
Neural Unbalanced Optimal Transport via Cycle-Consistent Semi-Couplings

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Comparing *unpaired* samples of a distribution or population taken at different points
2 in time is a fundamental task in many application domains where measuring pop-
3 ulations is destructive and cannot be done repeatedly on the same sample, such as
4 in single-cell biology. Optimal transport (OT) can solve this challenge by learning
5 an optimal coupling of samples across distributions from unpaired data. However,
6 the usual formulation of OT assumes conservation of mass, which is violated in
7 *unbalanced* scenarios in which the population size changes (e.g., cell proliferation
8 or death) between measurements. In this work, we introduce NUBOT, a *neural un-*
9 *balanced OT* formulation that relies on the formalism of *semi-couplings* to account
10 for creation and destruction of mass. To estimate such semi-couplings and general-
11 ize out-of-sample, we derive an efficient parameterization based on neural optimal
12 transport maps and propose a novel algorithmic scheme through a cycle-consistent
13 training procedure. We apply our method to the challenging task of forecasting
14 heterogeneous responses of multiple cancer cell lines to various drugs, where we
15 observe that by accurately modeling cell proliferation and death, our method yields
16 notable improvements over previous neural optimal transport methods.

17 1 Introduction

18 Modeling change is at the core of various problems in the natural sciences and is ideally done by
19 tracking particles over time. However, this is not always possible, as e.g. single-cell measurements
20 typically require to destroy the cells in the course of recording, making it impossible to measure
21 the same population more than once. In these situations, one must rely on comparing *different*
22 replicas of a population and, absent a natural identification of elements across the populations, infer
23 pairwise correspondences from data. Assuming molecular profiles of cells alter incrementally, recent
24 approaches have utilized optimal transport (OT) to tackle this problem (Schiebinger et al., 2019;
25 Bunne et al., 2022a; Tong et al., 2020). By returning a coupling between two measurements of
26 cell states, OT can solve that puzzle and reconstruct these incremental changes in cell states over
27 time. Despite these successes, the classic formulation of OT is ill-suited to model processes where
28 the population changes in *size*. This is the case in single-cell biology, where interventions typically
29 promote proliferation of certain cells and death of others, violating the assumption of conservation
30 of mass that the classic OT problem relies upon. Relaxing this assumption yields a generalized
31 formulation, known as the *unbalanced* OT (UBOT) problem.

32 In this work, propose a novel formulation of the unbalanced OT problem, that learns a parameterized
33 transport map that models the transformation between distributions (Fig. 1). We apply our proposed
34 method to the challenging task of predicting perturbation responses of single cells to multiple cancer
35 drugs, where our method successfully predicts cell proliferation and death.




Figure 1: **a.** A semi-coupling pair (γ_1, γ_2) consists of two couplings that together solve the *unbalanced* OT problem. Intuitively, γ_1 describes where mass goes as it leaves from μ , and γ_2 where it comes from as it arrives in ν . **b.** NUBOT parameterizes the semi-couplings (γ_1, γ_2) as the composition of reweighting functions η and ζ and the dual potentials f and g between the then *balanced* problem.

2 A Neural Unbalanced Optimal Transport Model

Relevant background on OT and the notation is summarized in Appendix §A. The method we propose weaves together a rigorous formulation of the UBOT problem based on semi-couplings (introduced below) with a practical and scalable OT mapping estimation method.

Semi-coupling formulation. To generalize OT for the unbalanced setting, Chizat et al. (2018b) introduced *semi-couplings* to allow for variations of mass, which are generalizations of couplings whereby only one of the projections coincides with a prescribed measure. Formally, the set of semi-couplings between measures μ and ν is defined as

$$\Gamma(\mu, \nu) \stackrel{\text{def.}}{=} \left\{ (\gamma_0, \gamma_1) \in (\mathcal{M}_+(\mathcal{X}^2))^2 : (\text{Proj}_1)_\# \gamma_0 = \mu, (\text{Proj}_2)_\# \gamma_1 = \nu \right\}. \quad (1)$$

Although this formulation lends itself to formal theoretical treatment, it has at least two limitations. First, it does not explicitly model a mapping between measures. Second, deriving a computational implementation of this problem is challenging by the very nature of the semi-couplings: being undetermined along one marginal makes it hard to model the space in (1).

Rebalancing with proxy measures. To turn the semi-coupling formulation of unbalanced OT into a computationally feasible method, we propose to conceptually break the problem into balanced and unbalanced subproblems, each tackling a different aspect of the difference between measures: feature transformation and mass rescaling. Specifically, we seek proxy measures $\tilde{\mu}$ and $\tilde{\nu}$ with equal mass (i.e., $\mu(\mathcal{X}) = \tilde{\nu}(\mathcal{X})$) across which to solve a *balanced* OT problem through a Monge/Brenier formulation. To decouple measure scaling from feature transformation, we propose to choose $\tilde{\mu}$ and $\tilde{\nu}$ as rescaled versions of μ and ν . Thus, formally, we seek $\tilde{\mu}, \tilde{\nu} \in \mathcal{M}_+(\mathcal{X})$ and $T, S : \mathcal{X} \rightarrow \mathcal{X}$ such that

$$\tilde{\mu} = \eta \cdot \mu, \quad \tilde{\nu} = \zeta \cdot \nu, \quad T_\# \tilde{\mu} = \tilde{\nu}, \quad S_\# \tilde{\nu} = \tilde{\mu}, \quad (2)$$

where $\eta, \zeta : \mathcal{X} \rightarrow \mathbb{R}^+$ are scalar fields, $\eta \cdot \mu$ denotes the measure with density $\eta(x) d\mu(x)$ (analogously for $\zeta \cdot \nu$), and T, S are a pair of forward/backward optimal transport maps between $\tilde{\mu}$ and $\tilde{\nu}$ (Fig. 1b).

Devising an optimization scheme to find all relevant components in (2) is challenging. In particular, it involves solving an OT problem whose marginals are not fixed, but that will change as the reweighting functionals η, ζ are updated. We propose an alternating minimization approach, whereby we alternatively solve for η, ζ (through an approximate scaling update) and T, S (through gradient updates on ICNN convex potentials, as described in Section A.1).

Updating rescaling functions. Given current estimates of η and T , we consider the UBOT problem (10) between $T_\#(\eta \cdot \mu) = T_\# \tilde{\mu}$ and ν . Although in general these two measures will not be balanced (hence why we need to use UBOT instead of usual OT), our goal is to eventually achieve this. To formalize this, let us use the shorthand notation $\pi_{\text{UB}}^*(\alpha, \beta) \stackrel{\text{def.}}{=} \text{argmin}_\pi \text{UBOT}(\pi; \alpha, \beta)$, where UBOT is defined in (11). For a fixed T , our goal is to find η such that $(\text{Proj}_1)_\# [\pi_{\text{UB}}^*(T_\#(\eta \cdot \mu), \nu)] = T_\#(\eta \cdot \mu)$. For the discrete setting (finite samples), this corresponds to finding a vector \mathbf{e} satisfying:

$$\sum_{j=1}^m [\Gamma]_{ij} = \mathbf{e} \odot \mathbf{u}, \quad \text{where } \Gamma = \text{argmin UBOT}(\mathbf{e} \odot \mathbf{u}, T(\mathbf{x}_i), \mathbf{v}, \mathbf{y}_j). \quad (3)$$

For a fixed T , the vector \mathbf{e}^* satisfying this system can be found via a fixed-point iteration. In practice, we approximate it instead with a single-step update using the solution to the unscaled problem:

$$\Gamma \leftarrow \text{UBOT}(\mathbf{u}, T(\mathbf{x}_i), \mathbf{v}, \mathbf{y}_j); \quad \mathbf{e} \leftarrow \Gamma \mathbf{1} \odot \mathbf{u};$$




Figure 2: Distributional fit of the predicted perturbed cell states to the observed perturbed cell states for each drug and timestep, measured by a weighted version of kernel MMD on a set of held-out cells.

71 which empirically provides a good approximation on the optimal \mathbf{e}^* but is significantly more efficient.
 72 For a given *backward map* S , we update ζ analogously (§ C). In order to be able to predict mass
 73 changes for new samples, we will use the discrete \mathbf{e} , \mathbf{z} to fit continuous versions of η , ζ via functions
 74 parameterized as neural networks trained to achieve $\eta(\mathbf{x}_i) \approx e_i \forall i \in \{1, \dots, n\}$ and $\zeta(\mathbf{y}_j) \approx z_j \forall j \in$
 75 $\{1, \dots, m\}$.

76 **Updating mappings.** For a fixed pair of η , ζ , we want T and S to be a pair of optimal OT maps
 77 between $\tilde{\mu}$ and $\tilde{\nu}$. Since these are guaranteed to be balanced due to the argument above, we can use a
 78 usual (balanced) OT formulation to find them. In particular, we use the formulation of (Makkuva
 79 et al., 2020) to fit them. That is, $T = \nabla g$ and $S = \nabla f$ for convex potentials f and g , parameterized
 80 as ICNNs with parameters θ_f and θ_g . The corresponding objective for these two potentials is:

$$\mathcal{L}(f, g) = \int_{\mathcal{X}} [f(\nabla g(x)) - \langle x, \nabla g(x) \rangle] \eta(x) d\mu(x) - \int f(y) \zeta(y) d\nu(y).$$

81 In the finite sample setting, this objective becomes:

$$\mathcal{L}(f, g) = \frac{1}{n} \sum_{i=1}^n e_i [f(\nabla g(\mathbf{x}_i)) - \langle \mathbf{x}_i, \nabla g(\mathbf{x}_i) \rangle] - \sum_{j=1}^m z_j f(\mathbf{y}_j). \quad (4)$$

82 The optimization procedure is summarized in Algorithm 1.

83 3 Evaluation

84 **Baselines.** To put NUBOT’s performance into perspective, we compare it to several baselines:
 85 CELLOT, UBOT GAN, IDENTITY and OBSERVED. Details can be found in the Appendix §F.2.

86 3.1 Single-Cell Perturbation Responses

87 Predicting responses of heterogeneous cell populations to perturbations (e.g. drugs) at the level
 88 of single cells is a crucial step towards deciphering underlying molecular processes. However,
 89 single-cell measurements typically require the destruction of cells in the course of recording, resulting
 90 in unaligned *snapshots* of cell populations before and after the perturbation. Using NUBOT and the
 91 considered baselines, we learn a map T that reconstructs how individual cells respond to a treatment.
 92 As cells can die or proliferate as a response to treatments, this is naturally an unbalanced OT problem.

93 The single-cell measurements used for this task were generated using the imaging technology 4i (Gut
 94 et al., 2018) over the course of 24 hours, resulting in three different unaligned snapshots ($t = 0h$,
 95 $t = 8h$ and $t = 24h$) for 25 drug treatments. The control cells, i.e., the source distribution μ , consists
 96 of cells taken from a mixture of melanoma cell lines at $t = 0h$ that are exposed to a dimethyl
 97 sulfoxide (DMSO) as a vehicle control. Further, we consider two different target populations ν
 98 capturing the perturbed populations after $t = 8h$ and $t = 24h$ of treatment, respectively. The cancer
 99 cell lines are characterized by the expression of mutually exclusive protein markers, i.e., one cell line
 100 strongly expresses a set of proteins detected by an antibody called MelA (MelA⁺ cell type), while
 101 the other is characterized by high levels of the protein Sox9 (Sox9⁺ cell type). As both cancer cell
 102 lines exhibit different sensitivities to the drugs (Raaijmakers et al., 2015), their proportion (Fig. 11)
 103 as well as the total cell counts (Fig. 13) vary over the time points. An evaluation of this cell line
 104 annotation can be found in Fig. 10 (8h) and Fig. 12 (24h). A description of the data can be found
 105 in § E.2, and details on the network architecture and hyperparameters in § F.3.




Figure 4: UMAP projections of the control cells for Ulixertinib at **a.** 8h and **b.** 24h. Cells are colored by the observed and predicted protein marker values (Ki67, MelA), and predicted weights. NUBOT correctly predicts weights ≥ 1 for proliferating cells in the MelA⁺ population (**a.** and **a.,** right panel), and increased levels of cell death in the Sox9⁺ population after 24h via weights ≤ 1 (**b.,** right panel), confirmed by the experimental observations (see Fig. 11).

106 **Results.** In terms of distributional fit of the predicted perturbed
 107 to the observed cells, NUBOT outperforms all baselines in almost
 108 all drug perturbations Fig. 2.

109 In the absence of a ground truth and in particular, we are required
 110 to base further analysis of NUBOT’s predictions on changes in cell
 111 count for each subpopulation (MelA⁺, Sox9⁺). Fig. 11 clearly
 112 shows that drug treatments lead to substantially different cell num-
 113 bers for each of the subpopulations compared to control. Weights
 114 predicted by NUBOT show a high correlation between observed cell
 115 counts of the two cell types and the sum of the predicted weights of
 116 the respective cell types after 8h of treatment for all drugs (Fig. 3).

117 The data further provides insights into biological processes such as
 118 apoptosis, a form of programmed cell death induced by enzymes
 119 called Caspases (CICasp3). While dead cells become invisible in the
 120 cell state space, *dying* cells are still present in the observed perturbed
 121 sample and can be recognized by high levels of CICasp3. Conversely,
 122 the protein Ki67 marks proliferating cells. Analyzing the correlation
 123 of CICasp3 and Ki67 intensity with the predicted weights provides an
 124 additional assessment of the biological meaningfulness of our results.
 125 For example, upon Ulixertinib treatment, the absolute cell counts
 126 show an increase in Sox9⁺ cells, and a decline of MelA⁺ cells at 24h
 127 (Fig. 11). Fig. 4 shows UMAP projections of the control cells at both
 128 time points, colored by the observed and predicted protein marker
 129 values and the predicted weights. At 8h, NUBOT predicts only little
 130 change in mass, but a few proliferative cells with high weights in ar-
 131 eas which are marked by high values of the proliferation marker Ki67.
 132 At 24h, our model predicts cell death in the Sox9⁺ (MelA⁻) cell type,
 133 and proliferation in the MelA⁺ cell type, which matches the observed
 134 changes in cell counts per cell type, seen in Fig. 11 in § D. We iden-
 135 tify similar results for Trametinib (Fig. 7), Ixazomib (Fig. 8), and Vin-
 136 desine (Fig. 9) which can be found in § D. These experiments demon-
 137 strate that NUBOT accurately predicts heterogeneous drug responses
 138 at the single-cell level, capturing both, cell proliferation and death.

139 4 Conclusion

140 This work presents a novel formulation of the unbalanced optimal transport problem that bridges
 141 two previously disjoint perspectives on the topic: a theoretical one based on semi-couplings and
 142 a practical one based on recent neural estimation of OT maps. The resulting algorithm, NUBOT, is
 143 scalable, efficient, and robust. Yet, it is effective at modeling processes that involve population growth
 144 or death. On the challenging single-cell perturbation task, NUBOT is able to successfully predict
 145 perturbed cell states, while explicitly modeling death and proliferation. This is an unprecedented
 146 achievement in the field of single-cell biology, which currently relies on the use of markers to
 147 approximate the survival state of cell population upon drug treatment. Thus, the application of
 148 NUBOT in the fields of drug discovery and personalized medicine could be of great implications,
 149 as it allows to identify cellular properties predictive of drug efficacy.

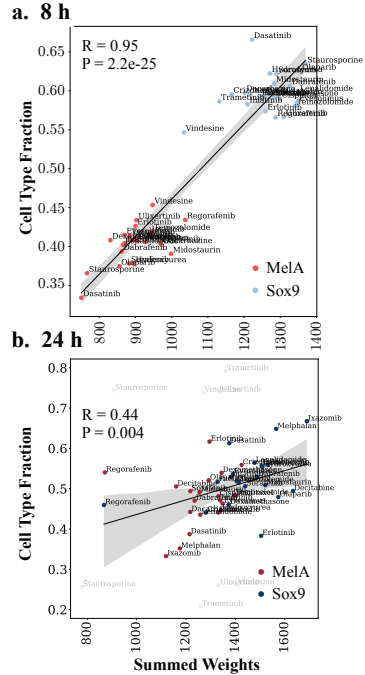


Figure 3: Given the ground truth on the known subpopulation (MelA⁺ (red) and Sox9⁺ (blue)) sizes for each drug, we analyze their level of correlation to our predicted weights after **a.** 8h and **b.** 24h. With increasing difficulty of the task and certain drugs completely removing both or one of the subpopulations, the level of correlation reduces from 8 to 24h.

References

- 150 David Alvarez-Melis, Yair Schiff, and Youssef Mroueh. Optimizing Functionals on the Space of Probabilities
151 with Input Convex Neural Networks. *Transactions on Machine Learning Research (TMLR)*, 2022.
- 153 Brandon Amos, Lei Xu, and J Zico Kolter. Input Convex Neural Networks. In *International Conference on*
154 *Machine Learning (ICML)*, volume 34, 2017.
- 155 Jean-David Benamou. Numerical resolution of an “unbalanced” mass transport problem. *ESAIM: Mathematical*
156 *Modeling and Numerical Analysis*, 37(5), 2003.
- 157 Jean-David Benamou, Guillaume Carlier, Marco Cuturi, Luca Nenna, and Gabriel Peyré. Iterative Bregman
158 Projections for Regularized Transportation Problems. *SIAM Journal on Scientific Computing*, 37(2), 2015.
- 159 Nicolas Bonneel and David Coeurjolly. SPOT: Sliced Partial Optimal Transport. *ACM Transactions on Graphics*
160 *(TOG)*, 38(4), 2019.
- 161 Yann Brenier. Décomposition polaire et réarrangement monotone des champs de vecteurs. *CR Acad. Sci. Paris*
162 *Sér. I Math.*, 305, 1987.
- 163 Charlotte Bunne, Stefan G Stark, Gabriele Gut, Jacobo Sarabia del Castillo, Kjong-Van Lehmann, Lucas
164 Pelkmans, Andreas Krause, and Gunnar Ratsch. Learning Single-Cell Perturbation Responses using Neural
165 Optimal Transport. *bioRxiv*, 2021.
- 166 Charlotte Bunne, Ya-Ping Hsieh, Marci Cuturi, and Andreas Krause. Recovering Stochastic Dynamics via
167 Gaussian Schrödinger Bridges. *arXiv Preprint arXiv:2202.05722*, 2022a.
- 168 Charlotte Bunne, Andreas Krause, and Marco Cuturi. Supervised Training of Conditional Monge Maps. In
169 *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, 2022b.
- 170 Charlotte Bunne, Laetitia Meng-Papaxanthos, Andreas Krause, and Marco Cuturi. Proximal Optimal Transport
171 Modeling of Population Dynamics. In *International Conference on Artificial Intelligence and Statistics*
172 *(AISTATS)*, volume 25, 2022c.
- 173 Luis A Caffarelli and Robert J McCann. Free boundaries in optimal transport and Monge-Àmpère obstacle
174 problems. *Annals of Mathematics*, 2010.
- 175 Laetitia Chapel, Mokhtar Z Alaya, and Gilles Gasso. Partial Optimal Transport with Applications on Positive-
176 Unlabeled Learning. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, 2020.
- 177 Laetitia Chapel, Rémi Flamary, Haoran Wu, Cédric Févotte, and Gilles Gasso. Unbalanced Optimal Transport
178 through Non-negative Penalized Linear Regression. *Advances in Neural Information Processing Systems*
179 *(NeurIPS)*, 34, 2021.
- 180 Yize Chen, Yuanyuan Shi, and Baosen Zhang. Optimal Control Via Neural Networks: A Convex Approach. In
181 *International Conference on Learning Representations (ICLR)*, 2019.
- 182 Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Scaling algorithms for
183 unbalanced optimal transport problems. *Mathematics of Computation*, 87(314), 2018a.
- 184 Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Unbalanced optimal transport:
185 Dynamic and Kantorovich formulations. *Journal of Functional Analysis*, 274(11), 2018b.
- 186 Marco Cuturi. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Advances in Neural*
187 *Information Processing Systems (NeurIPS)*, volume 26, 2013.
- 188 Atray Dixit, Oren Parnas, Biyu Li, Jenny Chen, Charles P Fulco, Livnat Jerby-Arnon, Nemanja D Marjanovic,
189 Danielle Dionne, Tyler Burks, Raktima Raychowdhury, et al. Perturb-Seq: Dissecting Molecular Circuits
190 with Scalable Single-Cell RNA Profiling of Pooled Genetic Screens. *Cell*, 167(7), 2016.
- 191 Debidatta Dwibedi, Yusuf Aytar, Jonathan Tompson, Pierre Sermanet, and Andrew Zisserman. Temporal
192 Cycle-Consistency Learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*,
193 2019.
- 194 Jiaojiao Fan, Shu Liu, Shaojun Ma, Yongxin Chen, and Haomin Zhou. Scalable Computation of Monge Maps
195 with General Costs. *arXiv preprint arXiv:2106.03812*, 2021a.
- 196 Jiaojiao Fan, Amirhossein Taghvaei, and Yongxin Chen. Scalable Computations of Wasserstein Barycenter via
197 Input Convex Neural Networks. In *International Conference on Machine Learning (ICML)*, 2021b.

- 198 Kilian Fatras, Thibault Sejourn, Remi Flamary, and Nicolas Courty. Unbalanced minibatch Optimal Transport;
199 applications to Domain Adaptation. In *International Conference on Machine Learning (ICML)*, 2021.
- 200 Alessio Figalli. The Optimal Partial Transport Problem. *Archive for Rational Mechanics and Analysis*, 195(2),
201 2010.
- 202 Charlie Frogner, Chiyuan Zhang, Hossein Mobahi, Mauricio Araya, and Tomaso A Poggio. Learning with a
203 Wasserstein loss. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 28, 2015.
- 204 Alexandre Gramfort, Gabriel Peyr, and Marco Cuturi. Fast Optimal Transport Averaging of Neuroimaging
205 Data. In *International Conference on Information Processing in Medical Imaging (IPMI)*. Springer, 2015.
- 206 Victoria A Green and Lucas Pelkmans. A systems survey of progressive host-cell reorganization during rotavirus
207 infection. *Cell Host & Microbe*, 20(1), 2016.
- 208 Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Scholkopf, and Alexander Smola. A kernel
209 two-sample test. *Journal of Machine Learning Research*, 13, 2012.
- 210 Qipeng Guo, Zhijing Jin, Ziyu Wang, Xipeng Qiu, Weinan Zhang, Jun Zhu, Zheng Zhang, and Wipf David. Fork
211 or Fail: Cycle-Consistent Training with Many-to-One Mappings. In *International Conference on Artificial
212 Intelligence and Statistics (AISTATS)*, 2021.
- 213 Gabriele Gut, Markus D. Herrmann, and Lucas Pelkmans. Multiplexed protein maps link subcellular organization
214 to cellular states. *Science*, 361(6401), 2018.
- 215 Ernst Hellinger. Neue Begrndung der Theorie quadratischer Formen von unendlichvielen Vernderlichen.
216 *Journal fr die reine und angewandte Mathematik*, 1909(136), 1909.
- 217 Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor
218 Darrell. CyCADA: Cycle-Consistent Adversarial Domain Adaptation. In *International Conference on
219 Machine Learning (ICML)*, 2018.
- 220 Chin-Wei Huang, Ricky T. Q. Chen, Christos Tsirigotis, and Aaron Courville. Convex Potential Flows: Universal
221 Probability Distributions with Optimal Transport and Convex Optimization. In *International Conference on
222 Learning Representations (ICLR)*, 2021.
- 223 YoonHaeng Hur, Wenxuan Guo, and Tengyuan Liang. Reversible Gromov-Monge Sampler for Simulation-Based
224 Inference. *arXiv preprint arXiv:2109.14090*, 2021.
- 225 Hicham Janati, Marco Cuturi, and Alexandre Gramfort. Spatio-Temporal Alignments: Optimal transport through
226 space and time. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 23,
227 2020.
- 228 Zdenek Kalal, Krystian Mikolajczyk, and Jiri Matas. Forward-Backward Error: Automatic Detection of Tracking
229 Failures. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2010.
- 230 L Kantorovich. On the transfer of masses (in Russian). In *Doklady Akademii Nauk*, volume 37, 1942.
- 231 Alexander Korotin, Vage Egiazarian, Arip Asadulaev, Alexander Safin, and Evgeny Burnaev. Wasserstein-2
232 generative networks. In *International Conference on Learning Representations*, 2020.
- 233 John Lee, Nicholas P Bertrand, and Christopher J Rozell. Parallel Unbalanced Optimal Transport Regularization
234 for Large Scale Imaging Problems. *arXiv preprint arXiv:1909.00149*, 2019.
- 235 Matthias Liero, Alexander Mielke, and Giuseppe Savar. Optimal Entropy-Transport problems and a new
236 Hellinger–Kantorovich distance between positive measures. *Inventiones Mathematicae*, 211(3), 2018.
- 237 Mohammad Lotfollahi, F Alexander Wolf, and Fabian J Theis. scGen predicts single-cell perturbation responses.
238 *Nature Methods*, 16(8), 2019.
- 239 Ashok Makkuva, Amirhossein Taghvaei, Sewoong Oh, and Jason Lee. Optimal transport mapping via input
240 convex neural networks. In *International Conference on Machine Learning (ICML)*, volume 37, 2020.
- 241 Facundo Mmoli and Tom Needham. Distance distributions and inverse problems for metric measure spaces.
242 *Studies in Applied Mathematics*, 2022.
- 243 Gaspard Monge. Mmoire sur la thorie des dblais et des remblais. *Histoire de l’Acadmie Royale des Sciences*,
244 pp. 666–704, 1781.

- 245 Ofir Pele and Michael Werman. Fast and robust Earth Mover’s Distances. In *International Conference on*
246 *Computer Vision (ICCV)*. IEEE, 2009.
- 247 Gabriel Peyré and Marco Cuturi. Computational Optimal Transport. *Foundations and Trends in Machine*
248 *Learning*, 11(5-6), 2019. ISSN 1935-8245.
- 249 Khiem Pham, Khang Le, Nhat Ho, Tung Pham, and Hung Bui. On Unbalanced Optimal Transport: An Analysis
250 of Sinkhorn Algorithm. In *International Conference on Machine Learning (ICML)*, 2020.
- 251 Marieke IG Raaijmakers, Daniel S Widmer, Melanie Maudrich, Tabea Koch, Alice Langer, Anna Flace, Claudia
252 Schnyder, Reinhard Dummer, and Mitchell P Levesque. A new live-cell biobank workflow efficiently recovers
253 heterogeneous melanoma cells from native biopsies. *Experimental Dermatology*, 24(5):377–380, 2015.
- 254 Jack Richter-Powell, Jonathan Lorraine, and Brandon Amos. Input convex gradient networks. *arXiv preprint*
255 *arXiv:2111.12187*, 2021.
- 256 Filippo Santambrogio. Optimal Transport for Applied Mathematicians. *Birkhäuser, NY*, 55(58-63):94, 2015.
- 257 Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. SuperGlue: Learning
258 Feature Matching with Graph Neural Networks. In *IEEE Conference on Computer Vision and Pattern*
259 *Recognition (CVPR)*, 2020.
- 260 Geoffrey Schiebinger, Jian Shu, Marcin Tabaka, Brian Cleary, Vidya Subramanian, Aryeh Solomon, Joshua
261 Gould, Siyan Liu, Stacie Lin, Peter Berube, et al. Optimal-Transport Analysis of Single-Cell Gene Expression
262 Identifies Developmental Trajectories in Reprogramming. *Cell*, 176(4), 2019.
- 263 Thibault Séjourné, François-Xavier Vialard, and Gabriel Peyré. Faster Unbalanced Optimal Transport: Trans-
264 lation invariant Sinkhorn and 1-D Frank-Wolfe. In *International Conference on Artificial Intelligence and*
265 *Statistics (AISTATS)*, 2022.
- 266 Daniel Sheldon, MA Elmohamed, and Dexter Kozen. Collective Inference on Markov Models for Modeling
267 Bird Migration. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 20, 2007.
- 268 Tianxiao Shen, Tao Lei, Regina Barzilay, and Tommi Jaakkola. Style Transfer from Non-Parallel Text by
269 Cross-Alignment. *Advances in Neural Information Processing Systems (NeurIPS)*, 30, 2017.
- 270 Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. TrajectoryNet: A
271 Dynamic Optimal Transport Network for Modeling Cellular Dynamics. In *International Conference on*
272 *Machine Learning (ICML)*, 2020.
- 273 Stefan Van der Walt, Johannes L Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D Warner, Neil
274 Yager, Emmanuelle Gouillart, and Tony Yu. scikit-image: image processing in python. *PeerJ*, 2:e453, 2014.
- 275 Karren D Yang and Caroline Uhler. Scalable Unbalanced Optimal Transport using Generative Adversarial
276 Networks. *International Conference on Learning Representations (ICLR)*, 2019.
- 277 Bo Yuan, Ciyue Shen, Augustin Luna, Anil Korkut, Debora S Marks, John Ingraham, and Chris Sander. CellBox:
278 interpretable machine learning for perturbation biology with application to the design of cancer combination
279 therapy. *Cell Systems*, 12(2), 2021.
- 280 Tianzuo Zhan, Giulia Ambrosi, Anna Maxi Wandmacher, Benedikt Rauscher, Johannes Betge, Niklas Rindtorff,
281 Ragna S Häussler, Isabel Hinsenkamp, Leonhard Bamberg, Bernd Hessling, et al. MEK inhibitors activate
282 Wnt signalling and induce stem cell plasticity in colorectal cancer. *Nature Communications*, 10(1), 2019.
- 283 Qiang Zhang, Tete Xiao, Alexei A Efros, Lerrel Pinto, and Xiaolong Wang. Learning Cross-Domain Correspon-
284 dence for Control with Dynamics Cycle-Consistency. *International Conference on Learning Representations*
285 *(ICLR)*, 2021.
- 286 Zhengxin Zhang, Youssef Mroueh, Ziv Goldfeld, and Bharath Sriperumbudur. Cycle Consistent Probability
287 Divergences Across Different Spaces. In *International Conference on Artificial Intelligence and Statistics*
288 *(AISTATS)*, 2022.
- 289 Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired Image-to-Image Translation using
290 Cycle-Consistent Adversarial Networks. In *IEEE Conference on Computer Vision and Pattern Recognition*
291 *(CVPR)*, 2017a.
- 292 Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A Efros, Oliver Wang, and Eli Shechtman.
293 Toward Multimodal Image-to-Image Translation. In 30, 2017b.

294 APPENDIX

295 A Background

296 A.1 Optimal Transport

297 For two probability measures μ, ν in $\mathcal{P}(\mathcal{X})$ with $\mathcal{X} = \mathbb{R}^d$ and a real-valued continuous cost function $c \in \mathcal{C}(\mathcal{X}^2)$,
 298 the optimal transport problem (Kantorovich, 1942) is defined as

$$\text{OT}(\mu, \nu) := \inf_{\gamma \in \Gamma(\mu, \nu)} \int_{\mathcal{X}^2} c(x, y) \gamma(dx, dy), \quad (5)$$

299 where $\Gamma(\mu, \nu) = \{\gamma \in \mathcal{M}_+(\mathcal{X}^2), \text{ s.t. } (\text{Proj}_1)_\# \gamma = \mu, (\text{Proj}_2)_\# \gamma = \nu\}$ is the set of couplings in the cone of
 300 nonnegative Radon measures $\mathcal{M}_+(\mathcal{X}^2)$ with respective marginals μ, ν . When instantiated on finite discrete
 301 measures, such as $\mu = \sum_{i=1}^n u_i \delta_{\mathbf{x}_i}$ and $\nu = \sum_{j=1}^m v_j \delta_{\mathbf{y}_j}$, with $\mathbf{u} \in \Sigma_n, \mathbf{v} \in \Sigma_m$ this problem translates to
 302 a linear program, which can be regularized using an entropy term (Peyré & Cuturi, 2019). For $\varepsilon \geq 0$, set

$$\text{OT}_\varepsilon(\mu, \nu) := \min_{\mathbf{P} \in U(\mathbf{u}, \mathbf{v})} \langle \mathbf{P}, [c(\mathbf{x}_i, \mathbf{y}_j)]_{ij} \rangle - \varepsilon H(\mathbf{P}), \quad (6)$$

303 where $H(\mathbf{P}) := -\sum_{ij} \mathbf{P}_{ij} (\log \mathbf{P}_{ij} - 1)$ and the polytope $U(\mathbf{u}, \mathbf{v})$ is the set of matrices
 304 $\{\mathbf{P} \in \mathbb{R}_+^{n \times m}, \mathbf{P} \mathbf{1}_m = \mathbf{u}, \mathbf{P}^\top \mathbf{1}_n = \mathbf{v}\}$. For clarity, we will sometimes write $\text{OT}_\varepsilon(\mathbf{u}, \mathbf{v}, \{\mathbf{x}_i\}, \{\mathbf{y}_j\})$.
 305 Notice that the definition above reduces to (5) when $\varepsilon = 0$. Setting $\varepsilon > 0$ yields a faster and differentiable
 306 proxy to approximate OT and allows fast numerical approximation via the Sinkhorn algorithm (Cuturi, 2013),
 307 but introduces a bias, since in general $\text{OT}_\varepsilon(\mu, \mu) \neq 0$.

308 **Neural optimal transport.** To parameterize (5) and allow to predict how a measure evolves from μ to ν ,
 309 we introduce an alternative formulation known as the Monge problem (1781) given by

$$\text{OT}_\varepsilon(\mu, \nu) = \inf_{T: T_\# \mu = \nu} \int_{\mathcal{X}} c(x, T(x)) d\mu(x), \quad (7)$$

310 with pushforward operator $\#$ and the optimal solution T^* known as the Monge map between μ and ν . When
 311 cost c is the quadratic Euclidean distance, i.e., $c = \|\cdot\|_2^2$, Brenier's theorem (1987) states that this Monge map
 312 is necessarily the gradient $\nabla \psi$ of a convex potential $\psi: \mathcal{X} \mapsto \mathbb{R}$ such that $\nabla \psi_\# \mu = \nu$, i.e., $T^*(x) = \nabla \psi(x)$.
 313 This connection has far-reaching impact and is a central component of recent neural optimal transport solvers
 314 (Makkuva et al., 2020; Bunne et al., 2022c; Alvarez-Melis et al., 2022; Korotin et al., 2020; Bunne et al., 2022b;
 315 Fan et al., 2021b). Instead of (indirectly) learning the Monge map T (Yang & Uhler, 2019; Fan et al., 2021a), it is
 316 sufficient to restrict the computational effort to learning a *good* convex potential $\nabla \theta$, parameterized via input con-
 317 vex neural networks (ICNN) (Amos et al., 2017), s.t. $\nabla \theta \psi_\# \mu = \nu$. Alternatively, parameterizations of such maps
 318 can be carried out via the dual formulation of (5) (Santambrogio, 2015, Proposition 1.11, Theorem 1.39), i.e.,

$$\text{OT}(\mu, \nu) := \sup_{\substack{f, g \in \mathcal{C}(\mathcal{X}) \\ f \oplus g \leq c}} \int f d\mu + \int g d\nu, \quad (8)$$

319 where the dual potentials f, g are continuous functions from \mathcal{X} to \mathbb{R} , and $f \oplus g \mapsto f(x) + g(x)$. Based on
 320 Brenier (1987), Makkuva et al. (2020) derive an approximate min-max optimization scheme parameterizing
 321 the duals f, g via two convex functions. The objective thereby reads

$$\text{OT}_n(\mu, \nu) = \sup_{f \text{ convex}} \inf_{g \text{ convex}} \underbrace{\frac{1}{2} \mathbb{E} [\|x\| + \|y\|]}_{\mathcal{C}_{\mu, \nu}} - \underbrace{\mathbb{E}_\mu [f(x)] - \mathbb{E}_\nu [\langle y, \nabla g(y) \rangle - f(\nabla g(y))]}_{\mathcal{V}_{\mu, \nu}(f, g)}. \quad (9)$$

322 When parameterizing f and g via a pair of ICNNs with parameters θ_f and θ_g , this neural OT scheme then allows
 323 to predict ν or μ via $\nabla g_{\theta_g \#} \mu$ or $\nabla f_{\theta_f \#} \nu$, respectively. We further discuss neural primal (Fan et al., 2021a; Yang
 324 & Uhler, 2019) and dual approaches (Makkuva et al., 2020; Korotin et al., 2020; Bunne et al., 2021) in §F.2.

325 A.2 Unbalanced Optimal Transport

326 A major constraint of problem (5) is its restriction to a pair of probability distributions μ and ν of equal mass.
 327 Unbalanced optimal transport (Benamou, 2003; Liero et al., 2018; Chizat et al., 2018b) lifts this requirement and
 328 allows a comparison between unnormalized measures, i.e., via

$$\inf_{\gamma \in \mathcal{M}_+(\mathcal{X}^2)} \int_{\mathcal{X}^2} c(x, y) \gamma(dx, dy) + \tau_1 \mathcal{D}_{f_1}((\text{Proj}_1)_\# \gamma | \mu) + \tau_2 \mathcal{D}_{f_2}((\text{Proj}_2)_\# \gamma | \nu), \quad (10)$$

329 with f -divergences D_{f_1} and D_{f_2} induced by f_1, f_2 , and parameters (τ_1, τ_2) controlling how much mass
 330 variations are penalized as opposed to transportation of the mass. When introducing an entropy regularization as
 331 in (6), the unbalanced OT problem between discrete measures \mathbf{u} and \mathbf{v} , i.e.,

$$\text{UBOT}(\mathbf{u}, \mathbf{v}) := \min_{\Gamma \in \mathbb{R}_+^{n \times m}} \langle \Gamma, [c(x_i, y_j)]_{ij} \rangle + \tau_1 \mathcal{D}_{f_1}(\Gamma \mathbb{1}_m \mid \mathbf{u}) + \tau_2 \mathcal{D}_{f_2}(\Gamma^\top \mathbb{1}_m \mid \mathbf{v}) - \varepsilon H(\Gamma), \quad (11)$$

332 can be efficiently solved via generalizations of the Sinkhorn algorithm (Chizat et al., 2018a; Cuturi, 2013;
 333 Benamou et al., 2015). We describe alternative formulations of the unbalanced OT problem in detail, review
 334 recent applications, and provide a broader literature review in the Appendix (§B.1).

335 B Related Work

336 In the following, we provide further information and review related literature on concepts discussed throughout
 337 this work.

338 B.1 Unbalanced Optimal Transport

339 Unbalanced optimal transport is a generalization of the classical OT formulation (5), and as such allows mass
 340 to be created and destroyed throughout the transport. This relaxation has found recent use cases in various
 341 domains ranging from biology (Schiebinger et al., 2019; Yang & Uhler, 2019), imaging (Lee et al., 2019), shape
 342 registration (Bonneeel & Coeurjolly, 2019), domain adaption (Fratras et al., 2021), positive-unlabeled learning
 343 (Chapel et al., 2020), to general machine learning (Janati et al., 2020; Frogner et al., 2015). Problem (10) provides
 344 a general framework of the unbalanced optimal transport problem, which can recover related notions introduced
 345 in the literature: Choosing for \mathcal{D}_f the Kullback-Leibler divergence, one recovers the so-called squared Hellinger
 346 distance. Alternatively, with $\mathcal{D}_f = \ell_2$ norm, we arrive at Benamou (2003), while an ℓ_1 norm retrieves a concept
 347 often referred to as partial OT (Figalli, 2010). The latter comprises approaches which do not rely on a relaxation
 348 of the marginal constraints as in (10). In particular, some strategies of partial OT expand the original problem by
 349 adding *virtual mass* to the marginals (Pele & Werman, 2009; Caffarelli & McCann, 2010; Gramfort et al., 2015),
 350 or by extending the OT map by *dummy* rows and columns (Sarlin et al., 2020) onto which excess mass can be trans-
 351 ported. A further review is provided in (Peyré & Cuturi, 2019, Chapter 10.2). Recent work has furthermore devel-
 352 oped alternative computational schemes (Chapel et al., 2021; Séjourné et al., 2022) as well as provided a computa-
 353 tional complexity analysis (Pham et al., 2020) of the generalized Sinkhorn algorithm solving entropic regularized
 354 unbalanced OT (Chizat et al., 2018a). Besides Yang & Uhler (2019), these approaches do not provide parameteri-
 355 zations of the unbalanced problem and allow for an out-of-sample generalization which we consider in this work.

356 B.2 Cycle-Consistent Learning

357 The principle of cycle-consistency has been widely used for learning bi-directional transformations between
 358 two domains of interest. Cycle-consistency thereby assumes that both the forward and backward mapping
 359 are roughly inverses of each other. In particular, given unaligned points $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, as well as maps
 360 $g : \mathcal{X} \mapsto \mathcal{Y}$ and $f : \mathcal{Y} \mapsto \mathcal{X}$, cycle-consistency reconstruction losses enforce $\|x - f(g(x))\|$ as well as
 361 $\|y - g(f(y))\|$ using some notion of distance $\|\cdot\|$, assuming that there exists such a ground truth bijection
 362 $g = f^{-1}$ and $f = g^{-1}$. The advantage of validating *good* matches by cycling between *unpaired* samples
 363 becomes evident through the numerous use cases to which cycle-consistency has been applied: Originally
 364 introduced within the field of computer vision (Kalal et al., 2010) and applied to image-to-image translation
 365 tasks (Zhu et al., 2017a), it has been quickly adapted to multi-modal problems (Zhu et al., 2017b), domain
 366 adaptation (Hoffman et al., 2018), and natural language processing (Shen et al., 2017). The original principle
 367 has been further generalized to settings requiring a many-to-one or surjective mapping between domains (Guo
 368 et al., 2021) via conditional variational autoencoders, dynamic notions of cycle-consistency (Zhang et al., 2021),
 369 or to time-varying applications (Dwibedi et al., 2019). These classical approaches enforce cycle-consistency
 370 by *explicitly* composing both maps and penalizing for any deviation from this bijection. In this work, we treat
 371 cycle-consistency differently. It is enforced implicitly by coupling the two distributions of interest through a
 372 sequence of reversible transformations: re-weighting, transforming, and re-weighting (Eq. (2) and Fig. 1).

373 Similarly to our work, Zhang et al. (2022) and Hur et al. (2021) establish a notion of cycle-consistency
 374 (reversibility) for a pair of pushforward operators to align two unpaired measures. Both methods rely on the
 375 Gromov-Monge distance (Mémoli & Needham, 2022), a divergence to compare probability distributions defined
 376 on different ambient spaces \mathcal{X} and \mathcal{Y} —a setting not considered in this work. They proceed by defining a
 377 reversible metric through replacing the single Monge map by a pair of two Monge maps, i.e., $f : \mathcal{X} \mapsto \mathcal{Y}$ and
 378 $g : \mathcal{Y} \mapsto \mathcal{X}$, minimizing the objective

$$\text{GM}(\mu, \nu) := \inf_{\substack{f: \mathcal{X} \mapsto \mathcal{Y}, f_{\#} \mu = \nu \\ g: \mathcal{Y} \mapsto \mathcal{X}, g_{\#} \nu = \mu}} \Delta_{\mathcal{X}}^p(f; \mu) + \Delta_{\mathcal{Y}}^p(g; \nu) + \Delta_{\mathcal{X}, \mathcal{Y}}^p(f, g; \mu, \nu), \quad (12)$$

$$\begin{aligned}\Delta_{\mathcal{X}}^p(f; \mu) &= (\mathbb{E} [|c_{\mathcal{X}}(x, x') - c_{\mathcal{Y}}(f(x), f(x'))|^p])^{\frac{1}{p}} \\ \Delta_{\mathcal{Y}}^p(g; \nu) &= (\mathbb{E} [|c_{\mathcal{X}}(y, y') - c_{\mathcal{Y}}(g(y), g(y'))|^p])^{\frac{1}{p}} \\ \Delta_{\mathcal{X}, \mathcal{Y}}^p(f, g; \mu, \nu) &= (\mathbb{E} [|c_{\mathcal{X}}(x, g(y)) - c_{\mathcal{Y}}(f(x), y)|^p])^{\frac{1}{p}}.\end{aligned}$$

380 Problem (12) shows similarities to the classical cycle-consistency objective of [Zhu et al. \(2017a\)](#), where cycle-
 381 consistency is indirectly enforced through $\Delta_{\mathcal{X}, \mathcal{Y}}^p$. [Zhang et al. \(2022\)](#) parameterize both Monge maps through
 382 neural networks in a similar fashion as done in ([Yang & Uhler, 2019](#); [Fan et al., 2021a](#)). Our approach differs
 383 from [Zhang et al. \(2022\)](#); [Hur et al. \(2021\)](#) as we model the problem through a single Monge map with duals f, g ,
 384 allowing us to map back-and-forth between measures μ and ν , and using a different parametrization approach
 385 (ICNNs). More importantly, the approaches presented by [Zhang et al. \(2022\)](#); [Hur et al. \(2021\)](#) do not generalize
 386 to the unbalanced case. While [Zhang et al. \(2022\)](#) proposed an unbalanced version of (12) by relaxing the
 387 marginals as done in [Chizat et al. \(2018a\)](#), they require the unbalanced sample sizes to be known (i.e., n and m
 388 need to be fixed). In our application of interest, particle counts of the target population are, however, not known
 389 *a priori*.

390 B.3 Convex Neural Architectures

391 Input convex neural networks ([Amos et al., 2017](#)) are a class of neural networks that approximate the family of
 392 convex functions ψ with parameters θ , i.e., whose outputs $\psi_{\theta}(x)$ are convex w.r.t. an input x . This property is
 393 realized by placing certain constraints on the networks parameters θ . More specifically, an ICNN is an L -layer
 394 feed-forward neural network, where each layer $l = \{0, \dots, L - 1\}$ is given by

$$z_{l+1} = \sigma_l(W_l^x x + W_l^z z_l + b_l) \text{ and } \psi_{\theta}(x) = z_L, \quad (13)$$

395 where σ_k are convex non-decreasing activation functions, and $\theta = \{W_l^x, W_l^z, b_l\}_{l=0}^{L-1}$ is the set of parameters,
 396 with all entries in W_l^z being non-negative and the convention that z_0 and W_0^z are 0. As mentioned above and
 397 through the connection established in § A, convex neural networks have been utilized to approximate Monge
 398 map T (7) via the convex Brenier potential ψ connected to the primal and dual optimal transport problem. In
 399 particular, it has been used to model convex dual functions ([Makkuva et al., 2020](#)) as well as normalizing flows
 400 derived from convex potentials ([Huang et al., 2021](#)). The expressivity and universal approximation properties
 401 of ICNNs have been further studied by [Chen et al. \(2019\)](#), who show that any convex function over a compact
 402 convex domain can be approximated in sup norm by an ICNN. To improve convergence and robustness of ICNNs
 403 —known to be notoriously difficult to train ([Richter-Powell et al., 2021](#))— different initialization schemes have
 404 been proposed: [Bunne et al. \(2022b\)](#) derive two initialization schemes ensuring that *upon initialization* $\nabla\psi$
 405 mimics an affine Monge map T mapping either the source measure onto itself (identity initialization) or providing
 406 a map between Gaussian approximations of measures μ and ν (Gaussian initialization). Further, [Korotin et al.](#)
 407 (2020) proposed to use quadratic layers as well as a pre-training pipeline to initialize ICNN parameters to encode
 408 an identity map.

409 B.4 Single-Cell Analysis

410 The problem of inferring correspondences across unpaired samples in biology has been traditionally tackled by
 411 relying on average and aggregate perturbation responses ([Green & Pelkmans, 2016](#); [Zhan et al., 2019](#); [Sheldon](#)
 412 [et al., 2007](#)) or by applying mechanistic or linear models ([Yuan et al., 2021](#); [Dixit et al., 2016](#)) in, potentially, a
 413 learned latent space ([Lotfollahi et al., 2019](#)). Cellular responses to treatments are, however, highly complex and
 414 heterogeneous. To effectively predict the drug response and capture such cellular heterogeneity, it is necessary
 415 to learn nonlinear maps describing such perturbation responses on the level of single cells.

416 C Additional Details of NUBOT

417 C.1 Algorithmic Details

418 The algorithmic scheme used to train NUBOT can be found in Algorithm 1.

419 **Backward direction UBOT.** For a given S , we choose ζ that ensures $(\text{Proj}_2)_{\#}[\pi_{\text{UB}}^*(S_{\#}(\zeta \cdot \nu), \mu)] =$
 420 $S_{\#}(\zeta \cdot \nu)$. For empirical measures, this yields the update:

$$\Gamma \leftarrow \text{UBOT}(\mathbf{v}, S(\mathbf{y}_j), \mathbf{u}, \mathbf{x}_i); \quad \mathbf{z} \leftarrow \Gamma \mathbf{1} \oslash \mathbf{v};$$

421 **Transforming new samples.** After learning f, g, η, ζ , we can use these functions to transform (map and
 422 rescale) new samples, i.e., beyond those used for optimization. For a given source datapoint \mathbf{x} with mass u , we
 423 transform it as $(\mathbf{x}, u) \mapsto (\nabla g(\mathbf{x}), \eta(\mathbf{x}) \cdot u \cdot \zeta(\nabla g(\mathbf{x}))^{-1})$. Analogously, target points can be mapped backed to
 424 the source domain using $(\mathbf{y}, v) \mapsto (\nabla f(\mathbf{y}), \zeta(\mathbf{y}) \cdot v \cdot \eta(\nabla f(\mathbf{y}))^{-1})$.

Algorithm 1 Neural Unbalanced Optimal Transport (NUBOT)

Input: f, g : ICNNs, initialized s.t. $\nabla g(x) \approx x$ and $\nabla f(y) \approx y$; η, ζ : NNs

```

1 for t in epochs do
2   Sample batch  $\{x_i\}_{i=1}^n \sim \mu$  and  $\{y_j\}_{j=1}^m \sim \nu$ 
3    $\hat{y} \leftarrow \nabla g(x)$ 
4    $\hat{x} \leftarrow \nabla f(y)$ 
5    $\Gamma_1 \leftarrow \text{unbalanced.sinkhorn}(\hat{y}, \frac{1}{n} \mathbf{1}_n, y, \frac{1}{m} \mathbf{1}_m)$ 
6    $e_i \leftarrow \frac{\sum_j \Gamma_{ij}}{\sum_{ij} \Gamma_{ij}} \cdot n$ 
7    $\Gamma_2 \leftarrow \text{unbalanced.sinkhorn}(\hat{x}, \frac{1}{m} \mathbf{1}_m, y, \frac{1}{n} \mathbf{1}_n)$ 
8    $z_i \leftarrow \frac{\sum_j \Gamma_{ij}}{\sum_{ij} \Gamma_{ij}} \cdot m$ 
9    $J(\theta_g, \theta_f) = \frac{1}{n} \sum_{i=1}^n e_i [f(\nabla g(x_i)) - \langle x_i, \nabla g(x_i) \rangle] - \frac{1}{m} \sum_{j=1}^m z_j f(y_j)$ 
10   $L_\eta(\theta_\eta) = \text{MSE}(\mathbf{e}, \eta(x))$ 
11   $L_\zeta(\theta_\zeta) = \text{MSE}(\mathbf{z}, \zeta(y))$ 
12  Update  $\theta_g$  to minimize  $J$ ,  $\theta_\eta$  to minimize  $L_\eta$ ,  $\theta_\zeta$  to minimize  $L_\zeta$ , and  $\theta_f$  to maximize  $J$ 

```




Figure 5: **Unbalanced sample mapping.** In all three scenarios (a,b,c), the source (gray) and target (blue) datasets share structure but have different shifts and per-cluster sampling proportions. Tasked with mapping from source to target, NUBOT and ubOT GAN predict the locations (middle pane, red) and weights (right pane) of the transported samples. The number next to the weights denotes the mean weights per cluster. While both methods map the samples to the correct location, NUBOT more accurately predicts the weights needed to match the target distribution, creating mass (dark blue) or destroying it (red) as needed.

425 C.2 Recovering semi-couplings.

426 Let us define $\tilde{\Gamma}_1 \stackrel{\text{def}}{=} \text{diag}(\mathbf{e}^{-1})^\top \Gamma_1$ and $\tilde{\Gamma}_2 \stackrel{\text{def}}{=} \text{diag}(\mathbf{z}^{-1})^\top \Gamma_2$, where Γ_1, Γ_2 are the solutions of the UBOT
 427 problems computed in Algorithm 1 (lines 7 and 9, respectively). It is easy to see that $(\tilde{\Gamma}_1, \tilde{\Gamma}_2^\top)$ is a valid pair of
 428 semi-couplings between μ and ν (cf. Eq. 1).

429 D Additional Experimental Results

430 D.1 Synthetic Data

431 Populations are often heterogeneous and consist of different subpopulations. Upon intervention, these sub-
 432 populations might exhibit heterogeneous responses. To simulate such heterogeneous intervention responses,
 433 we generate a dataset containing a two-dimensional mixture of Gaussians with three clusters in the source
 434 distribution μ . The target distribution ν consists of the same three clusters, but with different cluster proportions.
 435 Further, each particle has undergone a constant shift in space upon intervention. We consider three scenarios
 436 with increasing imbalance between the three clusters (see Fig. 5a-c). Table 1 shows the shares of the three
 437 clusters in the source and target distributions. In order to match the target distribution without transporting mass
 438 across non-corresponding clusters, the clusters have to be re-scaled with the factors presented in column 'True
 439 Scaling Factor'. The last two columns show the mean weights per cluster obtained by NUBOT and ubOT GAN,
 440 respectively. We evaluate NUBOT on the task of predicting the distributional shift from source to target, while at
 441 the same time correctly rescaling the clusters such that no mass is transported across non-corresponding clusters.

442 **Results.** The results (setup, predicted Monge maps and weights) are displayed in Fig. 5. Both NUBOT and
 443 ubOT GAN correctly map the points to the corresponding target clusters without transporting mass across
 444 clusters. NUBOT also accurately models the change in cluster sizes by predicting the correct weights for

445 each point. In contrast, UBOT GAN only captures the general trend of cluster growth and shrinkage, but does
 446 not learn the exact weights required to re-weight the cluster proportions appropriately. The exact setup and
 447 calculation of weights can be found in the §D (see Table 1). Fig. 6, shows the weighted MMD between the
 448 source distribution and the target distribution, confirming superior performance of NUBOT.




Figure 6: Distributional fit of the predicted samples to the target samples on synthetic data, measured by a weighted version of kernel MMD.

Table 1: Setup of the synthetic mixture of Gaussians dataset, showing the proportions of the three clusters in source and target distribution in three different settings (a., b., c.) as well as the required scaling factor per cluster needed to match the target without transporting points to non-corresponding clusters. The last two columns show the mean weights obtained by NUBOT and UBOT GAN.

	Cluster	Source Proportions (p)	Target Proportions (q)	True Scaling Factor (q/p)	Mean Weight NUBOT	Mean Weight UBOT GAN
a.	1	0.33	0.45	1.35	1.32	1.02
	2	0.33	0.45	1.35	1.36	0.99
	3	0.33	0.10	0.30	0.26	0.8
b.	1	0.33	0.70	2.10	2.07	1.18
	2	0.33	0.20	0.60	0.64	0.88
	3	0.33	0.10	0.30	0.29	0.81
c.	1	0.45	0.10	0.22	0.23	0.79
	2	0.45	0.45	1.00	0.98	0.94
	3	0.10	0.45	4.50	4.60	1.44

449 D.2 Single-Cell Perturbation Responses

450 As we lack ground truth for the correspondence of control and perturbed cells, we assess the biological
 451 meaningfulness of our predictions by comparing the weights to ClCasp3 and Ki67 intensity, the apoptosis and
 452 proliferation markers, respectively. Figures 7, 8 and 9 show UMAP projections computed on control cells for
 453 the drugs Trametinib, Ixazomib and Vindesine. In Figure 8 c., d., and Figure 9 c., d., regions of low predicted
 454 weights accurately correspond to regions of increased ClCasp3 intensity. Additionally, we compare predicted
 455 weights between the two cell types, and contrast them with observed cell counts.




Figure 7: UMAP projections computed on control cells for Trametinib at $t = 24h$. High predicted weights in the MeIA^+ cell type suggest proliferation, while the Sox9^+ population shows higher levels of cell death. This prediction is confirmed by the relative cell counts, where MeIA^+ cell counts increase and Sox9^+ counts decrease, demonstrating opposite response behaviors for each subpopulation, i.e., MeIA^+ cells show proliferation and Sox9^+ cells death.

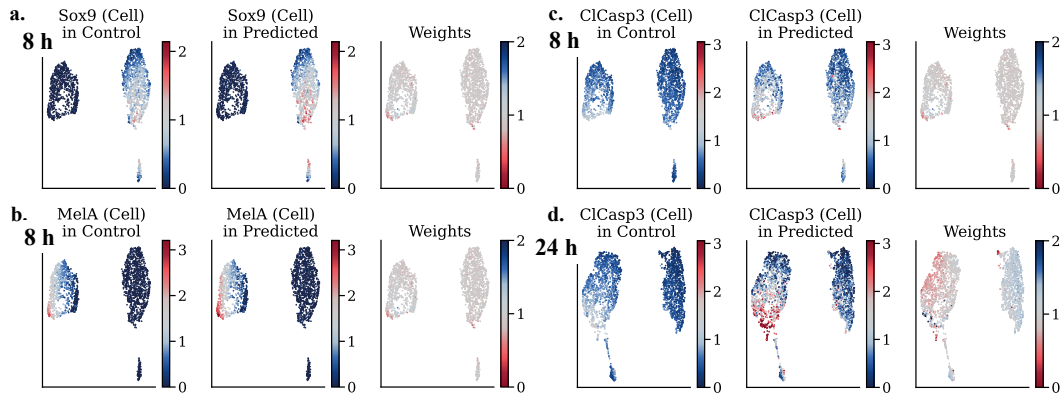


Figure 8: UMAP projections computed on control cells for Ixazomib for $t = 8h$, and $t = 24h$, colored by protein marker intensities **a.** MelA and **b.** Sox9, markers for the two subpopulations, as well as ClCasp3, a marker for cell death, at **c.** 8h and **d.** 24h. The UMAPs confirm the measured relative cell counts of each subpopulation. After 8h **a.-c.**, neither MelA⁺ nor Sox9⁺ cells are affected by the treatment, i.e., we mainly predict weights around 1. **d.** After 24h, we observe low weights in regions of high predicted apoptosis marker intensities (ClCasp3), especially at $t = 24h$, where the observed cell counts suggest death predominantly in the MelA⁺ cell cluster.

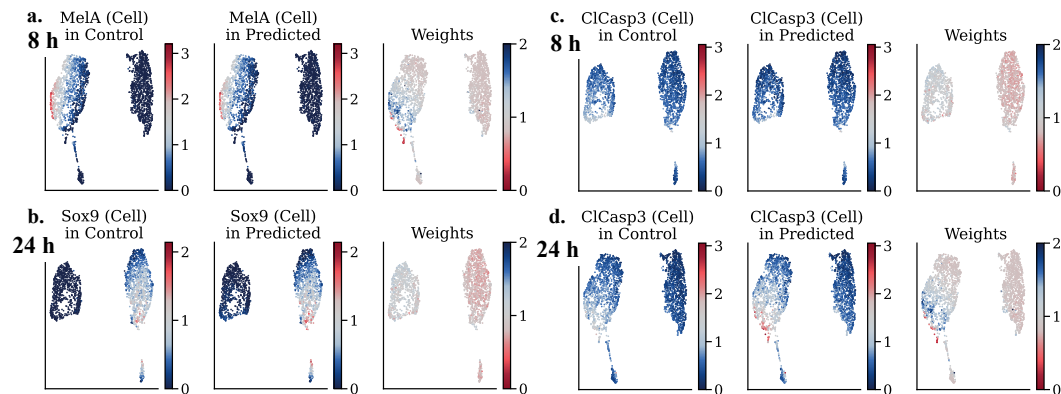


Figure 9: UMAP projections computed on control cells for Vindesine for $t = 8h$, and $t = 24h$, colored by protein marker intensities **a.** MelA and **b.** Sox9, markers for the two subpopulations, as well as ClCasp3, a marker for cell death, at **c.** 8h and **d.** 24h. The predicted weights (left) at **c.** 8h and **d.** 24h match the observed effects on each subpopulation, as initially only Sox9⁺ cells are affected by treatment with Vindesine, and only after 24h MelA⁺ cells show increased cell death.

456 E Datasets

457 We evaluate NUBOT on several tasks including synthetic data as well as perturbation responses of single cells.
 458 In both settings, we are provided with unpaired measures μ and ν and aim to recover map T which describes
 459 how source μ transforms into target ν . While in the synthetic data setting we are provided with a ground truth
 460 matching, this is not the case for the single-cell data as measuring a cell requires destroying it. In the following,
 461 we describe generation and characteristics of both datasets, as well as introduce additional biological insights
 462 allowing us to shed light on the learned matching T .

463 E.1 Synthetic Data

464 To evaluate NUBOT in a simple and low-dimensional setup with known ground-truth, we generate synthetic
 465 example: We model a source population with clear subpopulation structure through a mixture of Gaussians. Next,
 466 we generate a second (target) population aligned to the source population. We then simulate an intervention to
 467 which the subpopulations respond differently, including different levels of growth and death. Specifically, we
 468 generate batches of 400 samples with three clusters with different proportions before and after the intervention.
 469 Table 1 shows the proportions of the three clusters in the source and target distribution, as well as the required
 470 weight-factor and the obtained results from NUBOT and UBOT GAN.

Table 2: **Overview of all treatments and their inhibition type considered in this work.** Abbreviations PROTi (Proteasome inhibitor), DNASynthi (DNA synthesis inhibitor), panKi (pan kinase inhibitor), ImmuneMod. (Immune modulatory compound), MTDIsruptor (Microtubule disruptor), ApopInducer (Apoptosis inducer).

Drug Name	Inhibitor Type	Drug Name	Inhibitor Type
Ixazomib	PROTi	Olaparib	PARPi
Sorafenib	RAFi	Paclitaxel	MTDisruptor
Dabrafenib	BRAFi	Melphalan	Alkylator
Everolimus	mTORi	Regorafenib	panKi
Hydroxyurea	DNASynti	Vindesine	MTDisruptor
Midostaurin	panKi	Cisplatin	Alkalyting
Dexamethasone	ImmuneMod.	Ulixertinib	ERKi
Temozolomide	Alkylator	Staurosporine	ApopInducer
Decitabine	DNAMeti	Lenalidomide	ImmuneMod.
Dasatinib	SRCi-ABLi	Crizotinib	METi
Trametinib	MEKi	Imatinib	KITi-PDGFRi-ABLi
Erlotinib	EGFRi	Palbociclib	CDK4/6i
Dacarbazine	Alkylator		

471 E.2 Single-Cell Data

472 **Biological experiment.** The single-cell dataset used in this work was generated by the a multiplexed
473 microscopy technology called Iterative Indirect Immunofluorescence Imaging (4i) (Gut et al., 2018), which is
474 capable of measuring the abundance and localization of many proteins in cells. By iteratively adding, imaging and
475 removing fluorescently tagged antibodies, a multitude of protein markers is captured for each cell. Additionally,
476 cellular and morphological characteristics are extracted from microscopical images, such as the cell and nucleus
477 area and circularity. This spatially resolved phenotypic dataset is rich in molecular information and provides
478 insights into heterogeneous responses of thousands of cells to various drugs. Measuring different morphological
479 and signaling features captures pre-existing cell-to-cell variability which might influence perturbation effect,
480 resulting in various different responses. Some of these markers are of particular importance, as they provide
481 insights into the level of a cell’s growth or death as well as subpopulation identity. We utilized a mixture of
482 two melanoma tumor cell lines (M130219 and M130429) at a ratio of 1:1. The cell lines can be differentiated
483 by the mutually exclusive expression of marker proteins. The former is positive for Sox9, the latter for a set
484 of four proteins which are all recognized by and antibody called MelA (Raaijmakers et al., 2015). Cells were
485 seeded in a 384-well plate and incubated at 37C and 5% CO2 overnight. Next, the cells were exposed to multiple
486 cancer drugs and Dimethyl sulfoxide (DMSO) as a vehicle control for 8h and 24h after which the cells were
487 fixed and six cycles of 4i were performed TissueMAPS and the scikit-image library (Van der Walt et al., 2014)
488 were used to process and analyze the acquired images, perform feature extraction and quality control steps using
489 semi-supervised random forest classifiers.

490 **Data generation and processing.** Our datasets contain high-dimensional single-cell data of control and
491 drug-treated cells measured at two time points (8 and 24 hours). For both the 8h-dataset and the 24h-dataset, we
492 normalized the extracted intensity and morphological features by dividing each feature by its 75th percentile,
493 computed on the control cells. Additionally, values were transformed by a \log_{1p} function ($x \leftarrow \log(x + 1)$). In
494 total, our datasets consist of 48 features, of which 26 are protein marker intensities and the remaining 22 are
495 morphological features. For each treatment, we have measured between 2000 and 3000 cells. For training the
496 models, we perform a 80/20 train/test split. We trained all models on control and treated cells for each time step
497 and each drug separately. The considered drugs as well as their inhibition type can be found in Table 2.

498 **Cell type assignment.** We assigned M130219 and M130429 cells to the Sox9 and MelA cell types,
499 respectively, by first training a two component Gaussian mixture model on the features ‘intensity-cell-MelA-
500 mean’ and ‘intensity-nuclei-Sox9-mean’ of the control cells. Next, we used the aforementioned features and the
501 labels provided by the mixture model to train a nearest neighbor classifier, which we then used to predict the
502 cell type labels of the drug treated cells. The procedure was performed separately for the 8h- and 24h dataset.
503 Results of the classification can be found in Figure 10 and Figure 12 respectively.

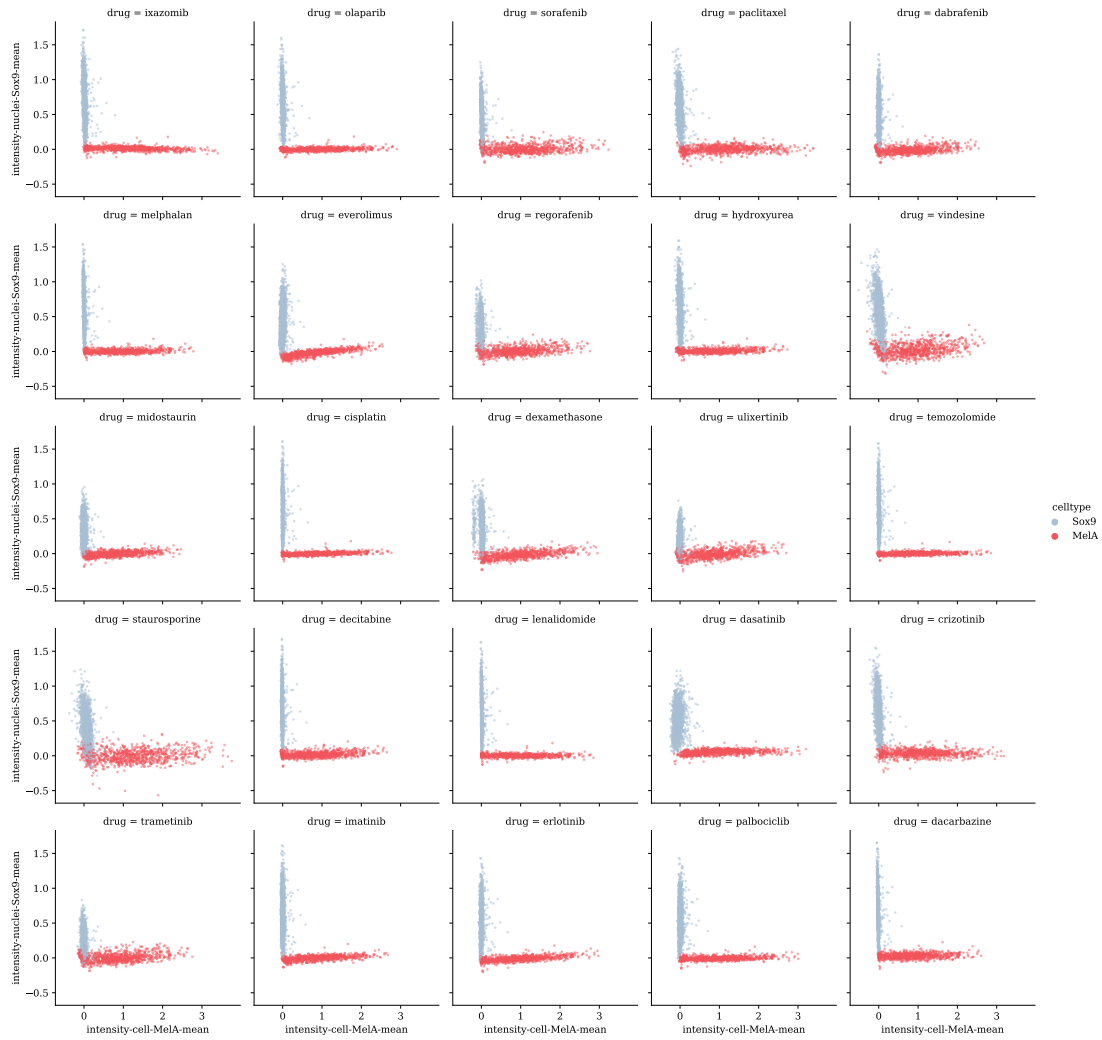


Figure 10: Classification of cells into cell types (MelA⁺, Sox9⁺) based on protein marker intensities of MelA and Sox9, for all drugs, at $t = 8h$ § E.2. Each tile represents one drug. MelA⁺ cells colored in red, Sox9⁺ in blue.

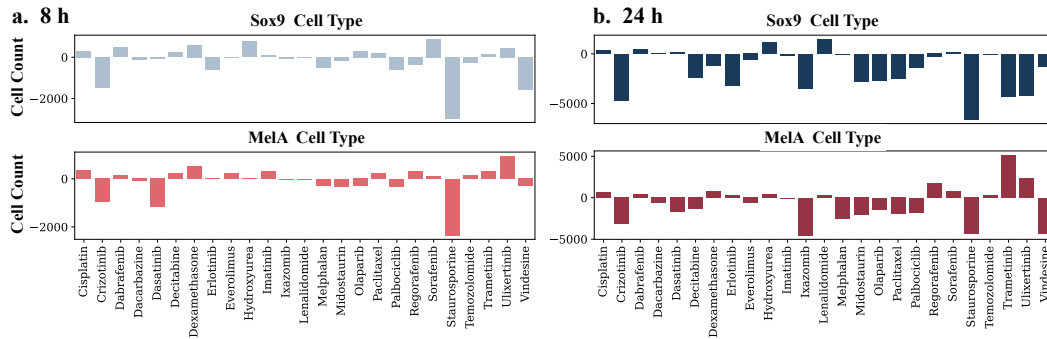


Figure 11: Drug treatment-induced change in cell counts in the two cell types compared to the cell count of the respective cell types in the control condition. **a.** Cell count change for cell types Sox9⁺ (top) and MelA⁺ (bottom) at $t = 8h$. **b.** Cell count change for cell types Sox9⁺ (top) and MelA⁺ (bottom) at $t = 24h$.

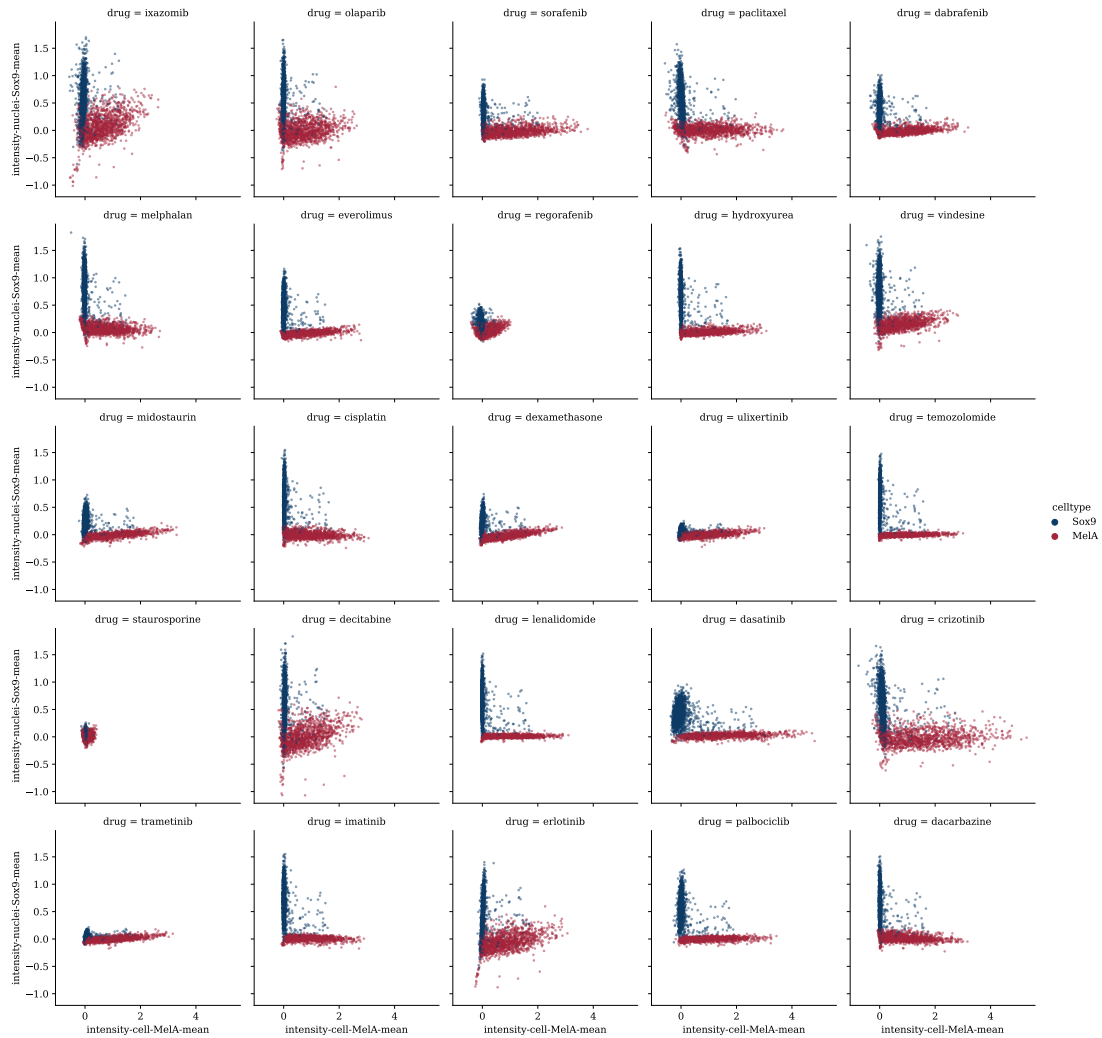


Figure 12: Classification of cells into cell types (MelA⁺, Sox9⁺) based on protein marker intensities of MelA and SOX9, for all drugs, at $t = 24h$ § E.2. MelA⁺ cells colored in red, Sox9⁺ in blue.




Figure 13: Observed cell counts of drug-treated cells normalized to control cell counts, per drug and time point. 8h treatment in light blue, 24h treatment in dark blue.

504 **F Experimental Details**

505 NUBOT consists of several modules and its performance is compared against several baselines. In the following,
 506 we provide additional background on experimental details, including a description of the evaluation metrics
 507 and baselines considered, as well as further information on the parameterization and hyperparameter choices
 508 made for NUBOT.

509 **F.1 Evaluation Metrics**

510 We evaluate our model by analyzing the distributional similarity between the predicted and observed perturbed
 511 distribution. For this, we compute the kernel maximum mean discrepancy (MMD) (Gretton et al., 2012). To take
 512 the mass variation into consideration, we compute a weighted version of MMD, by weighting each predicted
 513 point by its associated normalized weight.

514 **F.2 Baselines**

515 We compare NUBOT against several baselines. First, we consider a balanced neural optimal transport method
 516 CELLOT (Bunne et al., 2021). Further, we benchmark NUBOT against the current state-of-the-art UBOT
 517 GAN, an unbalanced OT formulation proposed by Yang & Uhler (2019). Additionally, we consider two naive
 518 baselines: IDENTITY, simulating the identity matching and modeling cell behavior in absence of a perturbation,
 519 and OBSERVED, a random permutation of the observed target samples and thus a *lower bound* when comparing
 520 predictions to observed cells on the distributional level. In the following, we briefly motivate and introduce each
 521 baseline.

522 **CELLOT.** By introducing reweighting functions η and ζ , NUBOT recovers a balanced problem parameterized
 523 by dual potentials f and g . An important ablation study to consider is thus to compare its performance to its
 524 balanced counterpart. Ignoring the fact that the original problem includes cell death and growth, and thus varying
 525 cell numbers, we apply ideas developed in Makuva et al. (2020); Bunne et al. (2021) and learn a balanced OT
 526 problem via duals f and g . These duals are parameterized by two ICNNs and optimized in objective (9) via an
 527 alternating min-max scheme.

528 **UBOT GAN.** Using (10), Yang & Uhler (2019) propose to model mass variation in unbalanced OT via a
 529 relaxation of the marginals. Similar to Fan et al. (2021a), Yang & Uhler (2019) reformulate the constrained
 530 Monge problem (7) as a saddle point problem with Lagrange multiplier h for the constraint $T_{\sharp}\mu = \nu$, i.e.,

$$\begin{aligned} \sup_h \inf_T \int_{\mathcal{X}} c(x, T(x))\mu(x)dx + \int_{\mathcal{X}} h(y) (\nu - T_{\sharp}\mu) dy \\ = \int_{\mathcal{X}} [c(x, T(x)) - h(T(x))]\mu(x)dx + \int_{\mathcal{X}} h(y)\nu(y)dy, \end{aligned}$$

531 parameterizing T and h via neural networks. To allow mass to be created and destroyed, Yang & Uhler (2019)
 532 introduce scaling factor $\xi : \mathcal{X} \rightarrow \mathbb{R}^+$, allowing to scale mass of each source point x_i . The optimal solution then
 533 needs to balance the cost of mass and the cost of transport, potentially measured through different cost functions
 534 $c_1 : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^+$ (cost of mass transport) and $c_2 : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ (cost of mass variation). Parameterizing the
 535 transport map T_{θ} , the scaling factors ξ_{ϕ} , and the penalty h_{ω} with neural networks, the resulting objective is

$$l(\theta, \phi, \omega) := \frac{1}{n} \sum_{i=0}^n [c_1(x_i, T_{\theta}(x_i))\xi_{\phi}(x_i) + c_2(\xi_{\phi}(x_i)) + \xi_{\phi}(x_i)h_{\omega}(T_{\theta}(x_i)) - \Psi^*(h_{\omega}(y_i))],$$

536 with Ψ^* approximating the divergence term of the relaxed marginal constraints (see (10)), and is optimized via
 537 alternating gradient updates.

538 **IDENTITY.** A trivial baseline is to compare the predictions to a map which does not model any perturbation
 539 effect. The IDENTITY baseline thus models an identity map and provides an *upper bound* on the overall
 540 performance, also considered in Bunne et al. (2021).

541 **OBSERVED.** In a similar fashion we might ask for a *lower bound* on the performance. As a ground truth
 542 matching is not available, we can construct a baseline for a comparison on a distributional level by comprising a
 543 different set of observed perturbed cells, which only vary from the true predictions up to experimental noise. The
 544 closer a method can approach the OBSERVED baseline, the more accurate it fits the perturbed cell population.

545 **F.3 Hyperparameters**

546 We parameterize the duals f and g using ICNNs with 4 hidden layers, each of size 64, using ReLU as activation
 547 function between the layers. We choose the identity initialization scheme introduced by Bunne et al. (2022b)

548 such that ∇g and ∇f resemble the identity function in the first training iteration. As suggested by [Makkuva](#)
549 [et al. \(2020\)](#), we relax the convexity constraint on ICNN g and instead penalize its negative weights W_i^z

$$R(\theta) = \lambda \sum_{W_i^z \in \theta} \|\max(-W_i^z, 0)\|_F^2.$$

550 The convexity constraint on ICNN f is enforced after each update by setting the negative weights of all $W_i^z \in \theta_f$
551 to zero. Duals g and f are trained with an alternating min-max scheme where each model is trained at the same
552 frequency. Further, both reweighting functions η and ζ are represented by a multi-layer perceptron (MLP) with
553 two hidden layers of size 64 for the single-cell and of size 32 for the synthetic dataset, with ReLU activation
554 functions. The final output is further passed through a softplus activation function as we do not assume negative
555 weights. For the unbalanced Sinkhorn algorithm, we choose an entropy regularization of $\varepsilon = 0.005$ and a
556 marginal relaxation penalty of 0.05. We use both Adam for pairs g and f as well as η and ζ with learning rate
557 10^{-4} and 10^{-3} as well as $\beta_1 = 0.5$ and $\beta_2 = 0.9$, respectively. We parameterize both baselines with networks
558 of similar size and follow the implementation proposed by [Yang & Uhler \(2019\)](#) and [Bunne et al. \(2021\)](#).