DISSECTING TUMOR TRANSCRIPTIONAL HETEROGENEITY FROM SINGLE-CELL RNA-SEQ DATA BY GENERALIZED BINARY COVARIANCE DECOMPOSITION

BY YUSHA LIU[1,a], PETER CARBONETTO[2,b], JASON WILLWERSCHEID[3,c], SCOTT A. OAKES[4,d], KAY F. MACLEOD[5,e] AND MATTHEW STEPHENS[6,f]

[1]*Department of Biostatistics, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA,* [a]*yushaliu@unc.edu*

[2]*Department of Human Genetics and the Research Computing Center, University of Chicago, Chicago, IL, USA,* [b]*pcarbo@uchicago.edu*

[3]*Department of Mathematics and Computer Science, Providence College, Providence, RI, USA,* [c]*jwillwer@providence.edu*

[4]*Department of Pathology, University of Chicago, Chicago, IL, USA,* [d]*soakes@uchicago.edu*

[5]*Ben May Department for Cancer Research, University of Chicago, Chicago, IL, USA,* [e]*kmacleod@uchicago.edu*

[6]*Departments of Statistics and Human Genetics, University of Chicago, Chicago, IL, USA,* [f]*mstephens@uchicago.edu*

Profiling tumors with single-cell RNA sequencing (scRNA-seq) has the potential to identify recurrent patterns of transcription variation related to cancer progression, and produce new therapeutically relevant insights. However, the presence of strong inter-tumor heterogeneity often obscures more subtle patterns that are shared across tumors, some of which may characterize clinically relevant disease subtypes. Here we introduce a new statistical method, generalized binary covariance decomposition (GBCD), to address this problem. We show that GBCD can help decompose transcriptional heterogeneity into interpretable components — including patient-specific, dataset-specific and shared components relevant to disease subtypes — and that, in the presence of strong inter-tumor heterogeneity, it can produce more interpretable results than existing methods. Applied to data from three studies on pancreatic cancer adenocarcinoma (PDAC), GBCD produces a refined characterization of existing tumor subtypes (e.g., classical vs. basal), and identifies a new gene expression program (GEP) that is prognostic of poor survival independent of established prognostic factors such as tumor stage and subtype. The new GEP is enriched for genes involved in a variety of stress responses, and suggests a potentially important role for the integrated stress response in PDAC development and prognosis.

Introduction

Profiling tumors with single-cell RNA sequencing technologies (scRNA-seq) has the potential to generate new insights into progression and metastasis of cancer as well as therapeutic response [1]. In particular, by analyzing patterns of transcriptional variation, studies can identify "gene expression programs" (GEPs) — i.e., sets of genes whose transcription tends to vary in a coordinated fashion [2–5]. Some GEPs may be patient-specific, whereas others may be shared across (subsets of) patients. Crucially, shared GEPs may characterize different molecular subtypes or cellular states, and so provide insights into disease etiology. Further, associating such GEPs with disease progression or therapeutic response may provide medically relevant insights.

A fundamental challenge to identifying shared GEPs is that tumors often exhibit substantial *inter-tumor* heterogeneity (differences in expression from one tumor or patient to another) [1, 6]. Thus, in scRNA-seq datasets involving multiple tumors, malignant cells typically cluster by patient, with little overlap between patients. This strong inter-tumor heterogeneity tends to obscure the more subtle, but nonetheless important, patterns of heterogeneity

shared among tumors. Indeed, as we illustrate in a simulation later, integrative analyses that first combine expression data from malignant cells across all tumors, and then apply off-the-shelf statistical methods (e.g., nonnegative matrix factorization, NMF [7–9]) to identify GEPs, tend to identify patient-specific GEPs, but often miss shared GEPs. Further, the combined analysis of multiple studies — which is desirable due to the typically small number of tumors in individual scRNA-seq studies — introduces potential dataset or batch effects on top of patient effects, further masking shared GEPs.

To address these issues, many recent studies have taken a *tumor-by-tumor* approach in which GEPs that capture *intra-tumor* heterogeneity (differences in expression among cells in a given tumor or patient) are identified separately in each tumor (e.g., by NMF), then compared across tumors to compile a candidate set of shared GEPs [10–16]. However, this strategy often fails to identify some of the more subtle shared GEPs. For example, when tumors are comprised of cells that are homogeneous in terms of molecular subtype identity, the tumor-by-tumor strategy tends to miss GEPs that distinguish among cancer subtypes (see for example the head and neck squamous cell carcinoma case study below). Moreover, this strategy does not benefit from the improved statistical power of an integrated analysis that combines information across tumors.

Another strategy, which tries to combine the benefits of integrated and tumor-by-tumor analyses, is to treat patient effects as "unwanted variation" to be removed [17–19] — a process sometimes known as "harmonization" — and then apply clustering, NMF, or another method for GEP identification to the harmonized data [13, 14, 20–23]. Many harmonization methods have recently been developed for scRNA-seq data [24–29]. While such methods can successfully harmonize datasets from non-diseased tissues [30], they are known to work best when cell types and states are present in comparable proportions across individuals, which is often not the case for tumor data [6]. A recent benchmark study [31] showed that Harmony [24], Conos [28] and Reciprocal PCA (implemented in Seurat 3 [26]) tend to "over-harmonize" cancer datasets, removing genuine variation between patients and thus failing to separate patients with different glioma types.

Therefore, there is an urgent need for improved statistical methods that can better dissect transcriptional heterogeneity in scRNA-seq data into both patient/dataset-specific and shared GEPs. Here we introduce such a method, based on a new approach to matrix factorization that we call "generalized binary covariance decomposition" (GBCD). We illustrate with simulations how GBCD addresses limitations of existing approaches and yields more interpretable results in the presence of strong inter-tumor heterogeneity. To demonstrate its potential to yield new biological insights we apply GBCD to scRNA-seq data from a study on head and neck squamous cell carcinoma (HNSCC), and to three studies on pancreatic cancer adenocarcinoma (PDAC). In HNSCC our approach captures subtype-related GEPs contributing to shared patterns of inter-tumor heterogeneity which were missed by a tumor-by-tumor NMF analysis [10]. In PDAC our analyses provide a refined characterization of previously-identified clinically relevant subtypes, and identify a novel GEP related to stress response pathways that is prognostic of poor survival even when accounting for established prognostic factors such as tumor stage and subtype.

Results

Methods overview. We assume that the combined scRNA-seq data from all tumors are contained in an $N \times J$ matrix $\boldsymbol{Y}$ of expression values with entries y_{ij}, where $i = 1, \ldots, N$ indexes cells and $j = 1, \ldots, J$ indexes genes. In typical applications, $\boldsymbol{Y}$ contains log-transformed pseudo-count-modified UMI counts ("log-pc counts" for brevity; see Section S.1.1.1 for details). Our method is "unsupervised" in that, unlike tumor-by-tumor and many harmonization approaches, it does not use information about which cell comes from which tumor.

Our method ultimately yields a decomposition of the data matrix $\boldsymbol{Y}$ into matrices $\boldsymbol{L}$ and $\boldsymbol{F}$ such that $\boldsymbol{Y} \approx \boldsymbol{L}\boldsymbol{F}^T$, or equivalently

$$y_{ij} \approx \sum_{k=1}^{K} l_{ik} f_{jk}. \tag{1}$$

Our goal is that the K components in (1) should be interpretable as GEPs, with l_{ik} representing the membership of cell i in GEP k, and f_{jk} representing the effect of GEP k on the expression of gene j. When y_{ij} are log-pc counts, each f_{jk} approximately represents the *log-fold change* (LFC) of gene j in a cell, i, with membership $l_{ik} = 1$ relative to a cell i' with membership $l_{i'k} = 0$. We refer to the f_{jk} values as LFCs, and we call $\boldsymbol{f}_k := (f_{1k}, \ldots, f_{Jk})^T$ the "signature" of GEP k.

Throughout, we adopt "GEP" as a term of art for the components in (1), but we emphasize that some components may not have a coherent biological function; for example, some components may capture technical effects, or effects on gene expression induced by copy number variation. Further, since our method does not know which cells come from which tumors, any patient-specific or batch effects would also appear as "GEPs". That is, our matrix decomposition takes an agnostic, unsupervised, approach to discovering components of transcriptional variation, which may be both technical and biological; determining which components are most biologically relevant is done by subsequent analyses of the gene signatures (see below).

We make two key assumptions about the matrix decomposition (1):

Assumption 1. Memberships l_{ik} are nonnegative, and often nearly binary (close to 0 or 1).

Assumption 2. The GEP signatures $\boldsymbol{f}_k$ are mutually orthogonal.

Assumption 1 is motivated by our particular interest in capturing discrete substructures such as patient effects and tumor subtypes, where membership among cells can be modeled as binary. However, since non-discrete structures may also be of biological interest, and can naturally arise from spatial or continuous cell processes, we implement this assumption as a "soft" assumption; that is, our modeling approach encourages, but does not require, binary membership values. We refer to this assumption as the *generalized binary assumption.*

Assumption 2 is motivated by our empirical observation that, without this assumption, a shared GEP can get absorbed into multiple patient-specific effects (namely, the effects corresponding to the patients who share that GEP). The orthogonality assumption helps to preserve the shared GEP because absorbing it into multiple patient-specific GEPs would make the signatures of the patient-specific effects dependent, and so non-orthogonal. The MNN method of [27] also makes an orthogonality assumption, and indeed we see in the simulation study that this helps MNN separate shared GEPs from patient effects. Similar assumptions have also been suggested in computer vision as a way to force matrix decomposition methods to learn factors that capture shared image features [32–34].

Note that (1) allows for modeling *overlapping structures* in the cell population, in that multiple nonzero l_{ik}'s are permitted for each cell i. For example, a given cell might have membership in four GEPs — a dataset-specific GEP, a patient-specific GEP, a tumor subtype GEP, and a cellular state GEP — and this configuration would be represented by four nonzero l_{ik}'s. This ability to model overlapping structures is a key feature that distinguishes our approach from clustering methods, which identify non-overlapping structures.

Combining (1) with Assumption 2 implies the following:

$$\boldsymbol{Y}\boldsymbol{Y}^T \approx \boldsymbol{L}\boldsymbol{F}^T\boldsymbol{F}\boldsymbol{L}^T = \boldsymbol{L}\boldsymbol{D}\boldsymbol{L}^T, \tag{2}$$

where $\boldsymbol{D} := \boldsymbol{F}^T\boldsymbol{F}$ is a diagonal matrix. In practice, our approach proceeds by first finding $\boldsymbol{L}$ and $\boldsymbol{D}$ in (2), and then estimating $\boldsymbol{F}$ based on (1) (at which point we do not force the

orthogonality assumption, so in this sense Assumption 2 is also a "soft" assumption.) The matrix $\boldsymbol{Y}\boldsymbol{Y}^T$ is closely related to the cell-by-cell covariance matrix, and so we refer to the decomposition (2) as *generalized binary covariance decomposition* (GBCD), and for brevity here we also use GBCD to refer to our overall approach. To implement GBCD we apply the empirical Bayes matrix factorization (EBMF) method [35] implemented in the R package `flashier` [36–38]. Conveniently, EBMF provides an automatic way to select the number of components K (Section S.1.3).

We assess the biological relevance of inferred GEPs by examining the top "driving genes" — that is, genes j in GEP k with largest magnitude of f_{jk} — and we perform gene set enrichment analyses [39] to identify biological processes enriched for the driving genes. While driving genes can be either up-regulated (positive f_{jk}) or down-regulated (negative f_{jk}), in our analyses we focus on up-regulated genes. This facilitates comparisons with previously derived expression programs which have usually focused on up-regulated genes.

Illustrative example. To illustrate the difference between GBCD and existing integrative strategies — including NMF methods, which are commonly used to identify structure in scRNA-seq data — we use a simple simulation designed to mimic patient-specific effects and different types of shared gene expression programs. Specifically we simulate scRNA-seq counts for $J = 10,000$ genes in $N = 3,200$ cells from 8 patients (400 cells per patient), with three different types of GEPs (Fig. 1A):

- Eight patient-specific GEPs: GEPs P1 to P8 are each uniquely expressed by a single patient, with binary memberships. These GEPs mimic the strong patient effects that are present in single-cell tumor data.
- Two subtype-related GEPs: GEPs S1 and S2 are respectively expressed in patients 1–4 and 5–8 with binary memberships. These GEPs mimic cancer subtype programs shared by multiple patients.
- One continuous GEP: this GEP is expressed in 600 cells randomly chosen from the 3,200 cells independently of patient identity, with non-zero membership values drawn from $\text{Unif}(0.4, 2)$. This GEP mimics some cellular process or activity whose expression intensity shows continuous variation across cells and may be unrelated to cancer subtype.

In short, the continuous GEP contributes to shared patterns of *intra-tumor* heterogeneity across patients whereas the subtype-related GEPs do not. Further, the subtype-related GEPs induce *hierarchical* relationships between patients, in the sense that patients 1–4 are transcriptionally more closely related to each other than patients 5–8; however, we note that the simulated data do not conform to a strict hierarchy due to the continuous GEP whose expression levels in cells are independent of patient or subtype identity.

Each of the 11 GEPs has a set of differentially expressed genes unique to that program, resulting in GEP signatures that are orthogonal to each other (satisfying Assumption 2). In order not to unfairly disadvantage the NMF-based approaches, the differentially expressed genes for each GEP are up-regulated relative to some baseline expression level by design, so that the true GEP signatures are non-negative. More details on data generation are provided in Section S.1.2.

We independently simulated 20 scRNA-seq datasets under this scenario. Results for a typical simulation are shown in Fig. 1A, Supplementary Fig. S1 and S2. In this dataset, the strong simulated patient effects — that is, the large expression differences between patients — are reflected in the *t*-SNE embedding [41], in which the primary structure is the clustering by patient (Fig. 1B).

We first compared GBCD with NMF applied to combined expression data across all patients ("combined NMF"). GBCD and both versions of combined NMF (applied to UMI or

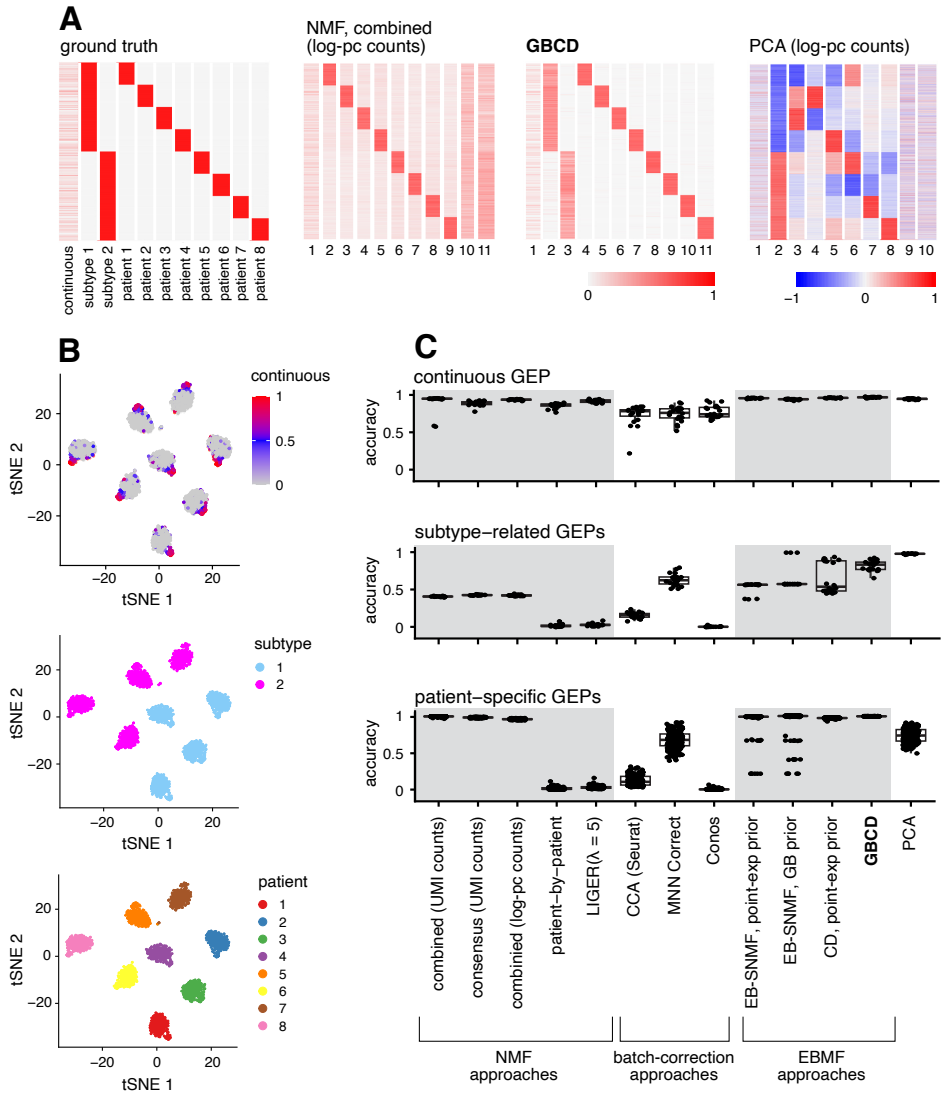

Fig. 1. Results of applying different methods to identify GEPs from simulated multi-patient scRNA-seq data. Panel A: the 3,200 × 11 membership matrix L used to simulate the data, and estimates of L obtained by applying NMF, PCA and GBCD to one of the 20 simulated datasets. Rows are cells, columns are GEPs. In these heatmaps, the columns were scaled separately so that the largest value in each column was 1. All methods identified one or more components that were strongly correlated with cellular detection rate [40]; these components were not included in the heatmaps. See Supplementary Fig. S1 for membership estimates produced by the other methods compared. Panel B: *t*-SNE visualizations of the same scRNA-seq data, in which cells are colored by membership to the continuous GEP, subtype GEPs, and patient-of-origin GEPs. Panel C: Performance of all methods in recovering the continuous, subtype-related and patient-specific GEPs in 20 simulated datasets. See Online Methods for a detailed description of the methods compared. For each true GEP, accuracy of a method's estimate was measured using the highest Pearson correlation between the true GEP membership and the estimated membership among all GEPs identified by the method.

log-pc counts) successfully identified patient-specific components (Fig. 1A, Supplementary Fig. S1), with GBCD providing corresponding membership estimates that were visually the least noisy. All three methods also effectively found the continuous GEP (Supplementary Fig. S2). Notably, GBCD was able to denoise the membership estimates of patient-specific

components without compromising the estimates of the non-zero membership values of the continuous component, highlighting the flexibility of the generalized binary assumption.

The primary difference among the three methods was that GBCD accurately identified the two subtype-related GEPs, whereas the combined NMF methods did not (Fig. 1A, Supplementary Fig. S1). This difference is traceable to the orthogonality assumption of GBCD; without an orthogonality constraint, the combined NMF methods can produce a solution where the subtype-related GEP signatures are absorbed into the patient-specific GEP signatures. That is, instead of estimating the 8 patient-specific components to have the "true" orthogonal gene signatures $\boldsymbol{f}_{\mathsf{P1}}, \ldots, \boldsymbol{f}_{\mathsf{P8}}$, the combined NMF methods estimated the patient-specific gene signatures as $\boldsymbol{f}_{\mathsf{P1}} + \boldsymbol{f}_{\mathsf{S1}}, \boldsymbol{f}_{\mathsf{P2}} + \boldsymbol{f}_{\mathsf{S1}}, \boldsymbol{f}_{\mathsf{P3}} + \boldsymbol{f}_{\mathsf{S1}}, \boldsymbol{f}_{\mathsf{P4}} + \boldsymbol{f}_{\mathsf{S1}}, \boldsymbol{f}_{\mathsf{P5}} + \boldsymbol{f}_{\mathsf{S2}}, \boldsymbol{f}_{\mathsf{P6}} + \boldsymbol{f}_{\mathsf{S2}}, \boldsymbol{f}_{\mathsf{P7}} + \boldsymbol{f}_{\mathsf{S2}}, \boldsymbol{f}_{\mathsf{P8}} + \boldsymbol{f}_{\mathsf{S2}}$. This decomposition of the expression data $\boldsymbol{Y}$ is less desirable, because it does not separate the shared components of expression from the patient-specific components, thus failing to recover the largely hierarchical relationships between patients. Since this solution does not not satisfy the orthogonality assumption (e.g., $\boldsymbol{f}_{\mathsf{P1}} + \boldsymbol{f}_{\mathsf{S1}}$ and $\boldsymbol{f}_{\mathsf{P2}} + \boldsymbol{f}_{\mathsf{S1}}$ are not orthogonal to each other because they share $\boldsymbol{f}_{\mathsf{S1}}$), it is ruled out by GBCD, which is thus encouraged to separate out the subtype-related GEPs.

To confirm that the differences shown above were systematic, rather than the chance outcome of a single simulation, we summarized the performance of these approaches to recover different types of GEP across the 20 replicate simulations. The results (Fig. 1C) confirmed that both GBCD and combined NMF consistently detected the patient-specific and shared continuous GEPs, whereas GBCD more reliably recovered the shared subtype-related GEPs.

We next compared GBCD with principal component analysis (PCA) applied to combined log-pc counts, another widely used approach for dimension reduction of scRNA-seq data. PCA successfully identified both the subtype-related and the continuous GEP, but it failed to identify the patient-specific GEPs (Fig. 1A). In place of patient-specific GEPs, PCA identified several principal components (PCs) that were loaded on cells belonging to a seemingly random subset of patients. This is problematic for interpretation because these PCs appear, falsely, to capture features in the data that are shared across patients. Part of the problem with PCA is that it constrains the PCs $\boldsymbol{l}_k$ to be orthogonal to each other, which is not satisfied when hierarchical relationships exist between patients (e.g., $\boldsymbol{l}_{\mathsf{S1}}$ and $\boldsymbol{l}_{\mathsf{P1}}$ are not orthogonal to each other). Analysis of results across all the 20 replicate simulations confirmed that PCA had substantially worse specificity for GEP discovery than GBCD, in the sense that a much larger proportion of the GEPs estimated by PCA (i.e., PCs) did not correlate with any of the 11 true GEPs (Supplementary Fig. S3).

Other strategies similarly failed to identify the shared subtype-related GEPs. (Details of these methods and their implementations are given in Sections S.1.3 and S.1.4.) Consensus NMF [5], which averages the solutions of combined NMF over multiple runs, failed to identify subtype effects because NMF could not detect such effects in the first place. Patient-by-patient NMF, which fits NMF to log-pc counts separately for each patient and then clusters the GEPs detected, failed to identify subtype effects because such effects do not vary across cells within a patient thus do not contribute to intra-tumor heterogeneity. A variation on our GBCD method, which replaces Assumption 1 with a generic sparsity assumption (point exponential priors), identified some subtype effects, but not as reliably as GBCD, and did not provide as clear a separation between cells that were active and inactive in each GEP having binary memberships. This finding demonstrates the importance of Assumption 1 in identifying discrete structures. Other variants of empirical Bayes matrix factorization that drop Assumption 2 tend, like NMF, to miss subtype-related GEPs. "Batch effect correction"-type methods that integrate scRNA-seq data from multiple samples, including LIGER [25], CCA implemented in Seurat [26], MNN Correct [27] and Conos [28], likewise failed to identify subtype effects, and also occasionally produced much worse membership estimates of the

continuous GEP, consistent with [31] which observed a tendency of these methods to over-correct cancer data. Among the batch-correction methods, MNN Correct was most successful in identifying components related to the subtype-related GEPs, and MNN Correct also makes a orthogonality assumption.

This simple example demonstrates that GBCD can produce qualitatively different results from NMF and other strategies, and highlights the potential of GBCD to identify GEPs that are shared across patients but do not necessarily lead to intra-tumor heterogeneity, which other methods may miss. However, we note that the representations obtained by existing methods here are not "incorrect". For example, the combined NMF solution is essentially equivalent to clustering cells by patient, which in some contexts could be the desired goal. In practical applications it may be helpful to run multiple methods to obtain different views of the structure present in the data, and GBCD may be seen as complementing existing analysis methods, rather than replacing them.

Head and neck squamous cell carcinoma data. To compare GBCD and NMF-based methods on real cancer data, we analyzed scRNA-seq data collected by Puram et al. [10] from primary tumors from 10 HNSCC patients and matching lymph node (LN) metastases from 5 of these patients. Puram et al. found that each of these 10 patients clearly mapped to a molecular subtype of HNSCC, whose signatures were previously defined by analysis of bulk expression data from 279 HNSCC tumors [42]. Here we demonstrate that GBCD can extract this molecular subtype information *de novo* from the single cell data alone, which other methods struggle to do.

We analyzed the $N = 2{,}176$ malignant cells, whose integration among tumors presents a greater challenge than non-malignant cells. As noted by Puram et al. [10], these cells demonstrate strong patient effects that are typical of cancer data; the major structure in the *t*-SNE visualization is the clustering of the cells by patient (Fig. 2A). Because of these strong patient effects, Puram et al. took a tumor-by-tumor analysis approach: they applied NMF separately to each tumor to identify GEPs, and then used clustering to identify similar GEPs from different tumors and create consensus GEPs that they called "meta-programs". Their analysis identified 6 meta-programs associated with cell cycle, stress, hypoxia, epithelial differentiation, and epithelial-mesenchymal transition or EMT (referred to by [10] as "partial EMT" because it demonstrated some features of classical EMT but lacked others). However, this tumor-by-tumor analysis of the single cell data failed to identify GEPs that classify the patients into their molecular subtypes. This illustrates a limitation of tumor-by-tumor analysis, which can only identify GEPs whose memberships vary sufficiently across cells within each tumor.

We applied GBCD to these data, estimating GEP memberships and signatures, and annotating the GEPs by their driving genes and gene set enrichment analysis, or GSEA (Fig. 2B, Table 1, Supplementary Fig. S18, Supplementary Tables S1–S2, Section S.1.6). In total, GBCD identified 19 GEPs (not including an "intercept" GEP on which all cells are loaded). Among the 19 GEPs, 12 were active in multiple patients and the remaining 7 were specific to an individual patient (Section S.1.8). Among the 12 shared GEPs, at least one GEP showed significant overlap with each of the 6 meta-programs from [10]. However, some of the shared GEPs were not strongly correlated with any of the 6 meta-programs; for example, GEP11 (enriched for genes involved in respiratory electron transport) and GEP12 (mRNA splicing). Most importantly, GBCD identified GEPs corresponding closely to the previously defined molecular subtypes: GEP1 was largely active only in cells from the 2 classical patients, and GEP2 was mainly active in cells from the 7 basal patients. Therefore, GEPs 1 and 2 were sufficient to accurately classify the 10 patients into the 3 subtypes. We emphasize that this analysis was entirely unsupervised; GBCD does not use the patient or subtype labels. Analysis of variance (ANOVA) of GEP memberships (Section S.1.8) for each of the 12 shared

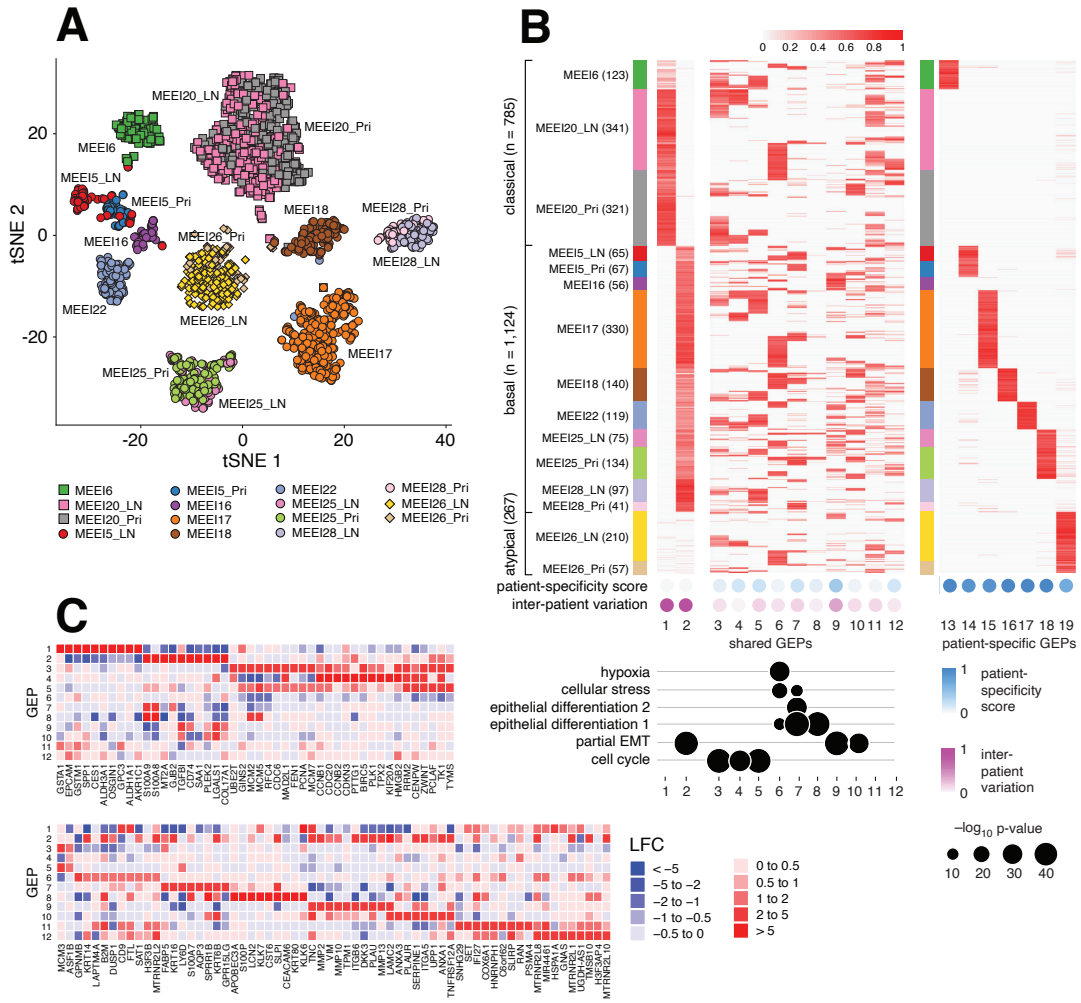

Fig. 2. GBCD analysis of HNSCC data. Panel A: *t*-SNE embedding from $N = 2{,}176$ scRNA-seq profiles in 10 primary tumors and 5 matched lymph node (LN) metastases. Cells are colored by patient of origin and tumor stage (primary tumor, LN metastasis), and shaped by tumor molecular subtype (square: classical; circle: basal; diamond: atypical). Panel B: Heatmap showing membership values of the $N = 2{,}176$ cells (rows) for the 19 GEPs (columns) identified by GBCD, in which cells are arranged top-to-bottom by tumor molecular subtype and patient, and GEPs are grouped left-to-right based on whether they are more shared across patients (GEP 1–12) or more patient-specific (GEP 13–19). For the heatmap, membership values were rescaled separately for each GEP so that the maximum membership for each GEP was always 1. Numbers in parentheses give the number of cells per tumor. Measures of patient-specific expression (relative to shared expression across patients) and inter-patient variation in expression (relative to intra-patient variation) are shown. See Section S.1.8 for definitions. Concordance between the gene signatures of shared GEPs and 6 previously identified meta-programs from a tumor-by-tumor analysis [10] is quantified by $-\log_{10}(p\text{-value})$, where *p*-values are produced from one-sided Wilcoxon rank-sum tests (Online Methods). Panel C: LFC estimates for selected genes in the shared GEPs.

GEPs revealed that GEPs 1 and 2 showed much greater inter-patient than intra-patient variation whereas the opposite was true for the other shared GEPs (Fig. 2B). This is consistent with the observation that tumor cells from an individual HNSCC patient are largely homogeneous in terms of molecular subtype identity [10].

TABLE 1

Top driving genes and enriched gene sets of shared GEPs identified in the HNSCC data. The left-hand column lists the genes with the largest increases in expression (LFC) in the GEP. The right-hand column lists the enriched MSigDB Hallmark and CP:Reactome gene sets [39, 43, 44] with Bonferroni-adjusted p-value < 0.05.

GEP	Top 10 driving genes	Top significantly enriched gene sets
1	*GSTA1, EPCAM, GSTM1, SPP1, CES1, ALDH3A1, OSGIN1, GPC3, ALDH1A1, AKR1C1*	Biological oxidations, Phase II—conjugation of compounds, Genes up-regulated through activation of mTORC1 complex, Metabolism of lipids, Phase I—functionalization of compounds
2	*S100A9, S100A8, MT2A, GJB2, TGFBI, CD74, SAA1, PLEK2, LGALS1, COL17A1*	Genes defining epithelial-mesenchymal transition, Genes up-regulated in response to IFNG, Interferon signaling, Interferon gamma signaling, Interferon alpha/beta signaling
3	*UBE2T, GINS2, MCM2, MCM5, RFC4, CDC6, MAD2L1, FEN1, PCNA, MCM7*	Genes encoding cell cycle related targets of E2F transcription factors, Cell cycle, Cell cycle (Mitotic), Cell cycle checkpoints, Genes involved in the G2/M checkpoint
4	*CCNB1, CDC20, CCNB2, CDKN3, PTTG1, BIRC5, PLK1, TPX2, KIF20A, HMGB2*	Genes involved in the G2/M checkpoint, Cell cycle, Cell cycle (mitotic), Mitotic prometaphase, Genes encoding cell cycle related targets of E2F transcription factors
5	*RRM2, CENPW, ZWINT, PCLAF, TK1, MCM5, TYMS, MCM3, MCM7, ASF1B*	Genes encoding cell cycle related targets of E2F transcription factors, Cell cycle, Cell cycle (mitotic), DNA strand elongation, Activation of ATR in response to replication stress
6	*GPNMB, LAPTM4A, KRT14, B2M, DUSP1, CD9, FTL, SAT1, H3F3B, MTRNR2L2*	Genes involved in p53 pathways and networks, Genes up-regulated in response to low oxygen levels (hypoxia), Genes regulated by NF-kB in response to TNF, Cell junction organization
7	*FABP5, KRT16, LY6D, S100A7, S100A8, AQP3, SPRR1B, KRT6B, GPR15LG, S100A9*	Formation of the cornified envelope, Keratinization, Developmental biology, Neutrophil degranulation, Genes defining late response to estrogen
8	*S100A9, APOBEC3A, S100P, LCN2, KLK7, CST6, SLPI, CEACAM6, KRT80, KLK6*	Formation of the cornified envelope, Keratinization, Antimicrobial peptides
9	*TNC, MMP2, VIM, MMP10, TPM1, ITGB6, DKK3, PLAU, TGFBI, MMP13*	Genes defining epithelial-mesenchymal transition, Extracellular matrix organization, Collagen formation, Assembly of collagen fibrils and other multimeric structures, ECM proteoglycans
10	*LAMC2, ANXA3, PLAUR, LGALS1, SERPINE1, ITGA5, UPP1, ANXA1, TNFRSF12A, TNC*	Genes defining epithelial-mesenchymal transition, Genes regulated by NF-kB in response to TNF, Hemostasis, Cell junction organization, Genes encoding components of apical junction complex
11	*SNHG29, SET, IFI27, B2M, COX6A1, HNRNPH1, C6orf62, SLIRP, RAN, HMGB1*	Respiratory electron transport, ATP synthesis by chemiosmotic coupling, The citric acid cycle and respiratory electron transport, Cellular response to chemical stress, A subgroup of genes regulated by MYC (version 1)
12	*MTRNR2L8, MTRNR2L2, MIR4461, HSPA1A, GNAS, MTRNR2L1, UGDH-AS1, TMSB10, H3F3AP4, MTRNR2L10*	mRNA Splicing, Processing of capped intron-containing pre-mRNA, mRNA Splicing—minor pathway, Processing of replication-dependent histone pre-mRNAs, RNA polymerase II transcription termination

Results of alternative analysis methods, including PCA and NMF-based methods, are summarized in Supplementary Fig. S5–S11. These results highlight behaviors that are qualitatively similar to those seen in the illustrative simulation above; for example, NMF captured the patient-specific programs, but failed to identify the molecular subtypes. PCA successfully identified the molecular subtypes, but it did not separate the patient-specific programs; instead, it reported multiple factors that appear to be random combinations of patient-specific effects. PCA was the only method other than GBCD to identify the molecular subtypes but

TABLE 2
PDAC scRNA-seq datasets analyzed. The datasets are labeled as "A", "B" and "C" in the results. The right-most column gives the number of high-quality malignant cells retained for analysis.

label	study	assay	specimens	*N*
A	Chan-Seng-Yue et al. [45]	10x 3' v2	13 primary, 1 metastatic	18,920
B	Peng et al. [46]	10x 3' v2	24 primary	10,936
C	Raghavan et al. [47]	Seq-Well	21 metastatic	5,814

otherwise produced results quite different to GBCD; indeed PCA failed to recapitulate any of the remaining 10 shared GEPs identified by GBCD (Supplementary Fig. S12). Taken together, these results illustrate the ability of GBCD to identify biologically meaningful expression programs explaining both inter-tumor and intra-tumor transcriptional heterogeneity.

Pancreatic ductal adenocarcinoma data. PDAC is highly lethal and demonstrates extensive heterogeneity in disease progression and treatment response among patients [45]. Many studies have attempted to define clinically relevant subtypes with the goal of personalizing treatment and improving outcomes [48]. An early study [49] applied NMF to microarray expression data and identified "classical" and "basal" subtypes. Later studies used newer technologies to refine these subtypes: Chan-Seng-Yue et al. [45] identified 4 subtypes — "classical-A", "classical-B", "basal-A" and "basal-B" — from an NMF analysis of bulk RNA-seq data in tumor specimens that underwent laser capture microdissection to improve tumor purity; Raghavan et al. [47] refined the basal and classical gene signatures derived from [49] using scRNA-seq data from metastatic patients; Hwang et al. [50] identified a classical and three basal-related expression programs ("basaloid", "squamoid" and "mesenchymal") from an NMF analysis of single-nucleus RNA sequencing data in primary tumors.

Understanding transcriptional heterogeneity in PDAC tumors is, therefore, a complex and evolving area. To examine whether GBCD can contribute new insights, we performed a *de novo* analysis of combined scRNA-seq data on PDAC tumors from three recent studies, which together include data on 35,670 malignant cells from 59 PDAC tumors (Table 2). Focusing on malignant cells prevents our analysis of tumor transcriptional heterogeneity from being influenced by a high degree of immune and stromal cell infiltration, which is a particular issue for bulk RNA-seq data analysis of PDAC tumors given their typically low tumor purity [45]. Further, since GBCD can effectively disentangle shared GEPs from patient- and study-specific effects, we expected our analysis to provide a deeper characterization of recurrent gene expression patterns by pooling tumors from multiple studies to achieve a larger sample size.

A 2-D embedding of the 35,670 scRNA-seq profiles by *t*-SNE shows that the cells cluster by patient, reflecting substantial inter-patient heterogeneity, and also by cohort, potentially reflecting (expected) technical and biological differences among studies (Fig. 3A). The substantial inter-patient and inter-study structure could obscure biologically interesting expression patterns recurring among the tumors. Applying GBCD to the combined expression data identified a total of 33 GEPs (not including an intercept GEP), whose memberships are shown in Fig. 3B. Consistent with the *t*-SNE analysis, a good number of GEPs identified by GBCD (GEPs 21–33) are each active predominantly in one patient, and thus mainly capture patient-specific effects, in part due to the influences of aneuploid copy number profiles that are common in most human tumors and can vary from patient to patient [51]. The patient-specific GEPs exhibit strong chromosomal structure (Fig. 3B, Supplementary Fig. S13; Online Methods), which suggests a link to copy number variation in tumor cells caused by aneuploidy.

Despite the strong inter-tumor expression heterogeneity, GBCD identified many factors (GEPs 1–20 in Fig. 3B) recurring in multiple tumors that represent potentially interesting

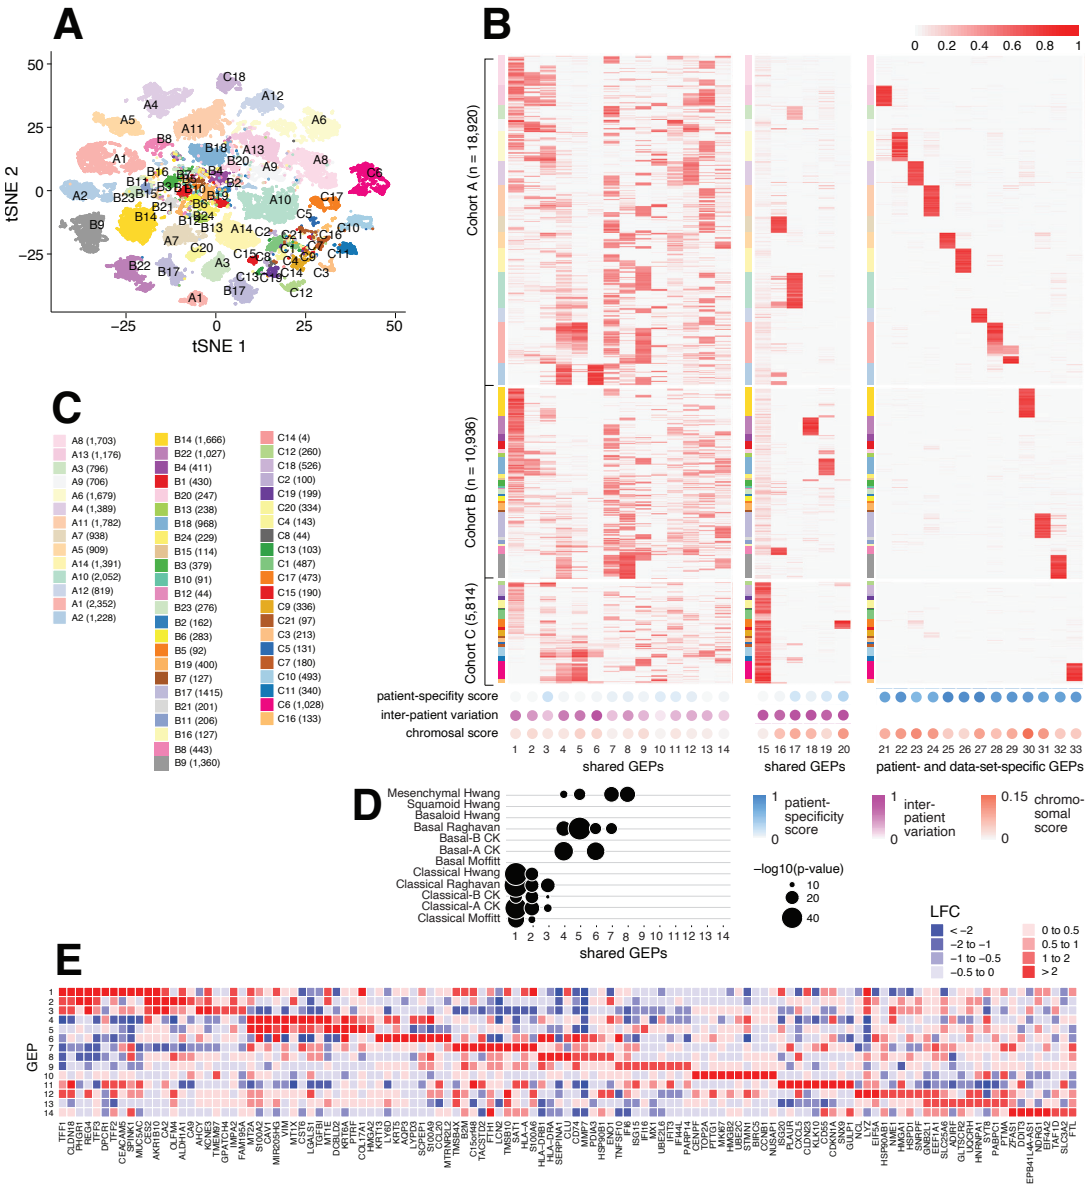

Fig. 3. GBCD analysis of PDAC data. Panel A shows the *t*-SNE embedding from scRNA-seq profiles of $N = 35{,}670$ malignant cells from 59 PDAC tumor samples (Table 2; Panel C). Cells are colored by patient of origin. Two patients (A1, B17) are labeled twice since the cells from those patients were split into two distinct clusters. Panel B is a heatmap showing membership values of the $N = 35{,}670$ cells (rows) in the 34 GEPs (columns) identified by GBCD. Cells are arranged top-to-bottom by study and patient. Within each study, patients are ordered by the proportion of cells expressing GEP1, which strongly correlates with the classical subtype of PDAC (Panel D). Measures of patient-specific expression (relative to shared expression across patients), inter-patient variation in expression (relative to intra-patient variation) and chromosomal structure of gene signature are shown. See Sections S.1.8 and S.1.9 for definitions. Panel D summarizes concordance between the GBCD-identified shared GEPs (GEPs 1–14) and literature-derived expression signatures for classical and basal PDAC subtypes, as quantified by $-\log_{10}(p\text{-value})$, where *p*-values are produced by one-sided Wilcoxon rank sum test (Online Methods). Literature-derived expression signatures include: Chan-Seng-Yue, Kim et al. [45] (CK), Hwang et al. [50], Moffitt et al. [49], Raghavan et al. [47]. Panel E: LFC estimates for selected genes in the shared GEPs.

components of transcriptional variation in PDAC tumors. To interpret the shared GEPs, we identified driving genes (most strongly up-regulated genes) and biological processes enriched for the driving genes of each GEP (Table 3, Supplementary Fig. S19, Supplementary Tables S3–S4, Section S.1.6), then we compared the driving genes with expression signatures from relevant PDAC literature [45, 47, 49, 50]. These analyses did not produce plausible interpretations for GEPs 15–20. GEP15 mainly captures dataset effects and distinguishes cohorts A and B from cohort C based on tumor stage (metastatic vs. primary) and sequencing platform (Seq-Well vs. 10x). GEPs 16–20 are each active in at most a few individuals and demonstrate much less within-patient than between-patient variation, as confirmed by ANOVA (Fig. 3B, Section S.1.8). We thus speculate that GEPs 16–20 capture patterns of copy number variation common to a few patients, which is supported by the fact that GEPs 16–20 display markedly stronger chromosomal structure than GEPs 1–14 (Fig. 3B, Supplementary Fig. S13).

A subset of shared GEPs (GEPs 1–8) align well with literature-derived classical or basal-like PDAC subtype signatures, confirming that these subtypes can be identified *de novo* from scRNA-seq data (Fig. 3D). The classical-subtype-associated programs (GEPs 1–3) are enriched for biological processes including O-linked glycosylation of mucins (GEP1), cholesterol homeostasis (GEP2), and MYC regulation (GEP3). GEPs 4–8 are associated with basal-related signatures, and are enriched for many different biological processes, some of them related. GEPs 4 and 5 are both enriched for genes involved in defining the EMT; other enriched processes include regulation of insulin-like growth factor transport (GEPs 4, 7, 8), extracellular matrix organization (GEP5), p53 pathway, cornification and keratinization (GEP6), neutrophil degranulation (GEPs 7, 8) and MHC-related processes (GEP8). Many of these enriched biological processes have also been found to be enriched in literature-derived subtype signatures [45, 47, 50, 52]. Notably, we observe at least 2-fold up-regulation of several previously identified marker genes of the classical subtype in at least one of the classical-related GEPs 1–3, including *AGR2*, *ANXA10*, *CLDN18*, *DDC*, *GATA6*, *LGALS4*, *REG4*, *TFF2*, and of several basal subtype marker genes in at least one of the basal-related GEPs 4–6, including *C16orf74*, *GPR87*, *KRT5*, *KRT6A*, *ITGA3*, *PTGES*, *S100A2* [49, 53, 54]. The classical-associated GEPs 1–2 and the basal-associated GEPs 4–6 are active in two largely exclusive subsets of cells, which sometimes co-occur in the same tumor (e.g., B17), consistent with previous findings that basal and classical programs can co-exist intratumorally [45, 55].

Despite the strong connection between GEPs 1–8 and literature-derived subtype signatures, these results do not simply recapitulate previous findings. Indeed, some previously identified signatures (e.g., the squamoid signatures from [50]) are not strongly correlated with any of our GEPs, and other previously identified signatures reflect some combination of our eight GEPs (Fig. 3D). Thus our GEPs 1–8 represent a new decomposition of classical- and basal-related transcriptional heterogeneity into eight components, with the potential to provide new insights into the disease (see the survival analysis below).

In contrast, GEPs 9–14 do not correlate with previously defined subtype signatures. These GEPs are enriched for interferon signaling (GEP9), cell cycle (GEP10), TNF-NFκB signaling (GEP11), respiratory electron transport (GEP12), translation (GEP13) and stress-induced pathways (GEP14).

A stress signaling program is strongly prognostic of poor survival. To assess the prognostic value of GEPs 1–14 in the context of previously identified subtype signatures, we scored each gene expression signature in bulk RNA-seq data of 391 resected primary PDAC tumors from 4 studies [52, 56–58] involving 260 deaths, and performed survival analysis (Section S.1.10). Specifically, we performed stepwise variable selection among all gene signatures, including the previously identified subtype signatures, in a Cox proportional hazards regression model with overall survival as the endpoint, adjusting for age, sex, and tumor stage [59]. The final

TABLE 3

Top driving genes and enriched gene sets of shared GEPs identified in PDAC data. The left-hand column lists the genes with the largest increases in expression (LFC) in the GEP. The right-hand column lists the enriched MSigDB Hallmark and CP:Reactome gene sets [39, 43, 44] with a Bonferroni-adjusted *p*-value < 0.05.

GEP	Top 10 driving genes	Top significantly enriched gene sets
1	*TFF1, CLDN18, PHGR1, REG4, TFF3, DPCR1, TFF2, CEACAM5, SPINK1, MUC5AC*	Cell-cell communication, O-linked glycosylation of mucins, Cell junction organization, O-linked glycosylation, Genes defining late response to estrogen
2	*REG4, CES2, AKR1B10, CA2, PHGR1, OLFM4, ALDH1A1, TFF3, CLDN18, CA9*	Metallothioneins bind metals, Genes involved in cholesterol homeostasis
3	*CA2, REG4, AHCY, AKR1B10, KCNE3, TMEM97, GPATCH4, IMPA2, CLDN18, FAM195A*	Mitochondrial translation, A subgroup of genes regulated by MYC (version 2), A subgroup of genes regulated by MYC (version 1), Translation, rRNA processing
4	*MT2A, S100A2, CAV1, MIR205HG, VIM, MT1X, CST6, LGALS1, TGFBI, MT1E*	Genes defining epithelial-mesenchymal transition, Regulation of insulin-like growth factor transport and uptake by insulin-like growth factor binding proteins, Metallothioneins bind metals, response to metal ions
5	*CAV1, DCBLD2, MT2A, S100A2, KRT6A, VIM, PTRF, COL17A1, MT1E, HMGA2*	Type I hemidesmosome assembly, Extracellular matrix organization, Genes defining epithelial-mesenchymal transition, Laminin interactions, Cell junction organization
6	*KRT13, LY6D, KRT5, AQP3, S100A2, LYPD3, SCPEP1, S100A9, CCL20, MTRNR2L2*	Genes involved in p53 pathways and networks, Genes regulated by NF-kB in response to TNF, Formation of the cornified envelope, Keratinization, Genes defining late response to estrogen
7	*TMSB4X, B2M, C15orf48, TACSTD2, FTH1, LCN2, TMSB10, SAT1, HLA-A, S100A6*	Neutrophil degranulation, Innate immune system, Hemostasis, Genes encoding components of the complement system, Regulation of insulin-like growth factor transport and uptake by insulin-like growth factor binding proteins
8	*HLA-DRB1, HLA-DRA, SERPINA1, CLU, CD74, SPINK1, MMP7, PDIA3, HSP90B1, ENO1*	Neutrophil degranulation, Innate immune system, Antigen presentation: folding, assembly and peptide loading of class I MHC, MHC class II antigen presentation, Regulation of insulin-like growth factor transport and uptake by insulin-like growth factor binding proteins
9	*TNFSF10, IFI6, ISG15, IFI16, MX1, UBE2L6, IFIT3, IFI44L, PARP14, CD74*	Genes up-regulated in response to alpha interferon proteins, Genes up-regulated in response to IFNG, Interferon signaling, Interferon alpha/beta signaling, Cytokine signaling in immune system
10	CENPF, TOP2A, PTTG1, MKI67, HMGB2, UBE2C, STMN1, BIRC5, CCNB1, NUSAP1	Genes encoding cell cycle related targets of E2F transcription factors, Cell cycle, Cell cycle (mitotic), Mitotic prometaphase, Genes involved in the G2/M checkpoint
11	*ISG20, PLAUR, CLDN18, CXCL5, CLDN23, KLK10, CD55, CDKN1A, SNX9, GULP1*	Genes regulated by NF-kB in response to TNF, Cytokine signaling in immune system, Interferon alpha/beta signaling, Tight junction interactions, Genes mediating apoptosis by activation of caspases
12	*TMSB10, NCL, LYZ, EIF5A,TFF2, HSP90AB1, NME1, HMGA1, HSPD1, SNRPF*	A subgroup of genes regulated by MYC (version 1), Respiratory electron transport, ATP synthesis by chemiosmotic coupling, Metabolism of RNA, Cellular response to chemical stress
13	*GNB2L1, EEF1A1, SLC25A6, ADIRF, GLTSCR2, UQCRH, HNRNPA1, SYT8, PABPC1, PTMA*	Eukaryotic translation initiation, Activation of mRNA upon binding of cap-binding complex and eIFs, Eukaryotic translation elongation, Translation
14	*ZFAS1, FTH1, DDIT3, EPB41L4A-AS1, NDRG1, EIF4A2, TAF1D, SLC3A2, GULP1, FTL*	Genes up-regulated in response to low oxygen levels (hypoxia), Genes up-regulated during unfolded protein response, A cellular stress response related to the endoplasmic reticulum, Genes up-regulated through activation of mTORC1 complex, Response of EIF2AK1 (HRI) to heme deficiency, Unfolded protein response (UPR)

model selects two of our GEPs, GEP4 and GEP14, as prognostic (at p-value < 0.01) but selects none of the previously identified subtype signatures (Fig. 4A). GEP4 is prognostic of poor survival (HR = 1.58, p-value = 2×10^{-10}) and is the only selected program related to subtype (basal), indicating that it is more strongly prognostic of survival than previously

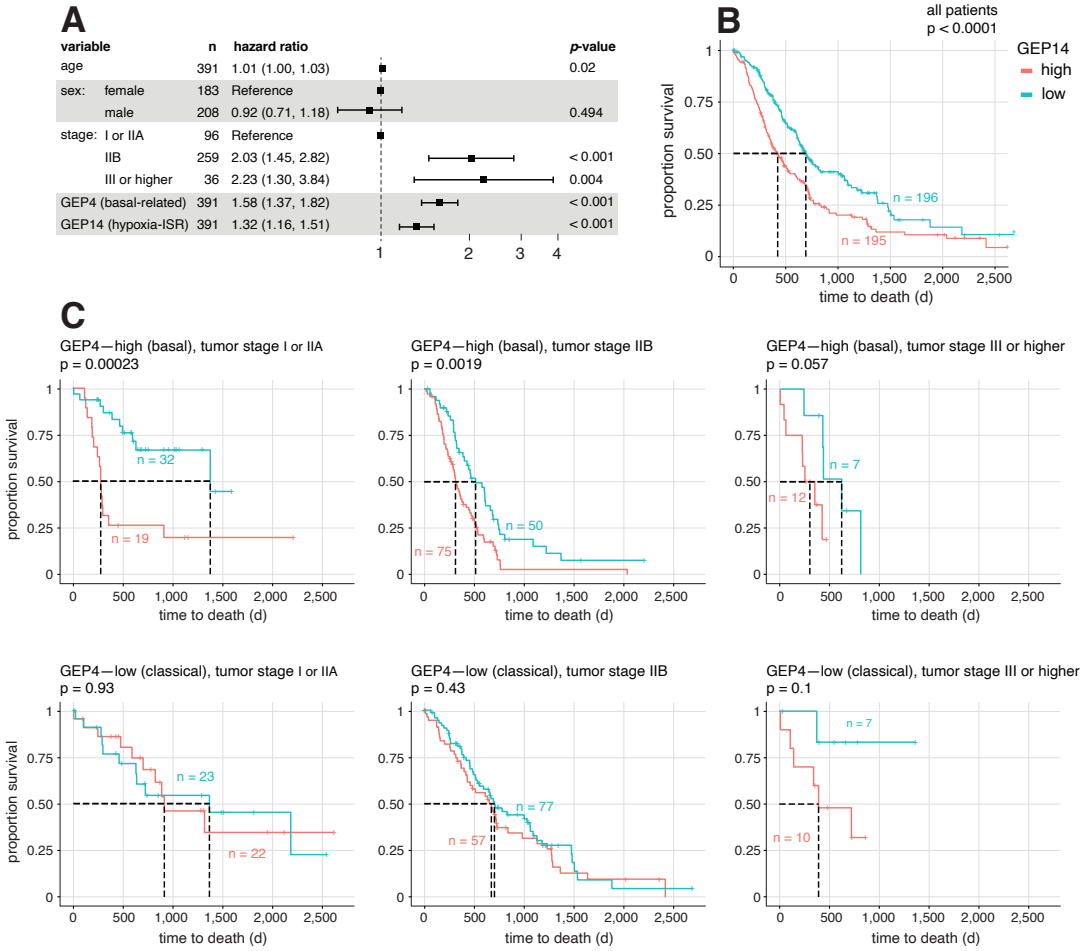

Fig. 4. Assessment of prognostic relevance of expression program signatures in bulk RNA-seq data from 391 primary PDAC tumors. Panel A shows the final model selected by stepwise variable selection in Cox regression analysis, with overall survival as endpoint and all gene expression signature scores as covariates to be selected. Age, sex and tumor stage are always included as covariates in the model. For each covariate included in the model, hazard ratio and its 95% confidence interval are presented. Panel B and C show the Kaplan-Meier plots comparing the overall survival between patients with relatively high or low levels of GEP14, together with two-sided log-rank test *p*-value. This comparison is done for all 391 patients in Panel B, and separately for subgroups of patients stratified by tumor subtype (dichotomized GEP4) and tumor stage in Panel C.

defined subtype signatures. GEP14 is also prognostic of poor survival (HR = 1.32, *p*-value = 5×10^{-5}), and represents a novel expression signature that is predictive of prognosis even when accounting for tumor subtype and tumor stage. Further analyses, stratified by tumor stage and classical/basal subtype, show that GEP14 is particularly prognostic of patient survival in earlier stage basal-like tumors (Fig. 4C).

Inspection of the top driving genes in GEP14 reveals interesting connections with previous literature on PDAC. The top driving gene is *ZFAS1*, a long non-coding RNA gene that has been shown to promote PDAC metastasis [60]. Other top driving genes include two genes related to iron metabolism, *FTH1* and *FTL*, both of which have been recently linked with PDAC progression arising from increased ferritinophagy [61, 62].

Enrichment analysis of a longer list of driving genes (Supplementary Fig. S19) highlights several gene sets related to stress response. These include genes up-regulated in response to

hypoxia (e.g., *BNIP3L, DDIT3, VEGFA*), in response to heme deficiency, and in response to misfolded protein accumulation in the endoplasmic reticulum. Notable driving genes related to stress response include activating transcription factor 4 *ATF4* — the main effector of the Integrated Stress Response [63] — and several of its validated targets, including *DDIT3/CHOP, HSPA5, MAP1LC3B, MTHFD2, SQSTM1* [64–66]. Among these, *SQSTM1* and *MAP1LC3B* are key modulators of autophagy [65, 67], which is known to promote PDAC progression and therapy resistance [68, 69], and *HSPA5* has been previously linked with PDAC progression through decreased ferroptosis [70, 71].

Computing requirements. GBCD builds on the general EBMF framework, and the fast model fitting algorithms implemented for EBMF makes it possible to apply GBCD to moderately sized data sets (e.g., the HNSCC data). However, it can be costly to explicitly compute the matrix $\boldsymbol{Y}\boldsymbol{Y}^T$ and then apply EBMF to that $N \times N$ (non-sparse) matrix for large data sets (e.g., the PDAC data). Therefore, we implemented the ability to apply EBMF to a matrix $\boldsymbol{Y}\boldsymbol{Y}^T$ *by performing computations using only the (sparse) matrix* $\boldsymbol{Y}$ (i.e., without ever explicitly forming the matrix $\boldsymbol{Y}\boldsymbol{Y}^T$), which is mathematically equivalent to the former but reduces computing costs when N is large and $\boldsymbol{Y}$ is sparse. This approach reduces the per-iteration computational complexity from $O(N^2K)$ to $O(SK)$, and reduces the memory usage from $O(N^2)$ to $O(NK + S)$, where S is the number of non-zero entries in $\boldsymbol{Y}$. For example, for the HNSCC data (2,176 cells, 17,113 genes), this approach reduced the running time and memory from 2 hours and 16 GB to 50 minutes and 2 GB; for the PDAC data (35,670 cells, 13,844 genes), this approach reduced computing time and memory from 50 hours and 200 GB to 21 hours and 48 GB. In this paper, we gave results for the approach that explicitly computes $\boldsymbol{Y}\boldsymbol{Y}^T$, but we note that results were highly consistent between the two approaches (Supplementary Fig. S17). Computing costs can also be reduced by optimizing the number of backfit iterations for estimating $\boldsymbol{L}$ (see Section S.1.11). We note that further improvements to the GBCD pipeline will likely be needed to analyze very large scRNA-seq datasets.

Discussion

In this paper, we introduced a new matrix factorization method, GBCD, and demonstrated its ability to provide new insights into transcriptional heterogeneity in scRNA-seq tumor data from multiple patients and studies. A key feature of GBCD is that it assumes expression programs to be orthogonal to one another, which helps avoid absorbing shared components of expression variability into patient-specific expression programs. When cells show strong inter-patient heterogeneity in their transcriptional profiles, as is typical of malignant cells, we demonstrated in both simulated and real data that GBCD can capture shared patterns of variation that existing MF methods miss. In contrast to NMF, GBCD does not constrain expression programs to be nonnegative, which allows it to model transcriptional repression of genes in a given program, as well as up-regulation.

In applications to scRNA-seq tumor data from PDAC and HNSCC patients, GBCD identifies GEPs that are shared across tumors and datasets, as well as GEPs that characterize patient and dataset effects. Some of the shared GEPs correspond to previously defined tumor subtypes or generic cellular processes such as cell cycle and immune activation, but others represent new findings. Most notably, we identified a new expression program (GEP14) that is an independent predictor of poor prognosis in patients with primary PDAC. GEP14 contains many genes related to stress response, including hypoxia. Hypoxia gene expression signatures have been previously related to worse PDAC outcomes and hypoxia is a promising target of innovative therapeutic strategies [72–76]. However, GEP14 is also associated with many other stress-response processes, including many genes that are induced by *ATF4* as part

of the Integrated Stress Response (ISR). The ISR is activated in response to stress, including hypoxia, by phosphorylation of the α subunit of eukaryotic initiation factor 2 (*eIF2*α) at serine 51, which suppresses global protein synthesis to reduce energy expenditure while enhancing translation of specific mRNAs, including *ATF4*. This then up-regulates the expression of many genes involved in stress adaptation [66]. Recently, *ATF4* has been shown to promote amino acid biosynthesis and tumor cell survival in PDAC [77] while *eIF2*α phosphorylation significantly correlates with tumor recurrence/metastasis and shorter overall survival [78]. Interestingly, the *ATF4* expression signature has been found to be a predictor of therapy resistance in other cancers. For example, it was revealed to be induced in melanoma cells that escaped RAF kinase inhibitor treatment in a recent scRNA-seq study [79]; tumor cells that persisted after treatment with BH3 mimetics similarly depended on *ATF4* and the ISR to survive [80]. The fact that GEP14 includes many genes induced during the ISR, and is an independent predictor of PDAC patient survival, suggests an important role for the ISR in PDAC progression, and raises the potential that pharmacological interventions targeting the ISR could provide useful therapeutic strategies. More work is needed to test these hypotheses.

Like other matrix factorization and clustering methods, GBCD is perhaps best viewed as a useful exploratory tool for identifying interesting structures in large datasets. And, like existing tools, GBCD has limitations that users should be aware of. For example, GBCD is solving a non-convex optimization problem, and so its solutions will typically depend on initialization, and it may be very difficult to find the global optimal solution. Indeed, although in our examples GBCD finds many GEPs corresponding to individual patient effects, there are some patients for which it finds no patient-specific program, and we believe this is likely a failure of the optimization routine to identify such programs. Applications of NMF sometimes try to ameliorate this issue by combining results across multiple runs with different initializations [5], and it is possible that similar strategies could be helpful for GBCD.

Also like other matrix factorization methods, GBCD requires user judgement in choosing K, the number of factors. The EBMF method that underlies GBCD requires users to set an upper bound, $K_{\max}$, on the number of factors, and then provides a method to automatically select K (up to $K_{\max}$). Thus, one possible strategy is to simply set $K_{\max}$ to be very large, and rely on this automatic method to select K. However, for several reasons — not least that computation increases with K — in practice users might want to more strictly limit K. In experimenting with various values of K in our analyses, we found that GEPs identified with a smaller K are generally also present when a larger K is used (Supplementary Fig. S16). Generally speaking then, a larger K allows for identifying finer structure in tumor transcriptional heterogeneity at the expense of higher computational cost, and the main downside to limiting K is that one may miss interesting structure in the data. In practice the value of K needed to capture most of the interesting structure will inevitably depend on the complexity of the data, and this will likely increase with the number of cohorts, patients and tumor cells analyzed.

Finally, we note that we do not expect GBCD to be uniformly superior to other methods, such as NMF, in all applications. The assumptions of GBCD, including orthogonality of GEPs, may be more appropriate in some applications than others. Since it may be difficult to know in advance whether one set of assumptions or another are more appropriate to a given dataset, it is common to apply multiple methods to see how the results differ. Our results here show that GBCD can produce results that are quite different from existing methods, and in so doing has the potential to provide new scientific insights.

URLs. The code and data resources for simulations and analysis of HNSCC and PDAC data are stored in the Zenodo repository `https://doi.org/10.5281/zenodo.8271036`. Other R packages used in this work include: gbcd (`https://github.com/`

stephenslab/gbcd), ebnm (https://github.com/stephenslab/ebnm), ashr (https://github.com/stephens999/ashr), flashier (https://github.com/willwerscheid/flashier), Seurat (https://satijalab.org/seurat/), fastTopics (https://github.com/stephenslab/fastTopics), NNLM (https://github.com/linxihui/NNLM), rliger (https://github.com/welch-lab/liger), cNMF (https://github.com/dylkot/cNMF), batchelor (https://bioconductor.org/packages/devel/bioc/html/batchelor.html), conos (https://github.com/kharchenkolab/conos), splatter (https://github.com/Oshlack/splatter), scran (https://doi.org/doi:10.18129/B9.bioc.scran) seqgendiff (https://github.com/dcgerard/seqgendiff), and infercnv (https://github.com/broadinstitute/infercnv).

Acknowledgements. This work was supported by NIH grants R01HG002585, DP2AI145100, U01AI160418, and a Gut Cell Atlas grant from The Leona M. and Harry B. Helmsley Charitable Trust. We thank the staff at the Research Computing Center at the University of Chicago for providing the high-performance computing resources used to implement the numerical experiments. This study was conducted using data provided by the Ontario Institute for Cancer Research which is supported by funding provided by the Government of Ontario. The views expressed in the publication are those of the author and do not necessarily reflect those of the Ontario Institute for Cancer Research or the Government of Ontario.

Author contributions. Y.L. and M.S. conceived the project. Y.L. and J.W. developed the method and implemented it in R under supervision of M.S. Y.L. performed the simulation study and analyses of real data. Y.L. and P.C. prepared the online code. Y.L., S.A.O. and K.F.M. interpreted the PDAC data analysis results with contributions from P.C., J.W. and M.S. All authors wrote the manuscript.

Competing interests. The authors declare no competing financial interests.

S.1. Online methods

S.1.1. GBCD analyses. In this section, we detail the steps that were taken to fit a GBCD model to a dataset and then obtain statistical quantities used in interpretation of the GBCD results.

S.1.1.1. *Preparation of scRNA-seq data.* In all our GBCD analyses, the initial data were combined into an $N \times J$ counts matrix $\boldsymbol{X}$ across tumors, cohorts, etc. In most cases, $\boldsymbol{X}$ was a matrix of unique molecular identifier (UMI) counts in which element x_{ij} was the UMI count of gene j in cell i. Here we describe the steps that were taken to convert the initial data $\boldsymbol{X}$ to the matrix $\boldsymbol{Y}$ that was then analyzed by GBCD.

Let s_i denote the "size factor" for cell i. For UMI count data, we took this size factor to be the total count over all genes, $s_i = \sum_{j=1}^{J} x_{ij}$. (We note that other size factor definitions [81–83] are possible, and in some settings may be preferred.) We scaled the data by the size factors, added a "pseudo-count" $c > 0$, then took the log:

$$y_{ij} = \log(c + \tilde{s}/s_i \times x_{ij}) - \log c, \tag{3}$$

in which $\tilde{s} := \text{median}\{s_1, \ldots, s_N\}$ denotes the median size factor. We subtracted $\log c$ from the log-transformed data, which ensured that all y_{ij}'s were non-negative. This also preserved sparsity in the original counts x_{ij}; that is, y_{ij} was zero if and only if x_{ij} was zero.

Regarding the choice of pseudo-count c, different pseudo-counts have been proposed (e.g., [84, 85]). Generally speaking, pseudo-counts that are too small overemphasize differences in normalized expression between zero and nonzero UMI counts, whereas pseudo-counts that are too large reduce differences in normalized expression [85]. We chose a reasonable "middle ground" value of $c = 0.1$; see [36, 86] for further discussion.

Unlike the other datasets, the HNSCC data were not UMI counts; rather, they were read counts produced by SMART-Seq2 [87]. Following [10], we defined the transformed counts as $y_{ij} = \log_2(1 + \text{TPM}_{ij}/10)$, where TPM_{ij} was the transcript-per-million (TPM) value [88, 89] for gene j in cell i. Note that this transformation is equivalent to (3) (up to a constant of proportionality) with $c = \tilde{s}/10^5$.

S.1.1.2. *Estimating GEP memberships.* In GBCD, we estimate the GEP membership matrix $\boldsymbol{L}$ in the matrix factorization (2) by decomposing $\boldsymbol{Y}\boldsymbol{Y}^T$ using the empirical Bayes matrix factorization (EBMF) framework [35]. In detail, we model $\boldsymbol{Y}\boldsymbol{Y}^T$ as

$$\begin{aligned} \boldsymbol{Y}\boldsymbol{Y}^T &= \boldsymbol{L}\boldsymbol{L}^T + \epsilon \boldsymbol{I}_N + \boldsymbol{E}, \\ e_{ij} &\overset{i.i.d.}{\sim} N(0, \varphi^2), \end{aligned} \tag{4}$$

where $\boldsymbol{I}_N$ is the $N \times N$ identity matrix, $\boldsymbol{E}$ is an $N \times N$ matrix with entries e_{ij}, and $\epsilon, \varphi^2 > 0$ are additional unknowns to be estimated.

If the rows of $\boldsymbol{Y}$ are centered, $\boldsymbol{Y}\boldsymbol{Y}^T$ is proportional to the sample covariance matrix that characterizes the covariation between the transcriptional profiles of every pair of cells. Thus, we refer to (4) as a "covariance decomposition" since it approximates $\boldsymbol{Y}\boldsymbol{Y}^T$ by $\boldsymbol{L}\boldsymbol{L}^T$, where $\boldsymbol{L}$ is an $N \times K$ nonnegative matrix.

GBCD assigns a "generalized binary" prior independently to each entry of $\boldsymbol{L}$,

$$l_{ik} \sim (1 - \pi_k^l)\delta_0 + \pi_k^l \, N_+(\mu_k, \sigma_k^2), \tag{5}$$

in which $N_+(\mu, \sigma)$ is a left-truncated normal distribution with mean μ and variance σ^2, with the truncation point at zero. The functional form of this prior is chosen to balance modelling flexibilty (ability to capture both discrete and continuous structures) and computational convenience. We call (5) a "generalized binary" (GB) prior because as $\sigma_k^2 \to 0$, the prior's support takes only two possible values, 0 or μ_k. A strictly binary prior — that is, (5) with $\sigma_k = 0$

— is suitable for modeling discrete structures in cell populations, such as distinct molecular subtypes and patient-specific effects, but cannot capture transcriptional variation of a continuous nature, such as biological processes with varying degrees of activity in cells. To get around this limitation, we have proposed the GB prior in which the ratio σ_k/μ_k is kept small by taking some pre-specified small value (see below), which encourages discrete structures while still allowing for continuous structures when they fit the data better. Rather than set the parameters π_k^l and μ_k by hand, we take an empirical Bayes approach and leverage the information across cells to adapt the priors to the data. The exact steps we take to estimate these parameters are described below.

Note that the diagonal matrix $\boldsymbol{D}$ from (2) does not appear in (4). This is because we assume that entries of $\boldsymbol{L}$ in (2) are *a priori* close to 0 or 1 (rather than 0 or μ_k) for ease of illustration, and so need to incorporate the scaling factors μ_k's into $\boldsymbol{D} = \operatorname{diag}\left(\mu_1^2, \ldots, \mu_K^2\right)$.

Our current EBMF implementation cannot fit $\boldsymbol{Y} \approx \boldsymbol{L}\boldsymbol{F}^T$ subject to $\boldsymbol{L} = \boldsymbol{F}$. Therefore, instead of fitting (4), in practice we fit a slightly relaxed matrix decomposition,

$$\boldsymbol{Y}\boldsymbol{Y}^T = \boldsymbol{L}\tilde{\boldsymbol{L}}^T + \epsilon \boldsymbol{I}_N + \boldsymbol{E}, \tag{6}$$

with independent GB priors (5) on $\boldsymbol{l}_k$ and $\tilde{\boldsymbol{l}}_k$. Clearly, (6) is the same as (4) when $\boldsymbol{L} = \tilde{\boldsymbol{L}}$. While we do not enforce $\boldsymbol{L}$ and $\tilde{\boldsymbol{L}}$ to be the same, in practice, we observe that fitting (6) always produces highly concordant $\boldsymbol{l}_k$ and $\tilde{\boldsymbol{l}}_k$ for the majority of components; in our implementation we only keep components k such that the Pearson correlation between $\boldsymbol{l}_k$ and $\tilde{\boldsymbol{l}}_k$ is greater than 0.8, and use $\boldsymbol{l}_k$'s as GEP membership estimates.

We now describe how to fit the EBMF model (6) with a GB prior (5). Thanks to the modularity of the model fitting algorithm, solving the EBMF problem for any prior family requires only solving a series of much simpler "empirical Bayes normal means" (EBNM) problems with the same prior family (see [35] for the connection between EBMF and EBNM). For an EBNM problem with a GB prior, we have observations $\boldsymbol{x} = (x_1, \ldots, x_n)^T$ that follow a multivariate Gaussian distribution with unknown means $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_n)^T$ and known covariance matrix $\operatorname{diag}\left(s_1^2, \ldots, s_n^2\right)$, where $\theta_1, \ldots, \theta_n$ are i.i.d. from a GB distribution. That is,

$$\boldsymbol{x} \mid \boldsymbol{\theta} \sim \mathrm{N}_n(\boldsymbol{\theta}, \operatorname{diag}(s_1^2, \ldots, s_n^2)), \tag{7}$$

$$\theta_1, \ldots, \theta_n \overset{i.i.d.}{\sim} (1-\pi)\delta_0 + \pi\, \mathrm{N}_+(\mu, \sigma^2). \tag{8}$$

Solving this EBNM problem involves the following two steps.

1. Estimate μ, σ^2, π by maximizing the marginal likelihood

$$\begin{aligned} L(\mu, \sigma^2, \pi; \boldsymbol{x}, \boldsymbol{s}) &:= \prod_{i=1}^n \left[(1-\pi)\, \mathrm{N}(x_i; 0, s_i^2) + \pi \int \mathrm{N}(x_i; \theta_i, s_i^2)\, \mathrm{N}_+(\theta_i; \mu, \sigma^2) d\theta_i \right] \\ &= \prod_{i=1}^n \left[(1-\pi)\, \mathrm{N}(x_i; 0, s_i^2) + \pi\, \mathrm{N}(x_i; \mu, \sigma^2 + s_i^2)\, \frac{\Phi(\tilde{\mu}_i/\tilde{\sigma}_i)}{\Phi(\mu/\sigma)} \right], \end{aligned} \tag{9}$$

where $\mathrm{N}(\cdot; \mu, \sigma)$ and $\mathrm{N}_+(\cdot; \mu, \sigma)$ respectively denote the probability density function of $\mathrm{N}(\mu, \sigma^2)$ and $\mathrm{N}_+(\mu, \sigma)$. $\Phi(\cdot)$ denotes the cumulative distribution function of a standard normal random variable; $\tilde{\mu}_i = \frac{s_i^2}{s_i^2+\sigma^2}\mu + \frac{\sigma^2}{s_i^2+\sigma^2}x_i$; and $\tilde{\sigma}_i^2 = (\frac{1}{\sigma^2} + \frac{1}{s_i^2})^{-1}$. In practice, we assume that $\sigma/\mu = \omega$ takes some pre-specified small value, e.g., $\omega = 0.02$, so we just need to maximize the marginal likelihood in (9) over μ and π. To facilitate solving this optimization problem, for each observation i, we define a latent indicator $z_i \in \{0, 1\}$ that is 1 if θ_i is drawn from $\mathrm{N}_+(\mu, \sigma^2)$ and 0 otherwise. With this data augmentation, we

maximize the complete data log-likelihood

(10)

$$\log L(\mu, \omega^2\mu^2, \pi; \boldsymbol{x}, \boldsymbol{s}, \boldsymbol{z})$$
$$= \sum_{i=1}^{n} \left[(1-z_i) \log\left((1-\pi)\,\mathrm{N}(x_i; 0, s_i^2)\right) + z_i \log\left(\pi\,\mathrm{N}(x_i; \mu, \omega^2\mu^2 + s_i^2)\,\frac{\Phi(\tilde{\mu}_i/\tilde{\sigma}_i)}{\Phi(1/\omega)}\right)\right],$$

using EM [90], which alternates between the following two steps at each iteration $t = 1, 2, \ldots$ until convergence.

- Given the current estimates of the parameters $\pi^{(t)}, \mu^{(t)}$, use Bayes theorem to calculate the conditional expectation of z_i, i.e., $\zeta_i^{(t)} := \mathbb{E}\left[z_i \mid \pi^{(t)}, \mu^{(t)}, \boldsymbol{x}, \boldsymbol{s}\right]$.
- Given the current $\zeta_i^{(t)}$, update the estimates of π, μ respectively by:

$$\pi^{(t+1)} = \underset{\pi \in [0,1]}{\operatorname{argmax}} \left\{ \sum_{i=1}^{n} \left[(1-\zeta_i^{(t)}) \log(1-\pi) + \zeta_i^{(t)} \log(\pi) \right] \right\} = \sum_{i=1}^{n} \zeta_i^{(t)}/n, \tag{11}$$

$$\mu^{(t+1)} = \underset{\mu \geq 0}{\operatorname{argmax}} \left\{ \sum_{i=1}^{n} \zeta_i^{(t)} \left[\log\left(\mathrm{N}(x_i; \mu, \omega^2\mu^2 + s_i^2)\right) + \log\left(\Phi(\tilde{\mu}_i/\tilde{\sigma}_i)\right) \right] \right\}. \tag{12}$$

The constrained one-dimensional optimization in (12) does not have a closed-form solution, so we solve it using the R function `optim` with method "L-BFGS-B". We denote the converged estimates of μ, π, ζ_i by $\hat{\mu}, \hat{\pi}, \hat{\zeta}_i$.

2. Compute the posterior distribution of $\boldsymbol{\theta}$ given $\hat{\mu}, \hat{\pi}, \hat{\zeta}_i$ by

$$p\left(\theta_i \mid \boldsymbol{x}, \boldsymbol{s}, \hat{\mu}, \hat{\pi}, \hat{\zeta}_i\right) = (1-\hat{\zeta}_i)\delta_0 + \hat{\zeta}_i\,\mathrm{N}_+(\hat{\tilde{\mu}}_i,\, \hat{\tilde{\sigma}}_i^2), \tag{13}$$

where $\hat{\tilde{\mu}}_i = \frac{s_i^2}{s_i^2+\omega^2\hat{\mu}^2}\hat{\mu} + \frac{\omega^2\hat{\mu}^2}{s_i^2+\omega^2\hat{\mu}^2}x_i$ and $\hat{\tilde{\sigma}}_i^2 = (\frac{1}{\omega^2\hat{\mu}^2} + \frac{1}{s_i^2})^{-1}$.

Finally, we fit the EBMF model (6) with GB prior (5) by passing the above EBNM solver as argument `ebnm.fn` to function `flash` in R package `flashier` [91]. More specifically, fitting (6) consists of the following 4 steps (i) – (iv). $\boldsymbol{L}$ and $\tilde{\boldsymbol{L}}$ are initialized by steps (i) – (iii), which are refined by a backfitting algorithm performed in step (iv).

(i). Add pairs of $\boldsymbol{l}_k/\tilde{\boldsymbol{l}}_k$ in a greedy manner, using the function `flash_greedy`, where independent point Laplace priors are placed on l_{ik} and $\tilde{l}_{ik}$.

(ii). Backfit existing pairs of $\boldsymbol{l}_k/\tilde{\boldsymbol{l}}_k$ (i.e., by updating each pair while keeping others fixed in a cyclical fashion), using the function `flash_backfit`, until convergence or completion of T_1 iterations.

(iii). Split each $\boldsymbol{l}_k$ into two nonnegative factors $\boldsymbol{l}_k^+$ and $\boldsymbol{l}_k^-$, where $l_{ik}^+ = \max\{l_{ik}, 0\}$ and $l_{ik}^- = \max\{-l_{ik}, 0\}$, so we have $\boldsymbol{l}_k = \boldsymbol{l}_k^+ - \boldsymbol{l}_k^-$. Do the same for $\tilde{\boldsymbol{l}}_k$ to obtain $\tilde{\boldsymbol{l}}_k^+$ and $\tilde{\boldsymbol{l}}_k^-$.

(iv). Backfit existing pairs of $\boldsymbol{l}_k^+/\tilde{\boldsymbol{l}}_k^+$ and $\boldsymbol{l}_k^-/\tilde{\boldsymbol{l}}_k^-$ using the function `flash_backfit`, where independent GB priors are placed on $l_{ik}^+, \tilde{l}_{ik}^+, l_{ik}^-, \tilde{l}_{ik}^-$, until convergence or completion of T_2 iterations.

S.1.1.3. *Estimating GEP signatures.* We estimate the GEP signature matrix $\boldsymbol{F}$ by fitting an EBMF model to $\boldsymbol{Y}$ with fixed $\boldsymbol{L}$, which is obtained via the covariance decomposition step in

S.1.1.2,

$$\boldsymbol{Y} = \boldsymbol{L}\boldsymbol{F}^T + \tilde{\boldsymbol{E}}, \tag{14}$$

$$\tilde{e}_{ij} \sim \mathrm{N}(0, \tilde{\varphi}_j^2), \tag{15}$$

where $\tilde{\boldsymbol{E}}$ is an $N \times J$ matrix of independent error terms $\tilde{e}_{ij}$, and $\tilde{\varphi}_j^2 > 0$ $(j = 1, \ldots, J)$ are additional unknowns to be estimated. We note that model (14) is a semi-nonnegative matrix factorization (SNMF) [92], because it places a nonnegativity constraint on $\boldsymbol{L}$ but not on $\boldsymbol{F}$. We also note that fitting (4) is, in principle, equivalent to fitting (14) with an orthonormal constraint on $\boldsymbol{F}$ (i.e., $\boldsymbol{F}^T\boldsymbol{F} = \boldsymbol{I}_K$), which can be seen by post-multiplying each side of (14) by its transpose.

For a GEP k, $\boldsymbol{f}_k$ characterizes the patterns of co-regulation of all genes. It is reasonable to assume that most genes do not exhibit differential expression in GEP k compared to other GEPs, so we put a point Laplace prior on f_{jk}, independent over j and k, which encourages shrinkage of small coefficients to 0 but avoids over-shrinkage of large coefficients.

$$f_{jk} \sim (1 - \pi_k^f)\delta_0 + \pi_k^f \,\mathrm{Laplace}(\lambda_k). \tag{16}$$

Similar to S.1.1.2, the component-specific hyper-parameters π_k^f, λ_k of the point Laplace prior are estimated from data by pooling information across all genes using an empirical Bayes method. For each k, the genes with largest (and positive) estimates of f_{jk} are interpreted as the genes driving the GEP k. The EBMF model (14) with point Laplace prior (16) is fit using `flashier`, where we call the function `flash` and set argument `ebnm.fn = ebnm::ebnm_point_laplace`. $\boldsymbol{F}^T$ is initialized by $\left(\boldsymbol{L}^T\boldsymbol{L}\right)^{-1}\boldsymbol{L}^T\boldsymbol{Y}$.

S.1.1.4. *Quantifying uncertainty in the GEP signatures.* Since mean field variational approximations, which are used in the EBMF framework, are known to underestimate uncertainty in posterior distributions [35], we performed additional calculations to obtain better uncertainty quantifications of the LFC estimates, which we then used to compute posterior z-scores and other posterior quantities for f_{jk}.

First, for each gene j, we fit a linear regression model in which the log-pc data for gene j, $\boldsymbol{y}_j = (y_{1j}, \ldots, y_{Nj})^T$, was the response variable, and the regressor matrix was the matrix of GEP membership estimates $\boldsymbol{L}$. We fit the linear regression models using the `lm` function in R. After fitting each gene-wise regression model, we extracted the point estimates $\hat{f}_{jk}$ and standard error estimates $\hat{s}_{jk}$ for the corresponding regression coefficients.

Next, we performed adaptive shrinkage, separately on the LFC estimates for each GEP k, to improve the estimation accuracy of the standard errors. To implement this step, we use the `ash` function from the `ashr` R package [93]. This adaptive shrinkage method takes as input a vector of effect estimates $\hat{f}_{1k}, \ldots, \hat{f}_{Jk}$ and associated standard errors $\hat{s}_{1k}, \ldots, \hat{s}_{Jk}$. The revised posterior mean estimates and standard errors returned by the adaptive shrinkage method are then used to calculate the final test statistics, including posterior z-scores (defined as the posterior mean divided by the posterior standard error [94]) and local false sign rates, or *lfsr* [93], which are then used in the volcano plots in the Supplementary Fig. S18 and S19.

Note that this approach to quantifying uncertainty of LFC estimates assumes that the membership estimates are "known", so does not fully account for uncertainty in the matrix factorization, but in practice it is an improvement over the uncertainty estimates provided by `flashier`.

S.1.2. Details for simulation design. We first simulated a UMI count matrix $\boldsymbol{X}^{\text{null}}$ containing scRNA-seq measurements of $J = 10,000$ genes in $N = 3,200$ cells using the Splatter framework [95],

$$x_{ij}^{\text{null}} \sim \text{Pois}(s_i 2^{\theta_{ij}}), \tag{17}$$

$$\boldsymbol{\Theta} = \mathbf{1}_N \boldsymbol{\alpha}^T, \tag{18}$$

where $\boldsymbol{\alpha}$ is a $J \times 1$ vector of gene-specific intercept terms that represent the relative expression levels for each gene. The parameters are generated from the Splat simulation model [95] using input parameters estimated from one PDAC dataset. Then we modified $\boldsymbol{X}^{\text{null}}$ using binomial thinning [96] to produce $\boldsymbol{X}$ such that

$$x_{ij} \sim \text{Pois}(s_i 2^{\tilde{\theta}_{ij}}), \tag{19}$$

$$\tilde{\boldsymbol{\Theta}} = \mathbf{1}_N \tilde{\boldsymbol{\alpha}}^T + \boldsymbol{L}\boldsymbol{F}^T, \tag{20}$$

where $\tilde{\boldsymbol{\alpha}}$ is a $J \times 1$ vector of modified gene-specific intercept terms due to thinning, and $\boldsymbol{L}$ and $\boldsymbol{F}$ are respectively $N \times K$ and $J \times K$ matrices representing cell memberships and gene signatures for $K = 11$ GEPs. Fig. 1A displays $\boldsymbol{L}$ for one simulation dataset.

Each GEP is characterized by up-regulation (relative to the "baseline" level of expression determined by $\tilde{\boldsymbol{\alpha}}$) of 75 to 500 non-overlapping genes, with a fold change having an average of 3 and ranging from 1.5 to 11 in magnitude. We repeated the above procedure to generate a total of 20 replicate datasets.

S.1.3. Implementation details for MF methods. Here we describe in detail how we ran the alternative MF methods that were implemented in the simulations.

- **Patient-by-patient NMF:** We fitted a NMF with 3 programs (an intercept term included) to the UMI counts, separately for each patient, by calling the `nnmf` function from the R package `NNLM` [97] with its default settings. Then we compared all programs by hierarchical clustering, using the fraction of top 50 genes (quantified by the Jaccard index [15]) that are shared between every pair of programs as the similarity measure. Following [10], we manually identified clusters of highly correlated programs, which were defined as "meta-programs". For a given meta-program, the cell-wise memberships of all patients are calculated as follows. If a patient is included in the associated cluster of programs, the cell-wise membership estimates come from the loadings of the corresponding program of that patient; otherwise the cell-wise membership estimates are set to zero.
- **Combined NMF (UMI counts):** We fitted a Poisson NMF to the UMI counts $\boldsymbol{X}$, which is equivalent to fitting a multinomial topic model [7, 98, 99], by calling the `fit_topic_model` function from the R package `fastTopics` [100] with its default settings.
- **Combined NMF (log-pc counts):** We fitted NMF with mean squared error loss function to $\boldsymbol{Y}$ by calling the `nnmf` function from the R package `NNLM` [97] with its default settings.
- **Consensus NMF (UMI counts):** We ran consensus NMF [5] with Frobenius loss to $\boldsymbol{X}$ using the Python package `cNMF` with its default settings.
- **PCA**: We performed PCA on the log-pc counts $\boldsymbol{Y}$ by calling the R function `svd`.
- **EB-SNMF, point-exp prior or GB prior:** we used the R package `flashier` to fit the model (14) to the log-pc counts $\boldsymbol{Y}$, with a point Laplace prior (16) on $\boldsymbol{F}$, and either a GB prior (5) or point-exponential prior on $\boldsymbol{L}$. No orthogonality constraint is placed on $\boldsymbol{F}$, which is the key difference from GBCD. More specifically, fitting EB-SNMF involves the following 4 steps. $\boldsymbol{L}$ and $\boldsymbol{F}$ are initialized by steps (i) – (iii), which are refined by a backfitting algorithm in step (iv).

(i). Add pairs of $\boldsymbol{l}_k/\boldsymbol{f}_k$ in a greedy manner, using the function `flash_greedy`, where independent point Laplace priors are placed on l_{ik} and f_{jk}.

(ii). Backfit existing pairs of $\boldsymbol{l}_k/\boldsymbol{f}_k$ (i.e., by updating each pair while keeping others fixed in a cyclical fashion), using the function `flash_backfit`, until convergence or completion of T_1 iterations.

(iii). Split each $\boldsymbol{l}_k$ into two nonnegative factors $\boldsymbol{l}_k^+$ and $\boldsymbol{l}_k^-$, where $l_{ik}^+ = \max\{l_{ik}, 0\}$ and $l_{ik}^- = \max\{-l_{ik}, 0\}$, so we have $\boldsymbol{l}_k = \boldsymbol{l}_k^+ - \boldsymbol{l}_k^-$.

(iv). Backfit existing pairs of $\boldsymbol{l}_k^+/\boldsymbol{f}_k$ and $\boldsymbol{l}_k^-/\boldsymbol{f}_k$ using the function `flash_backfit`, where independent GB (or point exponential) priors are placed on l_{ik}^+, l_{ik}^- and independent point Laplace priors are placed on f_{jk}, until convergence or completion of T_2 iterations.

- **CD, point-exp prior:** we ran a in which of GBCD where we replaced the GB prior (5) on $\boldsymbol{L}$ with a point exponential prior. The implementation is essentially the same as that of GBCD, as described in detail in Sec. S.1.1.

Choice of K in the simulation study. Determining the number of components K is generally a challenging task for any matrix factorization method. For NMF to UMI counts, we set K to its true value 11. For NMF to log-pc counts, we set $K = 12$ to allow the inclusion of an additional component that models the gene-specific baseline expression levels. For consensus NMF to UMI counts, K is determined from 8–12 based on the trade-off between stability and error. For PCA to log-pc counts, we set $K = 8, 10, 12$ to assess its performance when different numbers of principal components are included, and the PCA result shown in Fig. 1 was produced using $K = 10$. All EB-based methods, including GBCD, can perform automatic selection of K [35]. In particular, users of `flashier` just need to specify an upper bound $K_{\max}$ for K, among which the components whose prior is estimated to be a point mass at 0 are essentially removed from the model. We set $K_{\max} = 16$ to implement all EB-based methods.

S.1.4. Implementation details for harmonization methods. We implemented several widely used batch effect correction methods which were specifically developed to integrate scRNA-seq data from multiple sources. In particular, we used each of these methods to correct for patient effects in the simulated data, and assessed their performance in recovering the subtype-related and continuous GEPs which are shared across multiple patients. Here we describe in detail how these methods were implemented in the simulations.

LIGER. We implemented Liger [25] on the UMI counts using the R package `rliger` (version 1.0.0), following the suggested pipeline to integrate multiple scRNA-seq datasets as specified in the package vignette. We set the inner dimension of factorization `k=20` when calling the `optimizeALS` function. We assessed the impact of the regularization parameter λ, which controls the amount of penalization for patient-specific effects, on its performance (Supplementary Fig. S4). The Liger result shown in Fig. 1 was produced using the default value of $\lambda = 5$.

CCA (Seurat). We implemented Seurat 3 integration [26] on the log-pc counts using the R package `Seurat` (version 4.1.1), following the standard integration workflow described in the package vignette. We first identified the top 5000 highly variable genes separately for scRNA-seq data from each of the 8 patients using the `FindVariableFeatures` function. We then found a set of shared anchors across patients by canonical correlation analysis, using the `FindIntegrationAnchors` function with `anchor.features = 5000` and `dims = 1:50`, and performed data integration across patients based on these pre-computed anchors using the `IntegrateData` function with `dims = 1:50`.

MNN Correct. We implemented MNN Correct [27] on the log-pc counts by calling the `mnnCorrect` function in the R package `batchelor` (version 1.8.1).

Conos. We implemented Conos [28] on the log-pc counts using the R package `conos` (version 1.5.0). We first identified the top 8000 highly variable genes separately for scRNA-seq data from each of the 8 patients using the `FindVariableFeatures` function in the R package `Seurat` (version 4.1.1) (so that a sufficiently large number of genes can be used for the subsequent data integration step). We then constructed a Conos object using the `Conos$new` function, built the joint graph using the `buildGraph` function with `k = 15`, `k.self = 5`, and `n.odgenes = 5000`, and computed batch corrected gene expression data using the `correctGenes` function with `n.od.genes = 8000`.

We assessed the performance of Liger using the returned matrix `H.norm`, the nonnegative normalized cell loading estimates. For Seurat 3, MNN Correct and Conos, we applied NMF with $K = 12$ to the batch corrected log-pc counts returned by each method, and assessed their performance using the estimated program membership matrix $\boldsymbol{L}$ from NMF. Arguments of all functions called to implement these methods were kept at their default values unless otherwise specified above.

S.1.5. Tumor scRNA-seq datasets.

HNSCC data. The HNSCC dataset was generated using SMART-Seq2, and contained transcriptional profiles for 5,902 cells from 18 patients [10]. Cells were confidently classified as malignant or normal cells using several complementary approaches in this study. Our GBCD analysis focused on tumors with at least 50 sequenced malignant cells, which involved 2,176 malignant cells from 10 primary HNSCC tumors and 5 matched lymph node metastases from these patients. We removed genes which were expressed in fewer than 20 cells, mitochondrial genes, and ribosomal genes, leaving a total of 17,113 genes. Since SMART-Seq2 does not yield UMI counts, we directly used the normalized expression data from [10], which are obtained as $\log_2(\text{TPM}_{i,j}/10 + 1)$ where $\text{TPM}_{i,j}$ is the transcript-per-million value for gene j in cell i.

PDAC data. The three PDAC datasets come from recently published scRNA-seq studies on PDAC. Details about these datasets are provided in Table 2. For all cohorts, we excluded neuroendocrine tumors and analyzed the tumors which were confirmed to be PDAC by histology, and focused on analyzing malignant cells only. For cohorts B and C, cells were annotated as malignant or normal in the respective studies [46, 47]. For cohort A, such information was not available, so we distinguished tumor cells from normal cells in the surrounding tumor microenvironment based on transcriptional profiles. In particular, we clustered cells in cohort A by their gene expression patterns using the functions `FindNeighbors` and `FindClusters` from the R package `Seurat` (version 4.1.1), and determined the identity of each cell cluster using known marker genes. We also performed copy number variation (CNV) analysis on cohort A using inferCNV of the Trinity CTAT Project at `https://github.com/broadinstitute/inferCNV`, and found that the cells classified as malignant based on expression patterns exhibited significantly higher CNV levels, lending evidence to our identification of malignant cells. We further filtered out low-quality cells with a large fraction of mitochondrial counts ($> 10\%$) and genes which were expressed in fewer than 25 cells in any cohort, as well as mitochondrial genes and ribosomal genes, leaving a total of 13,844 genes measured in 35,670 cells. More specifically, with the filtering procedures described above, we focused our analysis on 18,920, 10,936 and 5,814 high-quality malignant cells respectively in cohort A, B and C, which initially contained 31,195, 41,986, and 23,042 cells respectively. We performed library size normalization and log-transformation on the UMI counts, as described in Sec. S.1.1.1.

S.1.6. Gene-set enrichment analysis. For each identified GEP in HNSCC and PDAC analyses, we performed gene-set enrichment analysis of the most strongly up-regulated genes, which we defined as having a fold change estimate of at least 1.5 and being among the top 2% ranked genes based on fold change estimates. We reported the significantly enriched gene sets from MSigDB Hallmark and CP:Reactome gene sets, which had a Bonferroni-corrected p-value < 0.05 from Fisher's exact test.

S.1.7. Measuring concordance between gene expression signatures. To assess the degree of overlap between an expression signature from the literature and a given GEP k identified by GBCD, we performed the Wilcoxon rank sum test to compare the distributions of f_{jk} between (a) the top 100 ranked genes of this gene signature and (b) a matched control gene-set of random genes (matched to have similar average expression levels as (a)). The matched control genes were selected following the strategy used in [10]: we grouped all J genes into 50 bins based on their average expression levels in all cells, then for each gene in (a), we randomly chose 100 genes from the same expression bin and added them to (b), so that (b) has a distribution of average expression levels comparable to that of (a) and is 100-fold larger. The concordance between this previously derived expression signature and our GEP k was then quantified by $-\log_{10}(p\text{-value})$, where the p-value is produced by an one-sided Wilcoxon rank sum test which tests whether the median of f_{jk} in (a) is smaller than that in (b).

S.1.8. Classification of GBCD-identified GEPs. We determined whether a given GEP k is patient-specific or shared among multiple patients as follows. For patient $m(=1,\ldots,M)$, we denote the proportion of cells whose membership in GEP k exceeds 0.05 (which can be considered active in this GEP) as prop_m. Letting $m' = \text{argmax}_{m=1,\ldots,M}\,\text{prop}_m$, we define the "patient-specificity score" υ_k for GEP k as

$$\upsilon_k = 1 - \frac{\max_{m \neq m'} \text{prop}_m}{\text{prop}_{m'}}. \tag{21}$$

υ_k takes a value between 0 and 1. A υ_k close to 1 means that GEP k is predominantly expressed in patient m', and a υ_k close to 0 means that another patient has a similar percentage of cells expressing GEP k compared to patient m'. In other words, υ_k reflects the degree of patient-specific expression of GEP k. GEP k is classified as "patient-specific" if $\upsilon_k > 0.5$, and "shared" if $\upsilon_k \leq 0.5$.

For each shared GEP k, we performed one-way ANOVA on the GEP membership values $\boldsymbol{l}_k$ with the patient identity of cells as the explanatory variable, and calculated the proportion of total variation explainable by patient identity. This proportion quantifies how strongly a shared GEP's memberships vary between patients as opposed to within patients. In the results, we refer to this measure as "inter-patient variation".

S.1.9. Quantifying the spatial structure of GEP signatures. For each GEP k identified by GBCD, we studied whether the gene signature demonstrates systematic patterns across chromosome locations, in the sense that genes located closer to one another in chromosomal base-pair position tend to show more similar f_{jk} than genes located farther away. To do so, we fit a nonparametric regression model with f_{jk} as the response and the base-pair position as the predictor to all genes j, separately on each autosomal chromosome, $l = 1,\ldots,22$, using the function `gam` from the R package `mgcv` [101], then we recorded the model deviance. We quantified the degree of spatial structure for a given GEP k as

$$\rho_k = \frac{\sum_{l=1}^{22} \text{model deviance}_{kl}}{\sum_{l=1}^{22} \text{null deviance}_{kl}}, \tag{22}$$

where the "null deviance" is based on a nonparametric regression model with a constant intercept term as the only predictor. A larger ρ_k indicates that a higher proportion of observed variation in f_{jk} can be explained by the chromosomal location of genes, which suggests greater spatial structure. In the results, we refer to ρ_k as the "chromosomal score" for GEP k.

S.1.10. Survival analysis on PDAC bulk RNA-seq data. PDAC bulk RNA-seq data from [52, 56, 58] were downloaded from cBioPortal [102, 103] along with clinical annotations including overall survival information (i.e., time to death or right censoring). PDAC data from [57] were provided by Ontario Institute for Cancer Research based on a data sharing agreement with University of Chicago. Only surgically resected primary tumors that were histologically confirmed to be PDAC with no evidence of metastasis at diagnosis were included in the survival analysis, yielding a total of 391 patients (TCGA [56], $n = 102$; CPTAC [58], $n = 87$; QCMG [52], $n = 70$; PanCuRx [57], $n = 132$). Gene expression levels are quantified by $\log_2$-transformed TPM values.

For each of GEPs 1-14 and literature-derived subtype signatures, we computed an expression score in each tumor by comparing the average expression level of the top 200 ranked genes of this gene signature relative to a control gene-set [10], using the `Seurat` function `AddModuleScore`, and then, separately in each study, normalized expression scores across all tumors in that study to have zero mean and unit standard deviation. Stepwise variable selection in a Cox proportional hazards regression model, with overall survival as the endpoint and age, sex, tumor stage and all the program expression scores as covariates, was performed using the R package `My.stepwise` in [59]. Among the covariates, age, sex, and tumor stage are always included in the model, whereas the program expression scores are selected by the stepwise variable selection procedure.

S.1.11. Optimization of the GBCD pipeline. As described in the **Computing requirements** section, our original GBCD implementation works with the (non-sparse) $N \times N$ matrix $\boldsymbol{Y}\boldsymbol{Y}^T$, which was used to generate all the results presented in the paper. However, this approach scales poorly with N, and therefore is unsuitable for large scRNA-seq data sets. We improved the GBCD implementation in the R package `gbcd`, which instead works directly with the (sparse) $N \times J$ matrix $\boldsymbol{Y}$ and achieves significant reductions in runtime and memory usage for large N.

To verify that the two approaches produced similar results, we ran GBCD on the PDAC data, which is a large and highly complex dataset and so should allow us to detect sensitivity of GBCD results to any subtle changes in the computational pipeline. Without explicitly computing $\boldsymbol{Y}\boldsymbol{Y}^T$, the GBCD model fitting completed in 21 hours and required 48 GB; when explicitly computing $\boldsymbol{Y}\boldsymbol{Y}^T$, the GBCD model fitting completed in 50 hours and required 200 GB of memory. While there were some small differences in the results obtained, overall the two sets of results were very consistent (Supplementary Fig. S17).

The most time-consuming steps in fitting GBCD are the steps (ii) and (iv) which involve running T_1 and T_2 backfit iterations, respectively (Sec. S.1.1.2). We also assessed the impact of T_1 and T_2 on GBCD results, again using the PDAC data (the results presented in the paper were produced with $T_1 = 500$ and $T_2 = 200$). Supplementary Fig. S17 shows the time spent on backfitting as a function of the number of iterations. The runtime per backfit iteration in step (iv) was much longer than that in step (ii). This was likely due to the fact that there were more factors to update in each iteration in step (iv) than in step (ii). While the choice of T_1 strongly affected a few GEPs (e.g., GEP2), the estimates of all GEPs stabilized at a T_2 setting of 25. While it is difficult to draw general conclusions, our experiments suggest that GBCD results are not sensitive to T_2, and can be set to a small number such as 25 to reduce runtime when N is large. Using the improved implementation on the PDAC data (i.e.,

working with $\boldsymbol{Y}$ rather then $\boldsymbol{Y}\boldsymbol{Y}^T$), setting $T_1 = 500$ and $T_2 = 25$ reduced computing time from 21 to 14 hours and achieved essentially the same results as our original implementation (which required 50 hours).

Computing environment. Computations with simulated and real datasets were performed in R 4.1.0 [104], linked to the OpenBLAS 0.3.13 optimized numerical libraries, on Linux machines (Scientific Linux 8) with Intel Xeon E5-2680v4 ("Broadwell") processors. Fitting GBCD to a single simulation dataset which contained expression data for 3,200 cells and 10,000 genes) required about 30 min and 2 GB of memory.

Data availability. Most of the single-cell and bulk RNA-sequencing data used in this study are publicly available and downloadable following the URLs provided in the papers that generated the corresponding datasets, except for the data produced by the PanCuRx Translational Research Initiative [45], which require controlled access through a data sharing agreement with the PanCuRx Data Access Committee at the Ontario Institute for Cancer Research. Data and results necessary to reproduce the figures in the paper are available online (see URLs).

Code availability. The R package `gbcd` implementing the GBCD method and the scripts for performing the simulations and the analyses of the HNSCC and PDAC data are available online (see URLs).

REFERENCES

[1] SUVÀ, M. L. and TIROSH, I. (2019). Single-cell RNA sequencing in cancer: lessons learned and emerging challenges. *Molecular Cell* **75** 7–12.

[2] GASCH, A. P., SPELLMAN, P. T., KAO, C. M., CARMEL-HAREL, O., EISEN, M. B., STORZ, G., BOTSTEIN, D. and BROWN, P. O. (2000). Genomic expression programs in the response of yeast cells to environmental changes. *Molecular Biology of the Cell* **11** 4241–4257.

[3] SEGAL, E., SHAPIRA, M., REGEV, A., PE'ER, D., BOTSTEIN, D., KOLLER, D. and FRIEDMAN, N. (2003). Module networks: identifying regulatory modules and their condition-specific regulators from gene expression data. *Nature Genetics* **34** 166–176.

[4] KRESO, A. and DICK, J. E. (2014). Evolution of the cancer stem cell model. *Cell Stem Cell* **14** 275–291.

[5] KOTLIAR, D., VERES, A., NAGY, M. A., TABRIZI, S., HODIS, E., MELTON, D. A. and SABETI, P. C. (2019). Identifying gene expression programs of cell-type identity and cellular activity with single-cell RNA-Seq. *Elife* **8** e43803.

[6] FAN, J., SLOWIKOWSKI, K. and ZHANG, F. (2020). Single-cell transcriptomics in cancer: computational challenges and opportunities. *Experimental and Molecular Medicine* **52** 1452–1465.

[7] LEE, D. D. and SEUNG, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature* **401** 788–791.

[8] LEE, D. D. and SEUNG, H. S. (2001). Algorithms for non-negative matrix factorization. In *Advances in Neural Information Processing Systems* **13** 556–562.

[9] GILLIS, N. (2021). *Nonnegative matrix factorization*. Society for Industrial and Applied Mathematics, Philadelphia, PA.

[10] PURAM, S. V., TIROSH, I., PARIKH, A. S., PATEL, A. P., YIZHAK, K., GILLESPIE, S. et al. (2017). Single-cell transcriptomic analysis of primary and metastatic tumor ecosystems in head and neck cancer. *Cell* **171** 1611–1624.

[11] FILBIN, M. G., TIROSH, I., HOVESTADT, V., SHAW, M. L., ESCALANTE, L. E., MATHEWSON, N. D. et al. (2018). Developmental and oncogenic programs in H3K27M gliomas dissected by single-cell RNA-seq. *Science* **360** 331–335.

[12] KINKER, G. S., GREENWALD, A. C., TAL, R., ORLOVA, Z., CUOCO, M. S., MCFARLAND, J. M. et al. (2020). Pan-cancer single-cell RNA-seq identifies recurring programs of cellular heterogeneity. *Nature Genetics* **52** 1208–1218.

[13] DONG, R., YANG, R., ZHAN, Y., LAI, H.-D., YE, C.-J., YAO, X.-Y. et al. (2020). Single-cell characterization of malignant phenotypes and developmental trajectories of adrenal neuroblastoma. *Cancer Cell* **38** 716–733.e6.

[14] GOJO, J., ENGLINGER, B., JIANG, L., HÜBNER, J. M., SHAW, M. L., HACK, O. A. et al. (2020). Single-cell RNA-seq reveals cellular hierarchies and impaired developmental trajectories in pediatric ependymoma. *Cancer Cell* **38** 44–59.

[15] GAVISH, A., TYLER, M., SIMKIN, D., KOVARSKY, D., CASTRO, L. N. G., HALDER, D. et al. (2021). The transcriptional hallmarks of intra-tumor heterogeneity across a thousand tumors. *bioRxiv* doi:10.1101/2021.12.19.473368.

[16] WU, S. Z., AL-ERYANI, G., RODEN, D. L., JUNANKAR, S., HARVEY, K., ANDERSSON, A. et al. (2021). A single-cell and spatially resolved atlas of human breast cancers. *Nature Genetics* **53** 1334–1347.

[17] LEEK, J. T. and STOREY, J. D. (2007). Capturing heterogeneity in gene expression studies by surrogate variable analysis. *PLoS Genetics* **3** e161.

[18] LEEK, J. T. and STOREY, J. D. (2008). A general framework for multiple testing dependence. *Proceedings of the National Academy of Sciences* **105** 18718–18723.

[19] LEEK, J. T. (2014). svaseq: removing batch effects and other unwanted noise from sequencing data. *Nucleic Acids Research* **42** e161.

[20] MONCADA, R., BARKLEY, D., WAGNER, F., CHIODIN, M., DEVLIN, J. C., BARON, M., HAJDU, C. H., SIMEONE, D. M. and YANAI, I. (2020). Integrating microarray-based spatial transcriptomics and single-cell RNA-seq reveals tissue architecture in pancreatic ductal adenocarcinomas. *Nature Biotechnology* **38** 333–342.

[21] LIANG, J., CHEN, Z., HUANG, Y., BI, G., BIAN, Y., JIN, X., ZHANG, H., SUI, Q., ZHAN, C. and WANG, Q. (2022). Signatures of malignant cells and novel therapeutic targets revealed by single-cell sequencing in lung adenocarcinoma. *Cancer Medicine* **11** 2244–2258.

[22] VELTEN, L., STORY, B. A., HERNÁNDEZ-MALMIERCA, P., RAFFEL, S., LEONCE, D. R., MILBANK, J., PAULSEN, M., DEMIR, A., SZU-TU, C., FRÖMEL, R. et al. (2021). Identification of leukemic and pre-leukemic stem cells by clonal tracking from single-cell transcriptomics. *Nature Communications* **12** 1366.

[23] XIE, X. P., LAKS, D. R., SUN, D., GANBOLD, M., WANG, Z., PEDRAZA, A. M., BALE, T., TABAR, V., BRENNAN, C., ZHOU, X. et al. (2022). Quiescent human glioblastoma cancer stem cells drive tumor initiation, expansion, and recurrence following chemotherapy. *Developmental Cell* **57** 32–46.

[24] KORSUNSKY, I., MILLARD, N., FAN, J., SLOWIKOWSKI, K., ZHANG, F., WEI, K., BAGLAENKO, Y., BRENNER, M., LOH, P.-R. and RAYCHAUDHURI, S. (2019). Fast, sensitive and accurate integration of single-cell data with Harmony. *Nature Methods* **16** 1289–1296.

[25] WELCH, J. D., KOZAREVA, V., FERREIRA, A., VANDERBURG, C., MARTIN, C. and MACOSKO, E. Z. (2019). Single-cell multi-omic integration compares and contrasts features of brain cell identity. *Cell* **177** 1873–1887.

[26] STUART, T., BUTLER, A., HOFFMAN, P., HAFEMEISTER, C., PAPALEXI, E., MAUCK, W. M., HAO, Y., STOECKIUS, M., SMIBERT, P. and SATIJA, R. (2019). Comprehensive integration of single-cell data. *Cell* **177** 1888–1902.

[27] HAGHVERDI, L., LUN, A. T., MORGAN, M. D. and MARIONI, J. C. (2018). Batch effects in single-cell RNA-sequencing data are corrected by matching mutual nearest neighbors. *Nature Biotechnology* **36** 421–427.

[28] BARKAS, N., PETUKHOV, V., NIKOLAEVA, D., LOZINSKY, Y., DEMHARTER, S., KHODOSEVICH, K. and KHARCHENKO, P. V. (2019). Joint analysis of heterogeneous single-cell RNA-seq dataset collections. *Nature Methods* **16** 695–698.

[29] HIE, B., BRYSON, B. and BERGER, B. (2019). Efficient integration of heterogeneous single-cell transcriptomes using Scanorama. *Nature Biotechnology* **37** 685–691.

[30] TRAN, H. T. N., ANG, K. S., CHEVRIER, M., ZHANG, X., LEE, N. Y. S., GOH, M. and CHEN, J. (2020). A benchmark of batch-effect correction methods for single-cell RNA sequencing data. *Genome Biology* **21** 1–32.

[31] RICHARDS, L. M., RIVERIN, M., MOHANRAJ, S., AYYADHURY, S., CROUCHER, D. C., DÍAZ-MEJÍA, J. J., COUTINHO, F. J., DIRKS, P. B. and PUGH, T. J. (2021). A comparison of data integration methods for single-cell RNA sequencing of cancer samples. *bioRxiv* 2021–08.

[32] SPRATLING, M. W. (2006). Learning iimage components for object recognition. *Journal of Machine Learning Research* **7** 793–815.

[33] LI, S. Z., HOU, X. W., ZHANG, H. J. and CHENG, Q. S. (2001). Learning spatially localized, parts-based representation. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* **1** 207–212.

[34] FENG, T., LI, S. Z., SHUM, H.-Y. and ZHANG, H. (2002). Local non-negative matrix factorization as a visual representation. In *Proceedings of the 2nd International Conference on Development and Learning* 178–186.

[35] WANG, W. and STEPHENS, M. (2021). Empirical Bayes matrix factorization. *Journal of Machine Learning Research* **22** 1–40.

[36] WILLWERSCHEID, J. (2021). Empirical Bayes matrix factorization: methods and applications, PhD thesis, University of Chicago.

[37] WILLWERSCHEID, J. and STEPHENS, M. (2021). ebnm: an R package for solving the empirical Bayes normal means problem using a variety of prior families. *arXiv* **2110.00152**.

[38] WILLWERSCHEID, J. (2022). flashier: empirical Bayes matrix factorization R package version 0.2.34.

[39] SUBRAMANIAN, A., TAMAYO, P., MOOTHA, V. K., MUKHERJEE, S., EBERT, B. L., GILLETTE, M. A., PAULOVICH, A., POMEROY, S. L., GOLUB, T. R., LANDER, E. S. et al. (2005). Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences* **102** 15545–15550.

[40] FINAK, G., MCDAVID, A., YAJIMA, M., DENG, J., GERSUK, V., SHALEK, A. K., SLICHTER, C. K., MILLER, H. W., MCELRATH, M. J., PRLIC, M., LINSLEY, P. S. and GOTTARDO, R. (2015). MAST: a flexible statistical framework for assessing transcriptional changes and characterizing heterogeneity in single-cell RNA sequencing data. *Genome Biology* **16** 1–13.

[41] VAN DER MAATEN, L. and HINTON, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research* **9**.

[42] CANCER GENOME ATLAS NETWORK (2015). Comprehensive genomic characterization of head and neck squamous cell carcinomas. *Nature* **517** 576.

[43] LIBERZON, A., BIRGER, C., THORVALDSDÓTTIR, H., GHANDI, M., MESIROV, J. P. and TAMAYO, P. (2015). The Molecular Signatures Database Hallmark Gene Set Collection. *Cell Systems* **1** 417–425.

[44] LIBERZON, A., SUBRAMANIAN, A., PINCHBACK, R., THORVALDSDOTTIR, H., TAMAYO, P. and MESIROV, J. P. (2011). Molecular signatures database (MSigDB) 3.0. *Bioinformatics* **27** 1739–1740.

[45] CHAN-SENG-YUE, M., KIM, J. C., WILSON, G. W., NG, K., FIGUEROA, E. F., O'KANE, G. M. et al. (2020). Transcription phenotypes of pancreatic cancer are driven by genomic events during tumor evolution. *Nature Genetics* **52** 231–240.

[46] PENG, J., SUN, B.-F., CHEN, C.-Y., ZHOU, J.-Y., CHEN, Y.-S., et al., etalZHAO, Y., YANG, Y.-G. and WU, W. (2019). Single-cell RNA-seq highlights intra-tumoral heterogeneity and malignant progression in pancreatic ductal adenocarcinoma. *Cell Research* **29** 725–738.

[47] RAGHAVAN, S., WINTER, P. S., NAVIA, A. W., WILLIAMS, H. L., et al., etalAGUIRRE, A. J. and SHALEK, A. K. (2021). Microenvironment drives cell state, plasticity, and drug response in pancreatic cancer. *Cell* **184** 6119–6137.

[48] ORTH, M., METZGER, P., GERUM, S., MAYERLE, J., SCHNEIDER, G., BELKA, C., SCHNURR, M. and LAUBER, K. (2019). Pancreatic ductal adenocarcinoma: biological hallmarks, current status, and future perspectives of combined modality treatment approaches. *Radiation Oncology* **14** 1–20.

[49] MOFFITT, R. A., MARAYATI, R., FLATE, E. L., VOLMAR, K. E., LOEZA, S. G. H., HOADLEY, K. A., RASHID, N. U., WILLIAMS, L. A., EATON, S. C., CHUNG, A. H. et al. (2015). Virtual microdissection identifies distinct tumor-and stroma-specific subtypes of pancreatic ductal adenocarcinoma. *Nature Genetics* **47** 1168–1178.

[50] HWANG, W. L., JAGADEESH, K. A., GUO, J. A., HOFFMAN, H. I., YADOLLAHPOUR, P., REEVES, J. W., MOHAN, R., DROKHLYANSKY, E., VAN WITTENBERGHE, N., ASHENBERG, O. et al. (2022). Single-nucleus and spatial transcriptome profiling of pancreatic cancer identifies multicellular dynamics associated with neoadjuvant treatment. *Nature Genetics* **54** 1178–1191.

[51] TAYLOR, A. M., SHIH, J., HA, G., GAO, G. F., ZHANG, X., BERGER, A. C., SCHUMACHER, S. E., WANG, C., HU, H., LIU, J. et al. (2018). Genomic and functional approaches to understanding cancer aneuploidy. *Cancer Cell* **33** 676–689.

[52] BAILEY, P., CHANG, D. K., NONES, K., JOHNS, A. L., PATCH, A.-M., GINGRAS, M.-C. et al. (2016). Genomic analyses identify molecular subtypes of pancreatic cancer. *Nature* **531** 47–52.

[53] RASHID, N. U., PENG, X. L., JIN, C., MOFFITT, R. A., VOLMAR, K. E., BELT, B. A., PANNI, R. Z., NYWENING, T. M., HERRERA, S. G., MOORE, K. J. et al. (2020). Purity independent subtyping of tumors (PurIST), a clinically robust, single-sample classifier for tumor subtyping in pancreatic cancer. *Clinical Cancer Research* **26** 82–92.

[54] MIYABAYASHI, K., BAKER, L. A., DESCHÊNES, A., TRAUB, B., CALIGIURI, G., PLENKER, D., ALAGESAN, B., BELLEAU, P., LI, S., KENDALL, J. et al. (2020). Intraductal transplantation models of human pancreatic ductal adenocarcinoma reveal progressive transition of molecular subtypes. *Cancer Discovery* **10** 1566–1589.

[55] HAN, J., DEPINHO, R. A. and MAITRA, A. (2021). Single-cell RNA sequencing in pancreatic cancer. *Nature Reviews Gastroenterology & Hepatology* **18** 451–452.

[56] RAPHAEL, B. J., HRUBAN, R. H., AGUIRRE, A. J., MOFFITT, R. A., YEH, J. J., STEWART, C., ROBERTSON, A. G., CHERNIACK, A. D., GUPTA, M., GETZ, G. et al. (2017). Integrated genomic characterization of pancreatic ductal adenocarcinoma. *Cancer Cell* **32** 185–203.

[57] AUNG, K. L., FISCHER, S. E., DENROCHE, R. E., JANG, G.-H., DODD, A., CREIGHTON, S. et al. (2018). Genomics-driven precision medicine for advanced pancreatic cancer: early results from the COMPASS trial. *Clinical Cancer Research* **24** 1344–1354.

[58] CAO, L., HUANG, C., ZHOU, D. C., BATHE, O. F., CHAN, D. W., HRUBAN, R. H., DING, L., ZHANG, B., ZHANG, H. and INVESTIGATORS, C. P. T. A. C. (2021). Proteogenomic characterization of pancreatic ductal adenocarcinoma. *Cell* **184** 5031–5052.

[59] HU, F.-C. (2017). My.stepwise: Stepwise Variable Selection Procedures for Regression Analysis R package version 0.1.0.

[60] LIU, J., ZHU, Y. and GE, C. (2020). LncRNA ZFAS1 promotes pancreatic adenocarcinoma metastasis via the RHOA/ROCK2 pathway by sponging miR-3924. *Cancer Cell International* **20** 1–15.

[61] SANTANA-CODINA, N., DEL REY, M. Q., KAPNER, K. S., ZHANG, H., GIKANDI, A., MALCOLM, C., POUPAULT, C., KULJANIN, M., JOHN, K. M., BIANCUR, D. E. et al. (2022). NCOA4-mediated ferritinophagy is a pancreatic cancer dependency via maintenance of iron bioavailability for iron–sulfur cluster proteins. *Cancer Discovery* **12** 2180–2197.

[62] MUKHOPADHYAY, S., ENCARNACIÓN-ROSADO, J., LIN, E. Y., SOHN, A. S., ZHANG, H., MANCIAS, J. D. and KIMMELMAN, A. C. (2023). Autophagy supports mitochondrial metabolism through the regulation of iron homeostasis in pancreatic cancer. *Science Advances* **9** eadf9284.

[63] PAKOS-ZEBRUCKA, K., KORYGA, I., MNICH, K., LJUJIC, M., SAMALI, A. and GORMAN, A. M. (2016). The integrated stress response. *EMBO Reports* **17** 1374–1395.

[64] LIU, L., WISE, D. R., DIEHL, J. A. and SIMON, M. C. (2008). Hypoxic reactive oxygen species regulate the integrated stress response and cell survival. *Journal of Biological Chemistry* **283** 31153–31162.

[65] RZYMSKI, T., MILANI, M., PIKE, L., BUFFA, F., MELLOR, H., WINCHESTER, L., PIRES, I., HAMMOND, E., RAGOUSSIS, I. and HARRIS, A. (2010). Regulation of autophagy by ATF4 in response to severe hypoxia. *Oncogene* **29** 4424–4435.

[66] WORTEL, I. M., VAN DER MEER, L. T., KILBERG, M. S. and VAN LEEUWEN, F. N. (2017). Surviving stress: modulation of ATF4-mediated stress responses in normal and malignant cells. *Trends in Endocrinology & Metabolism* **28** 794–806.

[67] B'CHIR, W., MAURIN, A.-C., CARRARO, V., AVEROUS, J., JOUSSE, C., MURANISHI, Y., PARRY, L., STEPIEN, G., FAFOURNOUX, P. and BRUHAT, A. (2013). The eIF2α/ATF4 pathway is essential for stress-induced autophagy gene expression. *Nucleic Acids Research* **41** 7683–7699.

[68] BRYANT, K. L., STALNECKER, C. A., ZEITOUNI, D., KLOMP, J. E., PENG, S., TIKUNOV, A. P., GUNDA, V., PIEROBON, M., WATERS, A. M., GEORGE, S. D. et al. (2019). Combination of ERK and autophagy inhibition as a treatment approach for pancreatic cancer. *Nature Medicine* **25** 628–640.

[69] KINSEY, C. G., CAMOLOTTO, S. A., BOESPFLUG, A. M., GUILLEN, K. P., FOTH, M., TRUONG, A., SCHUMAN, S. S., SHEA, J. E., SEIPP, M. T., YAP, J. T. et al. (2019). Protective autophagy elicited by RAF→MEK→ERK inhibition suggests a treatment strategy for RAS-driven cancers. *Nature Medicine* **25** 620–627.

[70] ZHU, S., ZHANG, Q., SUN, X., ZEH III, H. J., LOTZE, M. T., KANG, R. and TANG, D. (2017). HSPA5 regulates ferroptotic cell death in cancer cells. *Cancer Research* **77** 2064–2077.

[71] CHEN, D., FAN, Z., RAUH, M., BUCHFELDER, M., EYUPOGLU, I. and SAVASKAN, N. (2017). ATF4 promotes angiogenesis and neuronal cell death and confers ferroptosis in a xCT-dependent manner. *Oncogene* **36** 5593–5608.

[72] OLIVE, K. P., JACOBETZ, M. A., DAVIDSON, C. J., GOPINATHAN, A., MCINTYRE, D., HONESS, D., MADHU, B., GOLDGRABEN, M. A., CALDWELL, M. E., ALLARD, D. et al. (2009). Inhibition of Hedgehog signaling enhances delivery of chemotherapy in a mouse model of pancreatic cancer. *Science* **324** 1457–1461.

[73] SHAH, V. M., SHEPPARD, B. C., SEARS, R. C. and ALANI, A. W. (2020). Hypoxia: Friend or Foe for drug delivery in Pancreatic Cancer. *Cancer Letters* **492** 63–70.

[74] ENCARNACIÓN-ROSADO, J. and KIMMELMAN, A. C. (2021). Harnessing metabolic dependencies in pancreatic cancers. *Nature Reviews Gastroenterology & Hepatology* **18** 482–492.

[75] TAO, J., YANG, G., ZHOU, W., QIU, J., CHEN, G., LUO, W., ZHAO, F., YOU, L., ZHENG, L., ZHANG, T. et al. (2021). Targeting hypoxic tumor microenvironment in pancreatic cancer. *Journal of Hematology & Oncology* **14** 1–25.

[76] PARK, S. J., YOO, H. C., AHN, E., LUO, E., KIM, Y., SUNG, Y., YU, Y. C., KIM, K., MIN, D. S., LEE, H. S. et al. (2023). Enhanced glutaminolysis drives hypoxia-induced chemoresistance in pancreatic cancer. *Cancer Research* **83** 735–752.

[77] HALBROOK, C. J., THURSTON, G., BOYER, S., ANARAKI, C., JIMÉNEZ, J. A., MCCARTHY, A., STEELE, N. G., KERK, S. A., HONG, H. S., LIN, L. et al. (2022). Differential integrated stress response and asparagine production drive symbiosis and therapy resistance of pancreatic adenocarcinoma cells. *Nature Cancer* **3** 1386–1403.

[78] WANG, E. M., AKASAKA, H., ZHAO, J., VARADHACHARY, G. R., LEE, J. E., MAITRA, A., FLEMING, J. B., HUNG, M.-C., WANG, H. and KATZ, M. H. (2019). Expression and Clinical Significance of PERK and Phosphorylated EIF2α in Pancreatic Ductal Adenocarcinoma. *Pancreas* **48** 323.

[79] YANG, C., TIAN, C., HOFFMAN, T. E., JACOBSEN, N. K. and SPENCER, S. L. (2021). Melanoma subpopulations that rapidly escape MAPK pathway inhibition incur DNA damage and rely on stress signalling. *Nature Communications* **12** 1747.

[80] KALKAVAN, H., CHEN, M. J., CRAWFORD, J. C., QUARATO, G., FITZGERALD, P., TAIT, S. W., GODING, C. R. and GREEN, D. R. (2022). Sublethal cytochrome c release generates drug-tolerant persister cells. *Cell* **185** 3356–3374.

[81] LUN, A. T., BACH, K. and MARIONI, J. C. (2016). Pooling across cells to normalize single-cell RNA sequencing data with many zero counts. *Genome Biology* **17** 75.

[82] BULLARD, J. H., PURDOM, E., HANSEN, K. D. and DUDOIT, S. (2010). Evaluation of statistical methods for normalization and differential expression in mRNA-seq experiments. *BMC Bioinformatics* **11** 94.

[83] ANDERS, S. and HUBER, W. (2010). Differential expression analysis for sequence count data. *Genome Biology* **11** R106.

[84] LUN, A. (2018). Overcoming systematic errors caused by log-transformation of normalized single-cell RNA sequencing data. *bioRxiv*.

[85] TOWNES, F. W., HICKS, S. C., ARYEE, M. J. and IRIZARRY, R. A. (2019). Feature selection and dimension reduction for single-cell RNA-Seq based on a multinomial model. *Genome Biology* **20** 1–16.

[86] AHLMANN-ELTZE, C. and HUBER, W. (2022). Comparison of Transformations for Single-Cell RNA-seq Data. *bioRxiv*.

[87] PICELLI, S., FARIDANI, O. R., BJÖRKLUND, Å. K., WINBERG, G., SAGASSER, S. and SANDBERG, R. (2014). Full-length RNA-seq from single cells using Smart-seq2. *Nature Protocols* **9** 171–181.

[88] LI, B. and DEWEY, C. N. (2011). RSEM: accurate transcript quantification from RNA-Seq data with or without a reference genome. *BMC Bioinformatics* **12** 323.

[89] WAGNER, G. P., KIN, K. and LYNCH, V. J. (2012). Measurement of mRNA abundance using RNA-seq data: RPKM measure is inconsistent among samples. *Theory in Biosciences* **131** 281–285.

[90] DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society, Series B* **39** 1–22.

[91] WILLWERSCHEID, J. and STEPHENS, M. (2022). flashier: A faster and angrier package for EBMF R package version 0.2.27.

[92] DING, C. H., LI, T. and JORDAN, M. I. (2008). Convex and semi-nonnegative matrix factorizations. *IEEE transactions on pattern analysis and machine intelligence* **32** 45–55.

[93] STEPHENS, M. (2017). False discovery rates: a new deal. *Biostatistics* **18** 275–294.

[94] GELMAN, A., HILL, J. and YAJIMA, M. (2012). Why we (usually) don't have to worry about multiple comparisons. *Journal of research on educational effectiveness* **5** 189–211.

[95] ZAPPIA, L., PHIPSON, B. and OSHLACK, A. (2017). Splatter: simulation of single-cell RNA sequencing data. *Genome Biology* **18** 1–15.

[96] GERARD, D. (2020). Data-based RNA-seq simulations by binomial thinning. *BMC Bioinformatics* **21** 206.

[97] LIN, X. and BOUTROS, P. C. (2020). Optimization and expansion of non-negative matrix factorization. *BMC Bioinformatics* **21** 1–10.

[98] PRITCHARD, J. K., STEPHENS, M. and DONNELLY, P. (2000). Inference of population structure using multilocus genotype data. *Genetics* **155** 945–959.

[99] BLEI, D. M., NG, A. Y. and JORDAN, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research* **3** 993–1022.

[100] CARBONETTO, P., SARKAR, A., WANG, Z. and STEPHENS, M. (2021). Non-negative matrix factorization algorithms greatly improve topic model fits. *arXiv* **2105.13440**.

[101] WOOD, S. (2015). Package 'mgcv'. *R package version* **1** 729.

[102] CERAMI, E., GAO, J., DOGRUSOZ, U., GROSS, B. E., SUMER, S. O., AKSOY, B. A., JACOBSEN, A., BYRNE, C. J., HEUER, M. L., LARSSON, E. et al. (2012). The cBio cancer genomics portal: an open platform for exploring multidimensional cancer genomics data. *Cancer Discovery* **2** 401–404.

[103] GAO, J., AKSOY, B. A., DOGRUSOZ, U., DRESDNER, G., GROSS, B., SUMER, S. O., SUN, Y., JACOBSEN, A., SINHA, R., LARSSON, E. et al. (2013). Integrative analysis of complex cancer genomics and clinical profiles using the cBioPortal. *Science Signaling* **6** pl1–pl1.

[104] R CORE TEAM (2021). R: a language and environment for statistical computing. R Foundation for Statistical Computing.