56), which contain 34 and 9 different treatments, respectively. For more details on the datasets, see Online Methods. We benchmarked CELLOT against two autoencoder-based methods, SCGEN (40) and CAE (39). Due to the high dimensional nature of scRNA-seq data, we instead apply CELLOT on latent representations learned by an autoencoder. The marginal distributions for observed and predicted cell populations for two 4i treatments and two scRNAseq treatments are shown in **Fig. 2a, d**. Two features are selected for each perturbation and the complete set of marginals is shown in **Fig. S2, S3, S4, and S5**. While the autoencoder baselines tend to capture the mean of the treated cell population, they are less successful in matching all heterogeneous states of the perturbed population. More precisely, these methods do not capture higher moments of the perturbed population, as for example the standard deviation of phosphorylated AKT (pAKT) levels of Imatinib-treated 4i cells. Thus, these models tend to learn over-simplified perturbation effects, and are insufficient when aiming to understand heterogeneous rather than average cellular behaviors. CELLOT, on the other hand, is able to capture these higher moments, yielding accurate and nuanced predictions.

To quantify differences between the distributions of observed and predicted treated cell populations, we determine the maximum mean discrepancy (MMD) (22). Low values of MMD imply that all moments of two distributions are matched, and thus the entire distribution of perturbed cells is captured in fine detail, beyond the population average (see Online Methods for details). The MMDs between the predicted and observed populations for the selected perturbations are shown in **Fig. 2b, e**. For scRNA-seq data, MMD evaluations are computed using the top 50 marker genes. In addition to the autoencoder baselines, we include the trivial *identity* baseline that predicts treatment effects simply by returning the untreated states, as well as a theoretical lower bound, *observed*, comprising a different set of observed perturbed cells, thus only varying from the true predictions up to experimental noise. We find that CELLOT can approach the lower bound (*observed* setting), while the baseline methods often do not improve much over the *identity* setting.

Different evaluation metrics across all 34 4i therapies and 9 scRNAseq therapies are summarized in **Supplementary Tables S1 and S2**. Besides MMD, we additionally include the ℓ_2 (DS) metric that measures the distance between the observed and predicted *mean* drug effect, where these drug signatures are computed as the difference in means between the treated and untreated cell populations (see Online Methods). CELLOT outperforms the baselines in both metrics across all treatments, typically by one order of magnitude. We attribute the strong performance of CELLOT to its ability to learn a transport function that considers explicitly the data geometries of cell populations through the theory of optimal transport. This hypothesis is supported by the observation that the inter-feature correlation structure remains largely conserved between treated and untreated populations, thus depicting a setting where optimal transport approaches excel. For more information, see **Fig. S1**.

Finally, we computed UMAP (43) projections on a joint set of predicted and observed perturbed cells utilizing the full feature space, shown in **Fig. 2c, f**. We observe that the perturbed cell states inferred by CELLOT are well integrated with the observed perturbed cells. Again, both baselines do not recover the perturbed distribution in its entirety and thus the perturbed state of different subpopulations is not captured consistently.

CELLOT captures cell-to-cell variability to accurately predict drug responses of single cells

We co-cultured the aforementioned patient-derived melanoma cell lines (see Online Methods) at equal ratios and performed a boutique drug screen, during which we exposed cells 8h to a panel of 34 drugs. Using CELLOT, we learned transport maps for each of the drugs and predict the perturbed cell states for a perturbation from control (DMSO-treated) cells measured with

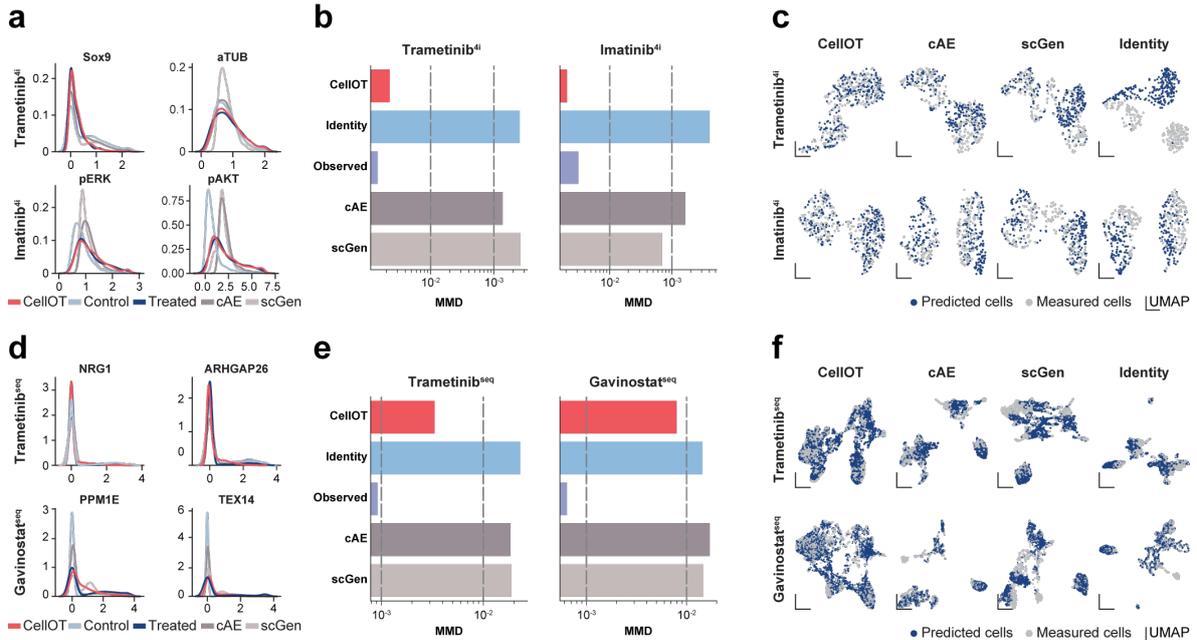


Figure 2: CELLOT outperforms current state-of-the-art methods on different data modalities. Marginal distribution of marker gene expression (x-axis) of cells profiled by **a)** 4i and **d)** scRNA. Observed control and treated states are shown in light and dark blue. CELLOT predictions are shown in red and baseline predictions are shown in gray. Joint UMAPs of observed treated cells and cells predicted by each model for **b)** 4i and **e)** scRNA data. Projections are computed on a joint set of cells, down-sampled such that the number of observed perturbed (gray) and predicted perturbed cells (blue) are equal. An identity coupling compares treated cells to untreated cells. We compare models based on the distributional distance MMD between observed perturbed and predicted perturbed cells, for **c)** 4i and **f)** scRNA data. The analysis is conducted for drugs Trametinib, Imatinib, and Gavinostat. 4i data was generated using cell lines M130219 and M130429 (Online Methods)

the 4i technology (**Fig. 3a**). Previous work (35) shows that phosphorylation levels of signaling kinases upon drug treatments are tightly linked to the cellular state. To assess whether this relationship was retained in predicted compared to observed perturbed cells, we analyzed the phosphorylation levels of extracellular signal-regulated kinases (pERK) using the transport maps learned by CELLOT on each drug. Using 750 predicted and 750 observed perturbed cells, we computed UMAP projections joint-wise from all features except pERK. **Fig. 3b** shows the predicted and observed population individually annotated with the respective pERK levels of each cell. We find the spatial organization of the two projections to look almost identical and that pERK levels had a highly comparable distribution across the cells of either class and all drug treatments. We furthermore supplemented this analysis by faithfully reconstructing unseen pERK levels of predicted cells from their local neighborhoods in multidimensional feature space (see **Fig. S6a, b** and **Online Methods**).

CELLOT enables cell state-aware drug profiling and disentangles subpopulation-specific drug effects

CELLOT allows us to profile the severity of drug perturbations on individual cellular states and understand how the cellular state of unperturbed cells determines drug response behavior. We can isolate the mode of action of each drug by computing the difference between the perturbed cells predicted by CELLOT and the untreated control cells, i.e., the *cost* of the optimal transport. A UMAP embedding of all cells color-coded by the treatment distinctly separates different treatments (**Fig. 3c** and **S6e**), all of which CELLOT is able to faithfully learn (**Tab. S2**). However, these distinct treatment embeddings are not present when accounting only for an average perturbation effect (**Fig. S6d**), indicating the importance of capturing the

cellular heterogeneity of drug responses.

Using Leiden clustering on the full feature set, we grouped unperturbed control cells in 12 cellular states (**Fig. 3d**, **Fig. S6g**, and Online Methods). Cellular states 1, 5, 6, 9, and 12 show high levels of MelA and no SOX9 and thus correspond to the melanocytic cell line M130429, whereas the SOX9⁺ and MelA⁻ states 2, 3, 4, 7, 8, 10, and 11 represent the mesenchymal cell line M130219 (see Online Methods). Overall, we find that M130429 cells have higher phosphorylation levels of the measured signaling kinases compared to M130219. The cluster structure derived from the control cells is further persistent in the perturbed states. More concretely, we find that a stereotypical spatial organization of cellular states is retained for the majority of the drugs, and cell states belonging to the same cell line cluster together (**Fig. S6f**).

Computing the difference between the control and treated state of each drug, i.e., what corresponds to the transport cost, allows us to further characterize a drug’s effect. High overall transport costs correspond to large feature value changes, i.e., a strong perturbation effect. An inducer of apoptosis, Staurosporine, shows high transport costs and thus substantial feature changes in all cellular states (**Fig. 3e**). Other drugs demonstrate less severe effects in the observed 8h incubation period. Nonetheless, we find that most drugs affect multiple cellular states and that individual cellular states have varying sensitivities to drugs in the drug panel. We further focus on perturbations introducing strong feature changes, mostly comprising proteasome inhibitors as single agents (Ixazomig and Carfilzomib) or as a combination treatment (Carfilzomib + Pomalidomide + Dexamethasone), a microtubule-stabilizing agent (Paclitaxel), the c-Met inhibitor Crizotinib, and Dasatinib, which is an ATP competitor for multiple tyrosine kinases such as c-KIT, and Bcr-Abl. We find all perturbations to increase levels of cleaved Caspase 3, an apoptosis marker, in various cellular states and in both cell lines (**Fig. S6k**), with the exception of Dasatinib, which specifically induced cell death in cellular states 5, 6, 9, and 19 associated to M130429 (**Fig. 3f**).

Previous work by Smith et al. (53) reports that M130429 cells reduce metabolic activity (a proxy of cell viability) upon treatment with inhibitors of MEK (MEKi) and RAF (RAFi), while M130219 cells are resistant to these inhibitors. When comparing the responses of the two cell lines to Trametinib (MEKi) and MLN2480 (panRAFi) in the MEK and PI3K pathway using pERK and pAKT as the respective readouts, we find that MEKi-sensitive M130429 cells down-regulate pAKT and pERK, whereas the MEKi-resistant M130219 cells only down-regulate pERK. Consistently, we also find that treatment with MLN2480 results in a similar differential drug response (**Fig. S6i**). This suggests that *decoupling* of the MEK and PI3K pathways may confer resistance to MEK and Raf inhibitors and constitute an adaptation to the escape of cancer therapy (36). We find further supporting evidence of pathway crosstalk alteration when we analyze pAKT and pERK levels upon treatment with a cocktail of Trametinib (MEKi) and Dabrafenib (BRAFi). In response to two drugs impinging on the MEK pathway, we observe pERK to be reduced in both cell lines but increased pAKT levels in the MEKi-resistant cell line M130219 (which resistance was acquired during pre-exposing a patient to MEKi) (**Fig. 3f**). This finding points towards a compensatory feedback mechanism acquired by M130219 during MEKi treatment by which inhibition of the MEK pathway (quantified as a reduction of pERK) would stimulate signaling through the PI3K pathway, possibly through activation of an upstream receptor kinase (13).

Our results on two co-cultured primary melanoma cell lines treated with various anti-cancer drugs show that CELLOT is capable of accurately capturing phenotypic heterogeneity in unperturbed populations of cells and accurately predict diverse drug responses by incorporating the underlying cell-to-cell variability.

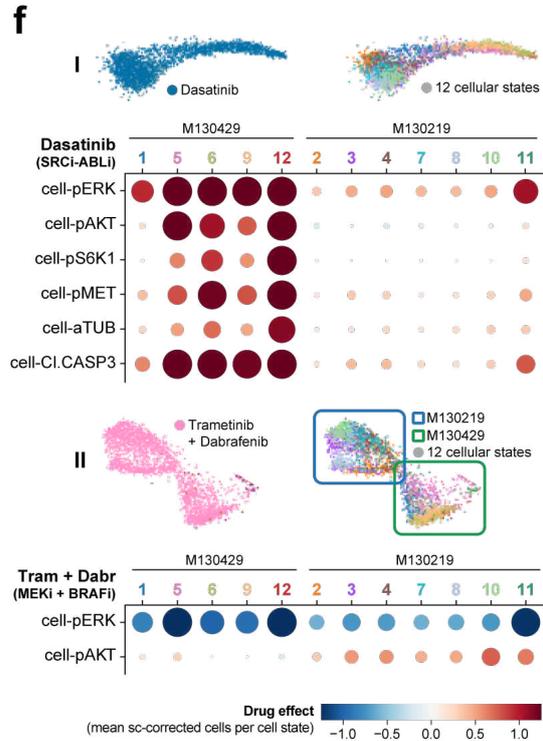
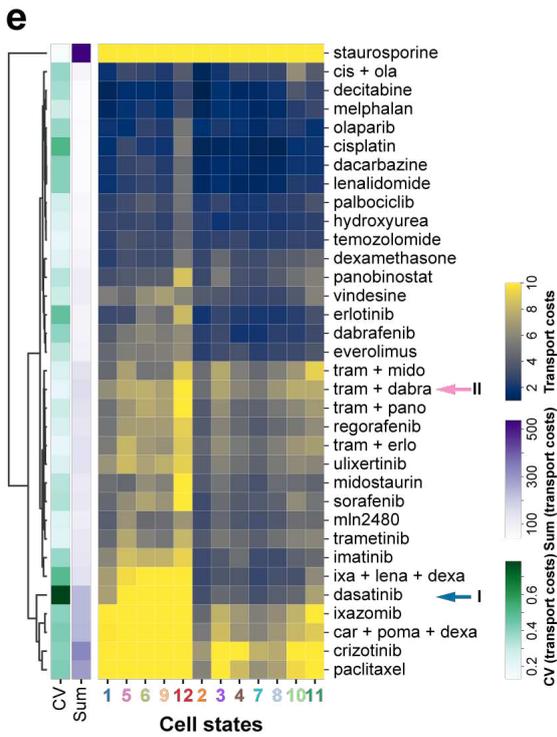
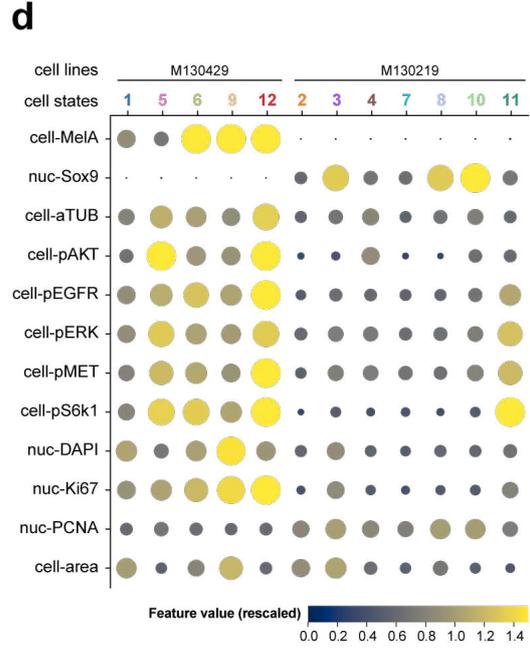
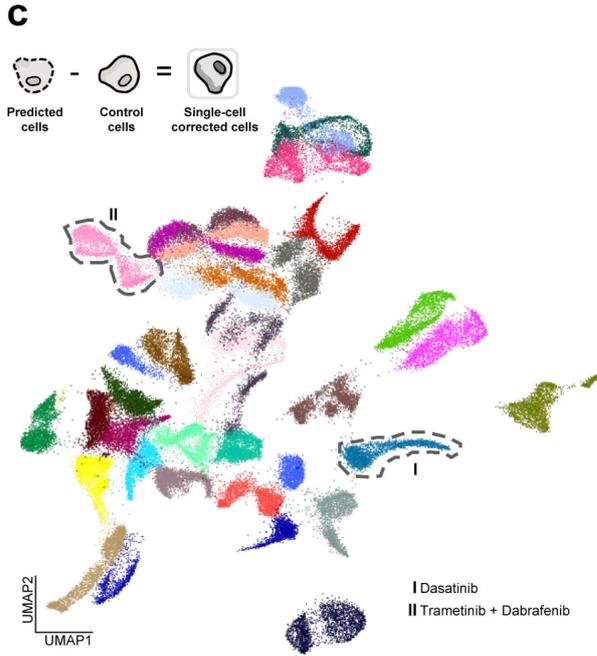
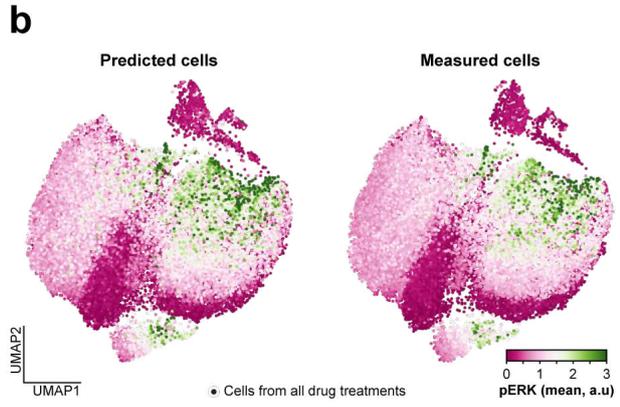
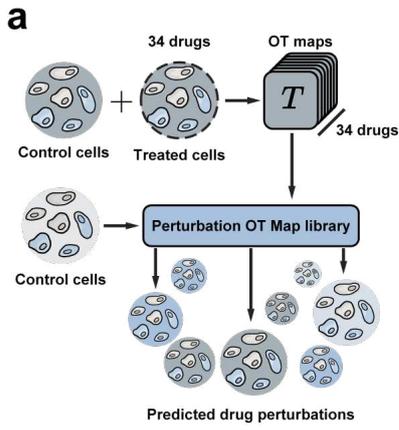


Figure 3: In-depth analysis of CELLOT’s predictions on unpaired 4i measurements of cells in control and those treated with different cancer drugs. **a)** CELLOT training and prediction setup. 34 CELLOT models were trained, one for each drug perturbation. Subsequently, each model was used to predict perturbed cells from a common set of previously unseen control cells. **b)** Projection of equal numbers of predicted and measured cells sampled from 34 perturbations in a jointly constructed UMAP. Each spot corresponds to a cell, color-coded for its measured or predicted pERK intensity, respectively, **c)** UMAP projection of single-cell perturbation effects using predicted cells. Each spot in the UMAP corresponds to a cell, color-coded for drug treatment it originates from (see **Fig. S6** for full legend). Single-cell perturbation effects were calculated by subtracting the feature profile of the original control cells from the predicted perturbation feature profile (see Online Methods). **d)** Cellular states identified in control cells using the full feature space and Leiden clustering (see Online Methods). Each column represents one of the identified cell state. Cell states are sorted in ascending order and based on their association to the cell lines M130219 and M130429 along the horizontal plot axis. A subset of features used to determine the cell states is plotted along the vertical plot axis (see **Fig. S6** for the full feature set). Size and color hue of the circles are scaled based on the feature value. **e)** Clustergram visualizing transport cost of individual drug treatments for each cell state (main heatmap, blue-yellow color scheme), the sum of transport costs (Sum) of all states per drug (first column left of the heatmap, purple), the coefficient of variation (CV) of the transport costs per drug (second column left of the heatmap, green) and the dendrogram based on the hierarchical clustering the drug’s cell state transport costs. Cell states are sorted as in **d**. **f)** In-depth view of cell state-specific responses to drug treatments. Top panel (I) Dasatinib treatment. Bottom panel (II) Trametinib + Dabrafenib treatment. Panel organization: top-left, condition-focused enlargement of UMAP projection from **c**. Top-right, same as top-left but color-coded for cell state assignment. Bottom, each column represents a cell state and rows represent highlighted features. ‘cell-’ stands for mean cell intensity. Circles are scaled based on drug effect, the larger the \pm effect the larger the circles. Negative values are encoded in hues of blue, positive values in red hues of the respective circles. 4i data was generated using cell lines M130219 and M130429 (Online Methods)

CELLOT accurately predicts single-cell responses out-of-sample for unseen patients

The accurate maps between molecular states before and after treatments learned by CELLOT contribute to a better understanding of the differences between cells that respond to certain drugs and cells that do not respond. This is crucial for inferring an incoming patient’s response to drugs and settings with high cell-to-cell variability. To make predictions on unseen patients, however, we need to demonstrate that the learned maps T model perturbation responses across different patients coherently and robustly, while still predicting personalized treatment outcomes for each patient instead of mere population averages. To test the generalization capacity of CELLOT in such an out-of-sample scenario, we use a peripheral blood mononuclear cells (PBMC) droplet scRNA-seq dataset. Kang et al. (30) characterize the cell type specificity and inter-individual variability of the response of eight lupus patients to IFN- β , a potent cytokine that induces genome-scale changes in immune cell transcriptional profiles. The dataset contains two pools, IFN- β -treated and control, prepared with the same number of cells from each individual, thus allowing us to investigate the cell fate differences between unperturbed and perturbed cells. In the following, we compare the performance of CELLOT and other baselines in an independent-and-identically-distributed (i.i.d.) setting, where models see cells from all patients, as well as in the out-of-sample (o.o.s.) setting, where models do not see cells from a specific holdout patient (see **Fig. 4a**).

As in the previous analysis, we evaluate how accurately CELLOT captures the change in the overall expression of different marker genes from control to IFN- β -treated cells and thus how well the predicted gene expression marginals are aligned with the treated population (**Fig. 4b**). Here, we consider the genes CXCL11, CCL2, and APOBEC3A, since they are connected with autoimmune diseases, including systemic lupus erythematosus (25, 45) and thus potential therapeutic targets in the management of patients with lupus and, likely, other interferonopathies (42, 48, 25, 42, 45, 18). These selected genes show a large change in expression from the control to the perturbed population, partially exhibiting a bimodal gene expression profile upon perturbation. In contrast to CELLOT, the baselines do not accurately predict these large transcriptomic shifts of these genes. An extended analysis of additional genes strongly affected by the IFN- β treatment can be found in the **Fig. S7, S8**.

All models, including CELLOT, show little performance drop when modeling the treatment

outcome on a new patient using the generalized perturbation model T_L trained on the patient cohort and using the control cells ρ_{c_z} of the unseen patient as input. This becomes evident when comparing the predicted population $\hat{\rho}_{k_z}$ with observations ρ_{k_z} using the MMD metric. **Fig. 4c** displays summary results in which each individual patient was considered for the holdout set. CELLOT outperforms previous baselines both in the i.i.d. and in the o.o.s. setting, while further showing a smaller performance drop when generalizing to the unseen patient. For more results, see **Fig. S9**. These results suggest that CELLOT not only accurately predicts treatment outcomes but also that the learned optimal transport maps correctly model the shift in the structures of the cellular subpopulation present in all patients, thus robustly performing out-of-sample.

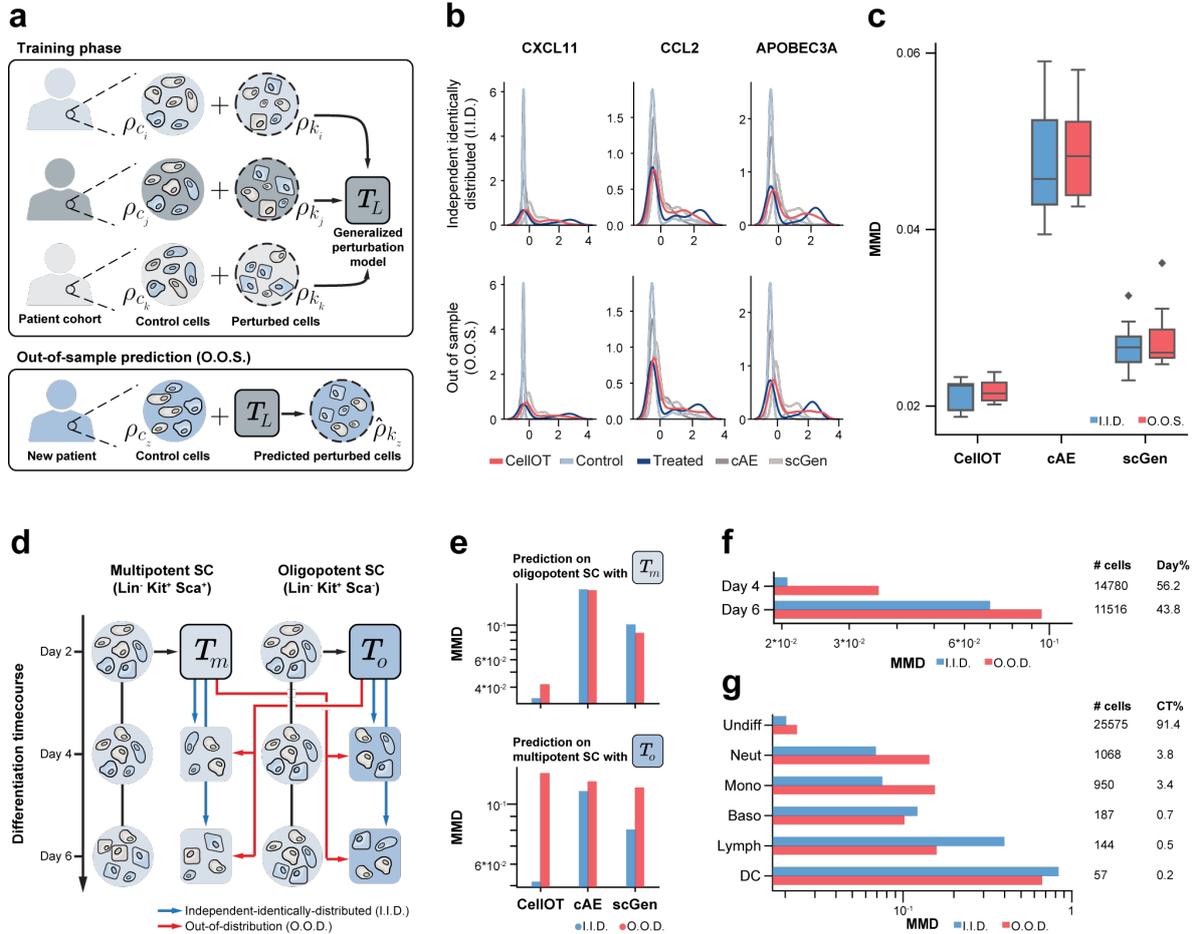


Figure 4: CELLOT generalizes to unseen patients and subpopulations. We consider the out-of-sample (o.o.s., **a-c**), and out-of-distribution (o.o.d., **d-g**) setting. **a**) Cells from eight lupus patients are measured in an untreated and IFN- β treated state (30). For each sample, we train two models, an o.o.s. model trained on cells from all other samples and an i.i.d. model trained with additional access to half of the cells in the holdout sample (not shown). **b**) Marginals for predicted cells of the holdout sample in the i.i.d. (top) and o.o.s. (bottom) setting. Predictions for both models are made on the same test set (not used for training the two models). **c**) MMD scores between the predicted distribution and the observed treated distribution across all holdout samples in the i.i.d. and o.o.s. settings. **d**) Cells from multipotent and oligopotent subpopulations are seeded and measured after 2, 4, and 6 days (64). We apply CELLOT to predict how cells from day 2 develop into the combined set of day 4 and 6 cells, when trained on only multipotent cells (T_m) or oligopotent cells (T_o). We then apply T_m to predict the o.o.d. oligopotent cells and T_o to predict the o.o.d. multipotent cells. Similar to the o.o.s. setting, i.i.d. models are trained that include half of the holdout subpopulation. **e**) MMD scores between the predicted and (observed) developed distributions for all models in both o.o.d. prediction task and i.i.d. prediction tasks (jointly for day 4 and 6 states). Performance of CELLOT, when predicting **f**) day 4 states and day 6 states **g**) for different cell types in each setting using T_m .

CELLOT generalizes perturbation responses out-of-distribution from multipotent populations to cells of lower potency

During developmental processes, stem and progenitor cells progress through a hierarchy of fate decisions, marked by a continuous differentiation of cells which refine their identity until reaching a functional end state. By tracking an initial cell population along the differentiation process, CELLOT allows us to recover individual molecular cell fate decisions and developmental trajectories. Here, the perturbation of the control population of progenitor cells is initiated by internal molecular factors driving developmental processes instead of some external factor.

Weinreb et al. (64) analyzed the fate potential of hematopoietic stem and progenitor cells (HSPCs), by tracking a broad class of oligopotent ($\text{Lin}^- \text{Kit}^+ \text{Sca}^-$) and multipotent ($\text{Lin}^- \text{Kit}^+ \text{Sca}^+$) progenitor cell subpopulations and observing samples on days 2, 4 and, 6 (**Fig. 4d**). Here, we test how well CELLOT and other baselines can learn the differentiation process of the cells observed on day 2 to the cells observed on days 4 and 6 (combined) and generalize from one subpopulation to another. We refer to the generalization task as out-of-distribution (o.o.d.), since unlike the o.o.s. setting, we expect the subpopulations to have more distinct responses. We learn two maps, where map T_o is trained exclusively on oligopotent cells, T_m on multipotent cells. I.i.d. versions of these maps are trained on both oligopotent and multipotent cells, such that each pair of i.i.d. and o.o.d. maps is evaluated on the same test set. Comparing the distributional distance between predicted and observed differentiated cell states using the MMD metric, CELLOT outperforms current state-of-the-art methods in this i.i.d. setting for both the oligopotent and the multipotent subsets (see **Fig. 4e**). Furthermore, while baselines struggle to perform in either o.o.d. setting, CELLOT is able to generalize its predictions in one direction, i.e., from multipotent cells to the oligopotent setting. In contrast to oligopotent cells, multipotent cells have a higher potency and thus can potentially differentiate into more cell types, and so we would expect T_m is more likely to generalize than T_o , trained on the less potent oligopotent cells. When predicting developmental perturbations on multipotent cells using T_o , the differentiated cell fates cannot be recovered.

CELLOT captures responses shared among cell types and different levels of temporal resolution

Depending on shared or individual receptors, signaling pathways, and regulatory networks, the response to perturbations and sensitivity to developmental factors or drugs might vary strongly between cell types or at different points in time. However, a robust perturbation model is required to capture the perturbation response across different cell types as well as at various levels of temporal resolutions. Therefore, we further investigate the difference in predictive performance at different time points and cell types. The transportation maps T_o and T_m were trained to learn the response from day 2 cells to day 4 and 6 together (**Fig. 4d**). **Fig. 4f** shows the accuracy of the modeled development of multipotent cells using map T_m individually for day 4 and day 6 cells, respectively. It is evident that CELLOT performs better when predicting developmental dynamics short range instead of states further away in time (further results in **Fig. S11**). This suggests a potential limitation for all of these methods, which might be unable to recover alignments over coarse time resolutions.

In addition, we further decompose the analysis to capture CELLOT’s consistency across different cell types. While the vast majority of cells on days 4 and 6 are still undifferentiated (undiff), some cells have evolved to neutrophils (neut), monocytes (mono), basophils (baso), lymphoid precursors (lymph), or dendritic cells (DC). As expected, the performance of CELLOT drops in terms of the MMD metric for those cell types that are only sparsely represented in the dataset (see **Fig. 4g**).

Discussion

In this work, we propose CELLOT, a framework to model single-cell perturbation responses from unpaired treated and untreated cell states using neural optimal transport. By adequately modeling the nature of the problem through the lens of optimal transport, CELLOT determines how perturbations affect cellular properties, reconstructs the most likely trajectory single cells take upon perturbation, and subsequently assists in a better understanding of driving factors of cell fate decision and cellular evasion mechanisms. CELLOT builds on the recent successes of optimal transport applications in single-cell biology (51, 37), by introducing a fully parameterized transport map that can be applied to incoming unseen samples. Previous methods (28, 67, 46) rely on an unconstrained parameterization of the *primal* optimal transport map, however the unconstrained nature of these models makes robust optimization challenging and results in reduced performance (41, Table 1). Instead, we learn the transformation of unperturbed to perturbed cell states through the *dual* optimal transport problem, parameterized via a pair of neural networks constrained to be convex (10, 41). These constraints are important inductive biases that facilitate learning and result in a reliable and easy-to-train framework, as evidenced by the consistently strong performance of CELLOT on several problems without the need for extensive hyperparameter tuning (see Online Methods). CELLOT infers the highly complex and nonlinear evolution of cell populations in response to perturbations without making strong simplifying assumptions on the nature of these dynamics. Unlike current approaches comprising autoencoder-based baselines (39, 40, 68), CELLOT does not necessarily rely on learning meaningful low-dimensional embeddings in which perturbations are modeled as linear shifts, thus faithfully capturing the heterogeneity of single-cell perturbation responses and accounting for high cell-to-cell variability. We confirm this advantage through experiments on single-cell responses to different drugs in cancer cell lines obtained with RNA-seq and spatially resolved 4i measurements, where CELLOT consistently outperforms (Fig. 2, Table. S1). Our evaluations went beyond the often used average treatment effect and correlation analysis across all cells; we analyzed marginals and computed MMD scores, a strong measure of how well predicted and observed distributions match.

Using CELLOT to perform cell-state-aware drug profiling enables us to quantify perturbation effects as a function of the underlying heterogeneity of the studied system, in our cases a co-culture of two melanoma cell lines with different sensitivities to drug treatments. In doing so, we *sharpen* the response profiles of the measured drugs and reveal cell-state-specific responses of multiple signaling pathway in relation to treatment history of the cell line donor. We find the signaling activity associated to the MEK and PI3k pathways to decouple in cells pre-exposed to MEK inhibitors, a known adaptation mechanism for therapy evasion in melanoma cells (36). This *pathway rewiring* is associated to alteration in the molecular feedback structure of cells from effectors to receptors (36, 60). Thus, combining CELLOT with a larger set of combination treatments, multiplexed imaging, and cellular systems reflective of disease adaptations may help us to elucidate the molecular mechanisms of signaling pathway evolution in the context of cancer therapy.

The results in Fig. 2-3 are based on predictions on cells from the same sample but that were not used for training (i.i.d. setting). The treatment effects can then be analyzed by scrutinizing the learned maps (cf. Fig. 3). However, for predicting the treatment effect in practice, it is much more relevant how well the learned maps generalize beyond samples used for training (o.o.s. setting) and to different sample compositions (o.o.d. setting). In Fig. 4, we therefore carefully analyze how well CELLOT and other methods generalize beyond settings containing cells derived from the same sample. Specifically, we study the challenging out-of-sample setting when predicting treatment responses in unseen lupus patients and out-of-distribution settings when inferring developmental trajectories on stem cells of lower potency. In both cases, CELLOT’s accuracy and precision are superior to current state-of-the-art methods (Fig. 4). Moreover, the predicted cell states after perturbation are still very close to the actually observed cell states. We consider these results as particularly promising, as it illustrates that accurate o.o.s. and o.o.d. predictions are indeed possible.

Nevertheless, we also observe that the predictive performance for CELLOT drops when perturbations are too strong, i.e., the cell distributions before and after perturbations are very different (Fig. 4f; a similar drop is observed for the other methods, see **Fig. S10**). The principle underlying the optimal transport theory is ideally suited for acute cellular perturbations during which single cells do not redistribute entirely and randomly in multidimensional measurement space, but typically only in a few dimensions, such that the overall correlation structure is preserved. While this modeling hypothesis is satisfied when perturbation responses are observed via regularly and frequently sampled snapshots, molecular transitions cannot be reconstructed when perturbation responses have progressed too far. For particularly strong or complicated perturbations, cellular multiplex profiles might change too drastically, violating OT assumptions and making it challenging to reconstruct the alignments between unperturbed and perturbed populations based on the *minimal effort* principle. In such settings, additional information is likely needed, for instance, a model of the underlying biology or models that integrate observations of multiple smaller time steps.

Despite of the stochastic nature of cell fate decisions and the fact that cellular dynamics are intrinsically noisy (65), CELLOT further models cell responses as deterministic trajectories. Approaches treating cell fate decisions as probabilistic events have previously allowed estimation of the full dynamical model to a greater extent than their deterministic counterparts (7). By connecting OT and stochastic difference equations, recent work (9) can build up on CELLOT to account for biological heteroscedasticity, at the cost of added model complexity and necessary simplifying assumptions.

Despite having provided a proof-of-concept of the capacity of CELLOT to model various chemical perturbations for different data modalities through an in-depth analysis of the nature of the learned mapping as well as a demonstration of its versatility in a broad class of applications, CELLOT’s generalization capacity has been evaluated on relatively small datasets. Crucially, large cohorts comprised of patients with different molecular profiles, such as cancer patients with various underlying genetics, could result in strongly heterogeneous treatment responses. It is evident that approaches to address these challenges could readily exploit the upcoming availability of large-scale patient cohort studies. The use of neural optimal transport to learn single-cell drug responses makes thus for an exciting avenue of future work, including its use to improve our understanding of cell therapies, study drug responses from patient samples, and better account for cell-to-cell variability in large-scale drug design efforts.

Contributions

G.G., K.L., L.P., A.K., and G.R. conceptualized the study. C.B. developed CELLOT, C.B. and S.G.S. implemented CELLOT, and S.G.S. and C.B. performed CELLOT experiments. G.G. and J.S.C. planned and performed 4i experiments. S.G.S., C.B., G.G., K.L., L.P., A.K., and G.R. analyzed results and generated figures. C.B., S.G.S., G.G., K.L., L.P., A.K., and G.R. wrote the manuscript. K.L., L.P., A.K., and G.R. supervised the study. M.L. provided reagents and cell lines, gave feedback on the manuscript and expert advice on the cell lines.

Acknowledgments

We are grateful to Hugo Yèche and Ximena Bonilla for their helpful comments, corrections, and discussions. C.B. and A.K. received funding from the Swiss National Science Foundation under the National Center of Competence in Research (NCCR) Catalysis under grant agreement 51NF40 180544. L.P. is supported by the European Research Council (ERC-2019-AdG-885579), the Swiss National Science Foundation (SNSF grant 310030_192622), the Chan Zuckerberg Initiative, and the University of Zurich. G.G. received funding from the Swiss National Science Foundation and InnoSuisse as part of the BRIDGE program, as well as from the University of Zurich through the BioEntrepreneur Fellowship. K.L. and S.G.S. were partially funded by ETH Zürich core funding (to G.R.) and from the Tumor Profiler Initiative (to G.R.).

Declaration of Interests

G.G. and L.P. have filed a patent on the 4i technology (patent WO2019207004A1).

Data Availability

Raw published data for the SciPlex 3 (56), lupus patients (30), and statefate dataset (64) are available from the Gene Expression Omnibus under accession codes GSM4150378, GSE96583, and GSE140802, respectively. We currently work on making the 4i melanoma dataset as well as all processed data available through a public data repository; we expect that it will be available for the revision of the manuscript. For now, the processed data is available for download here: <https://polybox.ethz.ch/index.php/s/RAykIMfD10qCJaM>.

Code Availability

CELLOT is written Python and uses standard Python libraries. The CELLOT library is available at <https://github.com/bunnech/cellot>.

Online Methods

1	Theoretical Background	14
1.1	Optimal Transport	14
1.2	Convex Neural Networks	14
1.3	Neural Optimal Transport	15
2	Model	15
2.1	Recovering Perturbation Effects via Neural Optimal Transport	15
2.2	Parametrization of the Optimal Transport Coupling	15
2.3	Predicting Perturbation Effects via CELLOT	16
3	Datasets and Preprocessing	16
3.1	Single-Cell Multiplex Data	16
3.2	Single-Cell RNA Sequencing Data	17
4	Training and Technical Details	17
4.1	Setup	17
4.2	Network Architectures	18
4.3	Hyperparameters	18
5	Evaluation	18
5.1	Metrics	18
5.2	Feature Selection	19
6	Materials	19
7	Experimental Details	20
7.1	In-Vitro Experiments	20
7.2	In-Silico Experiments	21

1 Theoretical Background

1.1 Optimal Transport

Optimal transport plays dual roles as it induces a mathematically well-characterized distance measure between distributions besides providing a geometry-based approach to realizing couplings between two probability distributions. Let $\mu = \sum_{i=1}^n a_i \delta_{x_i}$ and $\nu = \sum_{j=1}^m b_j \delta_{y_j}$ be two discrete probability measures in \mathbb{R}^d . The optimal transport (OT) problem (31) reads

$$W_2^2(\mu, \nu) = \inf_{\gamma \in \Gamma(\mu, \nu)} \int \|x - y\|^2 d\gamma(x, y), \quad (1)$$

where the polytope $\Gamma(a, b)$ is $\{\gamma \in \mathbb{R}_+^{n \times m}, \gamma \mathbf{1}_m = a, \gamma^\top \mathbf{1}_n = b\}$, describing the set of all couplings (or joint distributions) γ between μ and ν . The optimal transport plan γ thus corresponds to the coupling between two probability distributions minimizing the overall transportation cost. Computing optimal transport distances in (1) involves solving a linear program, and thus their computational cost is prohibitive for large-scale machine learning problems. Regularizing objective (1) with an entropy term results in significantly more efficient optimization (15),

$$W_2^{2,\varepsilon}(\mu, \nu) = \inf_{\gamma \in \Gamma(\mu, \nu)} \int \|x - y\|^2 d\gamma(x, y) - \varepsilon H(\gamma), \quad (2)$$

with entropy $H(\gamma) = -\sum_{ij} \gamma_{ij} (\log \gamma_{ij} - 1)$ and parameter ε controlling the strength of the regularization. W_2^ε is further differentiable w.r.t. its inputs and thus serves as a loss function in machine learning applications.

Problem (1) denotes the primal formulation for the Wasserstein-2 distance. The corresponding dual introduced by Kantorovich in 1942 is a constrained concave maximization problem defined as

$$W_2^2(\mu, \nu) = \sup_{(f,g) \in \Phi_c} \mathbb{E}_\mu[f(x)] + \mathbb{E}_\nu[g(y)], \quad (3)$$

where the set of admissible potentials is $\Phi_c := \{(f, g) \in L^1(\mu) \times L^1(\nu) : f(x) + g(y) \leq \frac{1}{2} \|x - y\|_2^2, \forall (x, y) d\mu \otimes d\nu \text{ a.e.}\}$ (62, Theorem 1.3). Villani (62, Theorem 2.9) further simplifies the dual problem (3) over the pair of functions (f, g) to

$$W_2^2(\mu, \nu) = \underbrace{\frac{1}{2} \mathbb{E} [\|x\|_2^2 + \|y\|_2^2]}_{c_{\mu, \nu}} - \inf_{f \in \tilde{\Phi}} \mathbb{E}_\mu[f(X)] + \mathbb{E}_\nu[f^*(Y)], \quad (4)$$

where $\tilde{\Phi}$ is the set of all convex functions in $L^1(d\mu) \times L^1(d\nu)$, $L^1(\mu) := \{f \text{ is measurable} \& \int f d\mu < \infty\}$, and $f^*(y) = \sup_x \langle x, y \rangle - f(x)$ is f 's convex conjugate. Villani (62, Theorem 2.9) then proves the existence of an optimal pair (f, f^*) of lower semi-continuous proper conjugate convex functions on \mathbb{R}^n minimizing (3).

1.2 Convex Neural Networks

In order to parameterize convex spaces such as $\tilde{\Phi}$ in (4), we need neural networks which are convex w.r.t. to their inputs. One example are input convex neural networks (ICNN) introduced by Amos et al. (2). ICNNs are based on fully-connected feed-forward networks that ensure convexity by placing constraints on their parameters. An ICNN with parameters $\theta = \{b_i, W_i^z, W_i^x\}$ represents a convex function $f(x; \theta)$ and, for a layer $i = 0 \dots L - 1$, is defined as

$$h_{i+1} = \sigma_i(W_i^x x + W_i^z h_i + b_i) \text{ and } f(x; \theta) = h_L, \quad (5)$$

where activation functions σ_i are convex and non-decreasing, and elements of all W_i^z are constrained to be nonnegative. Despite their constraints, ICNNs are able to parameterize a rich class of convex functions. In particular, Chen et al. (14) provide a theoretical analysis that any convex function over a convex domain can be approximated in sup norm by an ICNN. Huang et al. (26) further extend ICNNs from fully-connected feed-forward neural networks to convolutional neural architectures.

1.3 Neural Optimal Transport

Despite existing numerical approximations of the optimal transport distance and the corresponding optimal coupling (15, 3) (2), recent efforts have investigated neural network-based approaches as fast and scalable approximations to (1). Taghvaei and Jalali (58) consider solving (4) by parameterizing f with an ICNN and solving for f^* at each step, which has a high computational cost. Makkuva et al. (41) extend this work by approximating f^* with another ICNN g transforming the problem into a min-max optimization of two convex neural networks. Huang et al. (26) introduce a novel, OT-inspired parameterization of normalizing flows utilizing ICNNs. Korotin et al. (34) provide a detailed comparison of the current state of neural optimal transport solvers. Furthermore, convex neural architectures have been utilized to parameterize Wasserstein gradient flows (10, 1, 44).

2 Model

Recent high-throughput methods provide great insights on how cell populations respond to various perturbations on the level of individual cells. The provided data, however, is non-time-resolved and unaligned. Hence, snapshots taken of biological samples before and after perturbations do not provide information on single-cell trajectories. Perturbations might include the application of drugs affecting molecular functions in cells, or changes in the cellular environment causing shifts in biological signaling, thus impacting cells and their states in various ways. In the following, we describe our approach, which uncovers single-cell perturbation responses by predicting couplings between control and perturbed cell states. Hereby, let \mathcal{X} denote the biological data space spanned by cell morphology and gene expression features. We then treat a cell’s response to perturbation k as an evolution in a high-dimensional space of cell states $\mathbb{R}^d = \mathcal{X}$.

2.1 Recovering Perturbation Effects via Neural Optimal Transport

Given a dataset of n observations $\{x_1^c, \dots, x_n^c\}, x_i^c \in \mathcal{X}$ drawn from $\rho_c \in \mathcal{P}(\mathcal{X})$, the distribution of cells before applying a perturbation, we aim to learn the distribution of cells $\rho_k \in \mathcal{P}(\mathcal{X})$ upon some perturbation k , given a set of separate samples $\{x_1^k, \dots, x_m^k\}, x_i^k \in \mathcal{X}$.

Perturbation responses of cells are dynamic: after applying perturbation k , cell states evolve over time and thus can be modeled as a stochastic process on the cell data space. Despite this time-resolved nature of single-cell responses, we only have access to the distributions of cell states before, ρ_c , and after injecting perturbation k , ρ_k . We thus aim to understand the underlying stochastic process without access to time-resolved perturbation responses by uncovering the coupling γ between ρ_c and ρ_k . Given prior biological knowledge, we can assume that perturbations do not drastically or totally alter underlying cellular processes. We thus posit that the evolution of probability distributions of single-cells upon perturbation can be modeled via the mathematical theory of optimal transport. The coupling γ then corresponds to an optimal transport plan (1) between ρ_c and ρ_k .

Following Makkuva et al. (41), we infer the optimal coupling γ (1) between ρ_c and ρ_t . Thus, instead of computing a coupling individually for each pair of cell samples using existing solvers (15), we learn a parameterized optimal transport map using neural networks. The parameterized OT coupling then serves as a robust predictor for cellular distribution shifts upon perturbations on unseen samples $\{x_i^c\}_{i=1}^{n'} \sim \rho_c$, i.e., of another patient.

2.2 Parametrization of the Optimal Transport Coupling

Directly learning the optimal transport map in the primal (1) and dual (3) is notoriously difficult. Instead, Makkuva et al. build upon celebrated results by Knott and Smith (33) and Brenier (8), which relate the optimal solutions for the dual form (3) and the primal form (1),

to derive a min-max formulation replacing the convex conjugate in (4) (41, Theorem 3.3)

$$W_2^2(\rho_c, \rho_k) = \sup_{\substack{f \in \tilde{\Phi} \\ f^* \in L^1(\rho_k)}} \inf_{g \in \tilde{\Phi}} \mathcal{C}_{\rho_c, \rho_k} - \underbrace{\mathbb{E}_{\rho_c}[f(x)] - \mathbb{E}_{\rho_k}[\langle y, \nabla g(y) \rangle - f(\nabla g(y))]}_{\mathcal{V}_{\rho_c, \rho_k}(f, g)}. \quad (6)$$

We can further relax the constraint $g \in \tilde{\Phi}$ to $L^1(\rho_k)$, as a function $g \in L^1(\rho_k)$ minimizing (6) is convex and equal to f^* for any convex function f . In order to learn the resulting optimal transport, i.e., the solution of the minimization problem in (6), Makuva et al. (41) parameterize both dual variables f and g using input convex neural networks (§ 1.3) (2). The resulting approximate Wasserstein distance is thus defined as

$$\hat{W}_2^2(\rho_c, \rho_k) = \sup_{\phi} \inf_{\theta} \mathcal{C}_{\rho_c, \rho_k} - \mathcal{V}_{\rho_c, \rho_k}(f_{\phi}, g_{\theta}), \quad (7)$$

where θ and ϕ are the parameters of each ICNN. The resulting g_{θ}^* produces an approximate optimal transport plan $\gamma \approx (\nabla g_{\theta}^* \times Id)_{\#} \rho_c$.

2.3 Predicting Perturbation Effects via CELLOT

The framework described above allows us to recover couplings between control $\{x_1^c, \dots, x_n^c\}$ and perturbed cells $\{x_1^k, \dots, x_m^k\}$, giving insights into cellular response trajectories upon application of a perturbation k . Given a set of perturbations K , and sample access to the control distribution ρ_c as well as distributions ρ_k for each perturbation $k \in K$, CELLOT learns the optimal pair of dual potentials $(f_{\phi_k}^*, g_{\theta_k}^*)$ by solving equation 7. Given parametrizations of the convex potentials for each k , CELLOT then predicts the transformation of a control cell x_i^c upon perturbation k via $\hat{x}_i^k = \nabla g_{\theta_k}^*(x_i^c)$, i.e., samples following the predicted perturbed distribution $\hat{\rho}_k = (\nabla g_{\theta_k}^*)_{\#} \rho_c$. CELLOT thus provides a general approach to predict state trajectories on a single-cell level, as well as understand how heterogeneous subpopulation structures evolve under the impact of external factors.

3 Datasets and Preprocessing

3.1 Single-Cell Multiplex Data

Biologists have various powerful technologies at their disposal, capable of capturing multivariate single-cell measurements. High-content imaging, particularly when augmented by multiplexing abilities such as by Iterative Indirect Immunofluorescence Imaging (4i) (24), is ideally suited to study heterogeneous cell responses. With 4i, fluorescently labeled antibodies are iteratively hybridized, imaged, and removed from a sample to measure the abundance and localization of proteins and their modifications. Thus, 4i quickly generates large, spatially resolved phenotypic datasets rich in molecular information from thousands of treated and untreated (control) cells. Additionally to the multiplexed information 4i generates, information about cellular and nuclear morphology is routinely extracted from microscopy images (without the need for 4i) by image analysis algorithms (12).

Through multiplexing, 4i datasets are able to capture meaningful features related to both the treatment response heterogeneity (e.g., the phosphorylation or dephosphorylation of a kinase in a signaling pathway) and the pre-existing cell-to-cell variability (e.g., protein levels related to different cellular states or cell cycle phases) which may determine treatment response. Traditional high-content imaging datasets often need to compromise between features describing either the former or the latter and may thus struggle to provide sufficient information to pair treated and control cells accurately.

The cells were seeded in a 384-well plate, allowed to settle and adhere overnight. Drugs and Dimethyl sulfoxide as the vehicle control was added to the cells the next morning and incubated for 8 hours, after which the cells were fixed with Paraformaldehyde. Subsequently, 6 cycles of 4i were performed, for which the images were acquired with an automated high-content microscope. We utilized a mixture of two melanoma tumor cell lines (ratio 1:1) in order to image

a total of 97,748. For this, we consider two co-cultured primary melanoma cell lines (M130219 and M130429), which were derived from the same melanoma patient from different body sites. M130219 originates from a subcutaneous biopsy taken during treatment with Bimetinib (MEKi), whereas M130429 was derived from a bone autopsy one month after stopping said targeted therapy (47). Both cell lines share the same driver mutation (NRAS Q61R) but are phenotypically diverse. Consequently, the cell lines are also classed as two different Melanoma subtypes due to —amongst others— to differences in marker expression (47): the former a Mesenchymal subtype (SOX9+, MelA-), the latter a Melanocytic subtype (Sox9-, MelA+). 10,995 cells are imaged in the DMSO-treated control state and the rest are treated with one of 34 cancer therapies. Between 2,000 and 3,000 cells are profiled per treatment.

All image analysis steps were performed by our in-house platform called TissueMAPS (<https://github.com/TissueMAPS>). The steps included illumination correction (55), alignment of images from different acquisition cycles using Fast Fourier Transform (23), segmentation of nuclei and cell outlines (57), as well cellular and nuclear measurements of intensity and morphology features using the scikit-image library (61).

The extracted marker intensities and morphological features are then re-normalized to the same numerical scale by dividing each feature with its 75th percentile computed on control cells. Values are then transformed with a $\log_{1p}(x \leftarrow \log(x + 1))$ function. A total of 47 features are reported, 21 morphological features and 26 protein intensities.

3.2 Single-Cell RNA Sequencing Data

For statefate (64) and SciPlex 3 (56), raw counts were obtained from their GSA uploads. For each, the `scanpy` toolbox (66) was used to perform library size normalization, cell and gene filtering, and a \log_{1p} transformation. Highly variable genes are computed using the `scanpy`'s `highly_variable_genes` function. For the SciPlex 3 dataset we subset to the top 1k highly variable genes. Preprocessing for the lupus patients dataset (30) was inherited from (49), and we would like to thank the author for hosting this dataset.

4 Training and Technical Details

4.1 Setup

For SciPlex 3 and 4i experiments, we perform an 80/20 train test split. The split is performed on each drug and control condition independently to ensure that we train on 80% of each treatment. For the holdout samples in Kang et al. (30) and the holdout populations in (64), we train one model without cells under the holdout condition and train one model with half of the cells under the holdout condition. At evaluation time, we use the same set of cells to ensure that results are comparable. For scRNAseq datasets, we select hyperparameters for autoencoder models by doing a grid search over parameters summarized in the Table below and selecting the configuration that has the smallest reconstruction error over non-zero features when trained on the full SciPlex 3 dataset (i.e., trained on all drugs). We use the same configuration for all downstream tasks. The scRNA-seq dataset comprises more than 1,000 differentially expressed genes, typically assumed to lie in a low-dimensional manifold arising from the constraints of the underlying gene regulatory networks (40, 39). When applying CELLOT to scRNA embeddings, we use the same encoder that is used for SCGEN and embed gene expression data into a 50-dimensional latent space before applying CELLOT.. All models are trained for 250k iterations.

Parameter	Values	Selected
latent dimension	50, 100	100
num layers	2, 3	2
layer width	256, 512	512
dropout rate	0, 0.05, 0.1, 0.2	0
weight decay	0, 1e-5, 1e-3	1e-5
scheduler.step_size	10k, 50k, 100k	100k
scheduler.gamma	0.1, 0.25, 0.5, 0.9	0.5

4.2 Network Architectures

As suggested by Makuva et al. (41), we relax the convexity constraint on g_θ and instead penalize its negative weights W_l^h

$$R(\theta) = \lambda \sum_{W_l^h \in \theta} \left\| \max(-W_l^h, 0) \right\|_F^2. \quad (8)$$

The convexity constraint on f_ϕ is enforced after each update by setting the negative weights of all $W_l^h \in \phi$ to zero. Thus the full objective is

$$\max_{\phi: W_l^h \geq 0, \forall l} \min_{\theta} f_\phi(\nabla g_\theta(y)) - \langle y, \nabla g_\theta(y) \rangle - f_\phi(x) + \lambda R(\theta). \quad (9)$$

4.3 Hyperparameters

To learn the optimal transport maps, we use a batch size of 256, an ICNN architecture of 4 hidden layers of width 64, a learning rate of 0.0001 ($\beta_1 = 0.5$, $\beta_2 = 0.9$) and $\lambda=1$. f and g are learned in an iterative fashion. f is updated by fixing g and maximizing (9) with a single iteration. For each iteration, f is then fixed, and an inner loop of 10 updates is run to minimize g . To train all networks, we use the Adam optimizer (32).

5 Evaluation

5.1 Metrics

Since we lack access to the ground truth set of control and treatment observations on the single-cell level, we analyze the effectiveness of CELLOT using evaluations that operate on the level of the distribution of real and predicted perturbation states. Two metrics are considered, $\ell_2(\text{DS})$ and MMD.

$\ell_2(\text{DS})$ refers to the ℓ_2 -distance between the drug signatures (DS) computed on the true and predicted distributions, where signatures are computed as the difference in means between the distribution of perturbed states and control states. Let r_c be a set of observed untreated cells, r_k be a set of observed cells treated with perturbation k , and \hat{r}_k be the set of predictions made on r_c . The drug signature of perturbation is defined as:

$$DS(r_k, r_c) = \frac{1}{|r_k|} \sum_{x_i \in r_k} x_i - \frac{1}{|r_c|} \sum_{y_i \in r_c} y_i$$

We report the ℓ_2 distance between the observed signature $DS(r_k, r_c)$ and the predicted signature $DS(\hat{r}_k, r_c)$, which is a function of the difference in the means between the observed and predicted distributions.

MMD refers to the kernel maximum mean discrepancy (22), a metric to measure distances between distributions. Given two random variables x and y with distributions p and q , and a kernel function ϕ , Gretton et al. (22) define the squared MMD as:

$$MMD(p, q; \phi) = \mathbb{E}_{x, x'}[\phi(x, x')] + \mathbb{E}_{y, y'}[\phi(y, y')] - 2\mathbb{E}_{x, y}[\phi(x, y)]$$

We report an unbiased estimate of $MMD(r_k, \hat{r}_k)$ where the expectations are evaluated by averages over the cells in each set. The RBF kernel is employed, and as is usually done, the MMD reported is averaged over several length scales: 0.5, 0.1, 0.01, and 0.005.

5.2 Feature Selection

For 4i data (24), the above metrics are computed using all 47 features. Due to the high dimensionality of scRNA data, we report metrics using the top 50 marker genes. Marker genes are computed for each perturbation with the `scanpy` (66) function `rank_genes_groups`, using the untreated control cells as reference.

6 Materials

Cell Lines and Cell Culture Media Cell lines M130219 and M130429 are derived from the same human patient suffering from Melanoma cancer. M130219 originates from a subcutaneous biopsy, whereas M130429 originates from a bone biopsy. Cells were tested for the absence of mycoplasma before use and were gifted from the Levesque lab (University of Zürich/ University Hospital Zürich). Culture medium (CM) consists of 10% heat-inactivated Fetal Calf Serum (FCS), 1% Sodium Pyruvate, and 5% Glutamine in RPMI with 0.1mg/ml Anti/Anti. RPMI without L-Glutamine (Sigma Aldrich), Fetal Calf Serum (Gibco), Sodium Pyruvate (Gibco), Glutamine (Biochrome), Anti/Anti (Gibco).

Pharmacological Perturbations For a complete list of compounds, manufacturers and concentration see Supplementary Table S3. In general, compounds were stored at 5mM in dimethyl sulfoxide (DMSO) and diluted in three steps in CM to 5 μ M (0.5% DMSO) immediately before use on the cells. In the case of compound combinations, the final concentration of individual compounds was 5 μ M in CM (and 0.5% DMSO). (Aldrich Material ID: S990051-EA).

4i Blocking Solution (sBS) sBS consists of 1% Bovine Serum Albumine (BSA), and 150mM Maleimide in phosphate buffered saline (PBS). Maleimide is added to aqueous solution just before Blocking step in 4i protocol. BSA (Sigma Aldrich), Maleimide (Sigma Aldrich).

Conventional Blocking Solution (cBS) cBS consists of 1% Bovine Serum Albumine (BSA) (Sigma Aldrich) in phosphate buffered saline (PBS).

Imaging Buffer (IB) IB consists of 700mM N-Acetyl-Cysteine (NAC) in deionized water (dH2O) and 0.1M 4-(2-hydroxyethyl)-1-piperazineethanesulfonic acid (HEPES). Adjust to pH 7.4. NAC (Sigma Aldrich), HEPES (Sigma Aldrich).

Elution Buffer (EB) EB consists of 0.5M L-Glycine, 1.2M Urea, 3M Guanidinium chloride (GC), and 70mM TCEP-HCl (TCEP) in dH2O. Adjust to pH 2.5. L-Glycine (Sigma-Aldrich), Urea (Sigma-Aldrich), GC (Sigma-Aldrich), TCEP (Sigma-Aldrich).

DNA stain 4', 6-Diamidino-2-phenylindole (DAPI) at 0.4 μ g/mL in cBS. DAPI (Lifetechnologies)

Primary Antibodies The primary antibodies are listed below.

#	Name	Manufacturer	Catalogue # (Clone)	Species	Dilution (1/X)	4i Cycle
1	MelA Cocktail	Abcam	ab733	Mouse	400	1
2	Sox9	Abcam	ab185966	Rabbit	1000	1
3	pS6K1	Millipore	MABS82 (10G7.1)	Mouse	800	2
4	pAKT	Cell Signaling Technology	4060	Rabbit	600	2
5	PCNA	Abcam	ab139696	Chicken	4000	2
6	pEGFR	Cell Signaling Technology	2236	Mouse	500	3
7	pERK	Cell Signaling Technology	9101	Rabbit	500	3
8	Alpha-Tubulin	Millipore	MAB1864 (YL1/2)	Rat	8000	3
9	Ki67	Santa Cruz Biotechnology	sc-23900	Mouse	200	4
10	pMet	Cell Signaling Technology	3077	Rabbit	800	4
11	CD45	Abcam	ab187271	Mouse	400	5
12	Cleaved Caspase-3	Cell Signaling Technology	9664	Rabbit	500	5

Secondary Antibodies All secondary antibodies were diluted as listed below with cBS.

#	Name	Manufacturer	Catalogue #	Species	Dilution (1/X)
1	Anti-mouse AlexaFluor-488	Life Technologies (Invitrogen)	A-11029	Goat	400
2	anti-rabbit AlexaFluor-568	Life Technologies (Invitrogen)	A-11036	Goat	400
3	Anti-chicken AlexaFluor- 555	Life Technologies (Invitrogen)	A-32932	Goat	400
4	Anti-rat AlexaFluor- 555	Life Technologies (Invitrogen)	A-21434	Goat	400
5	Anti-rabbit AlexaFluor-647	Life Technologies (Invitrogen)	A-21245	Goat	400

7 Experimental Details

7.1 In-Vitro Experiments

Cell Culture Cells from both cell lines were cultured in Complete Medium at 37°C, 95% humidity and 5% CO₂. Per well 750 cells of each cell lines were seeded in a 384-well plate (Greiner, n^o781092 and lid n^o656191) and grown for 3 days in the above-mentioned conditions.

Pharmacological Perturbations Compounds were added to the cells using the Bravo liquid handling platform (Agilent Technologies) at the concentration specified in the Materials section. The cells were then incubated for 8h at 37°C, 95% humidity and 5% CO₂ prior to fixation.

Sample Preparation Sample preparation was performed as following. Cells were fixed in 4% Paraformaldehyde (Electron Microscopy Sciences) for 30min. Cells were then permeabilized with 0.5% Triton X-100 (Manufacturer) for 15 min. Fixation and permeabilization were performed at room temperature.

Iterative Indirect Immunofluorescence Imaging (4i) Each subsequent step was performed in sequence of mentioning and in every cycle of 4i. If not stated differently, all steps were performed at room temperature. (1) Antibody Elution. Sample was washed 4 times with dH₂O. Residual dH₂O was aspirated to minimal volume. Subsequent actions are repeated 3 times: EB was added to the sample and shaken at 100 rpm for 10 min. Then EB was aspirated to minimal volume. (2) Blocking. sBS was added to sample and shaken at 100 revolutions per minutes (rpm) for 1 hour. After 1h sample was washed 3 times with PBS. (3) Indirect immunofluorescence, primary antibody stain. Primary antibody solution was added to sample and shaken at 100 rpm for 2 hours. After 2 hours, the sample was washed 3 times with PBS. (4) Indirect immunofluorescence, secondary antibody stain. The secondary antibody solution was added to sample and shaken at 100 rpm for 2 hours. After 2 hours, the sample was washed 3 times with PBS. (5) Imaging. IB was added to sample and sample was imaged. Perform

step 1 to 5 until required plexity is achieved. All liquid dispensing and washing steps of the 4i protocol were performed using a Washer Dispenser EL406 (BioTek). Primary and secondary antibodies were dispensed using a Bravo liquid handling platform (Agilent Technologies).

Nucleus and Total Cell Staining Nuclei were stained using DSS during each 4i cycle by adding DAPI at the above specified concentration (Materials) to secondary antibody solution. Between step 4 and 5 of the last 4i cycles a cell staining was performed using AlexaFluor-647 NHS Ester (succinimidyl ester) (Invitrogen) for 5 minutes at a final concentration of $0.2\mu\text{g}/\text{mL}$ in 50mM carbonate-bicarbonate buffer pH 9.2. AlexaFluor-647 NHS Ester (Invitrogen, cat#A20006)

Microscopy An automated high-content microscope from GE Healthcare (IN Cell 6000) with an enhanced CSU-W1 spinning disk (Microlens-enhanced dual Nipkow disk confocal scanner, wide view type) was used in combination with a Nikon 40X (0.95 NA), Plan Apo, Correction Collar 0.11-0.23, CFI/Lambda, and Neo sCMOS cameras (Andor, $2,560 \times 2,160$ pixels) to acquire microscopy images. 7 by 7 images were acquired per well. 7 z-planes with a $1\mu\text{m}$ z-spacing were acquired per site and a maximum intensity projection was computed and used for subsequent image analysis. UV (406 nm), green (488 nm), red (568 nm) and far red (625 nm) signals were acquired sequentially.

7.2 In-Silico Experiments

Image Processing Image processing was done using TissueMAPS (TM): a cloud-based, interactive image processing and viewing tool developed by the Pelkmans Lab (<https://github.com/TissueMAPS>). As the first step during image processing, images were corrected for an illumination bias (4). Next, corrected images from different acquisition cycles from the same microscopy site were aligned as previously described (24). Finally, corrected and aligned images were used to generate pyramid views of the entire dataset, which were later used to train classifiers (see below).

Image Analysis and Feature Extraction Image analysis and feature extraction were performed using TM. Nuclei were segmented using DAPI signal of the first 4i cycle by applying Otsu thresholding and morphologically filling the identified objects (TM jterator modules: `threshold_otsu` & `separate_clumps`). Cell segmentation was performed using the AlexaFluor-647 NHS Ester (Sucs) signal acquired during the last 4i cycle by smoothing the Sucs signal, and adaptive thresholding (TM jterator modules: `smooth` & `segment_secondary`). Nucleus and cell morphology features were measured using TM jterator module `measure_morphology`. Prior to intensity feature extraction, all images were corrected for background signal by subtracting 120 pixel values from each pixel (TM jterator module: `rescale`). Intensity features were extracted for nucleus and cell objects using TM jterator module `measure_intensity`.

Semi-Supervised Classifiers and Data Clean-Up Cells tainted by artifacts related to sample preparation and image analysis (e.g., miss-segmentation, detachment during 4i procedure, fluorescent debris) were manually selected using TM’s graphical interface and used to train random forest classifiers to systematically exclude cells with similar artifacts from the dataset. Further, cells whose segmentation masks touched image boundaries were also excluded from the dataset.

Identification of Cell States Single-cell intensity and morphology features of DMSO-treated (control) cells, for which perturbation effects were predicted using CELLOT, were clustered using the Leiden algorithm (59) provided by the Python package `scanpy` (66, `scanpy.tl.leiden`) without customization of input parameters. Prior to Leiden clustering, a neighborhood graph was constructed for the Control cells using `scanpy.pp.neighbors` with the input parameter `_neighbors = 10` (no further customization of the input parameters).

UMAP Generation Uniform Manifold Approximation and Projection visualizations in Fig. 2, 3, S9, and S11 were generated using `scanpy`'s `scanpy.pl.umap` function preceded by `scanpy.pp.neighbors` (66).

3NN Cell Measurement The three nearest neighbor cells measurement (3NN) was calculated by identifying the three nearest cells of either measured or predicted cells in the population of measured cells using all features except pERK and then averaging their pERK value. The nearest neighbor search was performed for each drug condition separately.

Prediction Tasks in the I.I.D. Setting All marginals, UMAPs and metrics in figure 2 and tables S2 and S1 are computed using the unseen test set cells. UMAP projections are computed on the joint set of predicted and measured cells. The larger set is down samples such that their sizes are equal.

Prediction Tasks in the O.O.S. and O.O.D. Setting We test the ability of CELLOT to generalize to out-of-sample (o.o.s.) and out-of-distribution (o.o.d.) settings by predicting perturbation response on holdout samples and development trajectory on holdout cellular subpopulations. Results are reported in figure 4. To measure the drop in performance when switching to the o.o.d. or o.o.s. setting, for each holdout group, two models are trained, an i.i.d. and o.o.s./o.o.d. model. Both models are trained on all cells from the other groups, however, the i.i.d. model is additionally trained on half of the cells of the holdout group. Evaluations for both settings are done using the cells unseen by the i.i.d. model.

References

- [1] D. Alvarez-Melis, Y. Schiff, and Y. Mroueh. Optimizing Functionals on the Space of Probabilities with Input Convex Neural Networks. *arXiv Preprint*, 2021.
- [2] B. Amos, L. Xu, and J. Z. Kolter. Input Convex Neural Networks. In *International Conference on Machine Learning (ICML)*, volume 34, 2017.
- [3] G. Aude, M. Cuturi, G. Peyré, and F. Bach. Stochastic Optimization for Large-Scale Optimal Transport. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016.
- [4] N. Battich, T. Stoeger, and L. Pelkmans. Image-based transcriptomics in thousands of single human cells at single-molecule resolution. *Nature Methods*, 10(11), 2013.
- [5] N. Battich, T. Stoeger, and L. Pelkmans. Control of transcript variability in single mammalian cells. *Cell*, 163(7):1596–1610, 2015.
- [6] D. Berchtold, N. Battich, and L. Pelkmans. A systems-level study reveals regulators of membrane-less organelles in human cells. *Molecular cell*, 72(6):1035–1049, 2018.
- [7] V. Bergen, M. Lange, S. Peidli, F. A. Wolf, and F. J. Theis. Generalizing RNA velocity to transient cell states through dynamical modeling. *Nature Biotechnology*, 38(12), 2020.
- [8] Y. Brenier. Polar Factorization and Monotone Rearrangement of Vector-Valued Functions. *Communications on pure and applied mathematics*, 44(4):375–417, 1991.
- [9] C. Bunne, Y.-P. Hsieh, M. Cuturi, and A. Krause. Recovering Stochastic Dynamics via Gaussian Schrödinger Bridges. *arXiv Preprint*, 2022.
- [10] C. Bunne, L. Meng-Papaxanthos, A. Krause, and M. Cuturi. Proximal Optimal Transport Modeling of Population Dynamics. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2022.

- [11] Z. Cang and Q. Nie. Inferring spatial and signaling relationships between cells from single cell transcriptomic data. *Nature Communications*, 11(1), 2020.
- [12] A. E. Carpenter, T. R. Jones, M. R. Lamprecht, C. Clarke, I. H. Kang, O. Friman, D. A. Guertin, J. H. Chang, R. A. Lindquist, J. Moffat, et al. Cellprofiler: image analysis software for identifying and quantifying cell phenotypes. *Genome biology*, 7(10):1–11, 2006.
- [13] C. J. Caunt, M. J. Sale, P. D. Smith, and S. J. Cook. MEK1 and MEK2 inhibitors and cancer therapy: the long and winding road. *Nature Reviews Cancer*, 15(10), 2015.
- [14] Y. Chen, Y. Shi, and B. Zhang. Optimal Control Via Neural Networks: A Convex Approach. In *International Conference on Learning Representations (ICLR)*, 2019.
- [15] M. Cuturi. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 26, 2013.
- [16] P. Demetci, R. Santorella, B. Sandstede, W. S. Noble, and R. Singh. Gromov–Wasserstein Optimal Transport to Align Single-Cell Multi-Omics Data. *BioRxiv*, 2020.
- [17] A. Dixit, O. Parnas, B. Li, J. Chen, C. P. Fulco, L. Jerby-Arnon, N. D. Marjanovic, D. Dionne, T. Burks, R. Raychowdhury, et al. Perturb-Seq: Dissecting Molecular Circuits with Scalable Single-Cell RNA Profiling of Pooled Genetic Screens. *Cell*, 167(7):1853–1866, 2016.
- [18] J. Flier, D. M. Boorsma, P. J. van Beek, C. Nieboer, T. J. Stoof, R. Willemze, and C. P. Tensen. Differential expression of CXCR3 targeting chemokines CXCL10, CXCL9, and CXCL11 in different types of skin inflammation. *The Journal of Pathology: A Journal of the Pathological Society of Great Britain and Ireland*, 194(4), 2001.
- [19] C. J. Frangieh, J. C. Melms, P. I. Thakore, K. R. Geiger-Schuller, P. Ho, A. M. Luoma, B. Cleary, L. Jerby-Arnon, S. Malu, M. S. Cuoco, et al. Multimodal pooled perturb-cite-seq screens in patient models define mechanisms of cancer immune evasion. *Nature genetics*, 53(3):332–341, 2021.
- [20] F. Fröhlich, T. Kessler, D. Weindl, A. Shadrin, L. Schmiester, H. Hache, A. Muradyan, M. Schütte, J.-H. Lim, M. Heinig, et al. Efficient parameter estimation enables the prediction of drug response using a mechanistic pan-cancer pathway model. *Cell systems*, 7(6):567–579, 2018.
- [21] V. A. Green and L. Pelkmans. A systems survey of progressive host-cell reorganization during rotavirus infection. *Cell host & microbe*, 20(1):107–120, 2016.
- [22] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A Kernel Two-Sample Test. *Journal of Machine Learning Research*, 13(1), 2012.
- [23] M. Guizar-Sicairos, S. T. Thurman, and J. R. Fienup. Efficient subpixel image registration algorithms. *Optics letters*, 33(2):156–158, 2008.
- [24] G. Gut, M. D. Herrmann, and L. Pelkmans. Multiplexed protein maps link subcellular organization to cellular states. *Science*, 361(6401), 2018.
- [25] C. M. Hedrich and G. C. Tsokos. Epigenetic mechanisms in systemic lupus erythematosus and other autoimmune diseases. *Trends in Molecular Medicine*, 17(12), 2011.
- [26] C.-W. Huang, R. T. Q. Chen, C. Tsirigotis, and A. Courville. Convex Potential Flows: Universal Probability Distributions with Optimal Transport and Convex Optimization. In *International Conference on Learning Representations (ICLR)*, 2021.

- [27] G.-J. Huizing, G. Peyré, and L. Cantini. Optimal transport improves cell-cell similarity inference in single-cell omics data. *bioRxiv*, 2021.
- [28] L. Jacob, J. She, A. Almahairi, S. Rajeswar, and A. Courville. W2GAN: Recovering an Optimal Transport Map with a GAN. *arXiv Preprint*, 2018.
- [29] K. Kamimoto, C. M. Hoffmann, and S. A. Morris. CellOracle: Dissecting cell identity via network inference and in silico gene perturbation. *BioRxiv*, 2020.
- [30] H. M. Kang, M. Subramaniam, S. Targ, M. Nguyen, L. Maliskova, E. McCarthy, E. Wan, S. Wong, L. Byrnes, C. M. Lanata, et al. Multiplexed droplet single-cell rna-sequencing using natural genetic variation. *Nature Biotechnology*, 36(1), 2018.
- [31] L. Kantorovich. On the transfer of masses (in Russian). In *Doklady Akademii Nauk*, volume 37, 1942.
- [32] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*, 2014.
- [33] M. Knott and C. S. Smith. On the optimal mapping of distributions. *Journal of Optimization Theory and Applications*, 43(1), 1984.
- [34] A. Korotin, L. Li, A. Genevay, J. Solomon, A. Filippov, and E. Burnaev. Do Neural Optimal Transport Solvers Work? A Continuous Wasserstein-2 Benchmark. *arXiv Preprint*, 2021.
- [35] B. A. Kramer and L. Pelkmans. Cellular state determines the multimodal signaling response of single cells. *bioRxiv*, 2019.
- [36] E. Kun, Y. Tsang, C. Ng, D. Gershenson, and K. Wong. MEK inhibitor resistance mechanisms and recent developments in combination trials. *Cancer Treatment Reviews*, 92:102137, 2021.
- [37] H. Lavenant, S. Zhang, Y.-H. Kim, and G. Schiebinger. Towards a mathematical theory of trajectory inference. *arXiv preprint arXiv:2102.09204*, 2021.
- [38] P. Liberali, B. Snijder, and L. Pelkmans. A hierarchical map of regulatory genetic interactions in membrane trafficking. *Cell*, 157(6):1473–1487, 2014.
- [39] R. Lopez, J. Regier, M. B. Cole, M. I. Jordan, and N. Yosef. Deep generative modeling for single-cell transcriptomics. *Nature methods*, 15(12):1053–1058, 2018.
- [40] M. Lotfollahi, F. A. Wolf, and F. J. Theis. scGen predicts single-cell perturbation responses. *Nature Methods*, 16(8), 2019.
- [41] A. Makuva, A. Taghvaei, S. Oh, and J. Lee. Optimal transport mapping via input convex neural networks. In *International Conference on Machine Learning (ICML)*, volume 37, 2020.
- [42] A. Mathian, M. Hie, F. Cohen-Aubart, and Z. Amoura. Targeting interferons in systemic lupus erythematosus: current and future prospects. *Drugs*, 75(8), 2015.
- [43] L. McInnes, J. Healy, and J. Melville. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv Preprint*, 2018.
- [44] P. Mokrov, A. Korotin, L. Li, A. Genevay, J. Solomon, and E. Burnaev. Large-Scale Wasserstein Gradient Flows. *arXiv Preprint*, 2021.
- [45] D. Perez-Bercoff, H. Laude, M. Lemaire, O. Hunewald, V. Thiers, M. Vignuzzi, H. Blanc, A. Poli, Z. Amoura, V. Caval, et al. Sustained high expression of multiple APOBEC3 cytidine deaminases in systemic lupus erythematosus. *Scientific Reports*, 11(1), 2021.

- [46] N. Prasad, K. Yang, and C. Uhler. Optimal Transport using GANs for Lineage Tracing. *arXiv preprint arXiv:2007.12098*, 2020.
- [47] M. I. Raaijmakers, D. S. Widmer, M. Maudrich, T. Koch, A. Langer, A. Flace, C. Schnyder, R. Dummer, and M. P. Levesque. A new live-cell biobank workflow efficiently recovers heterogeneous melanoma cells from native biopsies. *Experimental Dermatology*, 24(5), 2015.
- [48] M. S. Rani, G. R. Foster, S. Leung, D. Leaman, G. R. Stark, and R. M. Ransohoff. Characterization of β -r1, a gene that is selectively induced by interferon β (ifn- β) compared with ifn- α . *Journal of Biological Chemistry*, 271(37), 1996.
- [49] S. Rybakov, M. Lotfollahi, F. J. Theis, and F. Alexander Wolf. Learning interpretable latent autoencoder representations with annotations of feature sets. *Cold Spring Harbor Laboratory*, page 2020.12.02.401182, Dec. 2020.
- [50] F. Santambrogio. Optimal Transport for Applied Mathematicians. *Birkäuser, NY*, 55 (58-63):94, 2015.
- [51] G. Schiebinger, J. Shu, M. Tabaka, B. Cleary, V. Subramanian, A. Solomon, J. Gould, S. Liu, S. Lin, P. Berube, et al. Optimal-Transport Analysis of Single-Cell Gene Expression Identifies Developmental Trajectories in Reprogramming. *Cell*, 176(4), 2019.
- [52] S. M. Shaffer, M. C. Dunagin, S. R. Torborg, E. A. Torre, B. Emert, C. Krepler, M. Beqiri, K. Sproesser, P. A. Brafford, M. Xiao, et al. Rare cell variability and drug-induced reprogramming as a mode of cancer drug resistance. *Nature*, 546(7658):431–435, 2017.
- [53] M. P. Smith, H. Brunton, E. J. Rowling, J. Ferguson, I. Arozarena, Z. Miskolczi, J. L. Lee, M. R. Girotti, R. Marais, M. P. Levesque, et al. Inhibiting drivers of non-mutational drug tolerance is a salvage strategy for targeted melanoma therapy. *Cancer Cell*, 29(3), 2016.
- [54] B. Snijder, R. Sacher, P. Rämö, E.-M. Damm, P. Liberali, and L. Pelkmans. Population context determines cell-to-cell variability in endocytosis and virus infection. *Nature*, 461 (7263):520–523, 2009.
- [55] B. Snijder, R. Sacher, P. Rämö, P. Liberali, K. Mench, N. Wolfrum, L. Burleigh, C. C. Scott, M. H. Verheije, J. Mercer, et al. Single-cell analysis of population context advances RNAi screening at multiple levels. *Molecular Systems Biology*, 8(1):579, 2012.
- [56] S. R. Srivatsan, J. L. McFaline-Figueroa, V. Ramani, L. Saunders, J. Cao, J. Packer, H. A. Pliner, D. L. Jackson, R. M. Daza, L. Christiansen, et al. Massively multiplex chemical transcriptomics at single-cell resolution. *Science*, 367(6473), 2020.
- [57] T. Stoeger, N. Battich, M. D. Herrmann, Y. Yakimovich, and L. Pelkmans. Computer vision for image-based transcriptomics. *Methods*, 85:44–53, 2015.
- [58] A. Taghvaei and A. Jalali. 2-Wasserstein Approximation via Restricted Convex Potentials with Application to Improved Training for GANs. *arXiv Preprint arXiv:1902.07197*, 2019.
- [59] V. A. Traag, L. Waltman, and N. J. Van Eck. From Louvain to Leiden: guaranteeing well-connected communities. *Scientific Reports*, 9(1), 2019.
- [60] A. B. Turke, Y. Song, C. Costa, R. Cook, C. L. Arteaga, J. M. Asara, and J. A. Engelman. MEK inhibition leads to PI3K/AKT activation by relieving a negative feedback on ERBB receptors. *Cancer Research*, 72(13), 2012.
- [61] S. Van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, and T. Yu. scikit-image: image processing in python. *PeerJ*, 2:e453, 2014.

- [62] C. Villani. *Topics in Optimal Transportation*, volume 58. American Mathematical Soc., 2003.
- [63] C. Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.
- [64] C. Weinreb, A. Rodriguez-Fraticelli, F. D. Camargo, and A. M. Klein. Lineage tracing on transcriptional landscapes links state to fate during differentiation. *Science*, 367(6479), 2020.
- [65] D. J. Wilkinson. Stochastic modelling for quantitative description of heterogeneous biological systems. *Nature Reviews Genetics*, 10(2), 2009.
- [66] F. A. Wolf, P. Angerer, and F. J. Theis. SCANPY: large-scale single-cell gene expression data analysis. *Genome Biology*, 19(1), 2018.
- [67] K. D. Yang and C. Uhler. Scalable Unbalanced Optimal Transport using Generative Adversarial Networks. In *International Conference on Learning Representations (ICLR)*, 2019.
- [68] K. D. Yang, K. Damodaran, S. Venkatachalapathy, A. C. Soylemezoglu, G. Shivashankar, and C. Uhler. Predicting cell lineages using autoencoders and optimal transport. *PLoS Computational Biology*, 16(4), 2020.
- [69] B. Yuan, C. Shen, A. Luna, A. Korkut, D. S. Marks, J. Ingraham, and C. Sander. CellBox: Interpretable Machine Learning for Perturbation Biology with Application to the Design of Cancer Combination Therapy. *Cell Systems*, 12(2), 2021.
- [70] S. Zhang, A. Afanassiev, L. Greenstreet, T. Matsumoto, and G. Schiebinger. Optimal transport analysis reveals trajectories in steady-state systems. *bioRxiv*, 2021.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [CellOTBunneetal.SupplementRC2.pdf](#)
- [CellOTBunneetal.TableS3CompoundListRC1.xls](#)