Comparing phenotypic manifolds with Kompot: Detecting differential abundance and gene expression at single-cell resolution

Dominik J. Otto[1,2,3], Erica Arriaga-Gomez[4], Elana Thieme[1,2,3,4,5], Ruijin Yang[1,2,3], Stanley C. Lee[4]+, Manu Setty[1,2,3]+

[1] Basic Sciences Division, Fred Hutchinson Cancer Center, Seattle WA

[2] Computational Biology Program, Public Health Sciences Division, Seattle WA

[3] Translational Data Science IRC, Fred Hutchinson Cancer Center, Seattle WA

[4] Translational Science and Therapeutics Division, Fred Hutchinson Cancer Center, Seattle WA

[5] Molecular and Cellular Biology Program, University of Washington, Seattle WA

+ Corresponding authors:

Manu Setty (msetty@fredhutch.org)

Stanley Lee (sclee3@fredhutch.org)

Abstract

Kompot is a statistical framework for holistic comparison of multi-condition single-cell datasets, supporting both differential abundance and differential expression. Differential abundance captures changes in how cells populate the phenotypic manifold across conditions, while differential expression identifies condition-specific changes in gene regulation that may be localized to particular regions of that manifold. Kompot models the distribution of cells and gene expression as continuous functions over a low-dimensional representation of cell states, enabling single-cell resolution inference with calibrated uncertainty estimates. Applying Kompot to aging murine bone marrow, we identified a continuum of shifts in hematopoietic stem cell and mature cell states, transcriptional remodeling of monocytes independent of compositional changes, and divergent regulation of oxidative stress response genes across cell types. By capturing both global and cell-state–specific effects of perturbation, Kompot reveals how aging reshapes cellular identity and regulatory programs across the hematopoietic landscape. This framework is broadly applicable to dissecting condition-specific effects in complex single-cell landscapes.

Introduction

Multi-condition single-cell datasets offer new opportunities to dissect how mutations or other perturbations affect both differentiation trajectories and disease phenotypes[1-4]. Systematic analyses of such datasets require characterization of two key aspects of condition-induced variability: differential abundance and differential expression.

Differential abundance refers to differences in the composition of cell states between two conditions (**Fig. 1A, B**). Even in the absence of overt phenotypic shifts such as loss or gain of broad cell types, changes in the frequency of specific cell states can provide early indicators of disease[5, 6] and response to treatment[7, 8]. These changes are often nuanced and cell-type specific, necessitating statistical approaches that rigorously assess significance while accounting for underlying biological variability.

Perturbations can also reshape the phenotypic landscape itself through differential expression, where gene expression changes that would be unlikely or impossible in unperturbed conditions can emerge[9] (**Fig. 1A, C**). These transcriptional responses reveal condition-specific molecular programs and regulatory networks that often manifest as subtle yet systematic changes across multiple states[9] and can frequently be masked by the inherent variability within biological systems.

Several computational approaches have been developed for comparative analysis of multi-condition single-cell data with differential abundance receiving greater attention[10-15]. A number of these methods rely on discretizations such as clustering cells into discrete groups to stabilize variability estimates. While facilitating interpretation, such discretizations lead to loss in resolution when changes are gradual or occur along continuous trajectories (**Supp. Fig. 1A**). Additionally, many existing methods lack principled significance measures, making it difficult to distinguish genuine biological signals from technical noise.

Here, we introduce Kompot, a novel computational framework that provides a continuous and statistically rigorous approach for differential abundance and differential expression testing using multi-condition single-cell data. Kompot leverages Gaussian Process (GP) modeling to estimate continuous functions across the cell-state space. This approach enables quantification of differential abundance and differential expression at single-cell resolution, while incorporating uncertainty estimates to provide statistically calibrated comparisons. By addressing both aspects of multi-condition analysis within a unified statistical framework, Kompot provides a comprehensive toolkit for extracting biological insights from increasingly complex single-cell datasets across a range of biological systems.

Results

Overview of the Kompot modeling framework

A single-cell RNA-seq dataset provides a snapshot of cellular states in a biological system, with the molecules measured representing a sample of the transcriptome of the cell. To characterize the system and compare it to similar systems under different conditions (such as mutations or other perturbations), we need to understand what this data tells us about the true composition of transcribed genes in a cell and the true distribution of cell states in the system. With Kompot we leverage high-dimensional, continuous cell-state representations and Bayesian Inference to learn both, the composition of the system and transcriptomes of each cell, enabling a holistic comparison of biological systems at single-cell resolution with uncertainty quantification. This comparison is divided into differential cell-state abundance that reflect changes in cell-state composition, and differential-expression that reflect changes in the transcriptome of cells in equivalent states.

The input to Kompot is a latent representation of co-embedded multi-condition single-cell data. This latent representation can incorporate batch effect correction[16, 17] to ensure that cell states considered equivalent share similar locations in that state space (**Supp. Fig. 1B**). From this foundation, Kompot employs diffusion maps[18] to derive a cell-state representation with biologically meaningful structure (**Fig. 1A**).

Diffusion maps are particularly effective for single-cell analysis as they capture the intrinsic geometry of the data while providing biologically meaningful measures of cell-to-cell distance[19, 20]. This distance metric supports both accurate local abundance estimation and principled information sharing when mapping from cell-state space to gene expression, enabling the statistical comparisons at the core of Kompot's differential testing framework.

Kompot differential abundance testing

To quantify differences in cell-state composition between conditions, Kompot first constructs condition-specific density functions on the diffusion map representation. For each condition, Kompot computes a separate continuous density function (**Fig. 1B, Supp. Fig. 2A-B**) using our Mellon framework[21]. Mellon leverages nearest-neighbor distances within a Gaussian Process model to estimate continuous density functions that capture the distribution of cells across the state space[21]. The continuous nature of these density estimates enables evaluation at any position within the cell-state space, including states predominantly observed in the alternate condition (**Supp. Fig. 2B**). By computing the ratio of density estimates at corresponding positions in the cell-state space, Kompot derives a measure of differential abundance that quantifies the fold-change in cell density between conditions at single-cell resolution (**Supp. Fig. 2C**).

While the ratio of density estimates provides a measure of fold-change, rigorous comparison between conditions requires statistical assessment of significance. Kompot addresses this challenge by computing uncertainty estimates for each density function. These uncertainty estimates account for two critical factors: limited data availability in regions of the cell-state space and inherent uncertainty through the sparse function representation. By incorporating these uncertainty estimates into the statistical framework, Kompot calculates posterior tail probabilities (PTPs), an alternative for p-values, for differential abundance at each cell state. This enables the generation of volcano plots for differential abundance testing where each point represents a single cell-state (**Fig. 1B**), allowing for threshold-based significance assessments. Thus, Kompot can identify subtle cell-state transitions and rare subpopulations that would likely be undetected by clustering-based approaches, providing insights into the continuous spectrum of cellular responses to perturbations (**Fig. 1D**).

Kompot differential expression testing

For differential expression testing between conditions, Kompot constructs condition-specific mappings from cell-states to gene expression space (**Fig. 1A, C**). These mappings are computed using Gaussian Processes (GPs)[22], by predicting log-transformed gene expression as a continuous function of the diffusion map coordinates. The non-linear relationship between diffusion maps and gene expression is effectively captured by GPs[22], essentially serving as a rigorous approach to gene expression imputation. The continuous nature of these mapping functions allows Kompot to estimate gene expression in one condition for cell states observed in another, producing counterfactual expression profiles that reflect how those same cell states would behave under different conditions (**Fig. 1C**). By computing the ratio between the actual expression in a cell and its counterfactual expression, Kompot quantifies gene expression-fold changes for every gene at single-cell resolution across the entire phenotypic landscape (**Fig. 1C**).

Given the large number of genes evaluated, we require a robust significance measure that allows us to identify the most relevant genes, accounting for both uncertainty in the gene expression function and statistical interdependence of imputed values. Kompot therefore employs the Mahalanobis distance, which measures the distance of a point from a multivariate normal distribution, considering both variance and covariance[23]. It effectively measures how many standard deviations a gene expression shift is away from expected variation, while accounting for the correlation structure between neighboring cell states,

making it an ideal significance measure for differential expression. Kompot therefore uses Mahalanobis distances as the score in a quasi-volcano plot (**Fig. 1C**) to identify significantly differentially expressed genes either in the full population of cells or in specific subsets (**Fig. 1D**).

Kompot detects cell-composition changes across conditions without relying on predefined populations

Aging is a compelling biological context for evaluating differential abundance methods, as it induces progressive and often subtle shifts in cellular composition rather than dramatic phenotypic transformations. Hematopoiesis has emerged as a particularly informative system for studying aging[24], characterized by declining stem cell function[25], reduced lymphocyte production[26], and increased susceptibility to myeloid malignancies[27]. To investigate these age-related changes, we generated a novel scRNA-seq dataset of murine bone marrow at 10, 63, and 103 weeks of age, representing young, middle-aged, and older timepoints (**Supp. Fig. 3A**). After quality control verification (**Supp. Fig. 3B, C**), we successfully identified most hematopoietic cell types across all ages (**Fig. 2A, B, Supp. Fig. 4**), providing us with a rich dataset to characterize cell-state changes and gene expression changes in aging hematopoiesis using Kompot.

We applied Kompot differential abundance testing on our aging hematopoiesis dataset to determine age-dependent composition changes (**Fig. 2C-E**). We first derived a unified cell-state representation for all cells using diffusion maps. We then calculated age-specific density functions using the subsets of cells from each timepoint (**Supp. Fig. 5A-B**). We evaluated these individual density functions at all cell-states in the dataset to derive condition-specific densities and associated uncertainties (**Supp. Fig. 5C-D**). This approach enabled pairwise comparisons that quantify age-related changes in cell-type composition (**Fig. 2C-D**). A key advantage of our framework is the ability to set thresholds based on both effect size and statistical significance at single-cell resolution, allowing for identification of cell states with significant differential abundance rather than relying solely on fold-change cutoffs (**Fig. 2C-E**).

Our analysis revealed significant expansions and contractions in distinct cell populations over time. Comparing young to old mice, we observed an expansion of the hematopoietic stem cell (HSC) compartment and a marked shift from naïve to memory cells in both B-cell and T-cell compartments (**Fig. 2F-G)**. The depletion of naïve B-cells and an enrichment of memory B-cells in older mice is a well-characterized phenomenon of hematopoietic aging[26]. This shift reflects fundamental changes in differentiation dynamics: naive B-cells rely on a continuous influx from hematopoietic progenitors, whereas memory B-cells accumulate over time due to antigen exposure. The observed decline in naïve B-cells reflects the reduced lymphoid differentiation potential of aged HSCs[28], while the expansion of memory B-cells, including age-associated B-cells, is driven by persistent immune activation[29]. Similarly, the T-cell compartment displayed a significant reduction in naïve T-cells with a concurrent increase in memory T-cells, a hallmark of immunosenescence that contributes to diminished adaptive immune responses in aging[30].

An increased accumulation of hematopoietic stem cells (HSCs) is the most significant change we observed in the progenitor cells (**Fig. 2F-G**). Progressive enrichment of HSCs with age is a well-documented phenomenon[31]. HSCs reside at the top of the hematopoietic hierarchy, maintaining blood production through self-renewal and differentiation into mature immune and blood cells. With aging, HSC function declines, leading to accumulation in the bone marrow while their ability to differentiate into different lineages becomes impaired[31]. While the accumulation of aged HSCs has been previously observed, it has traditionally been quantified using bulk measurements or clustering-based approaches. We took advantage of the single-cell resolution of Kompot to better characterize the age-induced change in HSCs. By comparing the differential abundance fold-change along pseudotime ordering[19], we observed that the expansion of HSC population is in fact a subtle shift in cell-state, rather than a global

168 increase in a predefined population (**Fig. 2H**). Specifically, HSCs in older mice progressively shift towards
169 a less differentiation-primed state, creating a continuum of altered stem cell states rather than a uniform
170 expansion (**Fig. 2H**). This subtle but consequential repositioning along the differentiation trajectory
171 provides a possible explanation for the functional decline in aged HSCs, revealing how changes in cell
172 abundance are accompanied by fundamental alterations in cellular function and developmental potential.

173 Comparing differential abundance results across ages, we observed that changes in mature cell
174 populations predominantly occurred between young and middle-aged mice (**Fig. 2D, Supp. Fig. 6A**),
175 while alterations in the HSC compartment were more pronounced between middle-aged and older mice
176 (**Fig. 2E, Supp. Fig. 6B**). This temporal separation suggests a sequential progression of aging, where
177 compositional shifts in differentiated cells precede fundamental changes in stem cell states. The early
178 alterations in mature immune cells may represent initial adaptive responses to aging, while the later HSC
179 changes potentially indicate more persistent reprogramming of the hematopoietic system.

180 Our results demonstrate the ability of Kompot to identify subtle but biologically meaningful abundance
181 changes between conditions. By combining a continuous representation of cell states with robust
182 statistical measures, Kompot enables precise detection of population-level changes while maintaining
183 single-cell resolution.

184

185 Kompot captures context-specific gene-expression shifts

186 We applied Kompot differential expression testing on our mouse aging hematopoiesis dataset to
187 characterize age-related transcriptional changes (**Fig. 3**). We first inferred age-specific mapping functions
188 from diffusion map coordinates to log-normalized gene expression values, using the subset of cells from
189 each timepoint. We then used these mappings to impute gene expression and associated uncertainty at
190 all cell-states across the dataset. This enabled us to compute gene-level fold changes by comparing
191 observed expression values to their counterfactual predictions under alternative age conditions. To
192 assess significance, we calculated Mahalanobis distances for each gene across specified groups of cells,
193 quantifying deviations from expected variation while accounting for the covariate structure among
194 neighboring cell states. We excluded T and NK cells for this analysis since do not differentiate in the bone
195 marrow (**Fig. 3H**).

196 We first applied our approach to characterize gene expression changes in HSCs, a population known to
197 undergo functional and transcriptional changes with age[31], by comparing HSCs from middle-aged to old-
198 aged mice (**Fig. 3A-B**). Aging in HSCs has been extensively studied, and the Hematopoietic Aging Atlas
199 has synthesized results from multiple datasets to identify genes consistently up- or downregulated with
200 age[32]. For each gene, this atlas assigns a consistency score that reflects the reproducibility of age-
201 associated expression changes across studies[32]. Kompot's differential expression results strongly
202 prioritized genes with the highest consistency scores in the Hematopoietic Aging Atlas, with genes
203 showing large fold changes and Mahalanobis distances highly enriched for reproducible aging signatures
204 (**Fig. 3C**, **Supp. Fig. 7**).

205 Our differential abundance analysis revealed that HSCs in older mice shift towards a less differentiated
206 state along the cell-state continuum (**Fig. 2H**). We next asked whether the genes differentially expressed
207 in aged HSCs reflected this shift. Indeed, we found that gene expression fold changes were structured
208 along the same continuum: genes exhibited larger fold changes in cell states that were more differentially
209 abundant (**Fig. 3D**). These results highlight the transcriptional changes in HSCs that accompany and
210 may underlie the shift toward a less differentiated state with age,

211 We next investigated age-related gene expression changes in mature cell types. We focused on changes
212 from young to old mice since the abundance changes in mature cell types were most pronounced in this
213 comparison (**Fig. 2D-F**). To distinguish expression changes from abundance-driven effects, we restricted

our analysis to cell types that *did not show significant compositional changes*. Monocytes exhibited clear transcriptional changes despite stable abundance (**Fig. 3E-F)**. Gene ontology analysis revealed that downregulated genes were enriched in core monocyte/macrophage functions such as defense response (**Fig. 3G**), consistent with their reduced functional capacity with age[33]. In contrast, upregulated genes were enriched for antigen processing and presentation pathways, driven by increased expression of MHC class II genes (**Fig. 3G**), which have been implicated in age-related immune remodeling[34]. Notably, our results recapitulate findings from a previous study that identified H2-Aa, H2-Ab1, H2-Eb1, Cd74, and Aw112010 as the top dysregulated genes in aging monocytes across both mice and humans, with only Igkc surpassing H2-Eb1 in our ranking[34]. The near-perfect agreement with experimentally validated gene sets underscores the power of our differential expression testing approach to detect biologically meaningful transcriptional shifts. Gene expression differences in other mature cell types were not as pronounced but did show a tendency for upregulation of MHC Class II genes similar to monocytes (**Supp. Fig. 8**).

Our use of Mahalanobis distance as a significance measure enables Kompot to identify genes that exhibit coordinated but opposing changes across subpopulations. For example, consider a gene that is upregulated in one cell type and downregulated in another. While the net fold change may be close to zero, the Mahalanobis distance remains high, reflecting a statistically significant deviation from expected variation if unaffected. This occurs because such a pattern is unlikely under the local covariance structure of the cell-state space, where neighboring states are typically expected to exhibit similar expression patterns. As a result, Kompot can detect structured differential expression that would be missed by conventional fold-change-based approaches.

To test this capability, we applied Kompot differential expression testing between young and old cells across all cell-types that differentiate in the bone marrow (**Fig. 3H-I**). The top genes with both high positive fold change and high Mahalanobis distance were strongly enriched for MHC class II genes, consistent with our monocyte-specific results and indicating that MHC class II upregulation is a general feature of hematopoietic aging (**Fig. 3J**). Likewise, genes with high Mahalanobis distance and strong negative fold change were consistently downregulated across multiple cell types (**Fig. 3J**). In contrast, genes ranked highly by Mahalanobis distance but with low absolute fold change exhibited divergent expression patterns across cell types, with upregulation in some and downregulation in others (**Fig. 3J**). To illustrate these distinctions, we visualized expression levels of representative genes from each category across the cell-state manifold (**Fig. 3K**). Genes selected by fold change showed broad, monotonic shifts in expression, whereas Mahalanobis-selected genes revealed distinct, cell-type-specific regulation (**Fig. 3K**).

We next performed gene ontology analysis to assess the biological relevance of genes with high Mahalanobis distance but low absolute fold change. Antioxidant activity emerged as the most significantly enriched term (**Fig. 3L**), a pathway known to be impacted by aging[35]. Notably, our results indicate that while antioxidant activity is broadly impacted across hematopoietic cell types, the direction of expression changes of the pathway genes varies across genes and cell types, suggesting a context-specific adaptation to aging (**Fig. 3M**). Oxidative stress is a known driver of HSC dysfunction, promoting DNA damage, apoptosis, and biased differentiation toward myeloid lineages[35]. The differential expression of antioxidant pathway genes observed in our study likely represents a compensatory mechanism to mitigate reactive oxygen species (ROS) accumulation in aged HSCs[35].

These results demonstrate that Mahalanobis distance captures biologically meaningful heterogeneity in gene expression and provides a powerful approach for aggregating structured expression differences across cell states without requiring prior clustering, making it especially suited for complex, continuous single-cell landscapes.

Differential Gene Expression in Mouse Embryonic Development

As a further test of Kompot's Mahalanobis Distance to detect expression differences across cell states without clustering, we applied Kompot differential expression to a single-cell RNA-seq dataset of mouse embryos at mid-gestation with wild-type and mutant *Tal1*[36] (**Supp. Fig. 9A, B**). *Tal1* is required for the generation of primitive erythroid, megakaryocyte, and myeloid lineages[36], but its effects on other lineages is less characterized. We therefore performed differential expression analysis between *Tal1* wildtype cells to *Tal1*$^{-/-}$ cells specifically within lineages not affected by abundance changes (**Supp. Fig. 9A, B**). As expected, hemoglobin genes showed the greatest changes in the mutant cells along with sex-associated genes such Xist (**Supp. Fig. 9C**). These signals likely reflect technical and experimental factors: hemoglobin transcripts contribute substantially to ambient RNA in single-cell data, and Tal1 mutant chimeras were generated exclusively from female donor cells, resulting in differential detection of imprinting-associated genes (**Supp. Fig. 9C**). After accounting for these confounders, we next checked whether Kompot highlights biologically meaningful differences in the remaining lineages.

We excluded the hemoglobin and imprinting-related genes and selected the top 500 differentially expressed genes ranked by Mahalanobis Distance for downstream analysis (**Supp. Fig. 9D**). Gene ontology analysis revealed significant enrichments for processes such as vasculature development and blood vessel development (**Supp. Fig. 9E**). Genes associated with these processes showed pronounced expression differences between wild-type and mutant cells, particularly within endothelial and hematoendothelial progenitor populations (**Supp. Fig. 9F)**, consistent with the known role of Tal1 in regulating these programs[36]. Further, the magnitude and direction of these changes varied across related cell states (**Supp. Fig. 9F, G**), suggesting a context-dependent role for Tal1 in modulating vasculature-associated gene expression.

Pathway enrichment analysis of the same gene set also highlighted effects of Tal1 loss beyond hematoendothelial populations (**Supp. Fig. 9H**). Genes associated with fluid shear stress pathways were broadly downregulated across multiple lineages (**Supp. Fig. 9I**), suggesting that blood flow during embryogenesis influences not only nutrient delivery but also the physical state of developing tissues. In addition, we observed widespread changes in cyclin gene expression following Tal1 knockout (**Supp. Fig. 9J**). Although Tal1 has been implicated in cell cycle regulation of hematoendothelial cells, these results suggest more indirect effects on cell cycle–associated programs across the embryo.

Rather than performing differential expression on a lineage-by-lineage or cluster-by-cluster basis, Kompot tests for structured expression differences across all cells simultaneously. This enables detection of both lineage-specific programs, such as vasculature-associated changes in endothelium, and broader patterns, such as altered cell cycle and shear-stress pathways across multiple lineages. Thus, these results illustrate how Kompot differential expression can recover biologically meaningful changes at multiple levels of organization without requiring predefined clusters.

Kompot reveals determinants of immunotherapy in advanced melanoma

We next used Kompot to examine how immunotherapy reshapes CD8 T cell states, leveraging its ability to quantify abundance and expression differences without imposing discrete cluster boundaries. Combination treatment with anti-LAG-3 (Relatlimab, Rela) and anti-PD-1 (Nivolumab, Nivo) has been shown to provide greater clinical efficacy than nivolumab alone in advanced melanoma[37]. Cillo et. al.[37] profiled CD8 T cells from patients in a clinical trial of these therapies using single-cell RNA-seq (**Fig. 4A**) and identified that combination therapy promotes an altered differentiation trajectory with increased cytotoxicity and exhaustion contributing to improved efficacy[37]. Using Kompot, we reanalyzed baseline and on-treatment samples within each therapy arm (**Fig. 4B**) to determine whether our cluster-free, continuous framework could recover complementary insights to those reported previously.

Kompot differential abundance analysis revealed significant but heterogeneous shifts across treatment arms relative to baseline (**Fig. 4C, Supp. Fig. 10**). In Nivo monotherapy, we observed a decrease in terminally exhausted cells, whereas Rela treatment led to an increase in TCR-activated exhausted states (**Fig. 4D, Supp. Fig. 10C**). The most prominent change under combination therapy was an expansion of activated cycling CD8 T cells (**Fig. 4D, Supp. Fig. 10C**). Given their central role in tumor cell killing[38], the expansion of activated cycling CD8 T cells represents a striking cellular correlate of the enhanced efficacy observed with combination therapy in advanced melanoma. We therefore next asked whether these therapy-specific abundance shifts detected by Kompot are mirrored in transcriptional programs across the CD8 T cell landscape.

T cell activation and exhaustion trajectories are continuous, producing related but heterogeneous states that complicate comparisons of expression programs using conventional approaches[39]. Kompot addresses this by quantifying covarying differences without clustering or discretization. Using Kompot differential expression, we identified the top 250 genes altered between baseline and treatment within each therapy arm (**Supp. Fig. 11A**). The union of these genes showed broad changes across therapies (**Fig. 4E**), with fold changes generally consistent across states within each arm (**Fig. 4F**). Nivo monotherapy more closely resembled combination therapy than Rela alone (**Fig. 4F**). Interestingly, activated cycling and effector states diverged the most, showing expression changes distinct from those induced by combination treatment (**Fig. 4F**). These results suggest that treatment-specific differences in efficacy may, in part, reflect cell-state–specific transcriptional responses within these key effector populations.

To further interpret these transcriptional differences, we performed gene ontology analysis of the top differentially expressed genes (**Supp. Fig. 11B**). This revealed significant enrichment for translation and T cell activation programs (**Supp. Fig. 11C**). T cell activation genes were upregulated in activated cycling cells under combination therapy but downregulated in the same states under Nivo monotherapy (**Fig. 4G-I, Supp. Fig. 11D**). Ribosomal genes displayed similarly opposing patterns (**Supp. Fig. 11E**). These results indicate that immunotherapy regimens elicit distinct and even opposing gene expression changes within the same effector populations, providing a mechanistic basis for the differential efficacy observed across treatments.

Together, these results show that Kompot provides an effective solution to the challenge of analyzing therapy-induced changes in continuous CD8 T cell landscapes. By jointly testing for abundance and expression differences without predefined clusters, Kompot recovered coherent, state-specific signals across overlapping populations. This underscores Kompot's utility for resolving condition-specific programs in single-cell datasets in which subtle, state-specific changes are distributed across continuous manifolds.

Resolving inter-patient variability with multi-sample differential testing

Multi-sample, multi-condition single-cell datasets present both an opportunity and a challenge: they capture shared therapeutic signals while also reflecting patient-to-patient variability. Most approaches trade off between robustness and resolution, either aggregating away heterogeneity or losing power across samples. Kompot addresses this by extending differential abundance and expression testing to the multi-sample setting, estimating condition-specific variance across samples at single-cell resolution and condition-specific covariance across samples at single-gene resolution (**Fig. 5A**). This enables statistically principled detection of both shared patterns and patient-specific differences in large single-cell cohorts.

We applied Kompot to a recently published single-cell atlas of high-risk neuroblastoma[40], comprising over 200,000 neoplastic cells from 22 patients profiled at diagnosis and after treatment (**Fig. 5B-E, Supp. Fig.**

352 **12A**). In this study, neoplastic cells were harmonized and assigned to six transcriptional states (**Fig. 5B,**
353 **C**). Broad therapy-associated shifts in state proportions between diagnosis and treatment were reported,
354 including increases in adrenergic (ADRN)-calcium, ADRN-dopaminergic, Interm-OXPHOS, and MES
355 states and decreases in ADRN-baseline and ADRN-proliferating states[40]. Using Kompot, we recovered
356 these reported shifts (**Fig. 5F, G**) and further resolved heterogeneity within ADRN-baseline and ADRN-
357 proliferating states: while most patients showed decreases after treatment, a subset—particularly those
358 with ALK mutations—displayed increases instead (**Fig. 5H, Supp. Fig. 12B**).

359 To test whether this patient-level heterogeneity reflects underlying biology, we categorized patients as
360 *Event Free* if they remained free of progression, relapse, or death and as *Event* otherwise (**Fig. 5E**). We
361 then ordered patients by the percentage of states showing significant differential abundance in ADRN-
362 baseline and ADRN-proliferating states (**Fig. 5H)**, the adrenergic programs that dominate at diagnosis
363 and whose persistence has been linked to poor outcomes[40]. We observed that patients with significant
364 reductions in these states were almost uniformly Event Free, with patient 14 as the sole exception.
365 Conversely, patients who either retained these states without significant changes in any cell state, or
366 exhibited increases in these states were all categorized as Event (**Fig. 5H)**. Moreover, all Event Free
367 patients displayed either reductions in these aggressive states or, in the case of patients 17, 20, and 22,
368 compensatory increases in alternative states (**Fig. 5H**), suggesting treatment-induced shifts away from
369 high-risk populations.

370 Kompot's multi-sample analysis therefore recapitulates broad, cohort-level trends while also resolving
371 patient-specific deviations that align with clinical outcomes. By capturing both shared and heterogeneous
372 abundance changes within a single unified framework, Kompot provides a general approach for
373 connecting variability in large-scale single-cell studies to biologically and clinically meaningful differences.

374

375 ## Discussion

376 Kompot provides a flexible statistical framework for quantifying condition-specific effects in single-cell
377 data, modeling both cell abundance and gene expression as continuous functions over the cell-state
378 manifold. By estimating condition-specific densities and expression profiles using Gaussian processes,
379 Kompot performs differential analysis at single-cell resolution while accounting for uncertainty and spatial
380 structure in the data. This approach supports biologically grounded interpretations of perturbation effects,
381 whether they manifest as global shifts in population structure, localized transcriptional responses, or
382 context-specific coordination between abundance and expression changes.

383 Applying Kompot to aging murine bone marrow revealed both well-established and previously
384 unrecognized features of hematopoietic aging. Kompot recapitulated known expansions in memory
385 lymphocytes and age-associated B cells, while resolving more subtle shifts in hematopoietic stem cell
386 states that do not correspond to discrete population changes. In the transcriptomic landscape, Kompot
387 uncovered gene expression programs that are either conserved or cell-type specific, including divergent
388 regulation of oxidative stress pathways. These findings demonstrate how modeling continuous
389 phenotypic structure supports more nuanced detection of biological change. Beyond aging, Kompot is
390 broadly applicable to studies of development, disease progression, and response to perturbation, where
391 condition-induced effects often unfold along continuous, high-dimensional trajectories.

392

393

394

Methods

Kompot framework

Kompot is a toolkit for comparing single-cell phenotypic manifolds across conditions. It provides frameworks for both differential abundance testing and differential gene expression testing at single-cell resolution without requiring predefined cell-types or clusters. The main assumption underlying Kompot is that both cell-state abundance and gene expression change continuously along cell-state space.

The input to Kompot is a latent representation of co-embedded multi-condition single-cell data (e.g. PCA, Harmony[16], scVI[41], LSI[42]). When required, batch-effect correction can be applied before co-embedding to ensure that condition-specific differences in the latent space are biologically meaningful rather than technical[17]. Using this latent representation as input, Kompot constructs diffusion maps to represent cell states, which provide biologically meaningful measures of distance between cells[18, 20, 43], Alternatively, Kompot can operate on any cell-state representation with a well-defined distance metric.

Differential Abundance Testing

Kompot differential abundance testing enables the detection of shifts in the abundance of cell-states across the phenotypic manifold. Using the co-embedded diffusion map representation, Kompot first constructs condition-specific density functions for each condition using Mellon[21]. These density functions describe how cells from each condition are distributed across the phenotypic manifold. Therefore, we can derive measures of differential abundance by comparing how these densities change between two conditions. Finally, since density functions are continuous, Kompot can perform these comparisons across the entire phenotypic landscape, enabling cluster- or cell-type–free differential abundance testing at single-cell resolution.

Our density quantification relies on Gaussian Process (**GP**) estimation implemented in Mellon[21]. While Mellon was previously introduced as a framework for continuous density modeling, here we extend it with an explicit uncertainty estimation at each cell state. This Bayesian inference procedure yields both a "best estimate" function that maps each point in the cell-state space, x_i to a log density value $\mu(x_i)$, and a quantification of uncertainty $\sigma(x_i)$ separately for each condition. Theoretical descriptions of the source of uncertainty and implementation are in **Supp. Notes 1**.

By comparing two conditions a and b, the posterior distribution of the change in density any cell state x_i is given by

$$\delta(x_i) \sim \mathcal{N}\left(\mu_b(x_i) - \mu_a(x_i), \sigma_a^2(x_i) + \sigma_b^2(x_i)\right).$$

Therefore, the density log-fold change between conditions is defined as $\mu_b(x_i) - \mu_a(x_i)$, with variance $\sigma_a^2(x_i) + \sigma_b^2(x_i)$.

Normalization for comparability across conditions

Mellon density estimates are in the units of log number of cells per volume and a result, direct comparisons of densities between conditions with different number of cells will not be accurate. Therefore, Kompot transforms these values into log fraction of cells per volume by subtracting the log total number of cells using

$$\mu'_a(x_i) = \mu_a(x_i) - \log N_a$$

where N_a represents the total number of cells measured for condition a. This normalization ensures that density values reflect the relative occupancy of the phenotypic manifold, independent of number of cells. For this to be valid, Kompot assumes that the intrinsic dimensionality of the manifold is estimated correctly. Mellon provides this estimate, and it is treated as constant across the dataset. Thus, the normalized local log-fold change representing abundance difference is

$$\Delta(x_i) = \mu_b(x_i) - \mu_a(x_i) - \log {}^{N_a}/_{N_b}\,.$$

Significance testing

To assess whether observed changes are statistically meaningful, Kompot computes a posterior tail probability (PTP). Intuitively, the PTP measures the probability that the true density difference has the opposite sign to the estimated change, $\Delta(x_i)$ i.e., if $\Delta(x_i)$ is positive, the PTP gives the probability that the true change is negative, and vice versa. Specifically, if the density change $\Delta(x_i)$ is positive, then we compute $\mathbb{P}(\delta(x_i) < 0)$, and $\mathbb{P}(\delta(x_i) > 0)$ when $\Delta(x_i)$ is negative. Since the posterior distribution is a normal distribution, the local Kompot PTP is given by:

$$PTP(x_i) = \Phi\left(\frac{-|\Delta(x_i)|}{\sqrt{\sigma_a^2(x_i) + \sigma_b^2(x_i)}}\right)$$

were Φ is the cumulative distribution function of the standard normal distribution. Thus, Kompot provides both an effect size (density log-fold change) and a principled measure of statistical significance (PTP) for every cell state.

Differential Expression Testing

Kompot differential expression testing enables the identification of gene expression changes between two conditions across the same manifold. The goal is to detect genes with significant differences between conditions, rather than between cell types or clusters within a single condition. Specifically, Kompot is designed to capture subtle, structured transcriptional shifts that vary across cell states, even when those shifts do not follow a consistent direction. Similar to differential abundance testing, Kompot first constructs condition-specific expression functions that map diffusion maps to observed expression values for each gene in each condition using Mellon[21]. This enables testing for gene expression changes across the entire phenotypic landscape without predefined clusters, by explicitly sharing information across neighboring cell states to detect structured patterns.

Similar to the density function, we leverage Mellon's GP implementation for gene expression functions (**Supp. Note 1**). For each condition, Kompot compute a gene-expression function that predicts an expression value for each gene at any location in the cell-state space, specific to that condition. These functions provide a smooth and continuous representation of gene expression across the manifold and allow for quantifying differential expression between conditions at any point in state space, effectively integrating all data points without requiring predefined cluster boundaries. The resulting gene-expression functions and uncertainty estimates can be used analogously to the cell-state density functions to compute local gene-expression differences, with local fold change and posterior tail probabilities, providing both effect size and uncertainty for each gene.

Beyond asking where in the manifold a gene is differentially expressed, we also want to know whether it shows consistent differences within a cell population or across the entire manifold. Unlike cell-state densities, GP-based expression functions are an imputation that already aggregate information across cells since they map state space to observed expression. Specifically, for gene g, the gene expression

values from function $f_g(x_i)$ at different cell states x_i are not independent measurements that can be aggregated directly. Kompot resolves this by explicitly modeling covariance across states using Mellon's uncertainty framework. Mellon provides not only variances for each predicted expression value, but also covariances that reflect the smoothness of gene expression across the manifold i.e., the uncertainty covers not only the variance for a function value $f(x_i)$ but also the covariance to other function values $f(x_j)$ based on the similarity of cell states x_i and x_j.

Structured DE testing with Mahalanobis distance

Formally, under our modelling assumption, the true gene expression values $f_g(x)$ at cell states $x = \{x_i\}_{i=1}^{m}$ are distributed according to a multivariate normal posterior

$$f_g(x) \sim \mathcal{N}\left(\mu^{(g)}(x), \Sigma^{(g)}(\mathrm{x})\right).$$

Therefore, difference in functions between the conditions a and b also follows a normal distribution:

$$\delta_g(x) \sim \mathcal{N}\left(\mu_b^{(g)}(x) - \mu_a^{(g)}(x), \Sigma_{\mathrm{b}}^{(g)}(\mathrm{x}) + \Sigma_{\mathrm{a}}^{(g)}(\mathrm{x})\right).$$

To summarize these correlated local differences into a single test statistic, Kompot uses the Mahalanobis distance (**Supp. Note 2**):

$$D(a,b) \quad = \sqrt{(\mu_a - \mu_b)^T (\Sigma_a + \Sigma_b)^{-1} (\mu_a - \mu_b)}$$

This statistic accounts for the covariance between nearby states, ensuring that consistent changes are emphasized while redundant information is down-weighted. As a result, Kompot can detect structured patterns of differential expression across the manifold, even when local changes differ in magnitude or direction.

Thresholds to identify significantly different genes

To enable robust identification of significantly different genes, Kompot generates an empirical null distribution of Mahalanobis distances by randomly shuffling gene expression values (default: 2000 genes). From the resulting distribution of null Mahalanobis distances, empirical p-values are computed and used to estimate false discovery rates (FDR) via the local FDR methods[44]. By default, Kompot applies with a cutoff of 0.05 on the local FDR, yielding a direct Mahalanobis distance threshold for declaring genes significantly differentially expressed (**Supp. Note 4**).

Scalability

Kompot inherits the scalability of the Mellon framework. Mellon demonstrated linear scaling in the number of cells through its sparse Gaussian process implementation, making inference feasible for datasets with millions of cells[21]. For differential expression, scaling is even faster due to its closed-form formulation: evaluating ~16,000 genes across 5,000 cells requires approximately one minute on a standard workstation, and runtime grows linearly with the number of genes since each is modeled by an independent Gaussian process.

The current bottleneck lies in the k-means clustering step used to generate landmark states. While k-means itself scales with complexity $O(t \times k \times d \times n)$ where t is the number of iterations, k the number of landmarks, d the dimensionality, and n the number of cells, this step can be easily replaced. All Kompot methods allow users to supply a custom set of landmarks. Since inference relies on covariance statistics, the exact landmark choice has little effect as long as they sufficiently cover the state space. Moreover,

nearly identical landmarks contribute negligibly to test statistics, as their strong covariance prevents redundant information from inflating significance.

Multi-Sample Framework

Kompot also supports multi-sample analyses by explicitly incorporating variability across replicates (**Supp. Note 3**). The density functions and gene-expression manifolds are always inferred jointly across all samples, which maximizes statistical power and prevents small or unbalanced samples from distorting the overall estimate. To capture heterogeneity between replicates, Kompot then incorporates an empirical variability terms that reflect differences within each condition. This combined uncertainty accounts for both biological variability (e.g., donor or tissue differences) and technical variability (e.g., sample preparation or sequencing batch effects). By integrating these sources of variability, Kompot provides a robust and conservative basis for identifying significant changes in both cell-state densities and gene-expression patterns.

Aging hematopoiesis single-cell RNA-seq data generation

Marrow collection

Bone marrow aspirate was collected using centrifugation from the femur, tibia, and hip of male C57BL/6 mice aged young (10 weeks, n=3), mid (63 weeks, n=2), and old (103 weeks, n=3). Whole marrow samples were washed twice in PBS supplemented with 10% FBS then were gently passed through a 40 micron strainer. The cell suspension was split 1:10 with one part set aside for FACS sorting of mature cells, and the rest for HSPC sorting.

HSPC processing

Cells intended for HSPC sorting were first subjected to MACS lineage depletion with BioLegend antibodies targeting the following mature lineage marker antigens: Ter119 (CAT# 79748), B220 (CAT# 103204), CD3e (CAT# 79751), CD4 (CAT# 100404), CD8a (CAT# 100704), Mac1/CD11b (CAT# 79749), Gr1 (CAT# 79750), FcεR1α (CAT# 134304), CD19 (CAT# 115504), NK1.1 (CAT# 108704), CD11c (CAT# 117304). Cells were incubated in lineage stain for 10 minutes, washed in MACS buffer (0.5% bovine serum albumin and 2 mM EDTA in PBS), then incubated another 15 minutes with anti-biotin MicroBeads (5 μL for every 10e7 cells, Miltenyi Biotec CAT# 130-090-485), washed in MACS buffer, and passed through a 40 micron filter immediately prior to processing with the AutoMACS Pro Separator. After MACS lineage depletion, samples were incubated for 10 minutes with TruStain FcX PLUS (1%, BioLegend CAT# 156604), then each sample was stained for 30 minutes with a cocktail of (1) a unique hashing antibody (BioLegend CAT#s 155831, 155833, 155835, 155837, 155839, 155841, 155843, or 155845) plus (2) the flow sorting antibodies DAPI, mCD45, anti-biotin and cKit plus (3) BioLegend CITE Total Seq antibodies targeting the following: CD135 (CAT# 135319), CD16/32 (CAT# 101345), Sca1 (CAT# 108149), CD127 (CAT# 135055), CD150 (CAT# 115951), CD48 (CAT# 103457), CD41 (CAT# 133941), CD34 (CAT# 128625). Samples were pooled then washed three times in Cell Staining Buffer (BioLegend CAT# 420201) then were sorted for DAPI (-), mCD45 (+), lineage (-), cKit (+).

Mature processing

Each of the 8 samples intended for Mature sorting were stained for 30 minutes with a cocktail of (1) a unique hashing antibody (BioLegend CAT#s 155831, 155833, 155835, 155837, 155839, 155841, 155843, or 155845) plus (2) the flow sorting antibodies DAPI, mCD45, and Ter119 plus (3) BioLegend CITE Total Seq antibodies targeting the following: CD4 (CAT# 100573), CD8a (CAT# 100783), Ly-6C (CAT# 128053), CD11b (CAT# 101273), Ly-6G (CAT# 127659), CD44 (CAT# 10371), CD19 (CAT#

115563), CD25 (CAT# 102067), B220 (CAT# 103271), CD115 (CAT# 135543), CD11c (CAT# 117359), CD62L (CAT# 104465), F4/80 (CAT# 123155), I-A/I-E (CAT# 107657), NK1.1 (CAT# 108763), CD3 (CAT# 100257), CD335 (CAT# 137641), CD27 (CAT# 124247), CCR2 (CAT# 150633), CXCR4 (CAT# 146521), CX3CR1 (CAT# 149045). Samples were pooled then washed three times in Cell Staining Buffer (BioLegend CAT# 420201) then were sorted for DAPI (-), Ter119 (-), mCD45 (+).

Sequencing

Single cell barcoded libraries were generated using chromium Next GEM Single Cell 3' Reagent Kits, Dual Index (v3.1 Chemistry) and were sequenced on an Illumina HiSeq.

Data Analysis

Aging hematopoiesis scRNA-seq

Preprocessing

CellRanger (v6.1.2) was used for demultiplexing and count matrix generation. We next performed quality control to remove low quality cells: First doublets were identified using Scrublet (v0.2.3)[45]. Suspected apoptotic cells were identified as any cell with >10% mitochondrial reads. Cells with 20% or more molecules coming from hemoglobin genes were identified as red blood cell contamination. Further, cells with abnormally low molecule count were additionally flagged. Finally, all cells failing these QC metrics were removed. Genes expressed in fewer than 10 cells were also filtered.

After QC, the dataset was preprocessed in Scanpy as follows. First, raw gene expression counts were normalized using Scanpy's normalize_total function so that all cells have the same total number of counts. Then the expression matrix was rescaled by adding a small pseudocount of 0.1 and taking the base-2 logarithm. The top 25% most highly variable genes were selected using scanpy's highly_variable_genes function with the cell_ranger flavor and batch_key=Compartment (the two compartments being HSPC and Mature). Cell cycle scores were computed using scanpy's score_genes_cell_cycle function, then the cell cycle effect was regressed out with Scanpy's regress_out. PCA was carried out on HVGs, then batch correction was performed across compartments with Harmony[16]. Neighborhood graphs were computed on the Harmony-corrected PCs, followed by UMAP and Leiden clustering. Cell typing was done manually using expression of marker genes gathered from public data sources

Mouse gastrulation data

Processed and annotated data comprising of wild type and *Tal1-/-* cells (56122 cells) from E8.5 mouse embryos were downloaded from ArrayExpress (E-MTAB-7325)[36]. Cell types that were exclusively present in the wild type embryos (>99) were excluded from analyses since our focus was on differential gene expression between Tal1 wild type and knock out cells. Pre-computed, batch corrected PCA was used input to compute diffusion maps with 50 diffusion components. Kompot differential expression analysis was performed with default parameters. Top 500 differentially expressed genes ranked by Mahalanobis Distance were used for all interpretive analyses.

Melanoma CD8 T cell landscape

Processed and annotated data was downloaded from GEO (GSE225063)[37]. The dataset consisted of 75160 cells spanning different CD8 T cell states from melanoma patients at baseline and following treatment with different ICI. All cells from baseline, week4 and week16 timepoints were used for analysis

607 (74644 cells)), with the union of week4 and week16 cells used as cells in the response condition. Pre-
608 computed scVI latent representation was used to compute diffusion maps with 10 components.

609 Kompot differential abundance testing was performed separately for Nivo, Rela and Combination
610 treatments to determine changes from baseline to response. Cell states with posterior tail probability <
611 1e-2 and absolute log fold change > 1 were identified to be significantly differentially abundant in each
612 treatment.

613 Kompot differential expression testing was performed separately for each treatment using genes detected
614 in at least 100 cells in the dataset (ls (15247 genes). Top 250 differentially expressed genes based on
615 Mahalanobis Distance were identified for each treatment condition and the union of these genes were
616 used interpretive analyses. STRING gene ontology analysis was performed with default parameters.

617

618 **Neuroblastoma cell atlas**

619 Processed and annotated data was downloaded from cellxgene[40]. This dataset consisted of 205153
620 neoplastic cells across 19 high-risk neuroblastoma patients at diagnosis and post-induction
621 chemotherapy. Pre-computed, batch corrected PCA was used as input to compute diffusion maps with
622 10 components. Kompot differential abundance testing was performed using the multi-sample mode with
623 default settings.

624

Data Availability

Raw single-cell data for the murine aging hematopoiesis will be deposited to GEO. Processed anndata objects with annotations are available to download from Zenodo[46].

Code Availability

Kompot is available as a Python module at https://github.com/settylab/Kompot. Jupyter notebooks detailing the usage of Kompot differential expression and differential abundance testing are available at https://github.com/settylab/kompot/tree/main/examples and documentation at https://kompot.readthedocs.io/en/latest/.

Acknowledgements

This study was supported by National Institute of General Medical Studies grant no. R35 GM147125 to M.S.; National Cancer Institute grant no. R01CA292932, Mark Foundation for Cancer Research Endeavor Award and Edward P. Evans Foundation Discovery Research Grant to S.C.L; Brotman Baty Institute Pilot Award and Translational Data Science IRC New Collaboration Award to M.S. and S.C.L; National Institute of General Medicine Studies grant no. T32GM136534 to E.T.; National Institutes of Health grant no. ORIP S10OD028685 to support high-performance computing at the Fred Hutchinson Cancer Center.

Competing Interests

The authors declare no competing interests.

Figure

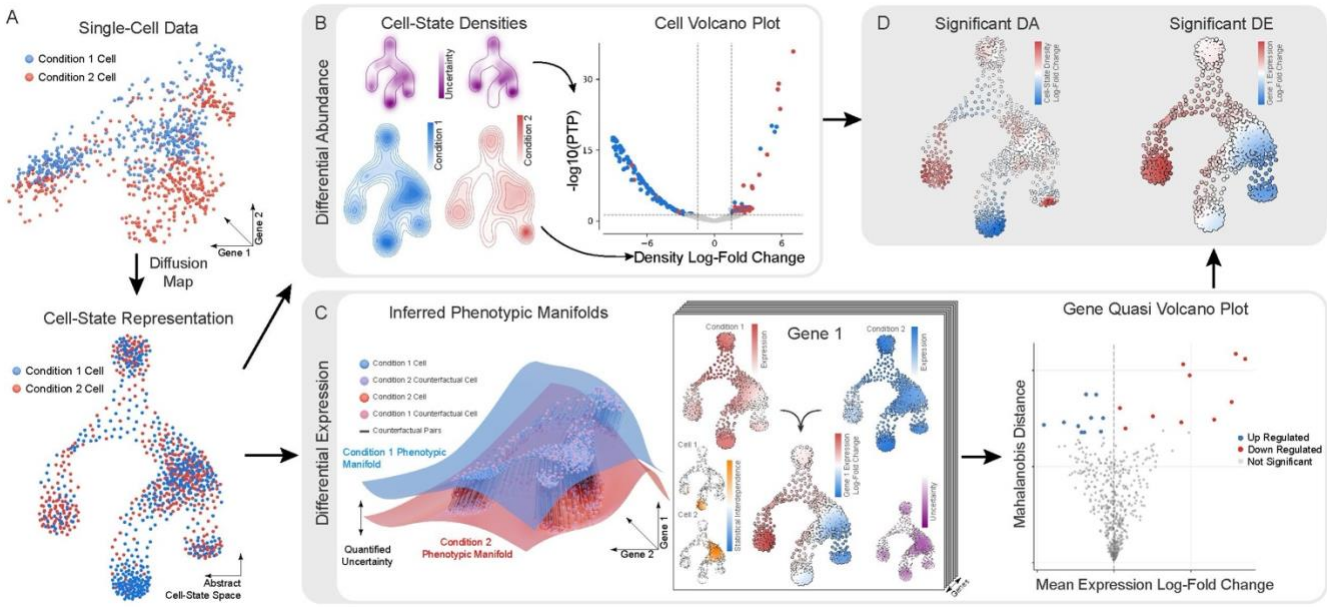

Figure 1: Overview of Kompot's framework for differential abundance and expression analysis.

(A) Illustration of high-dimensional single-cell transcriptomic data from two conditions, projected into a shared low-dimensional cell-state space using diffusion maps.

(B) Differential abundance analysis in Kompot. Condition-specific cell-state densities with associated uncertainty estimates are computed and compared to yield density log-fold changes. Significance is quantified using posterior tail probability (PTP), visualized in a quasi-volcano plot.

(C) Differential expression analysis. Left: schematic of inferred phenotypic manifolds for both conditions with counterfactual cell pairs illustrating expression comparison. Middle: gene-specific differences are illustrated for one example gene, showing condition-specific expression levels, statistical interdependence, and uncertainty. Right: A quasi-volcano plot summarizes gene-level results with the average log-fold change on the x-axis and Mahalanobis distance (accounting for variance and covariance) as a multivariate significance score on the y-axis.

(D) Visualization of differential abundance (left) and differential expression (right) across the cell-state space, colored by log-fold change magnitude.

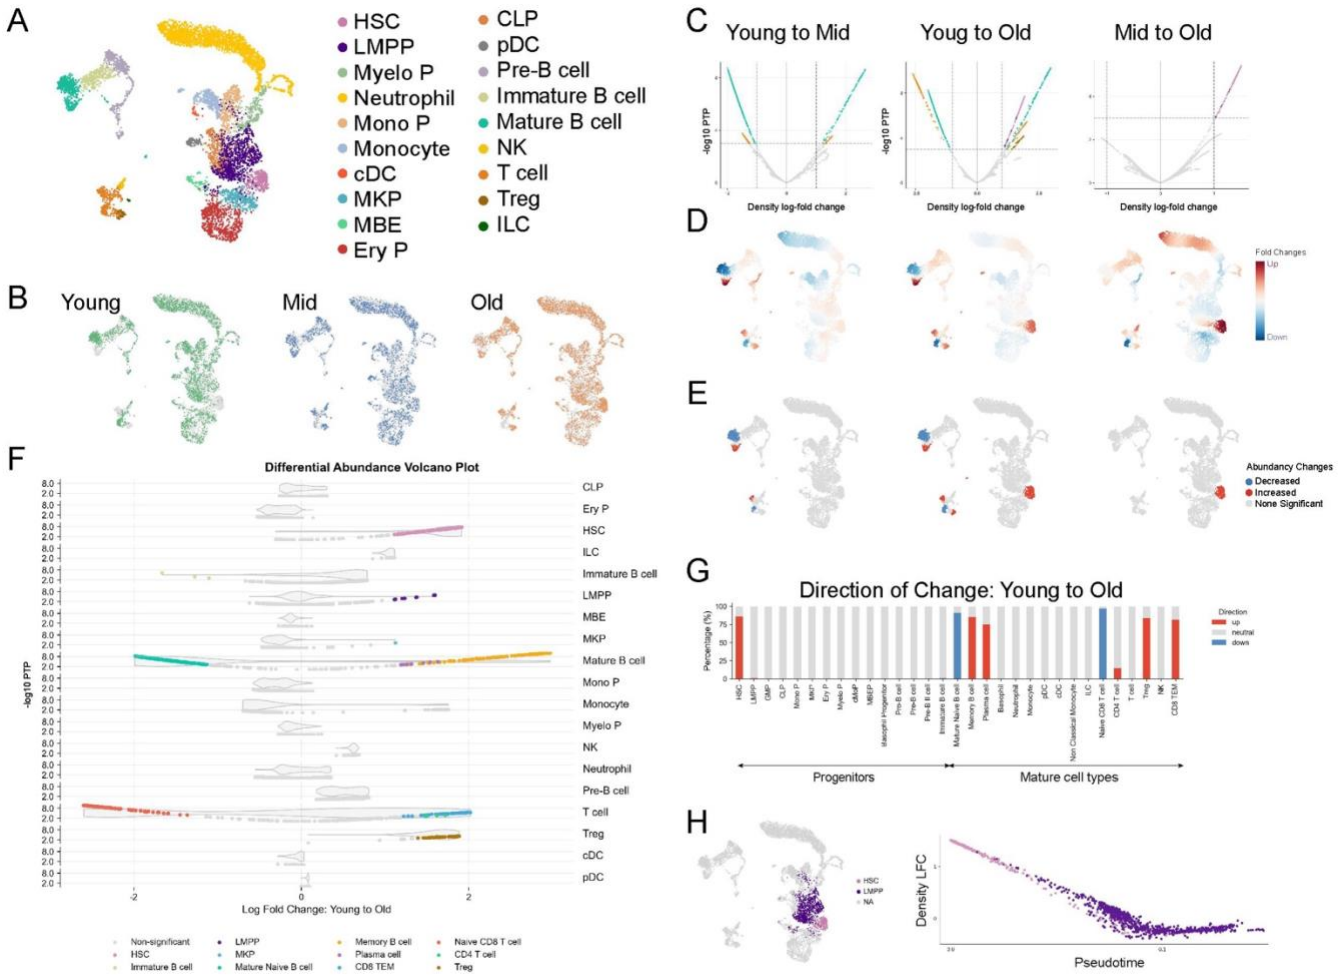

Figure 2: Kompot differential abundance analysis in aging hematopoiesis.

(A) UMAP projection of aging murine bone marrow scRNA-seq dataset, colored by medium-resolution cell types.

(B) UMAPs highlighting cells from mice grouped by age: Young (10 weeks), Mid (63 weeks), and Old (103 weeks).

(C) Differential abundance results for the Young to Mid (left), Young to Old (middle) and Mid to Old (right) comparisons. Quasi-cell volcano plot with density log-fold change on the x-axis and Posterior Tail Probabilities (PTP) on the y-axis. Significantly different cell-states (absolute log fold change > 1, PTP < 10^{-3}) are highlighted by their cell type annotations from (A).

(D) UMAPs displaying the estimated local density log-fold changes across the full joint cell-state space derived from all samples (Young, Mid, and Old) for comparisons in C.

(E) UMAPs highlighting the subset of cell states with statistically significant abundance changes, as defined by the thresholds shown in the volcano plots in C.

(F) Differential abundance volcano plots showing density log-fold changes across cells, stratified by major cell type and colored by subtype in the Young to Old comparison. y-axis indicates posterior tail probability. Violin plots in the background indicate abundance of cells at the respective density log-fold change in this group.

(G) Direction and magnitude of abundance changes per cell type in the Young to Old comparison. Bars show the fraction of significantly changing cell-states within each annotated cell type, color-coded by the direction of change. Cell types are ordered along a progenitor-to-mature axis.

(H) Kompot results for early hematopoietic progenitors. Left: UMAP highlighting HSC, and LMPP subsets. Right: Scatterplot of cells colored by density log-fold change along pseudotime (x-axis), showing dynamic abundance shifts during early differentiation.

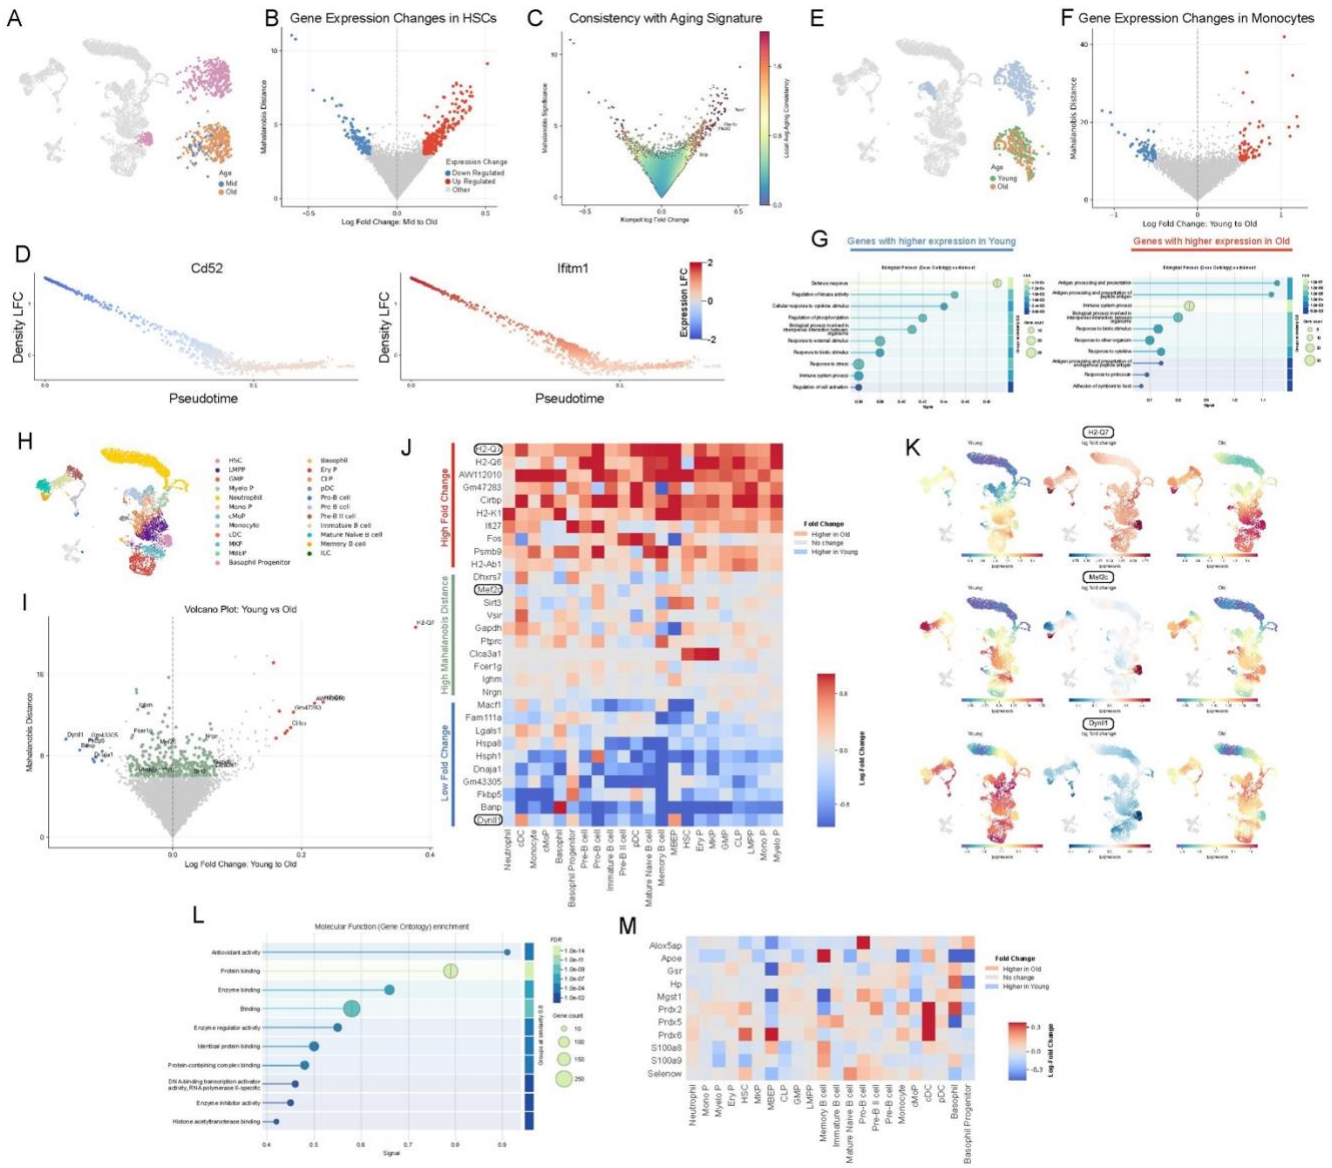

Figure 3: Kompot differential expression analysis in aging hematopoiesis.

(A) UMAP from Fig. 2A, with HSCs highlighted.

(B) Differential expression results for Mid to Old HSCs. Quasi-cell volcano plot with expression log-fold change on the x-axis and Mahalanobis distance on the y-axis. Significantly different genes (absolute log fold change > 0.15, Mahalanobis distance > 3) are highlighted.

(C) like (B) colored by weighted average of reported “Aging Consistency” using standard Gaussian Kernel after scaling x- and y-axis to standard deviation 1.

(D) Same plot as Fig. 2H, colored by Mid to Old HSC gene expression log fold change for a gene downregulated with age (Cd52, left) and an upregulated gene (Ifitm1, right).

(E) UMAP from Fig. 2A, with monocytes highlighted.

(F) Same as (B), for Young to Old monocytes.

(G) Gene ontology analysis from STRING[47] for downregulated genes (left) and upregulated genes (right) from Fig. 2F.

(H) UMAP for Fig. 2A, highlighting all cell types except T and NK cells

(I) Differential expression results for all cells in Fig. 2H for Young to Old comparisons. Top 10 genes with positive fold change and high Mahalanobis distance are shown in red, bottom 10 genes with negative fold change and high Mahalanobis distance are shown in blue, and genes with log absolute fold change but high Mahalanobis distance are shown in green (absolute log fold change < 0.07, Mahalanobis distance > 6). 20 genes are labeled.

(J) Heatmap showing Young to Old gene expression log fold changes across different cell types for genes highlighted in (I).

(K) Top: Imputed expression of gene H2-Q7 using young cells (left), using old cells (right) and fold change from young to old (middle). Middle, Bottom: Same as the top row for Mef2c and Dynll1. These genes serve as examples of genes with high positive fold change, high Mahalanobis but low absolute fold change, and high negative fold change respectively.

(L) Gene ontology analysis using genes with high Mahalanobis distance and low absolute fold change from (I).

(M) Heatmap showing Young to Old gene expression log fold changes across different cell types for genes involved in antioxidant activity.

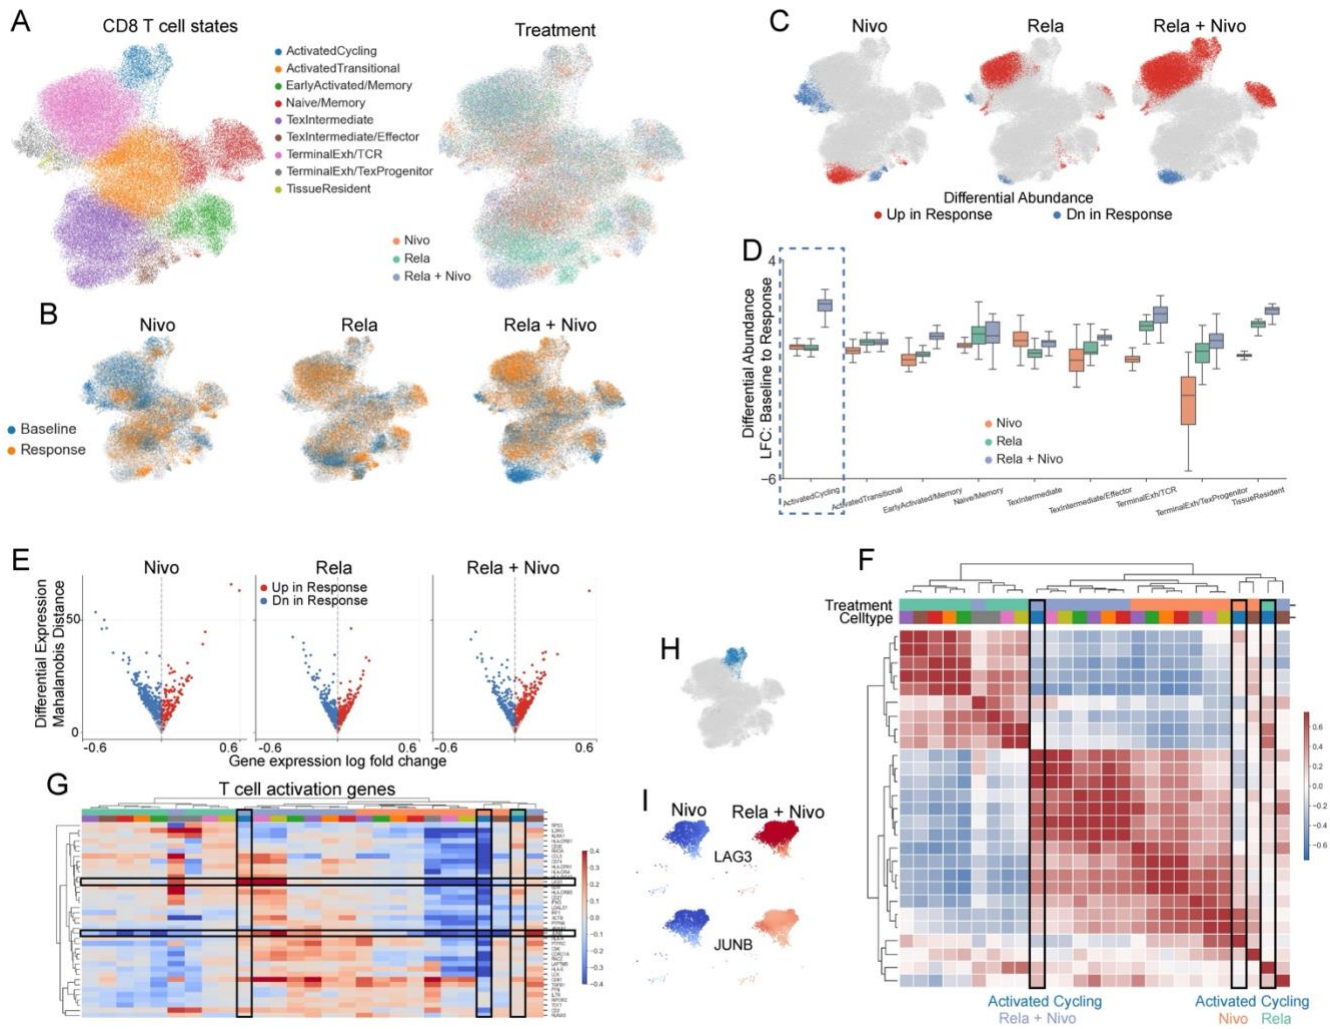

Figure 4: Kompot reveals determinants of immunotherapy in advanced melanoma

(A) UMAP highlighting the CD8 T cell state (left) landscape from immunotherapy recipients with advanced melanoma[37]. Right: Same UMAP as (left) with cells colored by treatment type.

(B) Same as (A) with cells colored by treatment condition for Nivolumab (anti-PD1, left), Relatlimab (anti-LAG-3, center) and combination of Nivolumab and Relatlimab (right).

(C) Kompot differential abundance results highlighting significantly altered cell states (p-value < 1e-3, absolute log-fold change > 1) for the three treatment conditions.

(D) Boxplot comparing the abundance fold changes for different CD8 T cell states across treatment conditions. Activated Cycling cells which have greater abundance in the combination treatment are highlighted.

(E) Kompot differential expression results for the three treatment conditions. The union of top 250 genes ranked by Mahalanobis Distance, across the three treatment conditions are highlighted in the pseudo-volcano plots.

(F) Heatmap showing the pairwise correlation between average fold change from baseline to response in different cell states and treatment conditions. Genes highlighted in (E) were used for computing correlations. Correlations of Activated cycling cells in the different treatment conditions are highlighted.

(G) Heatmap showing the log fold change from baseline to response conditions for T cell activation genes. These are a subset of genes in (E) selected based on gene ontology analysis. Columns of

744 activated cycling cells in the three treatment conditions are highlight. Columns are in the same order as
745 (F).

746 (H) Same UMAP as (A), with activated cycling T cells highlighted.

747 (I) Cells in (H) colored by gene expression fold change from baseline to response in Nivo (left) or the
748 combination treatment (right). Same genes show opposing expression changes with different
749 treatments.

750

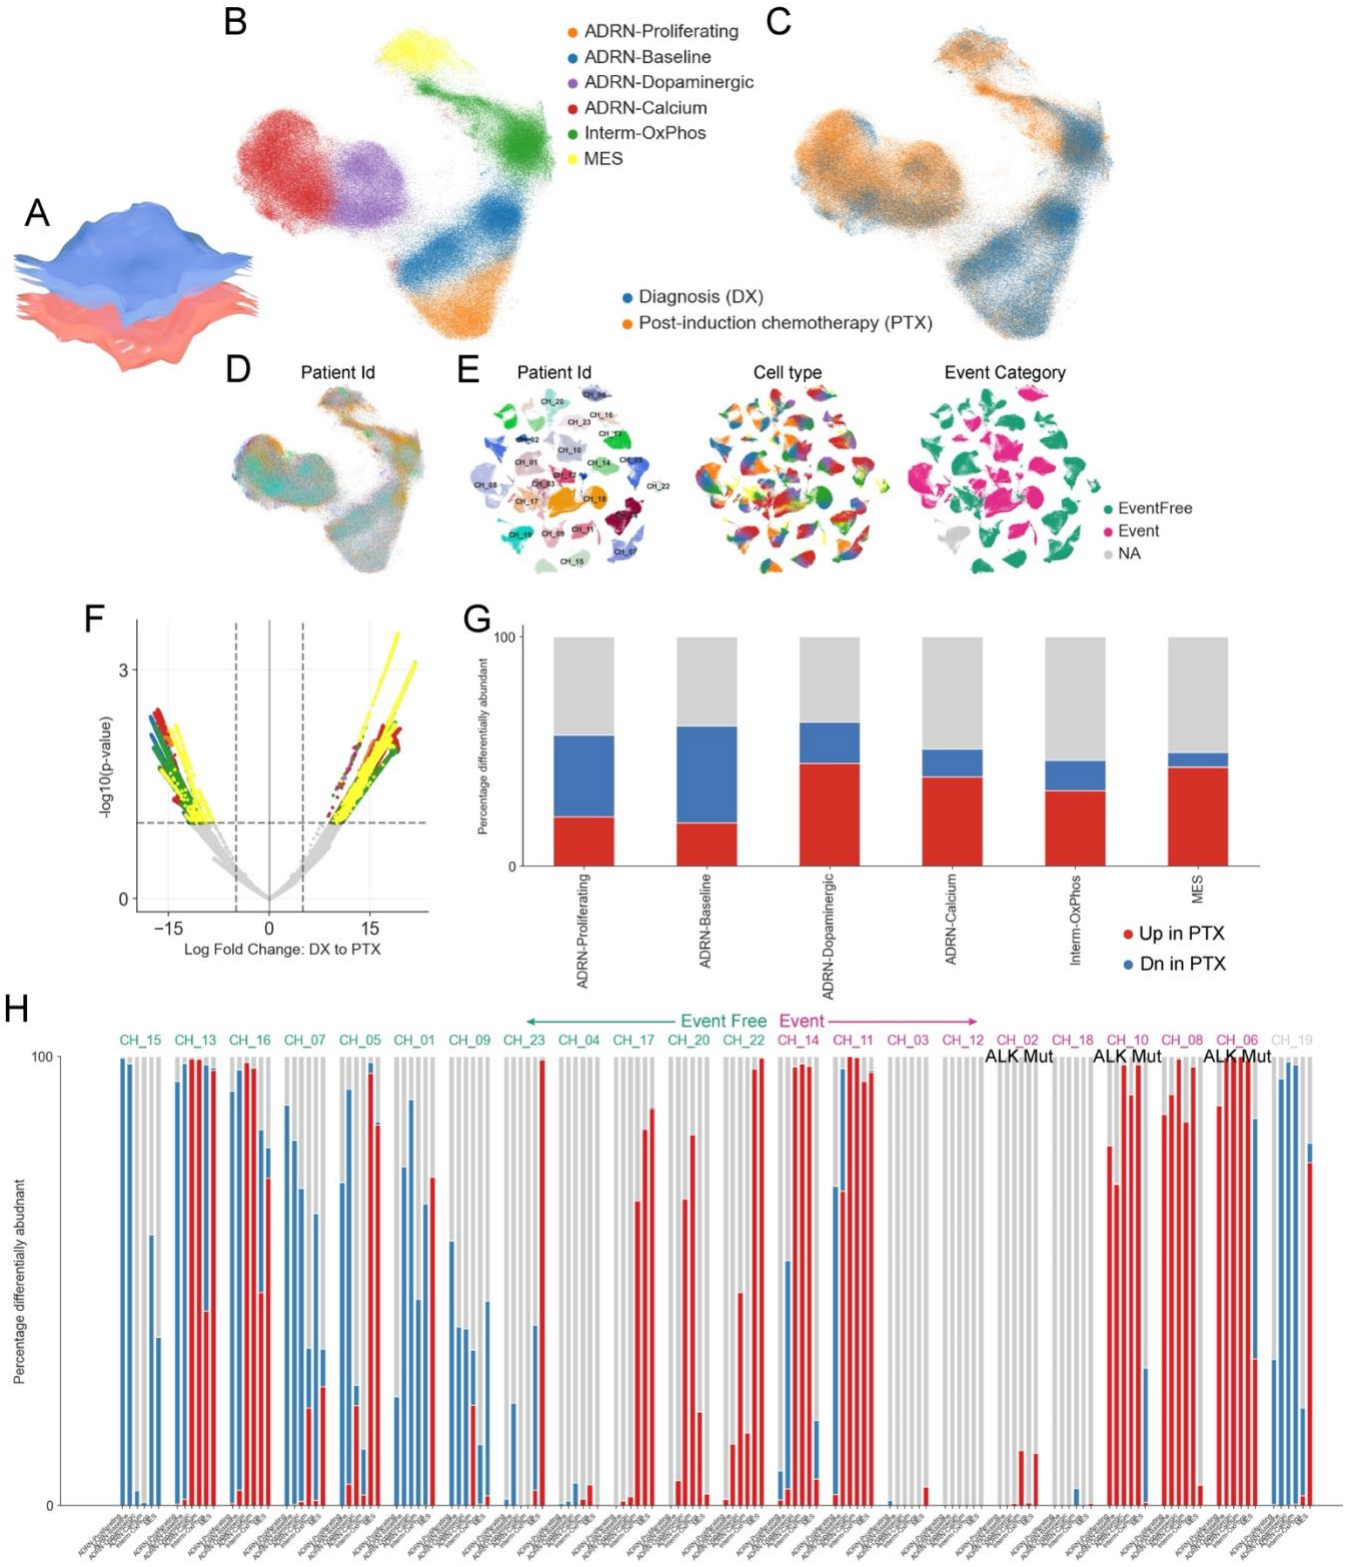

Figure 5: Inter-patient variability with multi-sample differential abundance testing

(A) Illustration of multi-sample differential abundance testing in Kompot. Variability is computed using density estimates for each sample within the condition for significance assessment.

(B) UMAP of neoplastic cells from high-risk neuroblastoma patients colored by cell-states[40].

(C-D) Same as (B), colored by treatment condition and patient id

(E) UMAP of the same cells in (B), prior to batch correction. Colored by patient id (left), cell type (middle) and patient event category (right).

(F) Kompot differential abundance volcano plot with significantly altered cell-states from diagnosis to post-induction chemotherapy (p-value < 5e-2, absolute log-fold change > 1) colored by cell types from (B).

(G) Direction and magnitude of abundance changes per cell type. Bars show the fraction of significantly changing cell-states within each annotated cell type, color-coded by the direction of change.

(H) Same as (G), but separately for each patient. Patients are ordered based on whether event free or suffered an adverse event. Within the two categories, patients are sorted by fraction of differentially abundant ADRN-proliferating and ADRN-Baseline cells. A greater reduction associates with event free survival.

References

1. Tabula Muris C. A single-cell transcriptomic atlas characterizes ageing tissues in the mouse. Nature. 2020;583(7817):590-5, PMC8240505.
2. Konturek-Ciesla A, Dhapola P, Zhang Q, Sawen P, Wan H, Karlsson G, et al. Temporal multimodal single-cell profiling of native hematopoiesis illuminates altered differentiation trajectories with age. Cell Rep. 2023;42(4):112304.
3. Herault L, Poplineau M, Mazuel A, Platet N, Remy E, Duprez E. Single-cell RNA-seq reveals a concomitant delay in differentiation and cell cycle of aged hematopoietic stem cells. BMC Biol. 2021;19(1):19, PMC7851934.
4. Kucinski I, Campos J, Barile M, Severi F, Bohin N, Moreira PN, et al. A time- and single-cell-resolved model of murine bone marrow hematopoiesis. Cell Stem Cell. 2023.
5. Gazestani V, Kamath T, Nadaf NM, Dougalis A, Burris SJ, Rooney B, et al. Early Alzheimer's disease pathology in human cortex involves transient cell states. Cell. 2023;186(20):4438-53 e23, PMC11107481.
6. Mayr CH, Simon LM, Leuschner G, Ansari M, Schniering J, Geyer PE, et al. Integrative analysis of cell state changes in lung fibrosis with peripheral protein biomarkers. EMBO Mol Med. 2021;13(4):e12871, PMC8033531.
7. Rehman SK, Haynes J, Collignon E, Brown KR, Wang Y, Nixon AML, et al. Colorectal Cancer Cells Enter a Diapause-like DTP State to Survive Chemotherapy. Cell. 2021;184(1):226-42 e21, PMC8437243.
8. van der Leun AM, Thommen DS, Schumacher TN. CD8(+) T cell states in human cancer: insights from single-cell analysis. Nat Rev Cancer. 2020;20(4):218-32, PMC7115982.
9. Replogle JM, Saunders RA, Pogson AN, Hussmann JA, Lenail A, Guna A, et al. Mapping information-rich genotype-phenotype landscapes with genome-scale Perturb-seq. Cell. 2022;185(14):2559-75 e28, PMC9380471.
10. Burkhardt DB, Stanley JS, 3rd, Tong A, Perdigoto AL, Gigante SA, Herold KC, et al. Quantifying the effect of experimental perturbations at single-cell resolution. Nat Biotechnol. 2021;39(5):619-29, PMC8122059.
11. Dann E, Henderson NC, Teichmann SA, Morgan MD, Marioni JC. Differential abundance testing on single-cell data using k-nearest neighbor graphs. Nat Biotechnol. 2022;40(2):245-53.
12. Yi H, Plotkin A, Stanley N. Benchmarking differential abundance methods for finding condition-specific prototypical cells in multi-sample single-cell datasets. Genome Biology. 2024;25(1):9.
13. Missarova A, Dann E, Rosen L, Satija R, Marioni J. Leveraging neighborhood representations of single-cell data to achieve sensitive DE testing with miloDE. Genome Biol. 2024;25(1):189, PMC11256449.
14. Boyeau P, Hong J, Gayoso A, Kim M, McFaline-Figueroa JL, Jordan MI, et al. Deep generative modeling of sample-level heterogeneity in single-cell genomics. bioRxiv. 2024:2022.10.04.510898.
15. Ahlmann-Eltze C, Huber W. Analysis of multi-condition single-cell data with latent embedding multivariate regression. Nat Genet. 2025;57(3):659-67, PMC11906359.
16. Korsunsky I, Millard N, Fan J, Slowikowski K, Zhang F, Wei K, et al. Fast, sensitive and accurate integration of single-cell data with Harmony. Nat Methods. 2019;16(12):1289-96, PMC6884693.
17. Luecken MD, Buttner M, Chaichoompu K, Danese A, Interlandi M, Mueller MF, et al. Benchmarking atlas-level data integration in single-cell genomics. Nat Methods. 2022;19(1):41-50, PMC8748196.
18. Coifman RR, Lafon S, Lee AB, Maggioni M, Nadler B, Warner F, et al. Geometric diffusions as a tool for harmonic analysis and structure definition of data: diffusion maps. Proc Natl Acad Sci U S A. 2005;102(21):7426-31, PMC1140422.
19. Setty M, Kiseliovas V, Levine J, Gayoso A, Mazutis L, Pe'er D. Characterization of cell fate probabilities in single-cell data with Palantir. Nat Biotechnol. 2019;37(4):451-60.
20. Haghverdi L, Buttner M, Wolf FA, Buettner F, Theis FJ. Diffusion pseudotime robustly reconstructs lineage branching. Nat Methods. 2016;13(10):845-8.
21. Otto D, Jordan C, Dury B, Dien C, Setty M. Quantifying Cell-State Densities in Single-Cell Phenotypic Landscapes using Mellon. Nature Methods. 2024, PMC10369887.
22. Rasmussen CE, Williams CKI. Gaussian Processes for Machine Learning2006.
23. PC M. Reprint of: Mahalanobis, P.C. (1936) "On the Generalised Distance in Statistics.". Sankhya A. 2018;80(1):1-7.
24. Jaiswal S, Ebert BL. Clonal hematopoiesis in human aging and disease. Science. 2019;366(6465), PMC8050831.
25. Chambers SM, Shaw CA, Gatza C, Fisk CJ, Donehower LA, Goodell MA. Aging hematopoietic stem cells decline in function and exhibit epigenetic dysregulation. PLoS Biol. 2007;5(8):e201, PMC1925137.

825 26. Li H, Cote P, Kuoch M, Ezike J, Frenis K, Afanassiev A, et al. The dynamics of hematopoiesis over the
826 human lifespan. Nat Methods. 2025;22(2):422-34, PMC11908799.
827 27. Appelbaum FR, Gundacker H, Head DR, Slovak ML, Willman CL, Godwin JE, et al. Age and acute
828 myeloid leukemia. Blood. 2006;107(9):3481-5, PMC1895766.
829 28. Florian MC, Klose M, Sacma M, Jablanovic J, Knudson L, Nattamai KJ, et al. Aging alters the epigenetic
830 asymmetry of HSC division. PLoS Biol. 2018;16(9):e2003389, PMC6168157.
831 29. Bulati M, Caruso C, Colonna-Romano G. From lymphopoiesis to plasma cells differentiation, the age-
832 related modifications of B cell compartment are influenced by "inflamm-ageing". Ageing Res Rev. 2017;36:125-
833 36.
834 30. Goronzy JJ, Weyand CM. Immune aging and autoimmunity. Cell Mol Life Sci. 2012;69(10):1615-23,
835 PMC4277694.
836 31. de Haan G, Lazare SS. Aging of hematopoietic stem cells. Blood. 2018;131(5):479-87.
837 32. Flohr Svendsen A, Yang D, Kim K, Lazare S, Skinder N, Zwart E, et al. A comprehensive transcriptome
838 signature of murine hematopoietic stem cell aging. Blood. 2021;138(6):439-51.
839 33. De Maeyer RPH, Chambers ES. The impact of ageing on monocytes and macrophages. Immunol Lett.
840 2021;230:1-10.
841 34. Barman PK, Shin JE, Lewis SA, Kang S, Wu D, Wang Y, et al. Production of MHCII-expressing classical
842 monocytes increases during aging in mice and humans. Aging Cell. 2022;21(10):e13701, PMC9577948.
843 35. Shao L, Li H, Pazhanisamy SK, Meng A, Wang Y, Zhou D. Reactive oxygen species and hematopoietic
844 stem cell senescence. Int J Hematol. 2011;94(1):24-32, PMC3390185.
845 36. Pijuan-Sala B, Griffiths JA, Guibentif C, Hiscock TW, Jawaid W, Calero-Nieto FJ, et al. A single-cell
846 molecular map of mouse gastrulation and early organogenesis. Nature. 2019;566(7745):490-5, PMC6522369.
847 37. Cillo AR, Cardello C, Shan F, Karapetyan L, Kunning S, Sander C, et al. Blockade of LAG-3 and PD-1
848 leads to co-expression of cytotoxic and exhaustion gene modules in CD8(+) T cells to promote antitumor
849 immunity. Cell. 2024;187(16):4373-88 e15, PMC11346583.
850 38. Wang W, Green M, Choi JE, Gijon M, Kennedy PD, Johnson JK, et al. CD8(+) T cells regulate tumour
851 ferroptosis during cancer immunotherapy. Nature. 2019;569(7755):270-4, PMC6533917.
852 39. Daniel B, Yost KE, Hsiung S, Sandor K, Xia Y, Qi Y, et al. Divergent clonal differentiation trajectories of T
853 cell exhaustion. Nat Immunol. 2022;23(11):1614-27, PMC11225711.
854 40. Yu W, Biyik-Sit R, Uzun Y, Chen CH, Thadi A, Sussman JH, et al. Longitudinal single-cell multiomic atlas
855 of high-risk neuroblastoma reveals chemotherapy-induced tumor microenvironment rewiring. Nat Genet.
856 2025;57(5):1142-54, PMC12081299.
857 41. Lopez R, Regier J, Cole MB, Jordan MI, Yosef N. Deep generative modeling for single-cell
858 transcriptomics. Nat Methods. 2018;15(12):1053-8, PMC6289068.
859 42. Granja JM, Corces MR, Pierce SE, Bagdatli ST, Choudhry H, Chang HY, et al. ArchR is a scalable
860 software package for integrative single-cell chromatin accessibility analysis. Nat Genet. 2021;53(3):403-11,
861 PMC8012210.
862 43. Setty M, Tadmor MD, Reich-Zeliger S, Angel O, Salame TM, Kathail P, et al. Wishbone identifies
863 bifurcating developmental trajectories from single-cell data. Nat Biotechnol. 2016;34(6):637-45, PMC4900897.
864 44. Efron B, Tibshirani R, Storey JD, Tusher V. Empirical Bayes Analysis of a Microarray Experiment. J Am
865 Stat Assoc. 2001;96(456):1151-60.
866 45. Wolock SL, Lopez R, Klein AM. Scrublet: Computational Identification of Cell Doublets in Single-Cell
867 Transcriptomic Data. Cell Syst. 2019;8(4):281-91 e9, PMC6625319.
868 46. Otto D, Arriaga-Gomez E, Thieme E, Yang R, Lee S, Setty M. Single-cell CITE-seq of murine bone
869 marrow across aging (Young, Mid, Old): Zenodo; 2025 [Available from: https://doi.org/10.5281/zenodo.15587767.
870 47. Szklarczyk D, Gable AL, Lyon D, Junge A, Wyder S, Huerta-Cepas J, et al. STRING v11: protein-protein
871 association networks with increased coverage, supporting functional discovery in genome-wide experimental
872 datasets. Nucleic Acids Res. 2019;47(D1):D607-D13, PMC6323986.

873

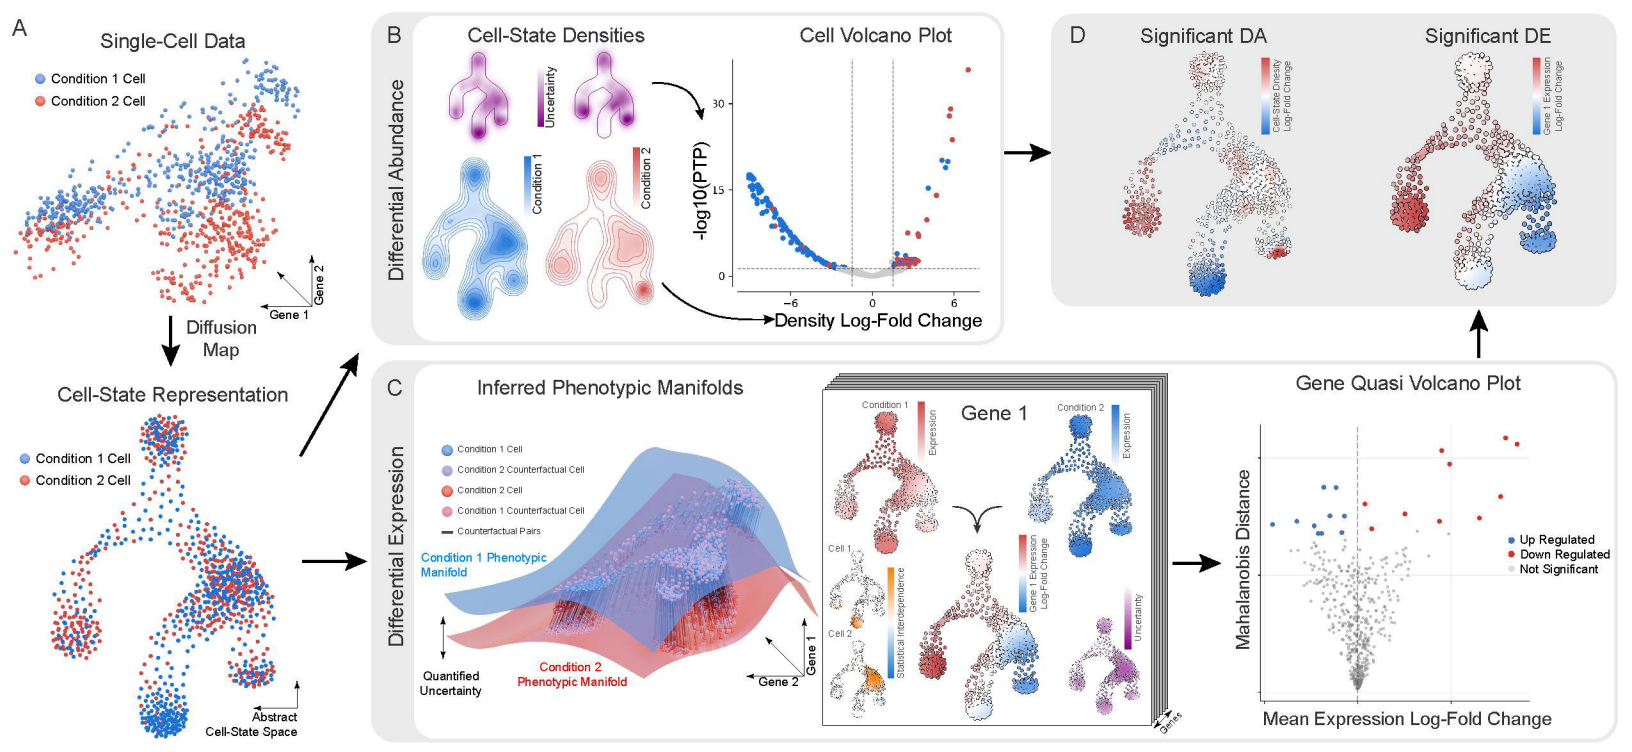
A
Single-Cell Data
Condition 1 Cell
Condition 2 Cell
Gene 1
Gene 2
Diffusion Map
Cell-State Representation
Condition 1 Cell
Condition 2 Cell
Abstract Cell-State Space
B
Differential Abundance
Cell-State Densities
Uncertainty
Condition 1
Condition 2
Cell Volcano Plot
-log10(PTP)
30
15
0
−6
0
6
Density Log-Fold Change
D
Significant DA
Cell-State Density Log-Fold Change
Significant DE
Gene 1 Expression Log-Fold Change
C
Differential Expression
Inferred Phenotypic Manifolds
Condition 1 Cell
Condition 2 Counterfactual Cell
Condition 2 Cell
Condition 1 Counterfactual Cell
Counterfactual Pairs
Condition 1 Phenotypic Manifold
Quantified Uncertainty
Condition 2 Phenotypic Manifold
Gene 2
Gene 1
Gene 1
Condition 1
Expression
Condition 2
Expression
Cell 1
Cell 2
Statistical Interdependence
Gene 1 Expression Log-Fold Change
Uncertainty
Genes
Gene Quasi Volcano Plot
Mahalanobis Distance
Up Regulated
Down Regulated
Not Significant
Mean Expression Log-Fold Change

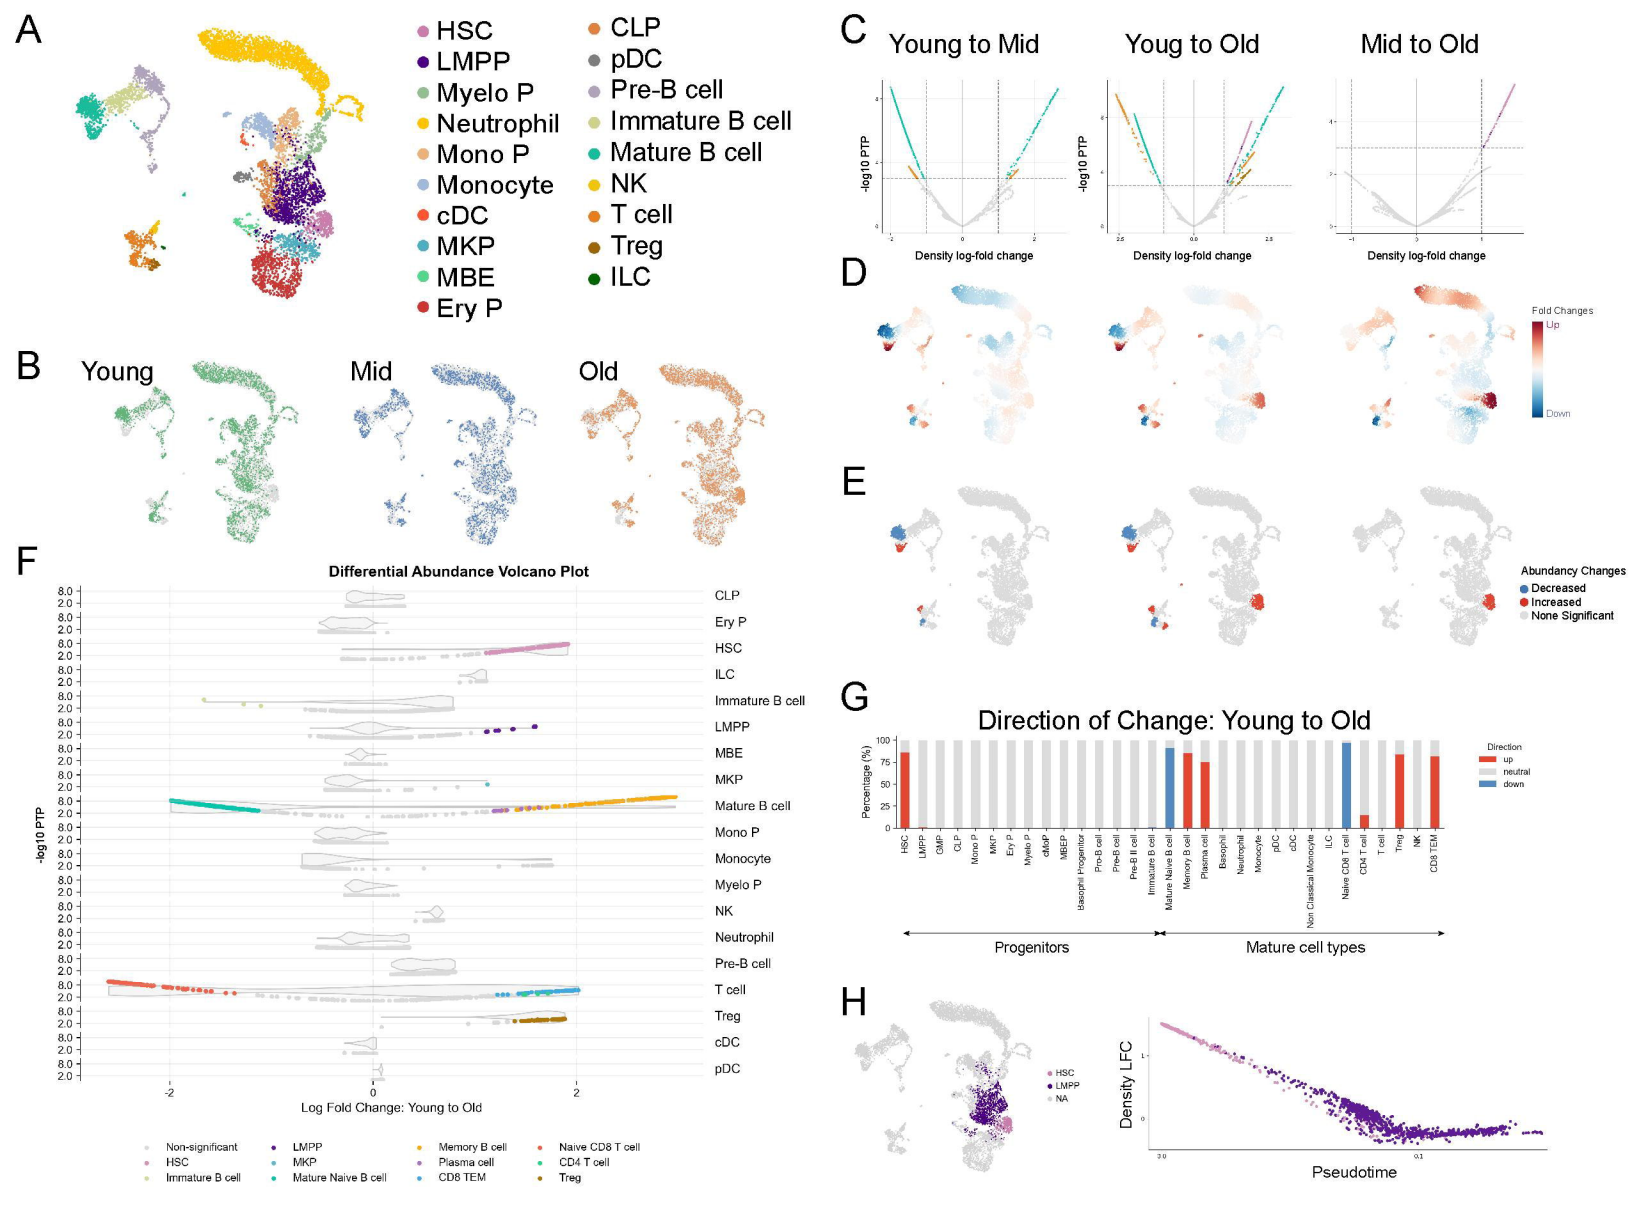
A
HSC
LMPP
Myelo P
Neutrophil
Mono P
Monocyte
cDC
MKP
MBE
Ery P
CLP
pDC
Pre-B cell
Immature B cell
Mature B cell
NK
T cell
Treg
ILC
B
Young
Mid
Old
C
Young to Mid
Youg to Old
Mid to Old
-log10 PTP
Density log-fold change
D
Fold Changes
Up
Down
E
Abundancy Changes
Decreased
Increased
None Significant
F
Differential Abundance Volcano Plot
-log10 PTP
Log Fold Change: Young to Old
Non-significant
HSC
Immature B cell
LMPP
MKP
Mature Naive B cell
Memory B cell
Plasma cell
CD8 TEM
Naive CD8 T cell
CD4 T cell
Treg
G
Direction of Change: Young to Old
Percentage (%)
Direction
up
neutral
down
Progenitors
Mature cell types
H
HSC
LMPP
NA
Density LFC
Pseudotime

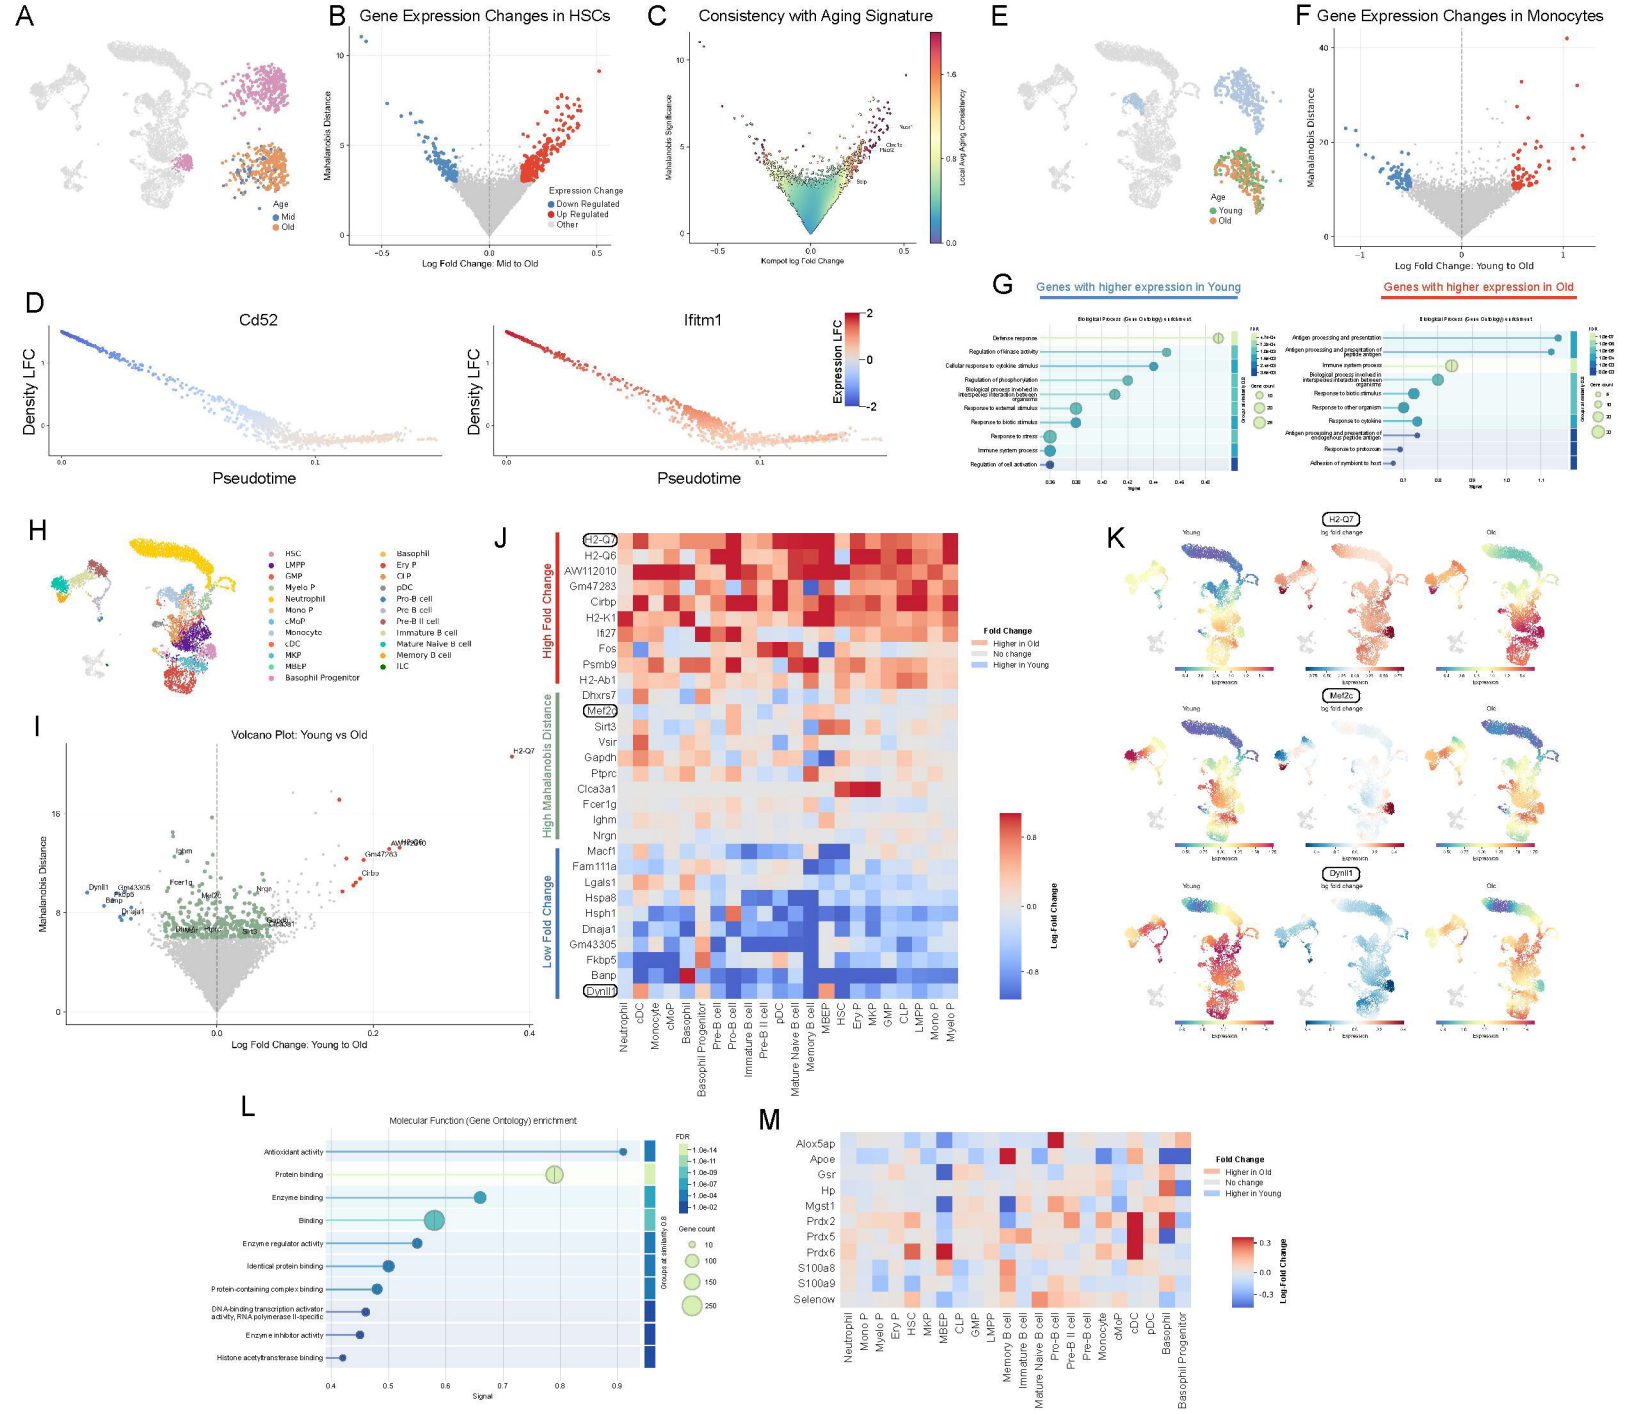
A
Age
Mid
Old
B
Gene Expression Changes in HSCs
Mahalanobis Distance
Log Fold Change: Mid to Old
Expression Change
Down Regulated
Up Regulated
Other
C
Consistency with Aging Signature
Mahalanobis Significance
Kompot log Fold Change
Local Avg Aging Consistency
E
Age
Young
Old
F
Gene Expression Changes in Monocytes
Mahalanobis Distance
Log Fold Change: Young to Old
D
Cd52
Density LFC
Pseudotime
Ifitm1
Density LFC
Pseudotime
Expression LFC
G
Genes with higher expression in Young
Genes with higher expression in Old
Biological Process (Gene Ontology) enrichment
Defense response
Regulation of kinase activity
Cellular response to cytokine stimulus
Regulation of phosphorylation
Biological process involved in interspecies interaction between organisms
Response to external stimulus
Response to biotic stimulus
Response to stress
Immune system process
Regulation of cell activation
Antigen processing and presentation
Antigen processing and presentation of peptide antigen
Immune system process
Biological process involved in interspecies interaction between organisms
Response to biotic stimulus
Response to other organism
Response to cytokine
Antigen processing and presentation of endogenous peptide antigen
Response to protozoan
Adhesion of symbiont to host
Signal
Gene count
H
HSC
LMPP
GMP
Myelo P
Neutrophil
Mono P
cMoP
Monocyte
cDC
MKP
MBEP
Basophil Progenitor
Basophil
Ery P
CLP
pDC
Pro-B cell
Pre B cell
Pre-B II cell
Immature B cell
Mature Naive B cell
Memory B cell
ILC
I
Volcano Plot: Young vs Old
Mahalanobis Distance
Log Fold Change: Young to Old
J
High Fold Change
High Mahalanobis Distance
Low Fold Change
H2-Q7
H2-Q6
AW112010
Gm47283
Cirbp
H2-K1
I627
Fos
Psmb9
H2-Ab1
Dhxrs7
Mef2c
Sirt3
Vsir
Gapdh
Ptprc
Clca3a1
Fcer1g
Ighm
Nrgn
Macf1
Fam111a
Lgals1
Hspa8
Hsph1
Dnaja1
Gm43305
Fkbp5
Banp
Dynll1
Fold Change
Higher in Old
No change
Higher in Young
Log Fold Change
Neutrophil
cDC
Monocyte
cMoP
Basophil
Basophil Progenitor
Pre-B cell
Pro-B cell
Immature B cell
Pre-B II cell
pDC
Mature Naive B cell
Memory B cell
MBEP
HSC
Ery P
MKP
GMP
CLP
LMPP
Mono P
Myelo P
K
H2-Q7
Mef2c
Dynll1
Young
Old
log fold change
expression
L
Molecular Function (Gene Ontology) enrichment
Antioxidant activity
Protein binding
Enzyme binding
Binding
Enzyme regulator activity
Identical protein binding
Protein-containing complex binding
DNA-binding transcription activator activity, RNA polymerase II-specific
Enzyme inhibitor activity
Histone acetyltransferase binding
FDR
Gene count
Groups at similarity 0.8
Signal
M
Alox5ap
Apoe
Gsr
Hp
Mgst1
Prdx2
Prdx5
Prdx6
S100a8
S100a9
Selenow
Fold Change
Higher in Old
No change
Higher in Young
Log Fold Change
Neutrophil
Mono P
Myelo P
Ery P
HSC
MKP
MBEP
CLP
GMP
LMPP
Memory B cell
Immature B cell
Mature Naive B cell
Pro-B cell
Pre-B II cell
Pre-B cell
Monocyte
cMoP
cDC
pDC
Basophil
Basophil Progenitor

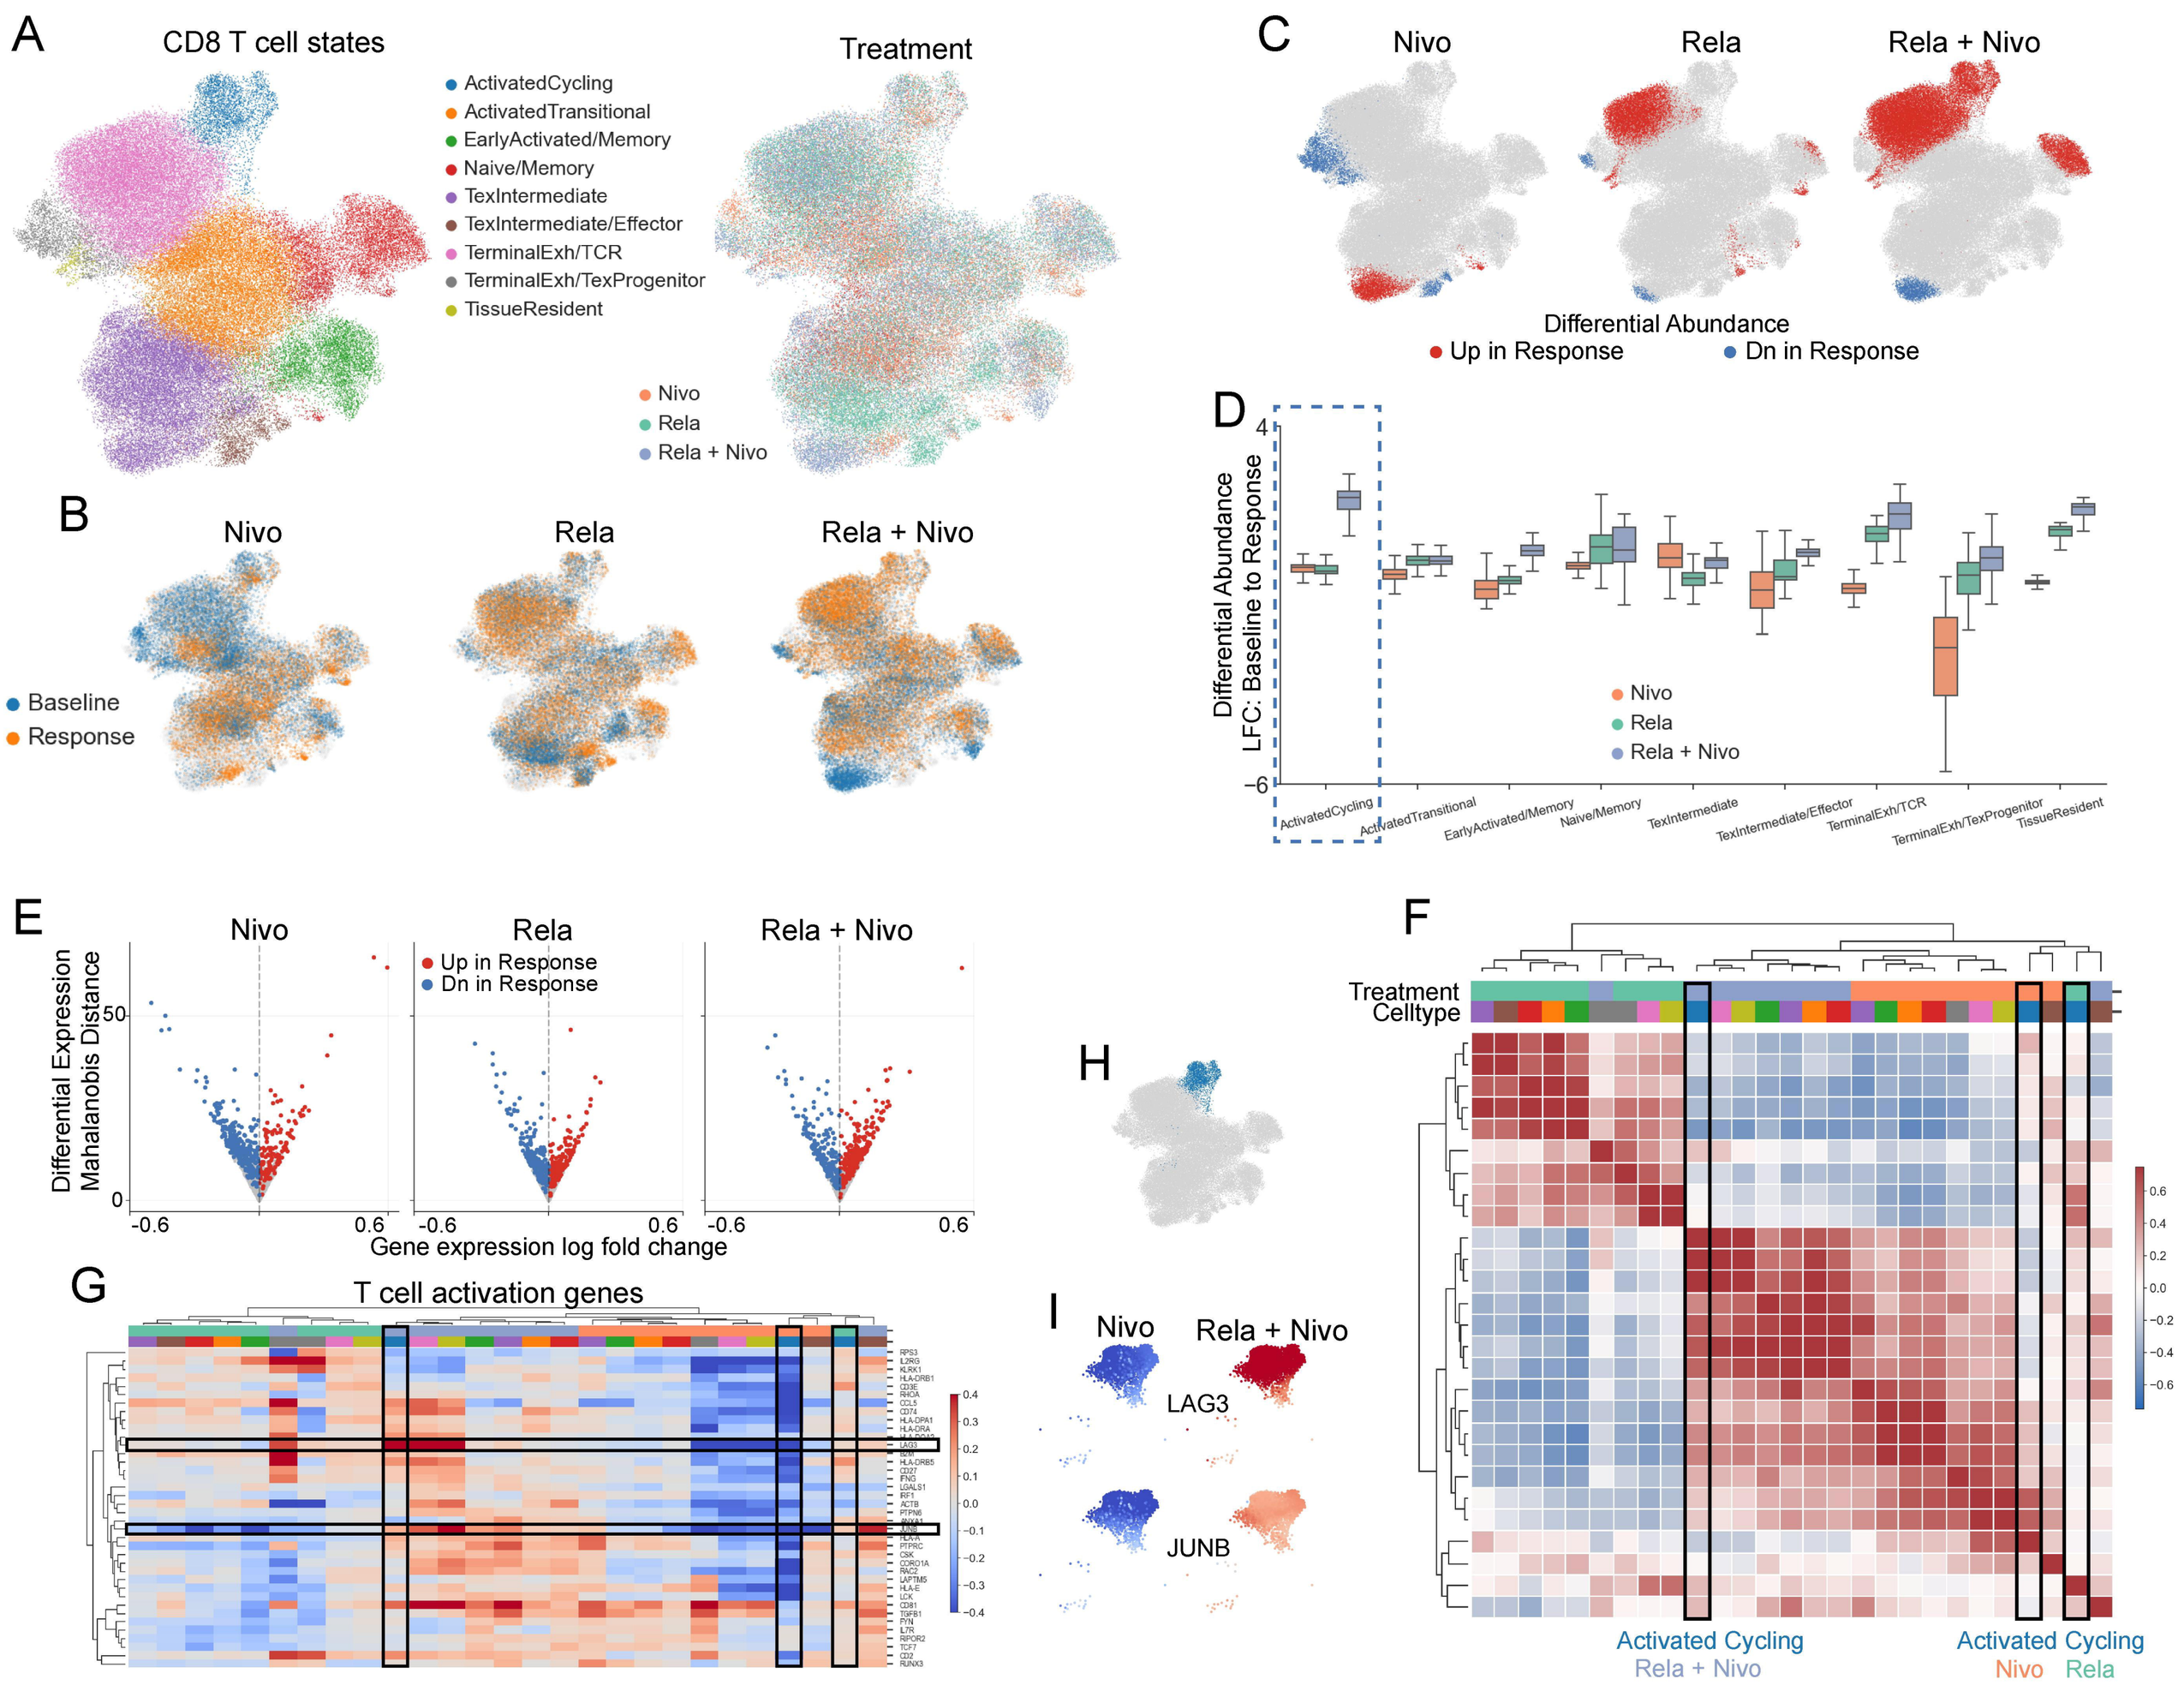
A
CD8 T cell states
ActivatedCycling
ActivatedTransitional
EarlyActivated/Memory
Naive/Memory
TexIntermediate
TexIntermediate/Effector
TerminalExh/TCR
TerminalExh/TexProgenitor
TissueResident
Treatment
Nivo
Rela
Rela + Nivo
B
Nivo
Rela
Rela + Nivo
Baseline
Response
C
Nivo
Rela
Rela + Nivo
Differential Abundance
Up in Response
Dn in Response
D
Differential Abundance
LFC: Baseline to Response
4
−6
Nivo
Rela
Rela + Nivo
ActivatedCycling
ActivatedTransitional
EarlyActivated/Memory
Naive/Memory
TexIntermediate
TexIntermediate/Effector
TerminalExh/TCR
TerminalExh/TexProgenitor
TissueResident
E
Nivo
Rela
Rela + Nivo
Up in Response
Dn in Response
Differential Expression
Mahalanobis Distance
50
0
-0.6
0.6
Gene expression log fold change
G
T cell activation genes
H
I
Nivo
Rela + Nivo
LAG3
JUNB
F
Treatment
Celltype
Activated Cycling
Rela + Nivo
Activated Cycling
Nivo
Rela

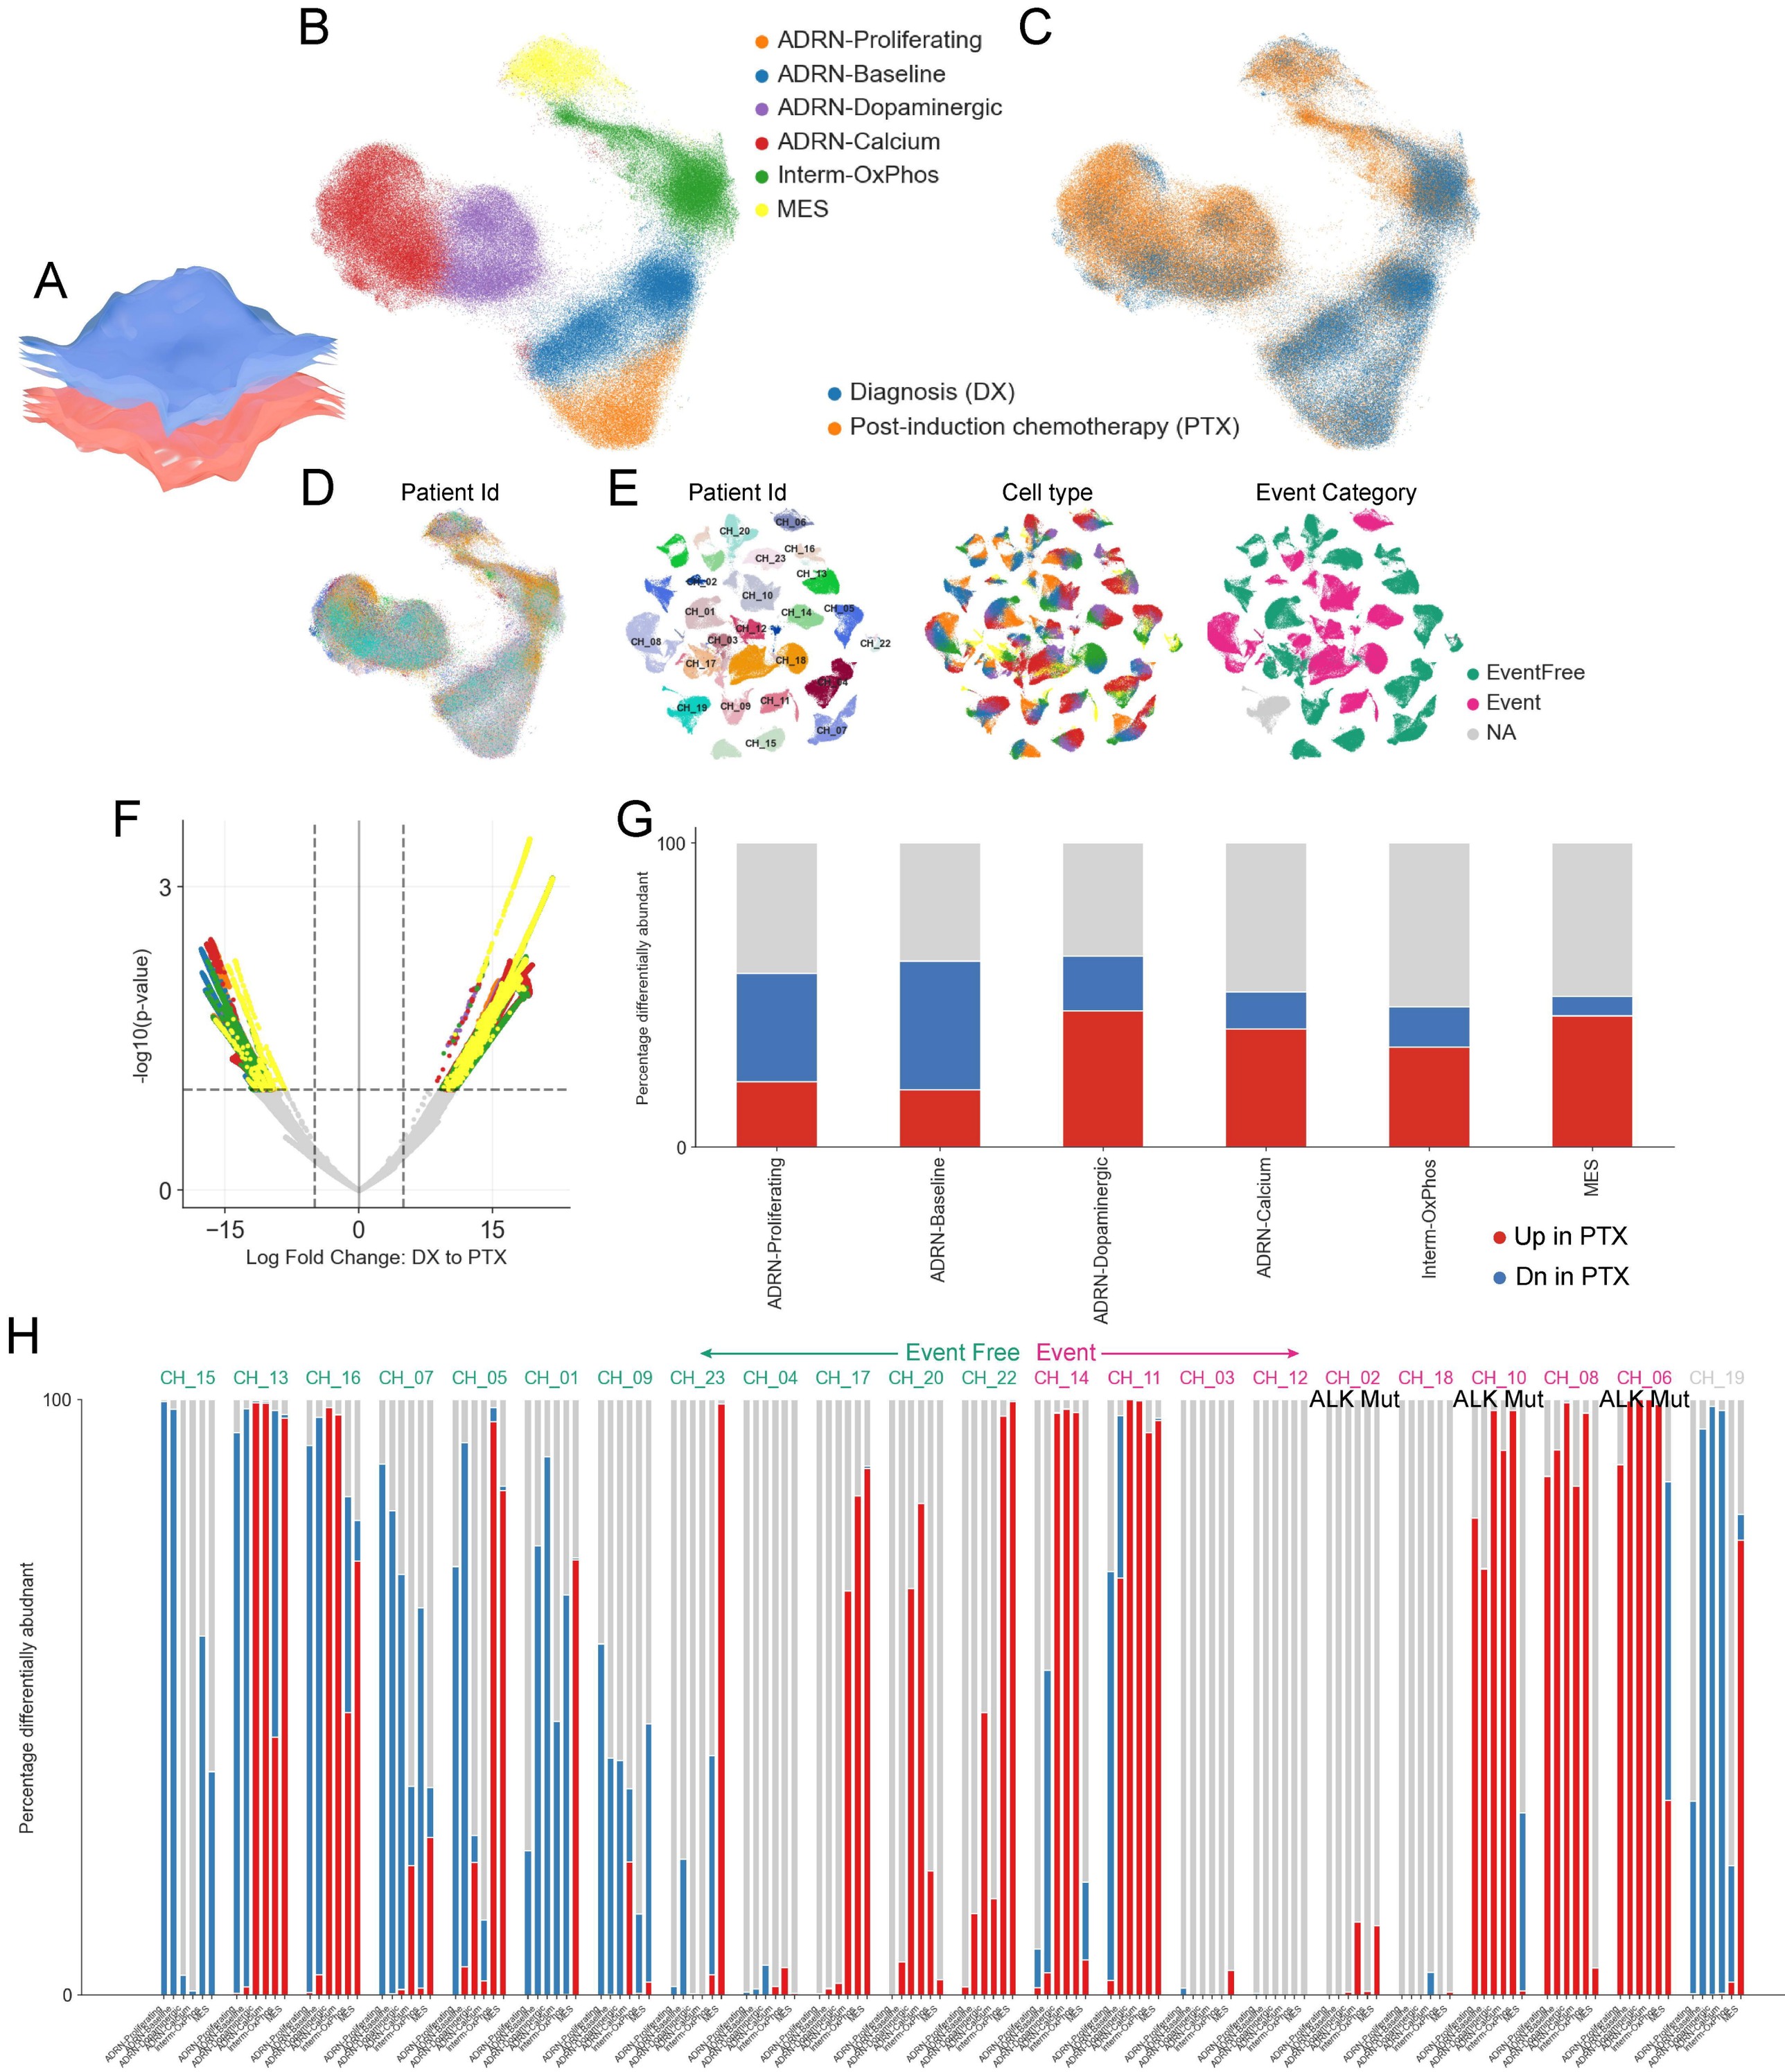
A
B
ADRN-Proliferating
ADRN-Baseline
ADRN-Dopaminergic
ADRN-Calcium
Interm-OxPhos
MES
C
Diagnosis (DX)
Post-induction chemotherapy (PTX)
D
Patient Id
E
Patient Id
Cell type
Event Category
EventFree
Event
NA
F
-log10(p-value)
Log Fold Change: DX to PTX
G
Percentage differentially abundant
Up in PTX
Dn in PTX
H
Event Free
Event
ALK Mut
ALK Mut
ALK Mut
Percentage differentially abundant