

Comprehensive Benchmarking and Integration of Tumor Microenvironment Cell Estimation Methods



Alejandro Jiménez-Sánchez, Oliver Cast, and Martin L. Miller

Abstract

Various computational approaches have been developed for estimating the relative abundance of different cell types in the tumor microenvironment (TME) using bulk tumor RNA data. However, a comprehensive comparison across diverse datasets that objectively evaluates the performance of these approaches has not been conducted. Here, we benchmarked seven widely used tools and gene sets and introduced Consensus^{TME}, a method that integrates gene sets from all the other methods for relative TME cell estimation of 18 cell types. We collected a comprehensive benchmark dataset consisting of pan-cancer data (DNA-derived purity, leukocyte methylation, and hematoxylin and eosin–derived lymphocyte counts) and cell-specific benchmark datasets (peripheral blood cells and tumor tissues).

Although none of the methods outperformed others in every benchmark, Consensus^{TME} ranked top three in all cancer-related benchmarks and was the best performing tool overall. We provide a Web resource to interactively explore the benchmark results and an objective evaluation to help researchers select the most robust and accurate method to further investigate the role of the TME in cancer (www.consensusTME.org).

Significance: This work shows an independent and comprehensive benchmarking of recently developed and widely used tumor microenvironment cell estimation methods based on bulk expression data and integrates the tools into a consensus approach.

Introduction

The tumor microenvironment (TME) plays an active role in tumor initiation, progression, metastasis, and treatment response (1). Thus, studying the TME is a central paradigm of cancer research. However, a great variety of stromal and immune cell types populate tumor tissues, and the complex interaction between these different components of the TME is still unclear. Traditionally, cells from the TME have been quantified using IHC, immunofluorescence, flow cytometry, and more recently using cytometry by time of flight (CyTOF) mass spectrometry. These methods, although accurate, are laborious, low throughput, and require preselected cellular markers, making their application in large number of samples and measurements challenging. Thus, the systematic application of these assays for comprehensively investigating the various different cell types in the TME in an unbiased manner is limited. Single-cell RNA sequencing (scRNA-

seq) has begun to fill this gap; however, scRNA-seq is expensive to apply on large patient cohorts, requires specific sample preparation, and cannot be applied to existing datasets, such as The Cancer Genome Atlas (TCGA), which consists of thousands of genomically profiled and clinically well-annotated tumor samples. The study of the different cell subpopulations of the TME in TCGA has become an important goal but also an important challenge for bioinformatics, because cell-type information identity is mixed in bulk tumor transcriptomics data.

Estimation of noncancerous cell proportions from bulk tumor samples can be performed using genomics data such as whole-exome sequencing, microarrays, RNA sequencing (RNA-seq), or DNA methylation data. During the last decade, multiple computational approaches have been developed intending to quantitatively or semiquantitatively calculate distinct TME cell-type population estimates (2). A variety of statistical frameworks and algorithmic procedures have been employed with each method using different benchmarking datasets (2). In general, two different algorithmic classes exist into which most TME cell estimation methods can be classified: regression-based deconvolution algorithms and gene set enrichment-based methods. Importantly, both classes rely on cell-type-specific markers that are selected according to prior knowledge. The deconvolution algorithms use signature matrices containing gene expression profiles of purified immune cells and impute to what degree each of the gene expression profile is represented in the bulk tumor gene expression profile. Gene set enrichment-based methods assign curated gene sets to represent cell types before computing enrichment scores as a function of the expression of the genes within each gene set. Both classes of computational approaches that intend to estimate TME cell content in tumors (i.e., TME cell estimation methods) require two components, a statistical framework (e.g.,

Cancer Research UK Cambridge Institute, University of Cambridge, Cambridge, United Kingdom.

Note: Supplementary data for this article are available at Cancer Research Online (<http://cancerres.aacrjournals.org/>).

A. Jiménez-Sánchez and O. Cast contributed equally to this article.

Corresponding Authors: Alejandro Jiménez-Sánchez, Cancer Research UK Cambridge Institute, University of Cambridge, Li Ka Shing Centre, Robinson Way, Cambridge CB2 0RE, United Kingdom. Phone: 646-238-0035; Fax: 44-0-1223-769-880; E-mail: ajs.scientia@google.com; and Martin L. Miller, martin.miller@cruk.cam.ac.uk

Cancer Res 2019;79:6238–46

doi: 10.1158/0008-5472.CAN-18-3560

©2019 American Association for Cancer Research.

regression or gene set enrichment) and a signature for the specific cell type of interest (e.g., a signature matrix or signature gene set).

Cell-type-specific estimation in the TME using bulk tumor data is a challenging task as certain stromal and immune cell populations are lowly abundant cell populations and is further convoluted because expression of particular genes is rarely unique to any particular cell type. Thus, there is not a straightforward solution for accurate TME cell estimation with various different gene signatures and statistical frameworks suggested as the optimal solution (2). One of the problems in the field is that each method has claimed to outperform others in their own benchmarking experiments (3, 4). Thus, the need for independent and more comprehensive benchmarks has become increasingly important (5, 6).

Here, we developed a consensus approach (Consensus^{TME}) for 18 different cell types that compiles cell-type-specific genes used by seven published gene sets or existing TME cell estimation methods: Bindea and colleagues gene sets (7), Davoli and colleagues gene sets (8), Danaher and colleagues gene sets (9), CIBERSORT (10), MCP-counter (11), TIMER (12), and xCell (13). We performed pan-cancer benchmarks using publicly available bulk genomic and transcriptomic data from TCGA and cell-type-specific benchmarks using "ground truth" data based on experiments carried out in each of the original manuscripts where available. The Consensus^{TME} approach, available as an R package, is evolvable by design allowing new gene signatures and algorithms to be incorporated and their performance compared with continuously updated benchmark datasets. Overall, the Consensus^{TME} approach provides a robust and improved method for the relative cell type estimation using bulk expression data of human tumor samples. We make all results accessible to the public (www.consensusTME.org) as well as provide a framework for incorporating new TME estimation methods and datasets.

Materials and Methods

Contact for resource sharing

Further information and requests for resources should be directed to and will be fulfilled by the lead contact, Martin L. Miller (martin.miller@cruk.cam.ac.uk).

Quantification and statistical analysis

Single-sample gene set enrichment analysis. Single-sample gene set enrichment analysis (ssGSEA; ref. 14), a modification of standard gene set enrichment analysis (15), was performed on RNA measurements for each sample using the GSVA package version 1.32.0 (16) in R version 3.6.0 with parameter: `method = "ssgsea"`. TME cell gene sets were obtained from previous publications or described as below (7–9).

Consensus^{TME}. To generate the Consensus^{TME} gene sets, we identified cell types for which signatures from at least two different sources were available, 18 cell types in total. To extract genes from the signature matrix "LM22" used by CIBERSORT, we first filtered out genes whose expression value was below 1.96 SD of the mean for each cell type. In addition, we collapsed activated and resting states for corresponding cell types. Once we had collected signature genes from Bindea and colleagues, Danaher and colleagues, Davoli and colleagues, CIBERSORT, MCP-Counter, and xCell, we created a unique union of the genes for each cell type. From this union of genes, a set of cell-type-specific genes was curated for

each of the TCGA cancer types. This was done using a similar approach to the TIMER algorithm where genes were included only if the expression of that gene has a negative correlation (Pearson's correlation < 0.2 , P value ≤ 0.05) with tumor purity (ABSOLUTE derived) for the corresponding cancer type (12, 17). The assumption behind this filter is that it removes immune and stromal cell genes, which may be aberrantly expressed by cancer cells in some cancer types, hereby ensuring that filtered gene sets likely represent the presence of the implicated immune cell while also leading to gene sets that account for tissue-specific variability in the immune response (18). The implementation of such an expression filter may on the other hand unintentionally remove immune- and stromal-specific genes that may indeed be cancer specific and hereby exclude interesting genes that are interesting to investigate in other contexts. However, in the current version, the implementation of the tumor-specific expression filter ensures that as additional signatures are incorporated into the Consensus^{TME} supersets, performance continues to increase. Finally, ssGSEA was employed to calculate normalized enrichment scale (NES) scores for each cell type as described previously. General immune scores for each tumor types were generated by combining the genes of the different immune cells into one gene set for each TCGA cancer type. The Consensus^{TME} method can be used through installation of the "ConsensusTME" R package via GitHub (<https://github.com/cansysbio/ConsensusTME/>).

Comparison statistical metrics. Concordance between computational estimates and ground truth values was measured using either Kendall's rank correlation coefficient or the multiple linear regression goodness-of-fit metrics: adjusted R^2 , Akaike information criterion (AIC), and Bayesian information criterion (BIC). AIC and BIC z-score values were calculated to compare across different tumor types in the comparisons because AIC and BIC values are unitless. Differences between groups of variables in the TIMER benchmark were identified using one-way ANOVA with Tukey honest significant differences *post hoc* tests. All statistical tests were adjusted for multiple testing using the Benjamini-Hochberg procedure to control for FDR.

TCGA immune estimations. TCGA RNA-seq data were collected from cBioPortal (19). Batch normalization had been applied and gene expression values calculated using the "RSEM" pipeline (20). Four existing TME cell estimation methods and three published gene sets were used alongside Consensus^{TME} to produce relative abundances of immune cell types per sample across 32 tumor types. For each method, a general immune score was also derived if it was not already provided, representing the total level of immune cell infiltration in each tumor sample; for the TME methods that were gene sets, this was done by collapsing the genes for each cell type together to form a new category and for the regression-based methods, this was done by summing the regression coefficients as per the method used by CIBERSORT algorithm with parameters: `absolute = TRUE`, `abs_method = 'sig.score'`.

TME cell estimation methods. TME estimation methods were used to estimate abundances of cell types from RNA expression profiles. The TME methods that were benchmarked were CIBERSORT (run in absolute mode; ref. 10), MCP-counter (11), TIMER (12), and xCell (13), as well as gene sets collected from three previous publications (Supplementary Table S1A; refs. 7–9).

Bindea and colleagues, Danaher and colleagues, and Davoli and colleagues gene sets. Gene sets provided by Bindea and colleagues, Danaher and colleagues, and Davoli and colleagues were used with ssGSEA to provide enrichment scores for each of the immune signatures (7–9). To generate general immune scores, genes selected for immune cells were combined into one gene set for each method independently.

xCell. The "xCell" R package (version 1.12) was used to generate immune estimates for the xCell method (13). A general immune estimation score is already generated by xCell.

MCP-counter. Estimations for the MCP-counter method were produced using the "MCPcounter" R package (version 1.1.0; ref. 11). Immune scores for this method were produced in a similar manner as the ssGSEA methods by creating a union of signature genes for each of the cell types. The "MCPcounter.estimate" function was altered to allow for the new signature.

CIBERSORT. CIBERSORT estimations were produced using the R source code, provided on request from the Web resource (10). CIBERSORT was run in "Absolute mode" (under beta development) using 100 permutations and quantile normalization disabled as recommended for RNA-seq data. Absolute scores representing the "overall immune content" are produced natively by the algorithm in absolute mode.

TIMER. TIMER estimations were produced using R source code, available from the Web resource (12). Immune scores for TIMER were produced as a sum of the coefficients for each cell type.

Purity score benchmark

Pan-cancer purity scores were downloaded from the NIH Genomic Data Commons (<https://gdc.cancer.gov/about-data/publications/PanCan-CellOfOrigin>; Supplementary Table S1B; ref. 21). Purity scores were generated using ABSOLUTE (17), which uses copy number, variant allele frequency, and tumor-specific karyotype data to calculate the cancer fraction of tumor samples. To benchmark the immune estimation methodologies using purity of samples, the immune scores were added to an independent stromal score, calculated through the use of ESTIMATE (version 1.0.13; ref. 22). The stromal score was added as this could negatively affect the performance of estimating RNA-based tumor purity for TME cell estimation methods, which estimate only immune cells compared with those that estimate both immune and stromal cells. ABSOLUTE's derived tumor purity and the different TME methods tumor purity scores were correlated independently for each tumor type.

Leukocyte fraction benchmark

Methylation-derived leukocyte fraction scores were downloaded from the NIH Genomic Data Commons (<https://gdc.cancer.gov/about-data/publications/PanCan-CellOfOrigin>; Supplementary Table S1B; ref. 21). Multiple linear regression was then employed using TME method cell-type estimates that were in the category of being leukocytes (Supplementary Table S1C) as explanatory variables and the methylation-derived leukocyte fraction as the response variable. Leukocyte methylation data were log transformed to meet the normality and heteroscedasticity assumptions of the model. Adjusted R^2 , AIC, and BIC metrics were calculated to compare the goodness of fit between the

methods while taking into consideration the number of variables included in the model.

Somatic single-nucleotide mutation data

Somatic single-nucleotide mutation data were downloaded from the Broad Institute GDAC Firehose (<https://gdac.broadinstitute.org/>).

Hematoxylin and eosin deep-learning lymphocyte fractions benchmark

Lymphocyte fractions were generated by Saltz and colleagues for 13 TCGA cancer types using deep-learning-based image analysis (Supplementary Table S1B; ref. 23). Multiple linear regression was applied in a similar manner as for the leukocyte methylation analysis, instead using a hyperbolic sine transformation of lymphocyte fraction as a response variable to meet normality and heteroscedasticity assumptions of the model. Models for each method were fitted using only method estimates of lymphocytes as explanatory variables (Supplementary Table S1C).

Independent cell-type-specific benchmarks

The benchmarking validation analyses for each of the methods were replicated, where possible, to match the parameters used in the original publications. Of the seven methods, there were four benchmarking datasets available; either online or provided by the authors. Each of the datasets contained samples with bulk gene expression values along with matched "ground truth" values (Supplementary Table S1B). The CIBERSORT benchmarking dataset, provided by the authors on request (10), consisted of flow cytometry values of different immune cell types from peripheral blood mononuclear cell (PBMCs) samples. The xCell benchmarking datasets, SDY311 and SDY420, were publicly available for download from ImmPort (24) and consisted of RNA-Seq and matching CyTOF quantification of immune cells from PBMC samples. The MCP-counter publication used gene expression profiles from Gene Expression Omnibus (accession number GSE39582) and IHC counts of CD3⁺, CD8⁺, and CD68⁺ cells (available on request from the authors; ref. 11). The TIMER benchmark consisted of hematoxylin and eosin (H&E) stained slides from TCGA bladder urothelial carcinoma study. A pathologist manually reviewed each of these slides to categorize each sample into one of three categorical levels for neutrophil abundance: "Low," "Medium," or "High"; estimations are available from the TIMER online resource (12). For all benchmarking experiments, except TIMER, concordance was measured using correlation between "ground truth" values and the immune estimations of each method. Because of the variation in the degree of specificity to which cell subsets were defined, summations of subsets were required to allow accurate comparisons in some cases (Supplementary Table S1D). For the TIMER benchmark, samples were grouped by low, medium, and high pathologic estimation and then TME method estimates were compared by ANOVA with Tukey *post hoc*.

Results

Consensus cell-type-specific gene supersets for estimating TME cell populations

Following the generation of large datasets of tumor genomic profiles such as TCGA and the International Cancer Genome

Consortium, various approaches to assessing TME cell populations have been developed, each using different algorithms, gene markers, and validation benchmarks (2). To build on the knowledge of cell-type-specific gene sets represented in the diversity of these TME cell estimation methods, we sought an integrative strategy that incorporates knowledge from existing signatures and statistical approaches. The Consensus^{TME} method integrates cell-type-specific gene markers from independent TME cell estimation methods and uses ssGSEA to compute relative TME cell-type- and tumor-specific enrichment scores from bulk expression data (Fig. 1A). The first step of the Consensus^{TME} approach is the selection of publicly available signatures or tools, which for the current version included four tools: CIBERSORT, TIMER, MCP-counter, xCell, and three gene sets curated and used by Bindea and colleagues, Danaher and colleagues, and Davoli and colleagues. Second, cell types are selected when at least two methods estimate their abundance. Third, we generated a gene set for each cell type by using the union of genes used by the methods to estimate that cell type. Fourth, we removed genes that correlate ($\rho > -0.2$) with tumor purity in a tumor-type-specific manner as applied in the TIMER method (see Materials and Methods; ref. 12). Fifth, for statistical framework, the ssGSEA approach was selected because gene set enrichment frameworks treat microarray and RNA-seq values in the same way, being based on the ranked genes rather than the mRNA transcript abundance values. In addition, ssGSEA outperforms other gene set enrichment calculations with the Consensus^{TME} gene sets (Supplementary Fig. S1A). Finally, the output from Consensus^{TME} is NES scores, which accounts for variations in gene set size (14) and captures the relative level of estimated abundance of specific cell types across samples. For example, in a sample with a high NES score for a specific cell type, the individual genes of the gene set would rank higher in a sorted list of the abundance level of all genes in the transcriptome compared with samples with low NES score of the same cell type. In sum, to improve performance, Consensus^{TME} aggregates cell-type-specific genes that have been independently considered relevant by different methods and estimates their relative abundance in a tumor-type-specific manner (Supplementary Fig. S1B).

Pan-cancer leukocyte and lymphocyte TCGA benchmarks

To benchmark the different methods in an objective and systematic manner, we used publicly available data from multiple tumor types comprising 9,142 samples in total. All immune estimation methods were utilized setting appropriate parameters but without optimization requiring changes to source code. This included running CIBERSORT absolute mode and TIMER using the appropriate cancer type for the dataset. For each of the Bindea and colleagues, Danaher and colleagues, and Davoli and colleagues gene sets, ssGSEA was applied with the same parameters as Consensus^{TME} (see Materials and Methods). First, to evaluate the ability of each TME cell estimation method to capture the overall amount of immune component in the TME, we calculated total immune scores independently with each method. Briefly, when immune scores were not calculated by default by a method, we customarily derived it by creating a unique union of all genes either for gene set enrichment methods and adding that as an additional signature or for regression-based approaches summing the coefficients from all cell types (see Materials and Methods). Because tumor purity accounts not only for immune cell infiltration but also for other stromal cells (e.g., fibroblasts and endothelial cells; ref. 25), we inferred stromal non-immune-related content of all

samples using ESTIMATE (22) and added this value to each of the TME estimation methods' immune scores to create an RNA-based immune- and stromal composition score (TME score). We found that all seven methods and Consensus^{TME} perform very similar to each other as estimated by the correlation (Kendall's correlation coefficient, τ) between the RNA-based TME score calculated from each of the methods and the DNA-derived tumor purity score based on ABSOLUTE (Fig. 1B; refs. 17, 21). As expected, in nearly all comparisons (all methods by all cancer types), there was a negative correlation between the TME score and the DNA-based tumor purity score. Across all cancer types (pan-cancer), we found that CIBERSORT, Consensus^{TME}, and Danaher performed as the top three TME estimation methods. Furthermore, using this specific correlation measure, all the methods performed well overall with many correlations being statistically significant, but there were some cancer types in which all methods performed poorly, for example, pancreatic adenocarcinoma. Variation in performance was largely independent of cancer cellularity, mutation load, leukocyte fraction, and sample size for any of the methods.

We further evaluated the performance of the TME estimation methods by using leukocyte fractions derived from DNA methylation data for 30 tumor types (see Materials and Methods; ref. 21). The leukocyte DNA methylation data, which are not a gold standard in itself, were used as an orthogonal inference for leukocyte infiltration across tumor types studied by the TCGA consortium. To assess the performance of each of the methods using leukocyte methylation data and to account for accuracy across multiple cell types, we fitted multiple linear regression models using the leukocyte fraction as a response variable and only cell-type estimates in the category of being leukocytes as explanatory variables for each method (Supplementary Table S1C). As there are many cell types classed as leukocytes, multiple linear regression was chosen as it allows us to assess how variation across all cell types in combination could explain the variation in the methylation-derived leukocyte fraction. Because different methods estimate different number of leukocytes, the coefficient of determination can be artificially increased by the number of variables in a model (i.e. overfitting). Thus, to more appropriately compare the models in an unbiased way, we used adjusted coefficients of determination (R^2), the AIC, and BIC for model comparison. These penalize model complexity (i.e., number of cell types used in the models) to varying degrees, with BIC putting the greatest emphasis on finding a parsimonious model. When comparing the R^2 of the different models, the best performing methods were Bindea, Consensus^{TME}, and Danaher (Fig. 1C). Similarly, AIC and BIC scores showed that Consensus^{TME}, Davoli, Danaher, and Bindea models perform better than the other methods (models with lower AIC and BIC values are preferred).

Similarly, we implemented multiple linear regression analysis using tumor-infiltrating lymphocyte counts derived from digitized H&E-stained images analyzed through a deep-learning convolutional neural network approach (23). The methods that showed better fit metrics were Consensus^{TME}, Davoli, Bindea, and Danaher (Fig. 1D). Together, these broad pan-cancer benchmarks show a variation in the performance of the different methods when compared with each other, and no single method consistently outperforms the others. There was also cancer-specific performance variation observed that should be taken into account when considering the appropriateness of using these TME estimation methods.

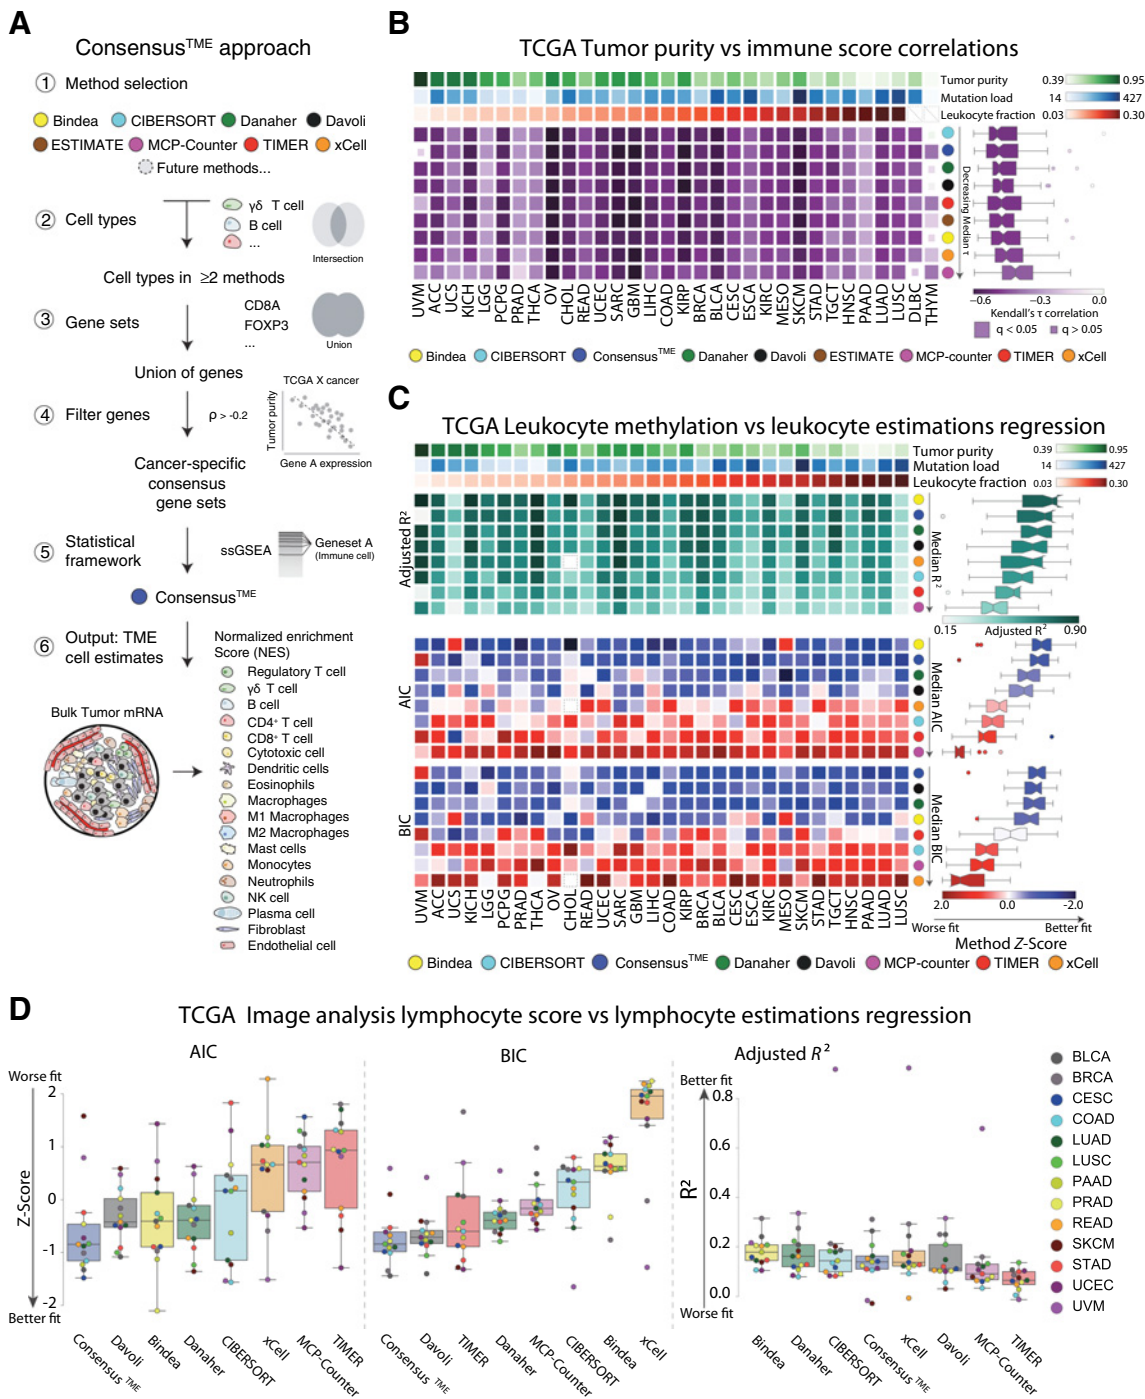


Figure 1. Benchmark of methods for estimating TME cell components using purity and leukocyte DNA methylation data from TCGA data. **A**, Consensus^{TME} development strategy (see Materials and Methods). **B**, Each heatmap square represents Kendall's correlation coefficients (τ) of DNA-derived ABSOLUTE tumor purity scores (17, 21) and RNA-derived TME score estimated by the different methods for the named TCGA cancer type. Box plots represent each method's performance across cancer types. **C**, Multiple linear regression models of leukocyte methylation scores as response variable (21) and RNA-derived leukocyte estimations as explanatory variables (Supplementary Table SIC). Column heatmaps are sorted according to median performance (top to bottom, decreasing performance). **D**, Multiple linear regression models of deep-learning H&E-derived lymphocyte counts as response variable (23) and RNA-derived lymphocyte estimations as explanatory variables (Supplementary Table SIC). Adjusted R^2 , AIC z-score, and BIC z-score were compared across models generated by each method's cell-type estimation. Lower AIC and BIC values represent a better goodness of fit penalizing the number of variables. Mutation load, median number of single-nucleotide variants (SNV). Tumor purity, DNA-derived ABSOLUTE cancer cell fraction. Leukocyte fraction, leukocyte cell fraction based on methylation data. Bar plots are sorted according to median correlation coefficient (left to right, decreasing performance).

Performance evaluation based on cell-type-specific datasets

All TME estimation methods tested, except Bindea and Davoli, performed their own independent benchmarks using experimental data with cell-type-specific measurements in the original publications. We collected benchmarking data for CIBERSORT, xCell, TIMER, and MCP-counter to carry out a side-by-side comparison. We used the CIBERSORT benchmark data that consisted of PBMCs of 20 healthy human subjects quantified by flow cytometry (10). Correlations between estimated immune cell types and the flow cytometry fractions showed that the best performing methods were MCP-counter, CIBERSORT, and xCell (Supplementary Fig. S1C). However, most of the correlations lacked statistical significance, and due to the different cell types estimated by the different methods, it is difficult to reach a conclusion. Similarly, the xCell benchmark dataset consisted of 16 PBMC leukocyte subsets from two different studies with 61 and 104 healthy human subjects each, where PBMCs fractions were measured using CyTOF (13). MCP-counter, CIBERSORT, and xCell were the methods that showed best performance in these PBMC benchmarks; however many cell types did not reach statistical significance (Supplementary Fig. S1C).

Finally, we used cancer-related benchmarks from TIMER (12) and MCP-counter (11). For TIMER's benchmark, 404 TCGA bladder cancer samples were analyzed by a pathologist who categorized them as low, medium, or high according to their neutrophil counts using H&E-stained slides. Performance was assessed by measuring the significance of difference between the computational estimates of samples in each category. Here, Consensus^{TME}, Bindea, and TIMER obtained the best separation between categories, but only Consensus^{TME} and Bindea separated significantly the three categories after multiple test correction (Fig. 2A). Interestingly, xCell, CIBERSORT, and MCP-counter did not differentiate the three categories significantly, whereas Davoli does not estimate neutrophils. For MCP-counter's benchmark, IHC digital quantification of CD3⁺ (T cells), CD8⁺ (CD8⁺ T cells), and CD68⁺ (monocytic lineage) cell densities was analyzed from 38 colorectal cancer samples. Correlations between the TME methods' estimations and the cellular fractions were computed (Fig. 2B). Consensus^{TME}, MCP-counter, and Danaher provided the best correlations, with Consensus^{TME} performing best on the three cell types evaluated.

When observing the rank of methods across all benchmarking experiments (Fig. 2C), no one method was shown to consistently outperform all others. However, the integrative Consensus^{TME} approach was in the top three for all cancer-based benchmarks and achieved the best mean rank of all TME estimation methods.

Discussion

With the recent generation of large publicly available molecular profiling of cancer samples, a variety of computational methods for analysis of cell components of the TME have been generated. In principle, the method of choice should be based on performance; however, the popularity and ease of use can also be reasons behind the specific method that researchers tend to select (5). In the case of TME cell estimation from bulk expression data, this problem is magnified by the lack of objective and independent benchmark analyses, because most methods use their own benchmarks that may introduce biases and reliance on one type of data. Here, we performed an unbiased and objective benchmarking exercise comparing seven of the most

widely used and recent TME estimation methods and also developed Consensus^{TME}, a gene set enrichment-based method that integrates cells types and genes from these seven different methods in order to generate a consensus gene superset for each cell type. The novelty of Consensus^{TME} is that it (a) draws on the current knowledge of cell-specific gene by integrating gene sets from publicly available methods, (b) provides a new TME cell estimation method that consistently performs well with the best mean rank across all benchmarks, and (c) will evolve and improve as new methods and data sets are published. We also provide a Web resource (www.consensustme.org) for evaluating current and future TME cell estimation methods across a diverse range of datasets allowing researchers to make informed decisions about the strengths and weaknesses of the different approaches to immune estimation.

The Consensus^{TME} approach is an evolvable method by design, meaning as new gene sets and signature matrices become available, they can be added to existing supersets and tested with already established benchmarks, thus potentially improving its performance as new methods and gene sets are developed. The "wisdom of the crowd" phenomenon, which the collective knowledge of a group supersedes that of individuals, is an approach that has previously been used to address complex problems in computational biology (26). With the complexity and diversity of the TME that also varies across different anatomical locations, the derivation of gene sets to capture this is a challenge that is well suited to this collaborative approach. Although this may lead to the inclusion of spurious genes in the supersets, the collection of multiple gene signatures can better capture the diversity in the transcriptome of immune cells in across different biological contexts. To combat this and to make gene sets specific for each cancer type, we employed a purity-based gene selection in Consensus^{TME}, similar to that used by TIMER (see Materials and Methods; ref. 12). Future versions may deploy alternative approaches such as gene exclusion or inclusion lists or weighting schemes, particularly as single-cell sequencing data of tumors become available, which will make it possible to derive cancer-specific immune and stromal gene signatures directly. To allow dissemination of Consensus^{TME}, we developed an easy to implement R package.

Currently, Consensus^{TME} leverages an ssGSEA statistical framework, but the compartmentalized approach allows for gene sets to be used with other statistical frameworks where appropriate. Although our benchmark of other available frameworks showed ssGSEA to be the most robust and accurate, future development on gene set enrichment algorithms may produce a new gold standard in the future (Supplementary Fig. S1A). Different frameworks can be especially important when considering the intended use of the output. Consensus^{TME} NES scores are imputed primarily for comparison of scores relatively across samples. Similar to another ssGSEA-based method, xCell, Consensus^{TME}, and NES scores also resemble absolute scores, which would allow comparison across cell types; however, currently this has not been benchmarked (13). Regression-based methods including CIBERSORT are better suited for solving this problem of comparing across cell types but require more careful interpretation when comparing across samples because outputs are fractional.

We performed pan-cancer benchmarks using orthogonal data types generated for TCGA samples. Although DNA-derived tumor purity scores correlated negatively with RNA-derived TME estimations for all methods, leukocyte methylation scores showed some

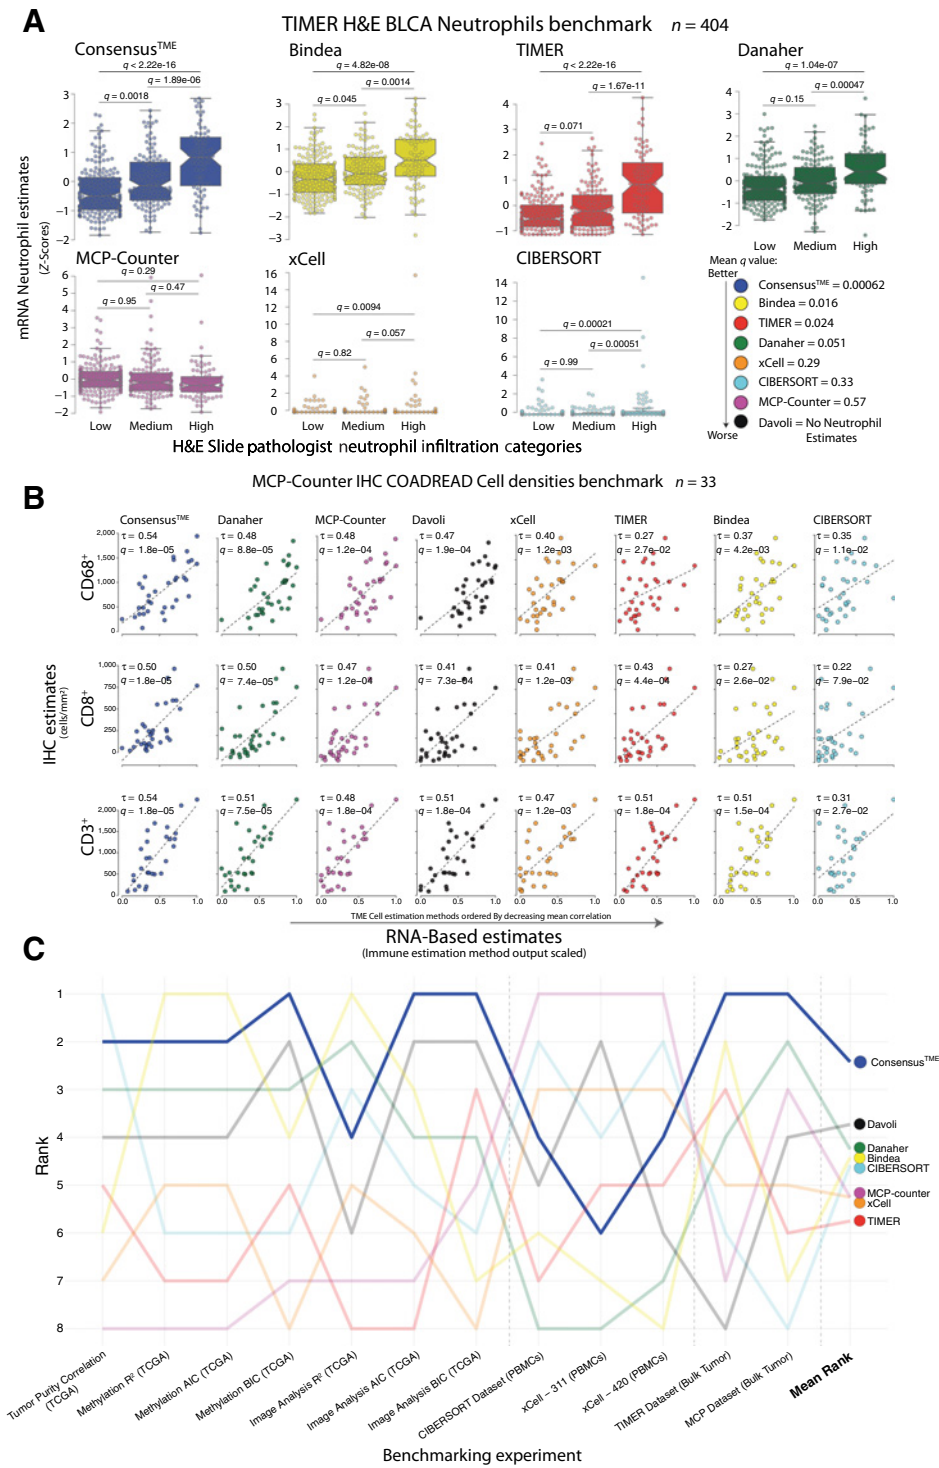


Figure 2. Cell-type-specific benchmark using original benchmarking datasets published by the individual methods. **A**, Comparison between low, medium, and high categories of BLCA ($n = 404$ samples) neutrophil H&E pathology counts from the TIMER benchmark data. One-way ANOVA with Tukey HSD *post hoc* tests were employed to calculate q values. **B**, Kendall's correlation coefficients (τ) of MCP-counter COADREAD IHC ($n = 38$ samples) cell densities (cell/mm²) against TME methods estimates for CD8⁺ T cells (CD8⁺), T cells (CD3⁺), and macrophage/monocytes (CD68⁺). Plots are sorted according to median correlation coefficient (left to right, decreasing performance). **C**, Overview of TME cell estimation methods across all benchmarking experiments, Consensus^{TME} method highlighted. Mean rank of cell estimations across all benchmarking experiments shown in final column.

variation across methods and a lack of correlation for some methods. Also, different tumor types showed varying levels of concordance with leukocyte estimations that were irrespective of features of the cancer type, such as average tumor cancer cell fraction, median mutation load, or average leukocyte fraction. The methylation-derived leukocyte fraction was generated by Hoadley and colleagues who compared the methylation patterns

of pure leukocyte cells and normal tissue methylation patterns before selecting regions of differential methylation between healthy or cancer cells and immune cells (21). This approach has the potential to be complementary to bulk RNA-based deconvolution (21, 27, 28). Lymphocyte deep-learning H&E quantifications provided a lower association with RNA-derived lymphocyte estimations; an observation that has been reported previously and

considered to be in part due to the RNA-derived estimates reflects more cell counts, whereas spatial image-derived estimates reflect the fraction of lymphocytes per area (23, 27). Thus, this benchmark should be used with caution due to the uncertainty of both RNA-derived and image-based derived lymphocyte estimates. However, similar performance results to the leukocyte methylation benchmark were observed.

Cell-type-specific benchmarking on the PBMC datasets showed that although MCP-counter and CIBERSORT performed best on PBMCs, in general a low number of significant correlations were obtained across methods. Because of the diversity of cells tested and estimated by the different methods, the introduction of potential biases was unavoidable, for example, a method predicting only highly abundant or transcriptionally distinct cell types would likely show good correlations but may not be as informative as a method providing estimates for a wider range of cell types. Thus, obtaining a concluding result out of the PBMC benchmarks was challenging, although the interactive portal can be used to compare how methods performed on a cell-by-cell basis. Moreover, given the gene expression of PBMCs in the circulation of healthy individuals is different to that of tumor-infiltrating lymphocytes, for the application of TME cell estimation using bulk RNA tumor data, PBMC benchmarks may be less informative (29, 30). In contrast, benchmarks using bulk tumor gene expression from both bladder and colorectal cancer tumors for neutrophils, CD3⁺, CD68⁺, and CD8⁺ cells showed significant associations for the majority of TME estimation methods, particularly Consensus^{TME}. Together, these benchmarks showed that although no one method consistently outperforms others, Consensus^{TME} was unique in ranking among the top three best performing methods in all cancer-related benchmarks. Our comprehensive benchmarking on diverse pan-cancer and cell-type-specific data complements recently published efforts in evaluating TME cell estimation methods focused mainly on specific cancer types (melanoma and ovarian) and scRNA-seq data (6).

In this work, we have provided a framework for a set of objective and unbiased benchmarks for immune and stromal cell estimation and deconvolution methods. Although each benchmark individually suffers from limitations, the combination of experiments allows a broad and informed assessment of the accuracy of current methods. The constant evolution within

the field of deconvolution led us to develop a shiny app for exploration of results (www.consensustme.org). Rather than these benchmarks providing a temporal snapshot of performance of TME cell estimation methods, this resource will ensure that new methods and datasets will be benchmarked and integrated into a common framework and the results can be explored in an interactive environment. Through use of this benchmarking methodology, we show that the Consensus^{TME} approach provides more accurate and robust cell-type estimates than any of the existing methods individually for the estimation of immune cell subtypes in the setting of bulk tumors.

Disclosure of Potential Conflicts of Interest

No potential conflicts of interest were disclosed.

Authors' Contributions

Conception and design: A. Jiménez-Sánchez, M.L. Miller

Development of methodology: A. Jiménez-Sánchez, O. Cast

Acquisition of data (provided animals, acquired and managed patients, provided facilities, etc.): A. Jiménez-Sánchez, O. Cast

Analysis and interpretation of data (e.g., statistical analysis, biostatistics, computational analysis): A. Jiménez-Sánchez, O. Cast

Writing, review, and/or revision of the manuscript: A. Jiménez-Sánchez, O. Cast, M.L. Miller

Administrative, technical, or material support (i.e., reporting or organizing data, constructing databases): O. Cast

Study supervision: A. Jiménez-Sánchez, M.L. Miller

Other (funding acquisition, resources): M.L. Miller

Acknowledgments

A. Jiménez-Sánchez was supported by a doctoral fellowship from the Cancer Research UK Cambridge Institute and the Mexican National Council of Science and Technology (CONACyT). O. Cast and M.L. Miller were supported by the Brown Performance Group, Innovation in Cancer Informatics Discovery Grant (BD523775). M.L. Miller was supported by Cancer Research UK core grant (C14303/A17197) and the Target Ovarian Cancer Translational Project Grant (Cambridge-1320 MM18). They also acknowledge the publishing authors of each of the TME methods for sharing their benchmarking datasets.

The costs of publication of this article were defrayed in part by the payment of page charges. This article must therefore be hereby marked *advertisement* in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

Received November 15, 2018; revised September 2, 2019; accepted October 17, 2019; published first October 22, 2019.

References

- Hanahan D, Weinberg RA. Hallmarks of cancer: the next generation. *Cell* 2011;144:646–74.
- Finotello F, Trajanoski Z. Quantifying tumor-infiltrating immune cells from transcriptomics data. *Cancer Immunol Immunother* 2018;67:1031–40.
- Li B, Liu JS, Liu XS. Revisit linear regression-based deconvolution methods for tumor gene expression data. *Genome Biol* 2017;18:127.
- Newman AM, Gentles AJ, Liu CL, Diehn M, Alizadeh AA. Data normalization considerations for digital tumor dissection. *Genome Biol* 2017;18:128.
- Zheng S. Benchmarking: contexts and details matter. *Genome Biol* 2017;18:129.
- Sturm G, Finotello F, Petitprez F, Zhang JD, Baumbach J, Fridman WH, et al. Comprehensive evaluation of transcriptome-based cell-type quantification methods for immuno-oncology. *Bioinformatics* 2019;35:i436–45.
- Bindea G, Mlecnik B, Tosolini M, Kirilovsky A, Waldner M, Obenauf AC, et al. Spatiotemporal dynamics of intratumoral immune cells reveal the immune landscape in human cancer. *Immunity* 2013;39:782–95.
- Davoli T, Uno H, Wooten EC, Elledge SJ. Tumor aneuploidy correlates with markers of immune evasion and with reduced response to immunotherapy. *Science* 2017;355. doi: 10.1126/science.aaf8399.
- Danaher P, Warren S, Dennis L, D'Amico L, White A, Disis ML, et al. Gene expression markers of tumor infiltrating leukocytes. *J Immunother Cancer* 2017;5:18.
- Newman AM, Liu CL, Green MR, Gentles AJ, Feng W, Xu Y, et al. Robust enumeration of cell subsets from tissue expression profiles. *Nat Methods* 2015;12:453–7.
- Becht E, Giraldo NA, Lacroix L, Buttard B, Elarouci N, Petitprez F, et al. Estimating the population abundance of tissue-infiltrating immune and stromal cell populations using gene expression. *Genome Biol* 2016;17:218.
- Li B, Severson E, Pignon J-C, Zhao H, Li T, Novak J, et al. Comprehensive analyses of tumor immunity: implications for cancer immunotherapy. *Genome Biol* 2016;17:174.
- Aran D, Hu Z, Butte AJ. xCell: digitally portraying the tissue cellular heterogeneity landscape. *Genome Biol* 2017;18:220.

14. Barbie DA, Tamayo P, Boehm JS, Kim SY, Moody SE, Dunn IF, et al. Systematic RNA interference reveals that oncogenic KRAS-driven cancers require TBK1. *Nature* 2009;462:108–12.
15. Subramanian A, Tamayo P, Mootha VK, Mukherjee S, Ebert BL, Gillette MA, et al. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proc Natl Acad Sci U S A* 2005;102:15545–50.
16. Hänzelmann S, Castelo R, Guinney J. GSEA: gene set variation analysis for microarray and RNA-seq data. *BMC Bioinform* 2013;14:7.
17. Carter SL, Cibulskis K, Helman E, McKenna A, Shen H, Zack T, et al. Absolute quantification of somatic DNA alterations in human cancer. *Nat Biotechnol* 2012;30:413–21.
18. Hu W, Pasare C. Location, location, location: tissue-specific regulation of immune responses. *J Leukocyte Biol* 2013;94:409–21.
19. Cerami E, Gao J, Dogrusoz U, Gross BE, Sumer SO, Aksoy BA, et al. The cBio cancer genomics portal: an open platform for exploring multidimensional cancer genomics data. *Cancer Discov* 2012;2:401–4.
20. Li B, Dewey CN. RSEM: accurate transcript quantification from RNA-Seq data with or without a reference genome. *BMC Bioinform* 2011;12:323.
21. Hoadley KA, Yau C, Hinoue T, Wolf DM, Lazar AJ, Drill E, et al. Cell-of-origin patterns dominate the molecular classification of 10,000 Tumors from 33 types of cancer. *Cell* 2018;173:291–304.
22. Yoshihara K, Shahmoradgoli M, Martínez E, Vegesna R, Kim H, Torres-García W, et al. Inferring tumour purity and stromal and immune cell admixture from expression data. *Nat Commun* 2013;4:2612.
23. Saltz J, Gupta R, Hou L, Kurc T, Singh P, Nguyen V, et al. Spatial organization and molecular correlation of tumor-infiltrating lymphocytes using deep learning on pathology images. *Cell Rep* 2018;23:181–93.
24. Bhattacharya S, Andorf S, Gomes L, Dunn P, Schaefer H, Pontius J, et al. ImmPort: disseminating data to the public for the future of immunology. *Immunol Res* 2014;58:234–9.
25. Binnewies M, Roberts EW, Kersten K, Chan V, Fearon DF, Merad M, et al. Understanding the tumor immune microenvironment (TIME) for effective therapy. *Nat Med* 2018;24:541–50.
26. Marbach D, Costello JC, Küffner R, Vega NM, Prill RJ, Camacho DM, et al. Wisdom of crowds for robust gene network inference. *Nat Methods* 2012;9:796–804.
27. Thorsson V, Gibbs DL, Brown SD, Wolf D, Bortone DS, Ou Yang T-H, et al. The immune landscape of cancer. *Immunity* 2018;48:812–30.
28. Chakravarty A, Furness A, Joshi K, Ghorani E, Ford K, Ward MJ, et al. Pan-cancer deconvolution of tumour composition using DNA methylation. *Nat Commun* 2018;9:3220.
29. Baine MJ, Chakraborty S, Smith LM, Mallya K, Sasson AR, Brand RE, et al. Transcriptional profiling of peripheral blood mononuclear cells in pancreatic cancer patients identifies novel genes with potential diagnostic utility. *PLoS One* 2011;6:e17014.
30. Sakai Y, Honda M, Fujinaga H, Tatsumi I, Mizukoshi E, Nakamoto Y, et al. Common transcriptional signature of tumor-infiltrating mononuclear inflammatory cells and peripheral blood mononuclear cells in hepatocellular carcinoma patients. *Cancer Res* 2008;68:10267–79.