Uncovering Predictive Gene and Cellular Signatures for Checkpoint Immunotherapy Response through Machine Learning Analysis of Immune Single-Cell RNA-seq Data

Asaf Pinhasi[1], Keren Yizhak[1,2,*]

[1]Department of Cell Biology and Cancer Science, The Ruth and Bruce Rappaport Faculty of Medicine, Technion - Israel Institute of Technology, Haifa 3525422, Israel

[2]The Taub Faculty of Computer Science, Technion - Israel Institute of Technology, Haifa 3200003, Israel

[*] Corresponding author: kyizhak@technion.ac.il

Abstract

Background

Immune checkpoint inhibitors have revolutionized cancer therapy by harnessing the body's immune system to eliminate tumor cells. However, only a subset of patients responds to treatment. Understanding why only some patients respond remains a critical challenge in cancer research due to the complexity and variability of cellular composition within the tumor microenvironment. Our study presents PRECISE (*Predicting therapy Response through Extraction of Cells and genes from Immune Single-cell Expression data*), a pipeline that leverages single-cell RNA sequencing data and machine learning techniques to predict ICI responses, while maintaining the richness of single-cell information and ensuring interpretability of the results.

Results

The PRECISE pipeline implements gene and cell-filtering approaches for optimizing treatment prediction. Utilizing the XGBoost algorithm for predicting patient response to ICI on a dataset of melanoma-infiltrated immune cell achieved an initial Area Under the Curve (AUC) of 0.84. This signal is further improved to 0.89 after Boruta feature selection, revealing an 11-gene predictive signature. Investigation of these genes through SHAP values identified various gene-pair interactions with non-linear conditional effects on predictions. Furthermore, a novel reinforcement learning framework implemented in PRECISE reveals non-predictive single cells that detracts the model's performance. Altogether, the identified gene- and cell-based signatures demonstrates high prediction power across independent datasets, including lung, breast, brain, and skin cancers.

Conclusions

Our approach demonstrates the potential of supervised machine learning and reinforcement learning to enhance the understanding of cancer immunity and improve the prediction of treatment responses using single-cell data.

Understanding patients' response to immune checkpoint inhibitors (ICIs) is a critical challenge in cancer research, given the complexity and variability of immune interactions within the tumor microenvironment. Our study leverages single-cell RNA sequencing data and aims to utilize machine learning techniques to predict immunotherapy responses, while maintaining the richness of single-cell information and ensuring interpretability of the results. Using a dataset of melanoma immune cells, we conducted thorough preprocessing and applied the XGBoost algorithm in a leave-one-out cross-validation fashion to predict sample response. By labeling cells according to their sample's response and aggregating predictions, we achieved an initial AUC score of 0.84. This score was improved to 0.89 with the application of Boruta feature

selection, which identified key predictive genes, leading to an 11-gene predictive signature. Further analysis pinpointed T cell clusters as significant contributors to immune response. Utilizing SHAP values provided deeper insights into gene behaviors, interactions, and their effects on the model's predictions. Additionally, a novel reinforcement learning model was developed for single-cell level prediction and characterization, allowing us to identify and analyze the most predictive cells for response and non-response. Our approach demonstrates the potential of sophisticated computational methods to enhance our understanding of cancer immunity and improve the prediction of treatment responses.

Introduction

Cancer immunotherapy, particularly immune checkpoint inhibitors (ICI), has revolutionized cancer treatment, offering durable responses in a subset of patients. ICI, such as anti-PD1 antibodies, work by reinvigorating exhausted T cells within the tumor microenvironment (TME), allowing the immune system to act effectively against tumor cells (1). Despite their success, the majority of patients fail to respond to these therapies (2), underscoring the need for reliable predictive biomarkers to identify patients who will likely benefit from treatment. Many efforts aim to decipher the behavior of the immune system against the tumor as a result of treatment (3). In particular, characterizing the immune system in relation to different ICI responses and finding robust biomarkers to differentiate between responders and non-responders has become a major goal in the field of cancer immunity, often focusing on different cells in the TME (4–6). Nevertheless, the complexity of immune interactions, and the context dependent behavior of the immune system to treatment are main obstacles in the endeavor for enhancing the efficacy of immunotherapy treatments.

67 Single-cell RNA sequencing (scRNA-seq) provides unparalleled insights into the cellular
68 heterogeneity and dynamic states within the TME, capturing the complexity of immune
69 responses at a granular level. However, despite the huge advances in the field of machine
70 learning (ML), integrating this high-dimensional data with advanced models and techniques to
71 predict therapeutic outcomes remains a challenge. Most existing approaches either oversimplify
72 the data, losing valuable single-cell resolution (7,8), or fail to provide interpretable results,
73 limiting their clinical applicability (9–11). Bridging this gap requires the development of
74 innovative computational methods that retain data richness while offering clear insights into the
75 biology underpinning the decision making of the models.

76 While machine learning in general has been widely adopted in the field of single-cell analysis
77 (12), relatively few studies use supervised learning to directly predict response to
78 immunotherapy. For example, Dong et al. used cell-type proportion extracted from scRNA-seq
79 data for training ML models to predict response to ICI (8). In a different study (7), Kang et al.
80 used pseudo-bulk expression calculated from scRNA-seq. Prediction models with similar aims to
81 ours have been developed, such as CloudPred (13), which uses differentiable machine learning
82 to predict lupus, and ScRAT (14), which uses neural networks to predict COVID-19 phenotypes.
83 However, these models have not been applied in the context of immunotherapy response.

84 Applying ML to scRNA-seq data raises a few challenges and open questions: for instance, should
85 all single cells from a responding patient be labeled as 'responders', or should we use pre-
86 existing knowledge on cell function to label cells as either ''favorable' or 'unfavorable'? How can
87 we infer the patient class from the predicted classes of the associated single cells? Should we
88 use all cells for prediction or perhaps a subset that is more predictive than others?

89 In this study we addressed these various questions by investigating scRNA-seq data of *CD45*$^+$
90 cells from a cohort of ICI-treated melanoma patients (15). We developed **PRECISE** (*Predicting*
91 *therapy Response through Extraction of Cells and genes from Immune Single-cell Expression*
92 *data*), a machine learning framework designed to predict immunotherapy response by and
93 extract gene- and cell-level insights from scRNA-seq data. Using XGBoost (eXtreme Gradient
94 Boosting) as the ML model (16), our basic methodology involved labeling cells based on their
95 sample's response, training models in a leave-one-out manner, and then aggregating predictions
96 to create a sample-level score. For feature selection (FS) we used Boruta (17) which provided
97 high performance in predicting patient response (AUC = 0.89). Following, we dissected the
98 contribution of selected genes using feature importance scores and the more advanced Shapley
99 Additive exPlanations (SHAP) values (18), identifying complex, non-linear contribution of single
100 genes and gene pairs to treatment response. Considering the cells-axis, we identified specific
101 cell-types that are more informative than others in predicting response. Moreover, we devised a
102 novel reinforcement learning model designed to identify single cells that are either predictive or
103 non-predictive of patient response (19). Applying these gene- and cell-based signatures to
104 different datasets, we managed to achieve high prediction power across different cancer types.
105 Overall, our findings highlight the potential of advanced computational approaches in enhancing
106 our understanding of cancer immunity and informing treatment.

107

108 **Results**

109 **Base Model Feature Selection and Per Cell Type Prediction**

110 We first utilized our previously published (15) single-cell RNA sequencing (scRNA-seq) dataset
111 that consists of 16,291 CD45$^+$ cells from 48 tumor biopsies taken from 32 stage IV metastatic

112 melanoma patients treated with ICI. 11 patients had longitudinal biopsies, and 20 patients had
113 one or two biopsies taken either at baseline or during treatment. Tumor samples were classified
114 according to radiologic evaluations as either progression/non-responder (NR; n = 31, including
115 stable disease (SD) and progressive disease (PD) samples) or regression/responder (R; n = 17,
116 including complete response (CR) and partial response (PR) samples). After the removal of non-
117 informative genes (Methods), we constructed an XGBoost-based machine-learning model and
118 investigated its power in predicting the patients' response to ICI (Figure 1A).

119 One of the main challenges in using supervised machine learning with single-cell data is the
120 imbalance between the vast amounts of single-cells and the typically small number of samples
121 from which these cells are taken. Naturally, the prediction target is often at the sample/patient
122 level, raising the question of how best to assign labels to individual cells. To address this, we
123 adopted a two-step approach: we trained the model and made predictions at the single-cell
124 level, and then aggregated these predictions to devise a sample score. Each cell was labeled
125 according to the response status of its sample of origin, such that cells from responding samples
126 were classified as 'responders' and those from non-responding samples as 'non-responders'.
127 After predicting the labels for all cells in a sample, we calculated the proportion of cells
128 predicted as 'responders,' and this proportion constituted the sample score. The logic for this
129 scoring could be described as a tug-of-war between 'good' and 'bad' cells. The more 'good' cells
130 a sample contains, the more likely it is to be pulled towards a favorable response. Considering
131 the small number of samples, training and testing was conducted in a leave-one-out (LOO)
132 cross-validation manner, with each fold training the model on the cells from all samples except
133 one, which was excluded for prediction. This approach also mitigates the stochastic effects often
134 encountered in standard cross-validation.

135 As our reference model, we included all 16,291 cells and all 10,082 genes that passed quality
136 control (Methods). The ROC AUC score for this analysis was 0.84, showing reasonably high
137 predictive accuracy (Figure 1B). To identify genes that are important for predicting patient
138 response to ICI, and can therefore serve as potential biomarkers, we used the feature
139 importance metric, which measures the contribution of a feature to the prediction of the model
140 based on how often it was used to split the data in the ensemble of trees. To quantify the
141 importance of each gene in the cohort we extracted the importance score from each of the 48
142 folds in the LOO simulation, and averaged them across all folds (Figure 1C). This analysis
143 highlighted several genes as important for predicting patient response, including *GAPDH*, *IFI6*,
144 *LILRB4*, *GZMH*, *STAT1*, *LGALS1* and *CD68* (Supplementary Table 2), all of which have already
145 been found to relate to ICI and tumor immunity (20–25). For example, *CD68*, a marker for
146 macrophages, was found to be associated with poorer clinical outcomes in multiple cancer types
147 in a pan-cancer analysis of expression data from TCGA and GTEx (20). However, it is generally
148 linked to M1 macrophages, which promote anti-tumor immunity (26). *STAT1* was found to
149 promote immunotherapy response in melanoma (21), and was found to be positively correlated
150 with *PD-L1* expression in ovarian cancer (22). *LILRB4*, has been found to strongly suppress tumor
151 immunity, with studies in murine tumor models demonstrating that its blockade can alleviate
152 this suppression, resulting in antitumor efficacy (23). Finally, *GAPDH*, receiving the highest
153 feature importance score, has been shown to be involved in cancer progression and immune
154 response regulation, particularly through its role in hypoxia-related pathways (24,25).

155 Examining the sample scores revealed that while most samples received scores consistent with
156 their response status, some samples were mis-scored. For example, sample Pre_P3, a non-
157 responder, received a high score, whereas sample Pre_P28, a responder, received a low score
158 (Supplementary Table 1). These discrepancies could be due to model performance, but they

159 might also indicate specific traits of these samples that warrant further investigation.
160 Specifically, sample Pre_P3 had a mutation in *B2M* limiting antigen presentation, demonstrating
161 how factors not shown in the transcription are causing the tumors to respond or resist to
162 treatment. In addition, sample Pre_P28 was unique in its high abundance of cells from the
163 myeloid lineage, emphasizing the importance of the specific cells used for prediction. Following,
164 and after establishing the prediction power of the base model that considers all genes and all
165 cells, we investigated the model's performance when using specific subsets of these entities.

166 **Improving model performance by exploring cell and gene subsets**

167 To improve model performance and study the contribution of different cell types and states, we
168 first assigned individual cells to their respective cell type. We used the classification from the
169 original paper which included 11 clusters determined by unsupervised clustering, and 13 cell
170 types or states determined in a supervised manner using known gene markers (15). We then
171 applied the XGBoost model to each of the different cellular groups separately. As expected, the
172 per cell-type/cluster prediction resulted in different ROC AUC scores for each of the groups
173 tested, testifying for their overall importance in predicting patient response. The results (Figure
174 1D, E) show that T cells subsets achieve substantially higher scores in comparison to the rest of
175 the cell-types (0.87-0.82), in concordant with their known role in immune response in general,
176 and specifically in melanoma (27). B cells (0.64), Macrophages (0.61) and other cell types were
177 significantly less predictive. While the quantity of cells in each cellular groups is different and
178 can potentially affect the accuracy of the model, this does not seem to be the dominant factor
179 since there are T cell subsets that have a comparable number of cells to other cell-types, but
180 produce much higher AUC scores (Figure 1D). Intriguingly, the model's accuracy was consistently
181 higher for the supervised cell types compared to the unsupervised clusters (Supplementary

182 Figure 2), suggesting that the supervised clusters may capture more biologically meaningful
183 signals and thus exhibit stronger predictive power.

184 We next moved to investigate if and to what extent feature selection would improve the base
185 model predictions. To this end we explored 5 main methods for evaluating feature contribution
186 - top highly variable (HVG), top differentially expressed, Lasso regression-based FS, simple top
187 feature importance (Supplementary Table 3), and Boruta (Methods). Importantly, in all feature
188 selection techniques applied, except for HVG method which is independent of response, we
189 assured that selection was made independently from the predicted sample, thus avoiding any
190 train-test leakage.

191 One of the main advantages of Boruta over all other feature selection techniques is the
192 automatic determination of the number of features (17). While it is possible to tune Boruta
193 hyperparameters to change the number of features, the method works well without any tuning,
194 and offers robust feature number instead of an arbitrary choice. Therefore, although other
195 feature selection techniques could offer slightly better accuracy for a specific choice of gene
196 number or hyperparameters, Boruta was the best method for consistent and robust
197 performance. The AUC scores produced by the other methods were mostly in the range 0.81-
198 0.85, depending on the method and the hyperparameter, while Boruta produced a score of 0.89
199 (Figure 1F), manifesting its superiority in this selection.

200 Finally, to account for the observed variability in performance achieved for different cellular
201 groups and gene subsets, we used Boruta feature selection to extract the top most contributing
202 features on the base model and per cell-type, getting insight into the most important cell type-
203 specific genes (Supplementary Table 4). To this end, we ran Boruta in the same manner as
204 before, but this time focusing on each cell-type and cluster. As expected, the accuracy across

205 most cell-types increased (Figure 1D,G). Exploring the top genes in each group (Figure 1H,I), we
206 found genes with an established relation to cancer immunity and activation. For example, In NK
207 cells, *FCGR3B*, known as *CD16b*, is involved in antibody-dependent cellular cytotoxicity (ADCC)
208 (28,29). It was suggested that NK cells mediate ADCC-driven cytotoxicity specifically against *PD-*
209 *L1*-positive cells, demonstrating the role of ADCC in enhancing the effectiveness of
210 immunotherapies, particularly those targeting *PD-L1* (30). In B cells, *CXCR4* is expressed on the
211 surface of mature B cells and interacts with *CXCL12* to recruit regulatory B cells to the tumor
212 site. These cells can inhibit T cell activity, contributing to an immunosuppressive tumor
213 environment. *HLA-DRB1* and *HLA-DP1*, both part of the MHC-II complex, were found in $CD4^+$
214 helper T cells to be of the greatest importance. In addition, $CD38^+HLA\text{-}DR^+$ expressing $CD4^+$ and
215 $CD8^+$ T cells were significantly expanded following ICI (immune checkpoint inhibitor) treatment
216 in HNSCC (head and neck squamous cell carcinoma) (31). *GAPDH*, *CD38*, *STAT1* and *IFI6* all have
217 well established roles in T cells activation and immunotherapy, and therefore they appear
218 robust in most T cells subtypes (24,25,32–35). In addition to these general markers, more
219 specific markers were extracted, such as *PRF1* in cytotoxic *CD8* T cell, a key component of the
220 cytotoxic machinery of activated T cells, directly involved in killing tumor cells (36).

221

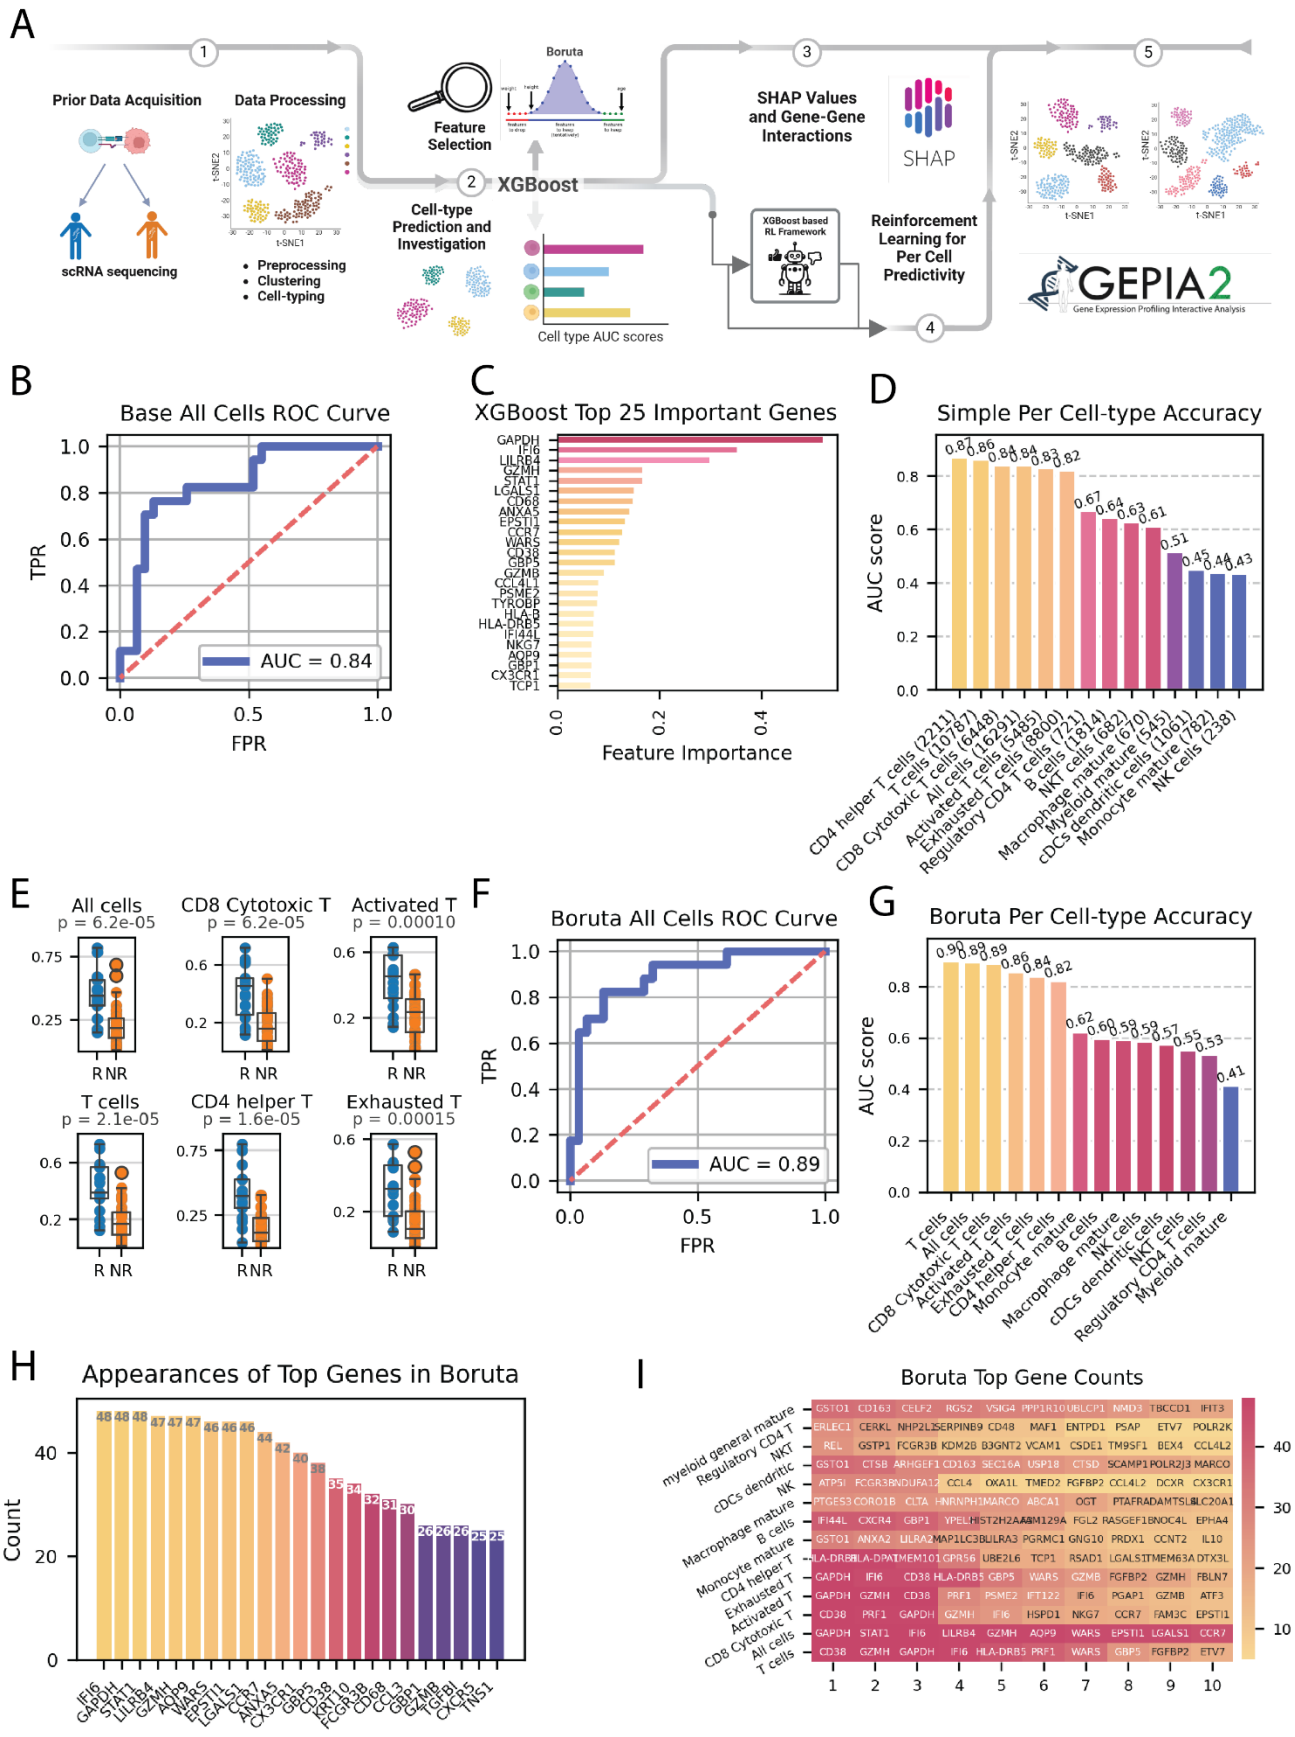
A
Prior Data Acquisition
scRNA sequencing
Data Processing
Preprocessing
Clustering
Cell-typing
Feature Selection
Boruta
XGBoost
Cell-type Prediction and Investigation
Cell type AUC scores
SHAP Values and Gene-Gene Interactions
SHAP
XGBoost based RL Framework
Reinforcement Learning for Per Cell Predictivity
GEPIA2
Gene Expression Profiling Interactive Analysis
B
Base All Cells ROC Curve
TPR
FPR
AUC = 0.84
C
XGBoost Top 25 Important Genes
Feature Importance
D
Simple Per Cell-type Accuracy
AUC score
E
All cells
p = 6.2e-05
CD8 Cytotoxic T
p = 6.2e-05
Activated T
p = 0.00010
T cells
p = 2.1e-05
CD4 helper T
p = 1.6e-05
Exhausted T
p = 0.00015
R NR
F
Boruta All Cells ROC Curve
TPR
FPR
AUC = 0.89
G
Boruta Per Cell-type Accuracy
AUC score
H
Appearances of Top Genes in Boruta
Count
I
Boruta Top Gene Counts

222

Figure 1: Prediction of ICI Response Using XGBoost and Boruta Feature Selection per Cell-Type

A Schematic workflow of the study. **B** ROC curve for the base model predicting response to ICI using all cells and genes in the cohort. **C** XGBoost feature importance bar-plot showing the top 30 most important genes. **D** AUC scores indicating the prediction accuracy of the base model across different immune cell types. **E** Box plots comparing the scores produced by base XGBoost between responders (R) and non-responders (NR) across top most accurate cell subtypes, with significant differences indicated by the Mann–Whitney U test. **F** AUC scores indicating the prediction accuracy of the Boruta selected model across different immune cell types. **G** ROC curve for the Boruta-selected model predicting response to ICI using all cells in the cohort. **H** Bar-plot of the number of occurrences of each gene in the Boruta selection across the LOO folds, showing the top most robust genes. **I** Heatmap displaying the top genes selected by Boruta for different immune cell types, showing number of occurrences of each gene for each cell type.

223

224 **SHAP Value Investigation Reveals Complex Gene Behaviors and Interactions**

225 While XGBoost's feature importance provides some insight into the most contributing genes, its
226 interpretability is limited. For instance, feature importance is non-directional, and thus does not
227 indicate whether a gene is more associated positively or negatively with response. To address
228 this, and considering that gene behavior may extend beyond simple linear or dichotomous
229 relationships, we employed SHAP (SHapley Additive exPlanations) values (18) (Methods). SHAP
230 values, derived from cooperative game theory, measure both the direction and magnitude of a
231 gene's effect on the predicted response (Figure 2 A,B), allowing visualization of how gene
232 expression influences it. Moreover, one of the strengths of SHAP values derives from its ability
233 to capture context dependent contributions of each gene, such that the SHAP value for a gene
234 can differ between two cells even if they share similar expression levels (Figure 2C).

235 After training an XGBoost model over all of the cells in the cohort, SHAP values were calculated
236 from that model, quantifying the contribution of each gene for each cell in the cohort. For
237 example, as illustrated in Figure 2C, *GAPDH*, and *STAT1*'s expressions are associated with non-

238 response. The directionality of the SHAP value is such that high SHAP values are associated with
239 response, and low SHAP values with non-response. However, the relationships of the two genes
240 differ: *GAPDH* exhibits a continuous relationship with a non-zero threshold separating
241 responders from non-responders, whereas *STAT1* shows an on-off behavior, with zero
242 expression linked to response and non-zero expression linked to non-response. Conversely,
243 other genes such as *HLA-G* (Supplementary Figure 3) were associated with response.
244 Additionally, SHAP reveals other, non-monotonic, gene behaviors. For example, *HSPA1A* (Figure
245 2C) shows a complex pattern where both very low and very high expression are associated with
246 response, while the middle range is more related to non-response. Another interesting example
247 is of *IFITM2* where zero expression is related to favorable response, the middle range has less
248 effect on the prediction and high *IFITM2* pushes the model towards non-response.

249 Going beyond the effects of single genes, we next investigated SHAP values for gene
250 interactions. Focusing on the top most important genes that were found in the FS section, we
251 examined their most interactive partners in predicting immunotherapy response (Figure 2D).
252 Notable pairs, showing a prominent conditional behavior in relation to SHAP, include *GAPDH-*
253 *STAT1*, *HLA-B-CCR7*, and *CD38-CCR5* (additional pairs are shown in Supplementary Figure 3). To
254 further simplify and enhance interpretability, we created decision trees for each gene pair.
255 These trees align well with interaction plots, quantifying expression thresholds that best
256 differentiate responses. As shown in Figure 2D, the interaction between *GAPDH* and *STAT1*
257 highlights that cells with low *GAPDH* (<11.1) and high *STAT1* (>2.45) are associated with non-
258 response, whereas low *GAPDH* with low *STAT1* is linked to response. High *GAPDH* is associated
259 with response regardless of *STAT1* expression. The *CCR7-HLA-B* pair show a different relation to
260 response, where the separation by *CCR7* is close to zero (<0.68), demonstrating an on-off
261 conditionality, where low *CCR7* is associated with non-response. In higher values of *CCR7*, *HLA-B*

262 best separates between outcomes with expression threshold of 11.3, where lower *HLA-B*
263 expression relates to response, and higher to non-response. Such interaction analyses constitute
264 a great way of modeling conditional effects of genes - relationships that are challenging to find
265 without advanced models like SHAP.

266 SHAP dependencies for the model that is trained on all cells, captures in part the cell-types and
267 sub cell-types expression patterns. Given that cells in different clusters exhibit distinct
268 expression profiles, a cluster associated with response or non-response is likely to reflect its
269 unique expression pattern in the dependence plot. For example, cluster G1 (B cells, n = 1445),
270 which is significantly more abundant in responders, shows a nearly 2.5-fold higher proportion of
271 low *GAPDH*, low *STAT1* cells compared to the entire cohort. In contrast, G6 and G11 (Exhaustion
272 clusters, n = 2222 and n = 1129) are significantly more abundant in non-responders.
273 Consistently, only a small fraction of these cells exhibits low *GAPDH* and *STAT1* expression, with
274 G6 showing a 3.3-fold lower abundance, and G11 a 20-fold lower abundance compared to its
275 proportion in the cohort. To explore intra-cell-type dependencies, we reduce the influence of
276 this effect by also focusing solely on T cells (Figure 2E).

277 The T cell-specific analysis revealed that in certain gene pairs, interactions became more
278 distinct. In particular, the separation of expression patterns was clearer in T cells for both *CCL5-*
279 *CD38* and *CCL4-HLA-B* (Figure 2E), highlighting the importance of these interactions within this
280 specific cell type. In the *CCL4-HLA-B*, the dependence is almost undetectable using the whole
281 cohort, but with the T cells this is one of the strongest dependencies. In other gene-pairs, such
282 as *GAPDH-STAT1*, the dependence is clearer in the whole cohort rather than in T cells, probably
283 due to the effect of cell-type bound expressions described above.

284 Interestingly, all the chemokines in this analysis—*CCL4*, *CCL5*, and *CCR7*— have an on-off
285 behavior. Non-zero expression, albeit low, of these chemokines is sufficient to dramatically alter
286 the behavior of other genes in relation to the immune response (Figure 2D,E). *CCL4* and *CCL5*
287 can promote anti-melanoma immune response (37), and are known to have a role in
288 recruitment of immune cells (38,39). They have also been found to promote tumor
289 development and progression by recruiting regulatory T cells to the tumor site (40). It is not
290 surprising then, that their presence or absence can significantly change the effect of other genes
291 on the immune response. This is especially true for *HLA-B*, which is critical for activating $CD8^+$ T
292 cells; the activation function can only be exerted following recruitment of these T cells,
293 facilitated by chemokines like *CCL4*.

294 To further elucidate the behavior of the model, we leveraged SHAP values' ability to deconvolve
295 the contributing factors to each prediction, allowing us to understand the model's choices. We
296 analyzed the model's performance to understand why certain samples were correctly or
297 incorrectly predicted. We focused on four samples: Pre_P2 and Pre_P35 that were correctly
298 predicted, and Pre_P3 and Pre_P28 that were misclassified. Using SHAP values, we trained the
299 XGBoost model on all cells except those from the target sample, mimicking the actual prediction
300 process. We then summed the contributions of all cells in each left-out-sample to obtain
301 aggregated scores for each gene (Methods). These scores, sorted by absolute value, reveal the
302 most critical genes for the model's predictions (Figure 2F). SHAP waterfall plots are commonly
303 used for interpreting a single prediction, but here, since each sample contains many prediction
304 targets (cells), the plots represent an aggregation of these values. Therefore, the scores for each
305 sample depict the average contribution of each gene to the prediction of all the cells in that
306 sample – showing a general trend rather than a single contribution value. For instance, Pre_P35
307 was classified correctly as a responder, primarily due to low *GAPDH* and HLA-B values, with

308 smaller contributions from several other genes. Pre_P2 was correctly classified as a non-
309 responder, mostly due to a high *GAPDH* expression, as well as a high *HLA-B* and *IFITM2*
310 expression and low *HLA-G* expression. Pre_P28 was misclassified as a non-responder due to the
311 high *HLA-C*, *TNFAIP3*, and *HLA*-B expression, coupled with low *HLA-G*. Conversely, Pre_P3 was
312 misclassified as a responder due to the average low expression of genes including *GAPDH*, *HLA-*
313 *B*, *HLA-C*, *FOS*, *STAT1*, *CD38*, *TSC22D3*, and *EPSTI1*. Although a low *HLA-G* score pushed the
314 model toward non-response, the cumulative effect of the rest of the genes was stronger.

315 Overall, the ability to interpret the model's performance and decision-making is crucial for
316 model improvement, gaining sample-level insights, and ensuring the model's applicability in a
317 clinical setting, which is the ultimate goal. Once the genes that affect the model’s prediction
318 either positively or negatively are identified, it is possible to examine the distribution of these
319 genes through the cohort and in the different cell-types.

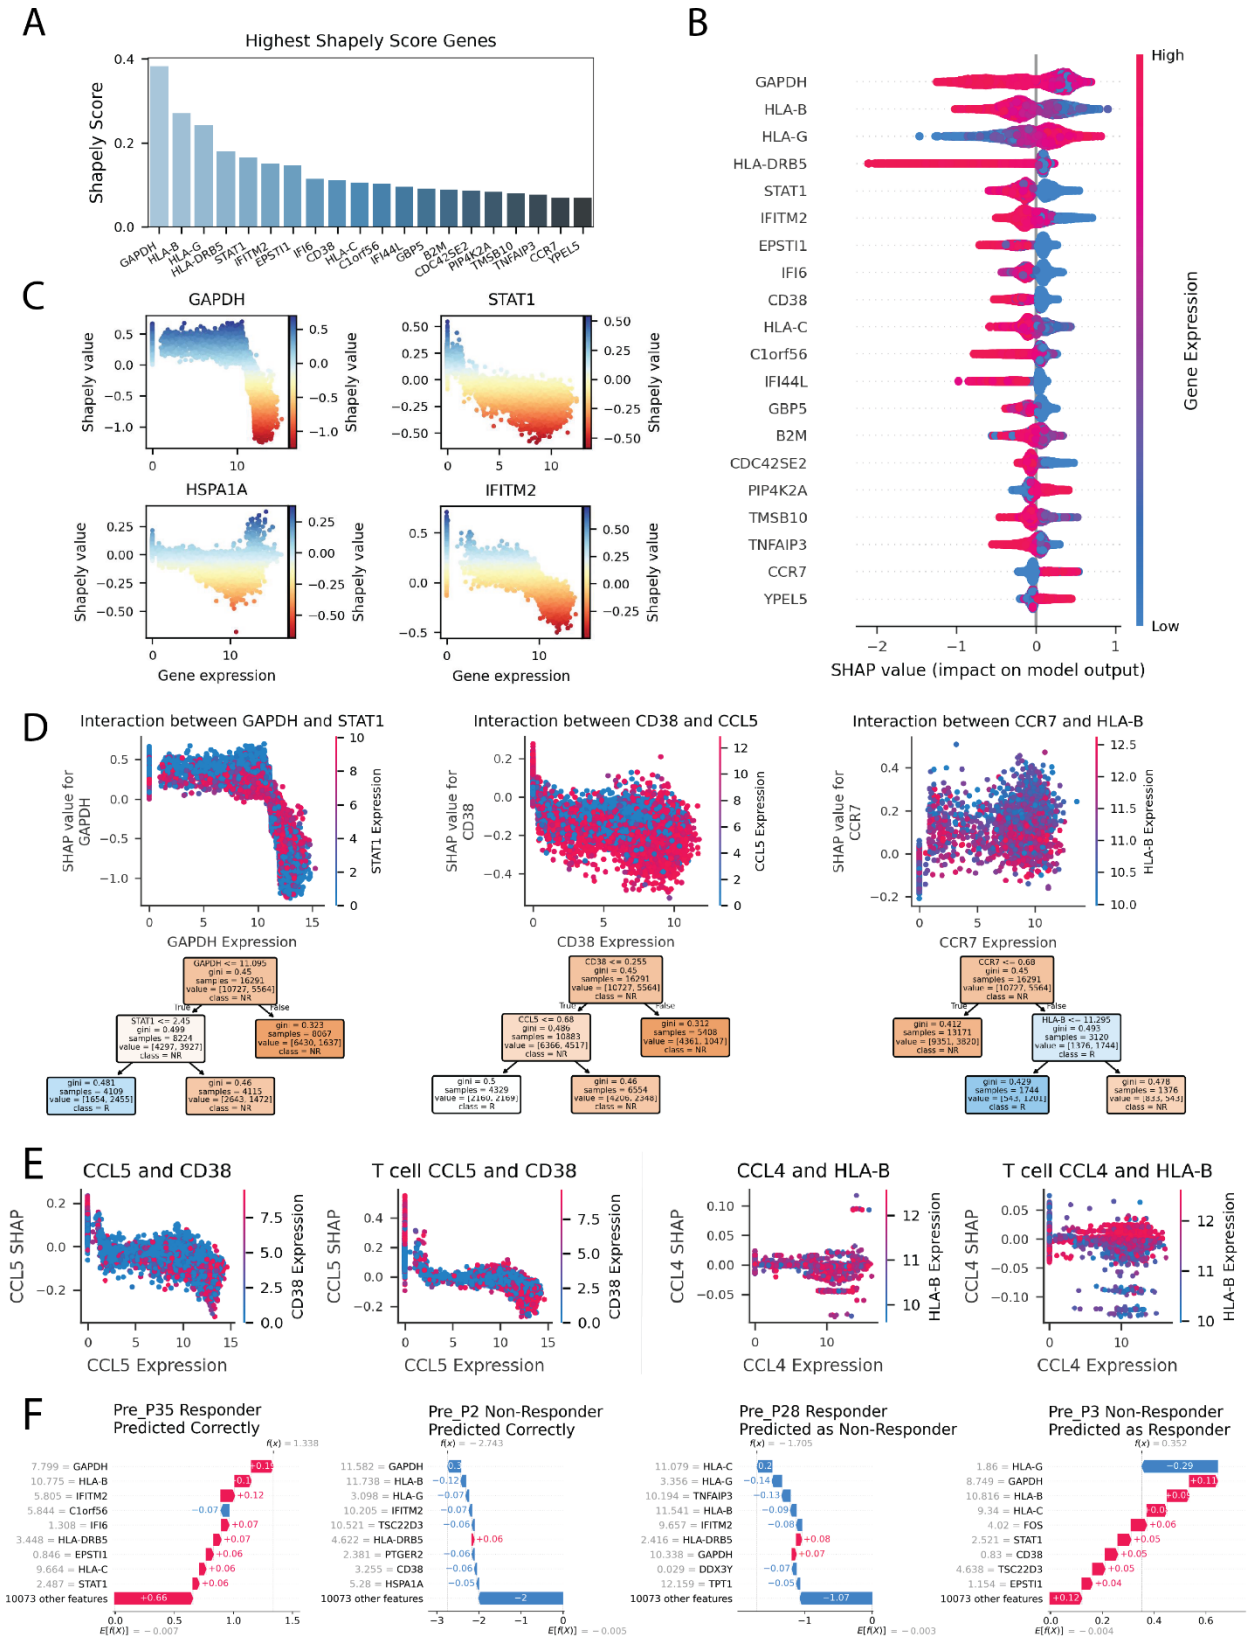
A
Highest Shapely Score Genes
Shapely Score
GAPDH
HLA-B
HLA-G
HLA-DRB5
STAT1
IFITM2
EPSTI1
IFI6
CD38
HLA-C
C1orf56
IFI44L
GBP5
B2M
CDC42SE2
PIP4K2A
TMSB10
TNFAIP3
CCR7
YPEL5
B
High
Low
Gene Expression
SHAP value (impact on model output)
C
GAPDH
STAT1
HSPA1A
IFITM2
Shapely value
Gene expression
D
Interaction between GAPDH and STAT1
SHAP value for GAPDH
GAPDH Expression
STAT1 Expression
Interaction between CD38 and CCL5
SHAP value for CD38
CD38 Expression
CCL5 Expression
Interaction between CCR7 and HLA-B
SHAP value for CCR7
CCR7 Expression
HLA-B Expression
E
CCL5 and CD38
T cell CCL5 and CD38
CCL5 SHAP
CCL5 Expression
CD38 Expression
CCL4 and HLA-B
T cell CCL4 and HLA-B
CCL4 SHAP
CCL4 Expression
HLA-B Expression
F
Pre_P35 Responder
Predicted Correctly
Pre_P2 Non-Responder
Predicted Correctly
Pre_P28 Responder
Predicted as Non-Responder
Pre_P3 Non-Responder
Predicted as Responder
10073 other features

320

Figure 2: SHAP Analysis of Gene-Gene Interactions in Immune Response Prediction

A 20 Top genes with highest absolute Shapley score identified in the model. **B** SHAP value summary plot depicting the impact of each of the most important genes on model output. Gene expression (shown by coloring) as a function of SHAP value (x-axis) showing the relation between expression pattern and model's prediction. **C** Scatter plots showing the relationship between gene expression and SHAP values for key genes (*GAPDH*, *STAT1*, *IFITM2*, and *HSPA1A*). **D** Interaction plots illustrating the SHAP value dependencies between specific gene pairs: *GAPDH* & *STAT1* (left), *CD38* & *CCL5* (middle), *HLA-B* & *CCR7* (right). Expression of the first gene shown as x-axis position, and the expression of the interacting gene shown as coloring; SHAP value shown as y-axis position. Decision trees beneath the plots show the simplified relation of the gene-pair conditional expression with the SHAP values, representing response to ICI. **E** Interaction plots comparing SHAP value dependencies between gene pairs trained on all cells vs T cells: *CCL5* & *CD38* (left), *CCL4* & *HLA-B* (right). Expression of the first gene shown as x-axis position, and the expression of the interacting gene shown as coloring; SHAP value shown as y-axis position. **F** Waterfall plots showing SHAP value contributions to individual model predictions (sample prediction). Four examples of patients predicted as responders or non-responders, with differences in gene expression patterns contributing to the outcome predictions. Pre_P28 – Responder predicted as Non-Responder, Pre_P3 – Non-Responder predicted as Responder, Pre_P2 – Non-Responder predicted correctly, Pre_P35 – Responder predicted correctly.

321

322 **Reinforcement Learning Enables Labeling Each Cell's Predictivity**

323 Our analysis revealed that certain subsets of T cells perform best in predicting patient response.
324 However, it remains unclear whether every cell within this group is necessary to achieve that
325 level of predictive accuracy, or if some cells may not contribute to the model's prediction and
326 even detract from its performance. Specifically, up until now we considered all the cells from a
327 responding sample as responders and all the cells from a non-responding sample as non-
328 responders. This approach represents a simplified model rather than a reflection of reality, as
329 the TME is heterogeneous, and contains synergistic, contradictory, and inert factors. For
330 example, while the proportion or exhausted T cells mostly associates with non-responding
331 samples in baseline, they do exist also in responders (41). To address that, we developed a novel

332 reinforcement leaning (RL) framework that is built upon the basic model. Our goal was to
333 quantify each cell's effect on model prediction and create a cell-level score by rewarding cells
334 with directional rewards when they are predicted correctly. By doing this, we move from a
335 coarse, cluster-level measure, to a fine-grained, cell-level score.

336 First, to get a continuous score, the XGBoost model was switched from a Classifier model to a
337 Regressor model. Moreover, to create the most accurate model possible, we only used the top
338 genes from the FS section. This not only improved running time but also ensured a better
339 distribution of the labels across samples. Initially, each cell from a responder sample is labeled 1,
340 and each cell from a non-responder is labeled -1. Then, the model iteratively updates the
341 continuous labels for each cell, ultimately assigning a score to each (Methods). Cells that are
342 consistently predicted in the same direction will receive repeated rewards and therefore will be
343 pushed to the edges of the distribution. In contrast, cells that are often predicted incorrectly will
344 be iteratively normalized without any reward, resulting in scores with a low absolute value.
345 These scores reflect the strength and direction of each cell's contribution to the response of the
346 sample (Figure 3A,B). Large positive and negative scores indicate strong predictive power of
347 positive and negative response, respectively, while close to zero scores indicate an overall weak
348 association of these cells with patient response.

349 Examining the score distribution across clusters and cell types, we categorized each cell into
350 three bins based on its score: 'non-response predictive' (< -0.5), 'non-predictive' (between -0.5
351 and 0.5), and 'response predictive' (> 0.5). Figure 3A illustrates the distribution of these bins
352 across the cohort. There is a clear association between the proportion and direction of
353 predictive cells within each cluster and the relation of the cluster to response (Figure 3C), as
354 determined by the original study (15). Specifically, clusters G3 (Monocytes/Macrophages), G6

355 (Exhausted CD8+ T-cells), and G11 (Lymphocytes exhausted/Cell cycle), which are prevalent in
356 non-responders, also had the highest proportion of non-response predictive cells. Conversely,
357 clusters G1 (B-cells) and G10 (Memory T-cells) that are associated with responders, were the top
358 clusters in terms of response predictive cells. This analysis also highlights clusters G9
359 (Exhausted/HS *CD8*$^+$ T-cells) and G7 (Regulatory T-cells), which have a relatively high prevalence
360 of non-response predictive cells, and cluster G5 (Lymphocytes), which contains more response
361 predictive cells, although with a lower proportion (see Supplementary Figure 5 for cell-type RL
362 distribution). This observation serves as a sanity check, reinforcing the validity of our model.

363 Notably, non-predictive cells are well distributed across most clusters and are prevalent
364 throughout the cohort. This might result from the arbitrary threshold used to bin the cells,
365 however, examining the distribution of the scores across the cohort shows that most of the non-
366 predictive cells are indeed close to zero and are therefore not affected by the specific threshold
367 used (Figure 3B). In addition, it is important to note that these cells have passed rigorous quality
368 control, so their non-predictive phenotype is not likely due to low-quality data but rather
369 suggests a different level of participation in immunity. Finally, the spreading of non-predictive
370 cells across different clusters and cell-types suggests that simply filtering cell-types or focusing
371 on a single cluster is not sufficient, emphasizing the potential need for a finer method to filter
372 cells and narrow the analysis to predictive cells.

373 Following, we examined the distribution of RL labels within each sample (Figure 3D). As
374 expected, the proportion of predictive cells in well-classified samples was much higher, while
375 misclassified samples had a higher proportion of non-predictive cells. Although likely reflecting
376 to some extent the accuracy of the model, this also highlights the heterogeneity of the TME
377 both within individual samples and across different samples.

378 To characterize these cells and understand what makes a cell predictive or non-predictive of
379 patient response, we performed differential expression analysis among the three bins, as shown
380 in Figure 3E. As expected, the genes used in the RL model appeared in the top differentiators
381 list, but other significant genes also emerged, suggesting that the RL model manages to
382 generalize and create biological groups. These include *PKM*, *PGAM1* and *ACTB* (42–44) in non-
383 response predictive, *TCF7*, *IL7R* in non-predictive, and *SELL*, *MS4A1* (45) in response predictive
384 cells. Although some biological genes were extracted as differentiators of non-predictive cells,
385 this group does not seem to have strong markers differentiating it from the rest of the cells.

386 Quantifying individual cell participation in immune response provides novel, granular insights
387 into the tumor microenvironment. This approach also offers an avenue for identifying and
388 filtering out non-contributory cells, refining predictions related to immunotherapy response.

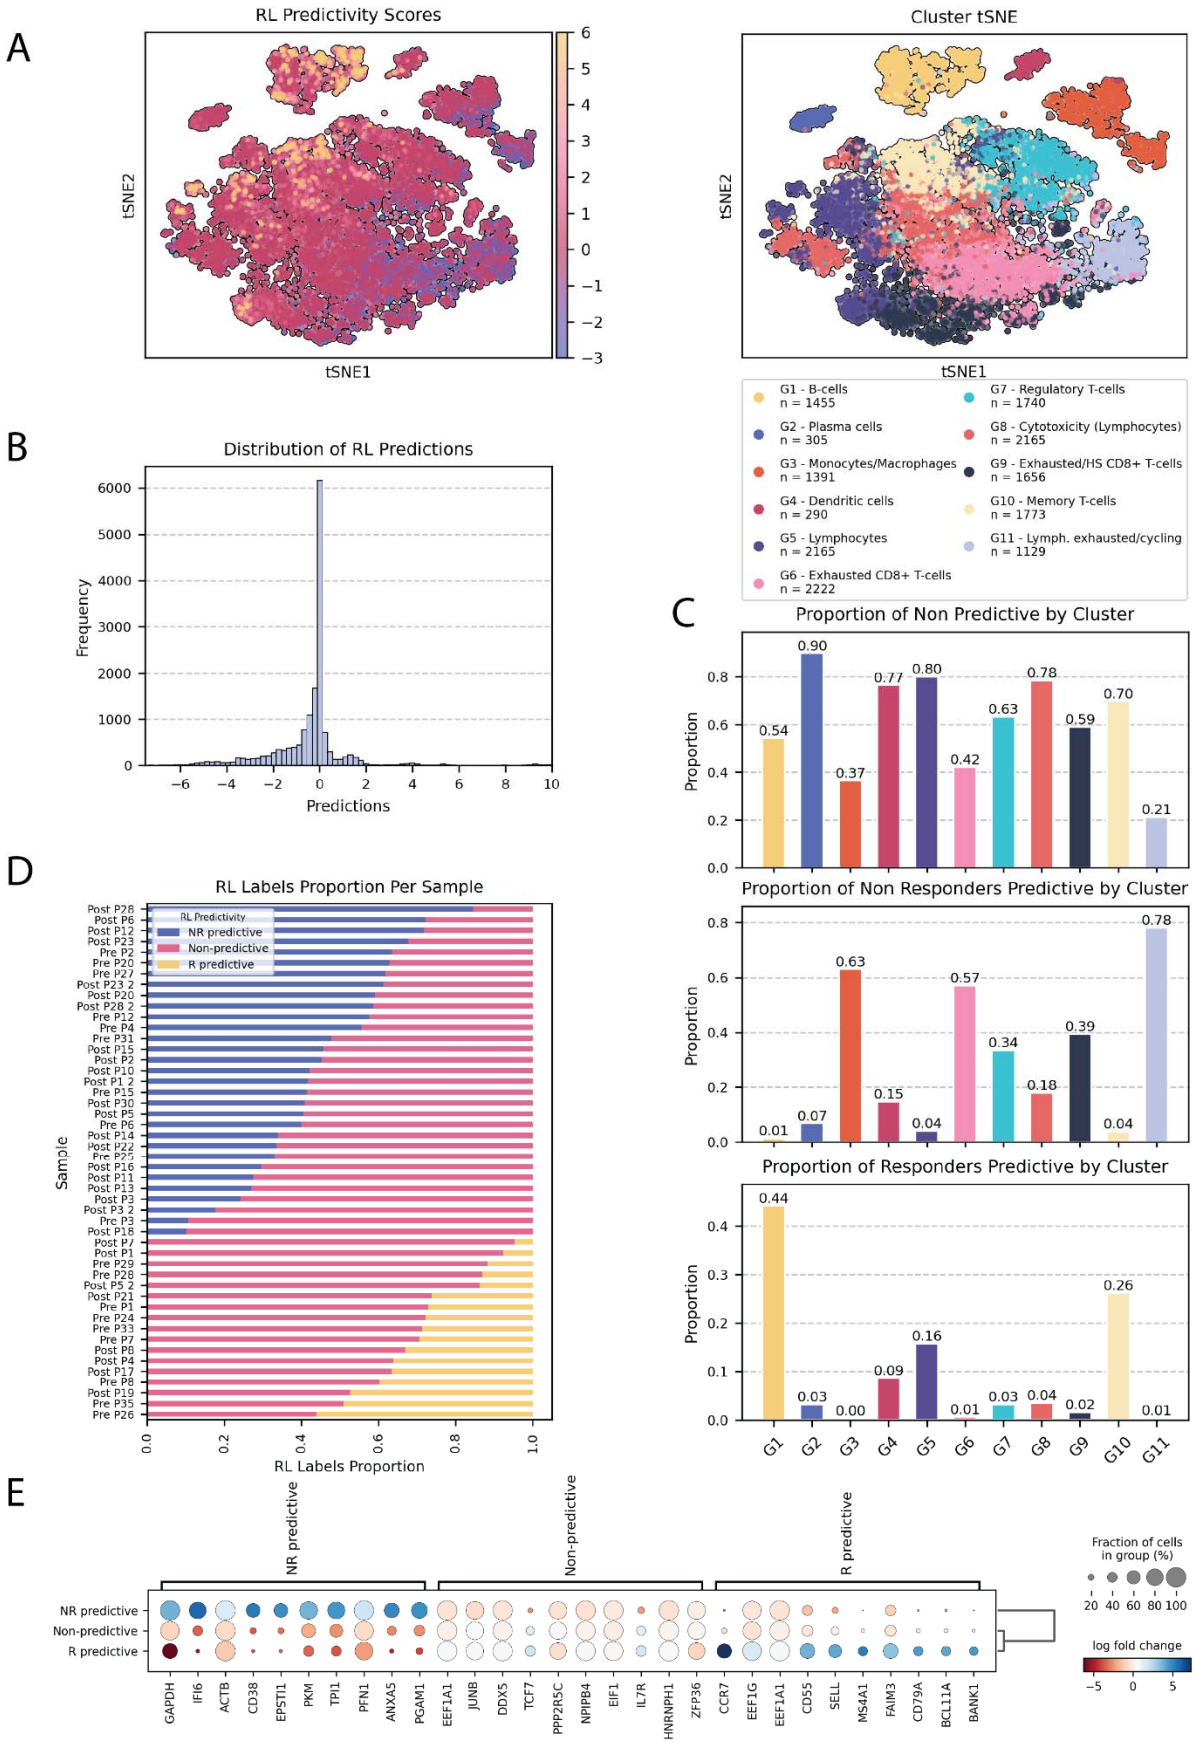
A
RL Predictivity Scores
tSNE2
tSNE1
Cluster tSNE
G1 - B-cells n = 1455
G2 - Plasma cells n = 305
G3 - Monocytes/Macrophages n = 1391
G4 - Dendritic cells n = 290
G5 - Lymphocytes n = 2165
G6 - Exhausted CD8+ T-cells n = 2222
G7 - Regulatory T-cells n = 1740
G8 - Cytotoxicity (Lymphocytes) n = 2165
G9 - Exhausted/HS CD8+ T-cells n = 1656
G10 - Memory T-cells n = 1773
G11 - Lymph. exhausted/cycling n = 1129
B
Distribution of RL Predictions
Frequency
Predictions
C
Proportion of Non Predictive by Cluster
Proportion of Non Responders Predictive by Cluster
Proportion of Responders Predictive by Cluster
Proportion
D
RL Labels Proportion Per Sample
RL Predictivity
NR predictive
Non-predictive
R predictive
Sample
RL Labels Proportion
E
NR predictive
Non-predictive
R predictive
GAPDH IFI6 ACTB CD38 EPSTI1 PKM TPI1 PFN1 ANXA5 PGAM1 EEF1A1 JUNB DDX5 TCF7 PPP2R5C NPIPB4 EIF1 IL7R HNRNPH1 ZFP36 CCR7 EEF1G EEF1A1 CD55 SELL MS4A1 FAIM3 CD79A BCL11A BANK1
Fraction of cells in group (%)
log fold change

389

Figure 3: RL Prediction power

A tSNE plots showing immune cell clusters (top) and RL prediction scores (bottom). Clusters are colored by group, and the RL prediction scores range from highly predictive to non-predictive (lower limit is determined as -3 and upper limit as 6 for visualization purposes). **B** Histogram of RL prediction distribution. **C** Bar plots showing the proportion of non-responders predictive (top), non-predictive (middle), responders predictive (bottom) cells by cluster. **D** Stacked bar plots representing the proportion of RL labels (R predictive, non-predictive, NR predictive) for each patient sample. Top 31 samples are Non-Responders and bottom 17 samples are Responders. **E** Dot plot of the top most differentially expressed genes between the three RL bins, showing log fold change of key genes in non-responders predictive (NR predictive), non-predictive, and responders predictive (R predictive) groups.

390

391 **Gene- and cell-based signatures predict patient response to ICI**

392 To evaluate the generalizability of our findings, we first constructed a gene signature
393 differentiating between responders and non-responders (Figure 4A). To this end, we extracted
394 the genes with the highest feature importance scores from each fold of the LOO (Methods).
395 Analyzing the intersection of genes from the 48 folds, we selected the top 11 genes that
396 consistently intersected across the folds: *GZMH*, *LGALS1*, *GBP5*, *HLA-DRB5*, *CCR7*, *IFI6*, *GAPDH*,
397 *HLA-B*, *EPSTI1*, *CD38*, and *STAT1*. Although somewhat arbitrary, the number of genes selected
398 was based on the minimum number of top genes needed from each fold to form the
399 intersection group: the top 11 genes required relatively few top genes from each fold, whereas
400 those further down the list required a larger number of intersecting genes, making them less
401 robust (Supplementary Figure 1). By examining the SHAP distribution for these genes'
402 expression, we determined that all except *CCR7* were negatively correlated with response.
403 Following, the score we devised was an equally weighted combined score of the genes in the
404 signature, with the weight of *CCR7* specifically negatively adjusted (Methods).

405 To mitigate biases arising from differences in cell-type abundance and considering that some
406 datasets contain only T cells, we focused exclusively on T cells for all of our validation datasets.
407 First, we evaluated the scores on the training melanoma cohort, which produced a high AUC
408 score of 0.93 (n = 48, p-value = 6.16×10^{-7}). Importantly, a significant difference was
409 observed when examining baseline and post-therapy samples separately (n = 19 and 29 , p-
410 values of $3.5 * 10^{-4}$ and $3.9 * 10^{-4}$, respectively).

411 To establish the robustness of this score, we analyzed additional publicly available datasets
412 having scRNA-seq data from patients treated with ICI (Figure 4A). In a triple-negative breast
413 cancer (TNBC) dataset, our predictor achieved an AUC of 0.95 (n = 13, p value = 0.003) (46), and
414 baseline samples were perfectly separated using the score receiving an AUC of 1 (n = 8, p-value
415 = 0.018). In a non-small cell lung cancer (NSCLC) dataset containing only post-treatment
416 samples, our score metric achieved an AUC of 0.76 (n = 57, p-value = $5.5 * 10^{-4}$) (47). Two
417 additional datasets showed a positive trend in association with the score, though they were not
418 as significant with this simple score: a glioblastoma dataset and a Basal-cell carcinoma (BCC)
419 dataset both receiving an AUC score or 0.68 (n = 11, p-value = 0.2 and n = 22, p-value = 0.08,
420 respectively) (48), (49). The BCC dataset showed a relatively better separation using only
421 baseline samples (AUC = 0.77), though due to the smaller sample size, this separation remained
422 insignificant (n = 11, p-value = 0.09).

423 Next, we applied our predictor to a dataset of HER2 positive, ER positive and TNBC patients
424 having annotations for clonal expansion phenotype (50). The dataset consists of two cohorts,
425 one of treatment naive samples, and the second of patients receiving neoadjuvant
426 chemotherapy prior to anti-PD1. In both cohorts, the score was significantly associated with
427 expansion, with the first cohort receiving an AUC of 0.81 and the second 0.77 (n=58, p-value =

428 $8.1 * 10^{-5}$, and n=22, p-value = 0.03). Combined, these datasets achieved an AUC score of 0.79
429 (n = 80, p-value = $2.8 * 10^{-5}$). This trend remained significant when separating baseline and
430 post-treatment samples. This finding suggests that samples with expanded clonal populations
431 are more likely to be non-responsive. This also aligns with a study in melanoma patients, which
432 showed that highly expanded clonotype families were predominantly distributed in cells with an
433 exhausted phenotype, characterized by decreased diversity of T cell receptors (TCRs) (51). Lastly,
434 we applied our signature to a dataset of breast cancer patients having a classification for
435 exhausted and non-exhausted TME (52). This data predominantly contained samples from
436 treatment-naive breast cancer patients. Our predictor achieved an AUC of 0.86 and significantly
437 differentiated between these two TME states (n = 14, p value = 0.01).

438 Following, we sought to validate the RL based filtration of non-predictive cells by constructing a
439 cell-filtration score (Methods). This score guided the removal of cells from each dataset, after
440 which we reapplied the 11-gene signature to the remaining cells and tested for prediction
441 accuracy (Figure 4A). While the genes identified in the differential expression analysis in the
442 previous section effectively differentiate response-predictive from non-response-predictive
443 cells, as aforementioned, non-predictive cells had less obvious distinguishing markers (Figure
444 3E). Since we aim to apply the score to remove non-predictive cells in new datasets, a more
445 sophisticated way of characterizing these cells is needed. To address this, we used a logistic
446 regression model to predict who are the non-predictive cells. Although this classifier cannot be
447 directly applied to other datasets due to differences in scale and potential batch effects, we can
448 derive a more generalizable score from it that can be applied to new data. The filtration score
449 was built according to the logistic regression classifier of non-predictive cells, using the
450 coefficients of the top 100 most important genes (Supplementary Table 5). A weighted mean
451 over all these genes determined the RL-based filtration-score (Methods).

452 To balance the removal of non-predictive cells while minimizing the risk of filtering out
453 predictive ones, we chose to remove 40% of the cells with the highest scores across all datasets.
454 The results were robust to the threshold choice between 15-60%: as expected, less filtration led
455 to smaller improvements, while more filtration resulted in greater improvements, up to a
456 certain point (Supplementary Table 6). In the melanoma dataset that was used for devising the
457 gene and cell scores (15), the addition of the cell-filtration improved the AUC score to over 0.94
458 (n = 48, p-value = 2.55×10^{-7}). Applying this filtering to the other datasets used for validation,
459 achieved an improvement in the prediction accuracy in five out of the six. Specifically, the
460 glioblastoma dataset showed a significant improvement in AUC, from 0.68 to over 0.9 (n = 11, p-
461 value = 0.03). Filtration in the BCC dataset produced an improved AUC of 0.7 (n = 22, p-value
462 = 0.06). The HER2 positive ER positive and TNBC dataset consisting of the two cohorts, saw an
463 increase in AUC to 0.8 in relation to expansion (n = 80, p-value = $1.3 * 10^{-5}$). The AUC of the
464 breast cancer cohort labeled with exhaustion increased to 0.88 (n = 14, p-value = 0.009). In the
465 TNBC (46) the AUC was improved from 0.95 to 0.98 (n = 13, p value = 0.0016). In the NSCLC
466 dataset the AUC score was reduced to 0.73.

467 **Gene signature predicts patients' survival**

468 Finally, we examined the specific genes that make up the 11-gene signature, using bulk RNA-seq
469 datasets to assess their association with patient overall survival (Figure 4B). We expected that
470 higher immune responsiveness would be associated with better survival, hence the signature
471 should differentiate survival rates in this direction. Of note, validation in bulk RNA sequencing is
472 challenging due to the mixture of different cell types. However, since survival analyses look at
473 outcomes across a cohort, they are less dependent on the precision of cell-type-specific
474 resolution. What matters is whether the overall expression of these immune genes, in the
475 context of the mixed population, correlates with better or worse outcomes for the patients.

476 Moreover, strong immune-related signals can still emerge because immune gene expression
477 often reflects the overall immune activity in the tumor microenvironment.

478 To examine this, we used the GEPIA2 online tool (53), and assessed the relation between each
479 of the 11 genes and survival. Using the median expression as a cutoff and using all cancer types
480 available in GEPIA (Methods), we found that 7 out of the 11 genes, *CCR7*, *STAT1*, *GBP5*, *LGALS1*,
481 *EPSTI1*, *IFI6* and *GAPDH*, were significantly differentiating the overall survival time in the
482 expected direction (after Bonferroni correction). The remaining 4 genes, *HLA-DRB5*, *HLA-B*,
483 *GZMH* and *CD38*, showed no significant association with survival in either direction (Figure 4B).
484 The fact that this signal comes from all cancer types and shows that patients with high
485 expression of these genes in their tumor exhibit an improved survival rate, strengthens our
486 hypothesis that this could indeed be the immune related activation.

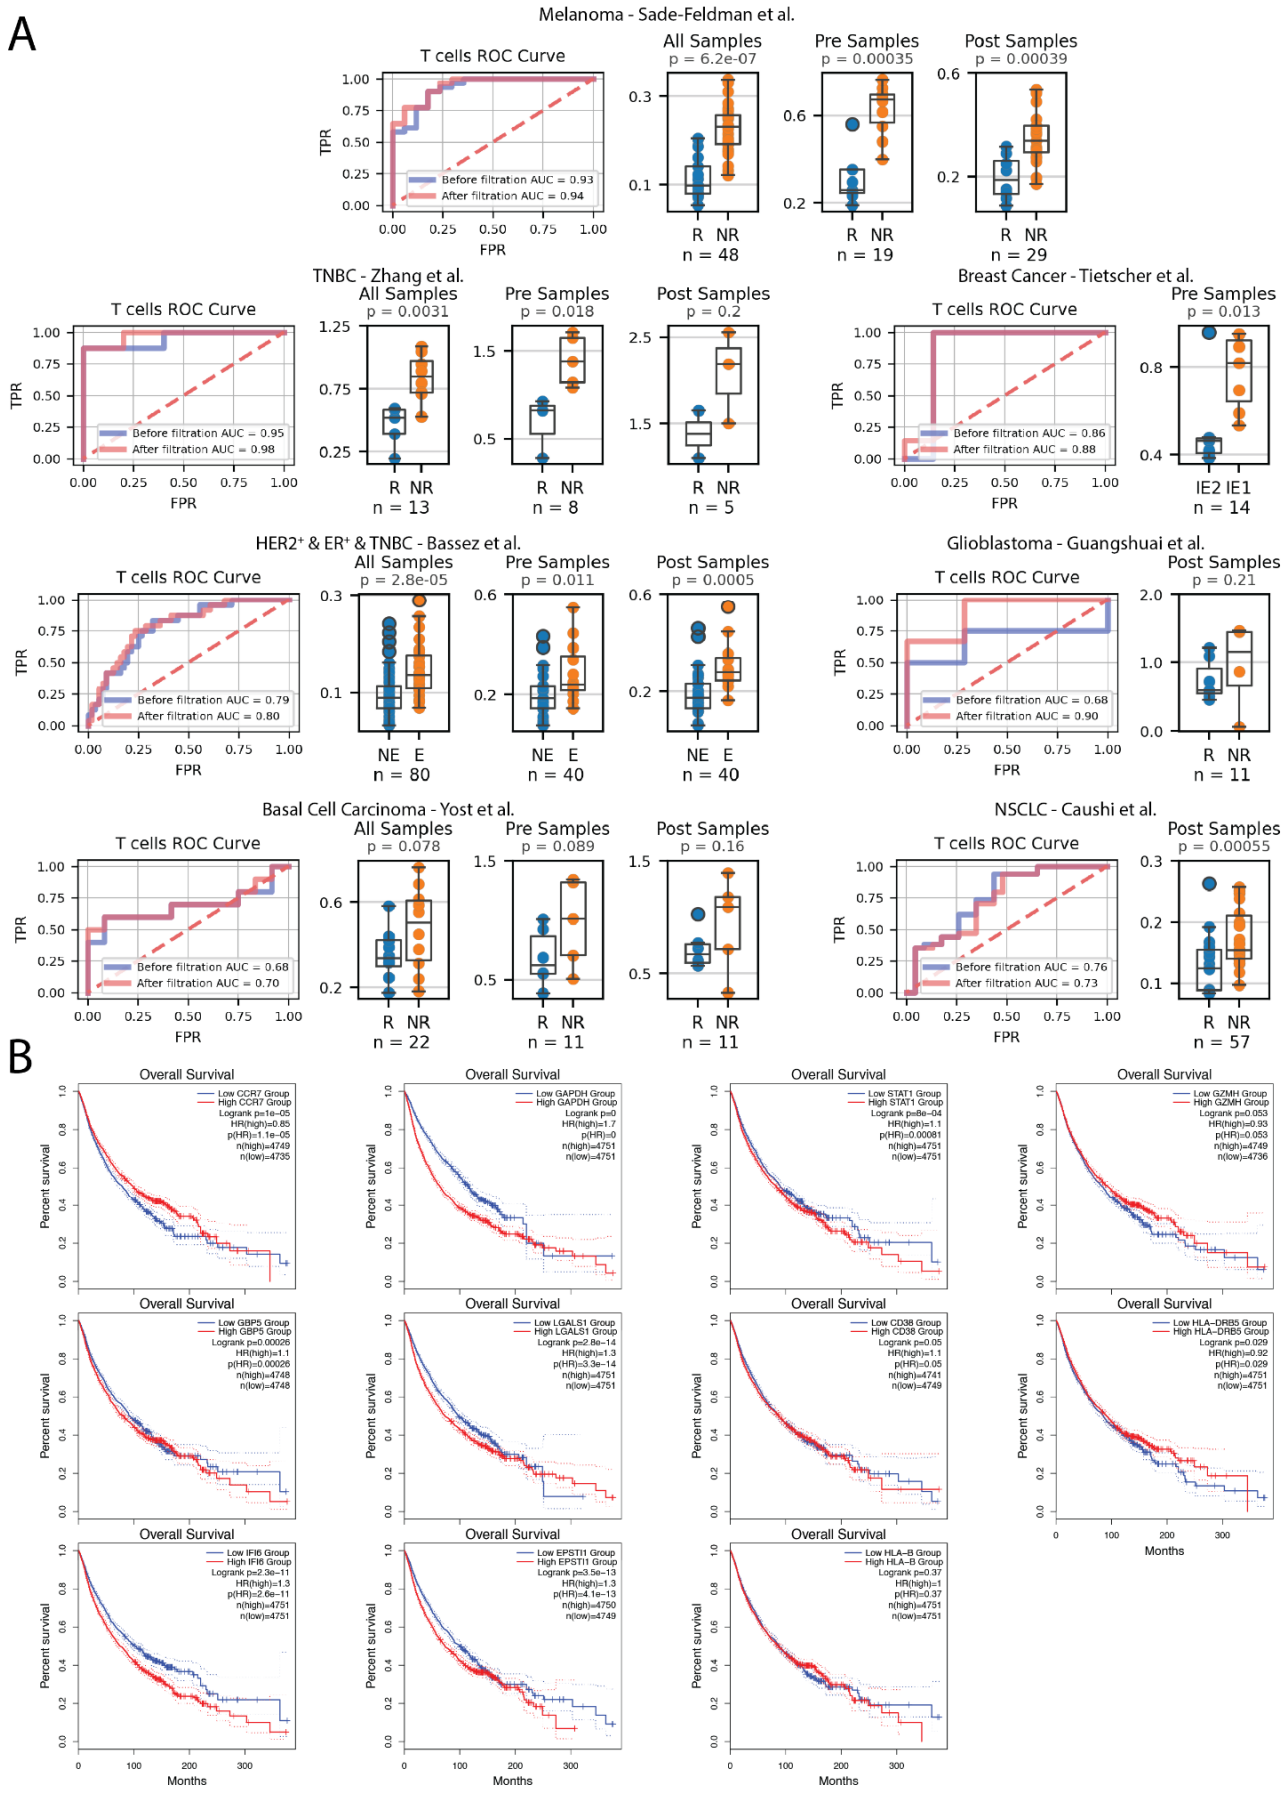
A
Melanoma - Sade-Feldman et al.
T cells ROC Curve
Before filtration AUC = 0.93
After filtration AUC = 0.94
All Samples
p = 6.2e-07
Pre Samples
p = 0.00035
Post Samples
p = 0.00039
R NR
n = 48
n = 19
n = 29
TNBC - Zhang et al.
Before filtration AUC = 0.95
After filtration AUC = 0.98
All Samples
p = 0.0031
Pre Samples
p = 0.018
Post Samples
p = 0.2
n = 13
n = 8
n = 5
Breast Cancer - Tietscher et al.
Before filtration AUC = 0.86
After filtration AUC = 0.88
Pre Samples
p = 0.013
IE2 IE1
n = 14
HER2⁺ & ER⁺ & TNBC - Bassez et al.
Before filtration AUC = 0.79
After filtration AUC = 0.80
All Samples
p = 2.8e-05
Pre Samples
p = 0.011
Post Samples
p = 0.0005
NE E
n = 80
n = 40
n = 40
Glioblastoma - Guangshuai et al.
Before filtration AUC = 0.68
After filtration AUC = 0.90
Post Samples
p = 0.21
n = 11
Basal Cell Carcinoma - Yost et al.
Before filtration AUC = 0.68
After filtration AUC = 0.70
All Samples
p = 0.078
Pre Samples
p = 0.089
Post Samples
p = 0.16
n = 22
n = 11
n = 11
NSCLC - Caushi et al.
Before filtration AUC = 0.76
After filtration AUC = 0.73
Post Samples
p = 0.00055
n = 57
TPR
FPR
B
Overall Survival
Percent survival
Months

487

Figure 4: T Cell Response and Overall Survival Analysis in Various Cancer Types

A T Cell ROC curves showing the predictive power of the 11-gene signature score combined with RL-based filtration in distinguishing responders from non-responders across multiple cancer types. The blue curve represents the ROC curve before filtration, and the red curve represents the ROC curve after filtration. Box plots, separated by pre- and post-treatment samples (where available), depict the 11-gene signature scores before RL filtration. The datasets included are: Melanoma – Sade-Feldman et al. (15), TNBC (Triple-Negative Breast Cancer) - Zhang et al. (46), NSCLC (Non-Small Cell Lung Cancer) – Caushi et al. (47), BCC (Basal Cell Carcinoma) – Yost et al. (48), Glioblastoma – Mei et al. (49), HER2+ & ER+ & TNBC - Bassez et al. (50), Breast Cancer - Tietscher et al. (52). **B** Kaplan-Meier survival analysis of the 11 genes in the signature using bulk RNA-seq data across various cancer types. The survival curves illustrate the association between high and low expression of each gene. Plots include survival for genes – GAPDH, CD38, CCR7, HLA-DRB5, STAT1, GZMH, LGALS1, IFI6, EPSTI1, HLA-G, GBP5.

Discussion

In this study, we developed PRECISE, a machine learning framework designed to predict response to ICI from single-cell RNA-seq data. Our aim was to provide a roadmap for how single-cell data can be used in prediction tasks, such as treatment response, while maintaining model interpretability. Machine learning models stand out as the only technology able to capture the complexity level seen in the TME and reflected in single-cell data. The main advantage of these supervised learning models is that they are trained directly on the target and adapt to the characteristics of the data, instead of merely searching for patterns linked to the target.

The ability to work simultaneously on two levels of complexity—sample-level and cell-level—is crucial for any model designed for single-cell interpretation. Our model enables exactly that: initially, by assigning each cell a label according to its sample; next, by combining cell-level predictions into a sample response score; later, by aggregating the single-cell effects on the model using SHAP and grouping them into sample contributions; and lastly, through the

502 reinforcement learning model, diffusing the sample responses back down to assign a predictivity
503 score to each cell.

504 Our model is highly flexible, both in terms of implementation and objective. Although we used
505 XGBoost, many other algorithms with sufficient capacity could handle the data complexity
506 equally effectively. Additionally, we opted to include all genes that passed preprocessing steps,
507 capturing the strongest signals across the data. However, this approach may overlook significant
508 subsets of the data. Future analyses could focus on specific gene groups (e.g., metabolic genes)
509 to explore different dimensions and revealing new insights.

510 Differences between cancer types, as well as the differences in prior treatments and medical
511 backgrounds, make transferability of results from one dataset to another very challenging.
512 Despite these challenges, the 11-gene signature showed a positive trend in all datasets,
513 indicating the robustness and generalizability of the model's results. The RL filtration score also
514 showed a positive trend in most datasets, improving prediction accuracy and effectively
515 removing noisy, non-predictive cells.

516 Focusing on the immune system is an excellent way to establish transferability between cancer
517 types, since while the cancers can greatly differ, the immune system retains common features.
518 This likely explains why our model's results could be generalized across different cancer types.
519 However, the model was only trained on one dataset. Training on multiple datasets together, or
520 transferring the model directly to new datasets was infeasible due to the large differences in
521 expression between cohorts and especially across sequencing technologies. This limits the
522 model's scalability and its ability to generalize. Future work should aim to address these
523 limitations, and the best current approach may be to use foundation models that integrate
524 single-cell data into a unified embedding (54,55) (interpretability should be kept in mind).

525 Incorporating multiple datasets from different cancer types will significantly enhance the
526 robustness of such models. Scaling up the size of the data, and number of samples used, will
527 open additional possibilities, such as adding higher-level models that aggregate scores from
528 different cell types to produce an overall sample score.

529 Another limitation is the use of both baseline and post-treatment samples. While the ultimate
530 objective of such models is the prediction of patient response before treatment, we included in
531 our training data both pre- and post-treatment samples to maintain as much data as possible for
532 the training step. Nonetheless, we found that the prediction accuracy was similar in the
533 validation datasets when estimated on pre-treatment samples alone, and in some cases even
534 produced better separation. Moreover, it should be noted that prediction on post-treatment
535 samples has a value in cases of advanced line ICI re-challenge, for patients not responding to
536 first line treatments.

537

538 **Conclusions**

539 This study demonstrates the potential of machine learning models trained on single-cell data to
540 predict immunotherapy responses and provide deep biological insights. Using PRECISE
541 (Predicting therapy Response through Extraction of Cells and genes from Immune Single-cell
542 Expression data), we offer a roadmap through which single-cell data can be used not only for
543 sample predictions, but also for understanding complex biological systems such as the TME.
544 Machine learning models in this context can go beyond mere prediction – we showed their
545 utility in capturing gene-gene interactions, identifying non-linear expression patterns, and
546 quantifying cell-type and single-cell participation in immune response. With sufficient data and

547 further development, these models could become exceedingly powerful tools. Expanding
548 training to larger, multi-dataset cohorts will be crucial for improving generalization.

549

550 **Methods**

551 **Data Collection and Preprocessing**

552 We utilized a publicly available scRNA-seq dataset containing immune cells from melanoma
553 patients, generated using Smart-seq technology (15,56). The dataset included 16,291 immune
554 cells from 48 samples. Following initial quality control done by the authors, preprocessing
555 primarily involved filtration of genes based on expression levels. Specifically, we excluded non-
556 coding genes, mitochondrial genes, and ribosomal protein genes (genes starting with 'MT-',
557 'RPS', 'RPL', 'MRP', and 'MTRNR'). Only genes expressed in at least 3% of cells were retained,
558 resulting in 10,082 genes. Cells were labeled according to their sample's response to ICI therapy,
559 with cells from responding samples labeled as responders and those from non-responding
560 samples labeled as non-responders.

561 Additional datasets were collected and preprocessed for validation, including droplet-based
562 datasets containing read counts (46–50). In the droplet-based datasets, expression levels were
563 first normalized to a standard sum of 10,000 reads per cell and then log-transformed. It is
564 important to note that the per-cell read coverage in these datasets was often higher than
565 10,000. Only samples with defined clinical annotations (response or clonal expansion) were
566 included in the analysis. We considered only patients treated with either immune checkpoint
567 inhibitors (ICI) or a combination of ICI and chemotherapy. Additionally, a dataset of a non-ICI
568 cohort containing mostly treatment-naive patients with annotations for exhausted and non-
569 exhausted tumor microenvironment (TME) was included (52).

570

571 **Cell-Type Assignment and Unsupervised Clustering**

572 In our previously published melanoma dataset (15), we used the original classification of T cells
573 that was based both on gene markers and manual curation. Eleven unsupervised clusters and
574 supervised cell-type assignments of the cells were obtained from the original article. In most
575 other datasets, T cells were either extracted using flow cytometry, or were identified *in silico*
576 using computational approaches in the paper (47–50,52) generating the data. The TNBC dataset
577 (46) was the only dataset without cell-type annotations. Identification of T cells in this dataset
578 was done by unsupervised clustering, followed by differential expression on these clusters and
579 determination of T cells according to known markers.

580 **Base Machine Learning Model**

581 Training was conducted in a leave-one-out (LOO) cross-validation manner, where the model was
582 trained on all samples except one, and predictions were made for the left-out sample. The
583 model was trained and used for prediction at the single-cell level, learning to predict the label
584 associated with the sample's response. After predicting the labels of all cells in the cohort, we
585 calculated the proportion of cells predicted as responders in each sample, forming the sample
586 score. Accuracy of the prediction was evaluated using the Receiver Operating Characteristic
587 (ROC) Area Under the Curve (AUC) score between the sample scores and the sample response
588 labels, making the accuracy assessment robust to the threshold choice. From each model fold,
589 feature importance scores were extracted for each gene using the model's built-in
590 'feature_importance_' attribute, which quantifies each gene's contribution to the model's
591 prediction.

592 **XGBoost Classifier Supervised Predictive Model**

593 An XGBoost model, known for its balance between accuracy and interpretability, was trained to
594 predict the response label of each cell (16). The XGBoost model was built using the XGBoost
595 Python package. The classifier was trained using the default parameters for XGBoost, except for
596 a learning rate (LR) of 0.2, a max_depth of 7, and an alpha of 1. The model's performance was
597 not sensitive to these parameters. No hyperparameter tuning was required. The objective
598 function was set to 'binary:logistic'. An important parameter included was the
599 'scale_pos_weight', set to the ratio of the number of negative cells to positive cells, addressing
600 potential bias towards responders or non-responders in the data. Although accuracy was high
601 without it, including this parameter ensures better generalization for more biased data.

602 **Per Cluster/Cell-Type Prediction**

603 Cellular groups we're extracted from the original paper (15), and included 11 clusters
604 determined by unsupervised clustering, and 13 cell types or states determined using known
605 gene markers. To examine cluster- and cell type-specific predictions, we ran the classification
606 model with LOO on each cellular group, and obtained an AUC score for each cluster
607 independently. Samples were excluded from the leave-one-out process if they didn't contain
608 any cells from the specific cellular group, and the accuracy scores were adjusted accordingly.

609 **Boruta Feature Selection**

610 We used the boruta_py from the Python Boruta package. Boruta is a wrapper algorithm that
611 iteratively removes features that are statistically less important than random probes (shadow
612 features). We used Boruta to extract only the confirmed genes from its output in a leave-one-
613 out fashion. We set alpha to 0.02, which produced a small, robust number of features. The
614 chosen features were saved for interpretation and used for LOO prediction, ensuring no train-
615 test leakage.

616 **Lasso Regression Based Feature Selection**

617 Lasso regression, a technique that performs both variable selection and regularization, was used
618 for feature selection. Lasso zeroes the coefficients of non-contributing features. Using different
619 values of the regularization parameter alpha, we performed LOO predictions with Lasso-based
620 feature selection. Although results were reasonable, they were consistently lower than those
621 obtained with Boruta and required handpicking of the alpha value. Feature selection was done
622 after scaling the genes, using Lasso from Python's sklearn package (with max_iter set to 500).

623 **Simple Feature Selections**

624 We also examined the performance of simple feature selection methods, such as differentiating
625 genes and the most highly variable genes. Differential expression analysis between responding
626 and non-responding samples was conducted within the LOO using either Wilcoxon ranksum test,
627 t-test (scanpy), or Fisher's exact tests with significance values of 0.05, 0.01 and 0.001. Highly
628 variable genes were extracted using scanpy's function on the top 1000, 2000, 4000 and 8000 top
629 highly variable genes. The Wilcoxon ranksum test produced the best results among these
630 methods but still yielded consistently lower performance compared to Boruta and resulted in a
631 much higher number of genes.

632 **SHAP Analysis**

633 SHAP (SHapley Additive exPlanations) values were used to interpret the model's predictions.
634 SHAP values explain the impact of each feature (gene) on the model's output, providing insights
635 into gene importance and interactions. We applied SHAP using the SHAP Python package on the
636 genes: [*'GZMH', 'LGALS1', 'GBP5', 'HLA-DRB5', 'CCR7', 'IFI6', 'GAPDH', 'HLA-B', 'EPSTI1', 'CD38',*
637 *'STAT1'*]. SHAP automatically identified gene interactions and interpret the mean contributions

638 of genes to a sample's prediction. We averaged the SHAP scores for each gene across the
639 sample to generate sample- level Waterfall plots (Figure 2, Supplementary Figure 4).

640 **Decision Trees for SHAP Interaction Analysis**

641 To model the pairwise conditional behaviors of genes as extracted by SHAP analysis, we used
642 decision trees. Each time, we insert the expression pattern of two genes extracted by SHAP to
643 the decision tree model. We used the DecisionTreeClassifier from python's sklearn with
644 max_depth set to 2. This limit on the depth helped exhibit a simple relation between the two
645 genes.

646 **Reinforcement Learning Framework**

647 To quantify the directional predictivity of each cell, we constructed a reinforcement learning
648 framework based on XGBoost, but this time we used a regression model. The initial labels were
649 set to 1 for cells from responders and -1 for cells from non-responders. Iteratively, we updated
650 the labels of the cells if they were predicted correctly, weighted by the prediction value. Each
651 iteration normalized the absolute value of the labels to keep their scale consistent. The
652 pseudocode below illustrates this process:

653 <u>Initialization:</u>

654 1. Initialize responder cells with label 1 and non-responder cells with label -1

655 2. Set the list of chosen features to those extracted with feature selection

656 3. Set the learning rate (eta = 0.01) and the number of iterations (200)

657 <u>Iteration Process:</u>

658 For each iteration:

1. For each sample:

a. Train the model on all samples except the current sample.

b. Predict the labels for the current sample.

c. for each cell in sample:

- If the sample is a non-responder and the prediction value is negative,

or if the sample is responder and the prediction value is positive:

- Update the label: label = label +eta*prediction

- Else:

- Do not update the label.

2. Normalize the absolute updated value of the labels to maintain their scale.

Return:

Return the updated labels.

The labels were then classified into 'Responder Predictive' for $label > 0.5$, 'Non-Responder Predictive' for $label < -0.5$, and 'Non-Predictive' for $-0.5 < label < 0.5$. These groups were then characterized with differential expression analysis.

Application of 11-gene Signature on External Datasets

678 We created a pseudo-bulk expression matrix by averaging the expression of each gene across all
679 T cells in each sample. To normalize the differences in the scale of expression among different
680 genes, we divided the expression of each gene in a sample by the total expression of that gene
681 across all samples. Formally,

$$E[sample, gene] = \frac{P[sample, gene]}{\sum_{s\ in\ samples}\ P[s, gene]}$$

682 Where P is the pseudo-bulk expression matrix with the samples as rows and genes as columns,
683 and E is the new, scaled matrix. The normalization term in the denominator sums the expression
684 of a given gene over all samples in the dataset.

685 After calculating the scaled pseudo-bulk, the score for each sample was calculated as follows:

$$score(sample)$$
$$= sum(E[sample, GZMH], E[sample, LGALS1], E[sample, GBP5], E[sample, HLA$$
$$- DRB5], E[sample, IFI6], E[sample, GAPDH], E[sample, HLA$$
$$- B], E[sample, EPSTI1], E[sample, CD38], E[STAT1\ ], - E[CCR7])$$

686 This score was then used to differentiate between the different responses in the samples.

687

688 **Cell Filtration Score Calculation and Application**

689 The filtration score was calculated based on a logistic regression model. The model was trained
690 to differentiate non-predictive cells from the rest of the cells (response predictive and non-
691 response predictive) using the python LogisticRegression algorithm from sklearn (57), with the
692 max_iter parameter increased to 500 to assure convergence. The 100 genes with the largest
693 absolute coefficient values were extracted from the model and used to create a weighted score.

694 Given the vector of the weights extracted by the logistic regressor and the expression matrix E
695 normalized for each gene's expression:

$$\underline{W} = w[gene1], w[gene2], \ldots, w[gene100]$$

$$\underline{E_{cell}} = E[cell, gene1], E[cell, gene2], \ldots, E[cell, gene100]$$

696 The score for a specific cell is calculated by the dot product of these two vectors

$$score(cell) = \underline{W} \cdot \underline{E_{cell}}$$

697 The cells with the largest scores are most likely to be non-predictive. The cells with the top 40%
698 scores were filtered out in all datasets.

699 **Bulk RNA-seq Survival Analysis**

700 GEPIA2, an online tool for analyzing RNA sequencing expression data, was used to examine the
701 relation of the expression of the 11 genes to overall survival (53). Using the median expression
702 as a cutoff, this produced Kaplan-Meier survival curves for each gene, over all available cancer
703 types in GEPIA2. These include: ACC, BLCA, BRCA, CESC, CHOL, COAD, DLBC, ESCA, GBM, HNSC,
704 KICH, KIRC, KIRP, LAML, LGG, LIHC, LUAD, LUSC, MESO, OV, PAAD, PCPG, PRAD, READ, SARC,
705 SKCM, STAD, TGCT, THCA, THYM, UCEC, UCS and UVM.

706

707 **Abbreviations**

708 **scRNA-seq:** Single-cell RNA sequencing

709 **TME:** Tumor microenvironment

710 **ICI:** Immune checkpoint inhibitors

711 **LOO:** Leave-one-out

712 **ROC:** Receiver Operating Characteristic

713 **AUC:** Area Under the Curve

714 **SHAP:** SHapley Additive exPlanations

715 **RL:** Reinforcement learning

716 **FS:** Feature selection

717 **HNSCC:** Head and neck squamous cell carcinoma

718 **BCC:** Basal-cell carcinoma

719 **NSCLC:** Non-small cell lung cancer

720 **TNBC:** Triple-negative breast cancer

721 **ADCC:** Antibody-dependent cellular cytotoxicity

722 **ML:** Machine learning

723 **NR:** Non-responder/Non-responsive

724 **R:** Responder/Responsive

725 **LR:** Learning rate

726 **HVG:** Highly variable genes

727 **TCRs:** T cell receptors

728

729 **Declarations**

730 **Ethics approval and consent to participate**

731 Not applicable.

732 **Consent for publication**

733 Not applicable.

734 **Availability of Data and Materials**

735 All data used in this study is published and publicly available. The scRNA-seq data for the
736 melanoma training cohort used in this paper is available through accession number GEO:
737 GSE120575 (15). TNBC scRNA-seq data was downloaded from GSE169246 (46). NSCLC scRNA-
738 seq data was downloaded from GSE176021 (47). Glioblastoma scRNA-seq was downloaded from
739 Figshare (https://doi.org/10.6084/m9.figshare.22434341) (49). BCC was attained from
740 GSE123813 (48). HER2 positive, ER positive and TNBC scRNA-seq data was accessed from the
741 Lambrecht's lab website (https://lambrechtslab.sites.vib.be/en/single-cell) (50). Breast cancer
742 scRNA-seq data classified by exhaustion was obtained from ArrayExpress database at EMBL-EBI
743 under E-MTAB-10607 (52). The methods described in this manuscript have been implemented in
744 a tool called PRECISE, which is freely available at the GitHub repository. The repository also
745 includes a tutorial for its use: https://github.com/yizhak-lab-ccg/PRECISE-.

746 **Competing interests**

747 The authors declare that they have no competing interests.

748 **Author's Contribution**

749 K.Y. supervised the project and guided it throughout. K.Y. and A.P. jointly wrote the main
750 manuscript text. A.P. prepared the figures and performed the analysis. Both authors reviewed
751 and approved the final manuscript.

752 **Acknowledgements**

753 We would like to thank Prof. Gad Getz, Dr. Moshe Sade-Feldman, Dr. Liron Zisman, Elad Zisman,
754 Ofir Shorer, and Sapir Levin for fruitful discussions and helpful comments on the manuscript.
755 Figure 1A was created with BioRender.com using a paid license.

756 **Funding**

757 This work was supported by the Ministry of Science and Technology (2032895), by the Israel
758 Science Foundation (3614/19), and by the Israel Cancer Research Fund (2032965 23-204-RCDA).
759 This work received additional support from the Ruth and Bruce Rappaport Technion Integrated
760 Cancer Center (RTICC).

761

762 **References**

763 1. Pardoll DM. The blockade of immune checkpoints in cancer immunotherapy. Nat Rev
764 Cancer. 2012 Apr;12(4):252–64.

765 2. Sharma P, Hu-Lieskovan S, Wargo JA, Ribas A. Primary, Adaptive, and Acquired Resistance to
766 Cancer Immunotherapy. Cell. 2017 Feb 9;168(4):707–23.

767 3. Fares CM, Van Allen EM, Drake CG, Allison JP, Hu-Lieskovan S. Mechanisms of Resistance to
768 Immune Checkpoint Blockade: Why Does Checkpoint Inhibitor Immunotherapy Not Work

769 for All Patients? Am Soc Clin Oncol Educ Book Am Soc Clin Oncol Annu Meet. 2019
770 Jan;39:147–64.

771 4. Ganesan S, Mehnert J. Biomarkers for Response to Immune Checkpoint Blockade. Annu Rev
772 Cancer Biol. 2020 Mar 4;4(Volume 4, 2020):331–51.

773 5. Toor SM, Sasidharan Nair V, Decock J, Elkord E. Immune checkpoints in the tumor
774 microenvironment. Semin Cancer Biol. 2020 Oct 1;65:1–12.

775 6. Wu T, Dai Y. Tumor microenvironment and therapeutic response. Cancer Lett. 2017 Feb
776 28;387:61–8.

777 7. Kang Y, Vijay S, Gujral TS. Deep neural network modeling identifies biomarkers of response
778 to immune-checkpoint therapy. iScience. 2022 May;25(5):104228.

779 8. Dong Y, Chen Z, Yang F, Wei J, Huang J, Long X. Prediction of immunotherapy
780 responsiveness in melanoma through single-cell sequencing-based characterization of the
781 tumor immune microenvironment. Transl Oncol. 2024 May 1;43:101910.

782 9. Liu R, Dollinger E, Nie Q. Machine Learning of Single Cell Transcriptomic Data From anti-PD-
783 1 Responders and Non-responders Reveals Distinct Resistance Mechanisms in Skin Cancers
784 and PDAC. Front Genet. 2022 Feb 1;12:806457.

785 10. Erfanian N, Heydari AA, Feriz AM, Iañez P, Derakhshani A, Ghasemigol M, et al. Deep
786 learning applications in single-cell genomics and transcriptomics data analysis. Biomed
787 Pharmacother. 2023 Sep 1;165:115077.

788 11. Ma Q, Xu D. Deep learning shapes single-cell data analysis. Nat Rev Mol Cell Biol. 2022
789 May;23(5):303–4.

790 12. Raimundo F, Meng-Papaxanthos L, Vallot C, Vert JP. Machine learning for single-cell
791 genomics data analysis. Curr Opin Syst Biol. 2021 Jun 1;26:64–71.

792 13. He B, Thomson M, Subramaniam M, Perez R, Ye CJ, Zou J. CloudPred: Predicting Patient
793 Phenotypes From Single-cell RNA-seq. arXiv; 2021 [cited 2024 Nov 13]. Available from:
794 http://arxiv.org/abs/2110.07069

795 14. Mao Y, Lin YY, Wong NKY, Volik S, Sar F, Collins C, et al. Phenotype prediction from single-
796 cell RNA-seq data using attention-based neural networks. Bioinformatics. 2024 Feb
797 1;40(2):btae067.

798 15. Sade-Feldman M, Yizhak K, Bjorgaard SL, Ray JP, de Boer CG, Jenkins RW, et al. Defining T
799 cell states associated with response to checkpoint immunotherapy in melanoma. Cell. 2018
800 Nov 1;175(4):998-1013.e20.

801 16. Chen T, Guestrin C. XGBoost: A Scalable Tree Boosting System. In: Proceedings of the 22nd
802 ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. San
803 Francisco California USA: ACM; 2016 [cited 2024 Aug 28]. p. 785–94. Available from:
804 https://dl.acm.org/doi/10.1145/2939672.2939785

805 17. Kursa MB, Jankowski A, Rudnicki WR. Boruta – A System for Feature Selection. Fundam
806 Informaticae. 2010;101(4):271–85.

807 18. Lundberg S, Lee SI. A Unified Approach to Interpreting Model Predictions. arXiv; 2017 [cited
808 2024 Aug 28]. Available from: http://arxiv.org/abs/1705.07874

809 19. Sutton RS, Barto AG. Reinforcement Learning: An Introduction.

810 20. Zhang J, Li S, Liu F, Yang K. Role of CD68 in tumor immunity and prognosis prediction in pan-
811 cancer. Sci Rep. 2022 May 12;12(1):7844.

812 21. Horowitch B, Lee DY, Ding M, Martinez-Morilla S, Aung TN, Ouerghi F, et al. Subsets of
813 interferon signaling predict response to immune checkpoint blockade in melanoma
814 patients. Clin Cancer Res Off J Am Assoc Cancer Res. 2023 Aug 1;29(15):2908–18.

815 22. Liu F, Liu J, Zhang J, Shi J, Gui L, Xu G. Expression of STAT1 is positively correlated with PD-L1
816 in human ovarian cancer. Cancer Biol Ther. 2020 Oct 2;21(10):963–71.

817 23. Sharma N, Atolagbe OT, Ge Z, Allison JP. LILRB4 suppresses immunity in solid tumors and is
818 a potential target for immunotherapy. J Exp Med. 2021 May 11;218(7):e20201811.

819 24. GAPDH: A common housekeeping gene with an oncogenic role in pan-cancer - PMC. [cited
820 2024 Aug 28]. Available from: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10470192/

821 25. Peng C, Ye H, Yi Z. GAPDH: unveiling its impact as a key hypoxia-related player in head and
822 neck squamous cell carcinoma tumor progression, prognosis, and therapeutic potential. Am
823 J Cancer Res. 2023 Dec 15;13(12):5846–67.

824 26. M1 Macrophage and M1/M2 ratio defined by transcriptomic signatures resemble only part
825 of their conventional clinical characteristics in breast cancer | Scientific Reports. [cited 2024
826 Nov 4]. Available from: https://www.nature.com/articles/s41598-020-73624-w

827 27. Ahmed H, Mahmud AR, Siddiquee MohdFR, Shahriar A, Biswas P, Shimul MdEK, et al. Role
828 of T cells in cancer immunotherapy: Opportunities and challenges. Cancer Pathog Ther.
829 2022 Dec 20;1(2):116–26.

830 28. Meng F, Zhang S, Xie J, Zhou Y, Wu Q, Lu B, et al. Leveraging CD16 fusion receptors to
831 remodel the immune response for enhancing anti-tumor immunotherapy in iPSC-derived NK
832 cells. J Hematol OncolJ Hematol Oncol. 2023 Jun 14;16(1):62.

833 29. Aguilar OA, Gonzalez-Hinojosa MDR, Arakawa-Hoyt JS, Millan AJ, Gotthardt D, Nabekura T,
834 et al. The CD16 and CD32b Fc-gamma receptors regulate antibody-mediated responses in
835 mouse natural killer cells. J Leukoc Biol. 2023 Jan 10;113(1):27–40.

836 30. Park JE, Kim SE, Keam B, Park HR, Kim S, Kim M, et al. Anti-tumor effects of NK cells and
837 anti-PD-L1 antibody with antibody-dependent cellular cytotoxicity in PD-L1-positive cancer
838 cell lines. J Immunother Cancer. 2020 Aug 1;8(2):e000873.

839 31. Luoma AM, Suo S, Wang Y, Gunasti L, Porter CBM, Nabilsi N, et al. Tissue-resident memory
840 and circulating T cells are early responders to pre-surgical cancer immunotherapy. Cell.
841 2022 Aug 4;185(16):2918-2935.e29.

842 32. Chen L, Diao L, Yang Y, Yi X, Rodriguez BL, Li Y, et al. CD38-Mediated Immunosuppression as
843 a Mechanism of Tumor Cell Escape from PD-1/PD-L1 Blockade. Cancer Discov. 2018 Sep
844 4;8(9):1156–75.

845 33. Revach OY, Cicerchia AM, Shorer O, Petrova B, Anderson S, Park J, et al. Disrupting CD38-
846 driven T cell dysfunction restores sensitivity to cancer immunotherapy. bioRxiv. 2024 Mar
847 26;2024.02.12.579184.

848 34. Nakayama Y, Mimura K, Tamaki T, Shiraishi K, Kua LF, Koh V, et al. Phospho-STAT1
849 expression as a potential biomarker for anti-PD-1/anti-PD-L1 immunotherapy for breast
850 cancer. Int J Oncol. 2019 Apr 10;54(6):2030–8.

851 35. Liu J, Chen H, Qiao G, Zhang JT, Zhang S, Zhu C, et al. PLEK2 and IFI6, representing
852 mesenchymal and immune-suppressive microenvironment, predicts resistance to
853 neoadjuvant immunotherapy in esophageal squamous cell carcinoma. Cancer Immunol
854 Immunother CII. 2022 Sep 19;72(4):881–93.

855 36. Osińska I, Popko K, Demkow U. Perforin: an important player in immune response. Cent-Eur
856 J Immunol. 2014;39(1):109–15.

857 37. Vilgelm AE, Richmond A. Chemokines Modulate Immune Surveillance in Tumorigenesis,
858 Metastasis, and Response to Immunotherapy. Front Immunol. 2019 Feb 27;10:333.

859 38. Griffith JW, Sokol CL, Luster AD. Chemokines and Chemokine Receptors: Positioning Cells for
860 Host Defense and Immunity. Annu Rev Immunol. 2014 Mar 21;32(Volume 32, 2014):659–
861 702.

862 39. Honey K. CCL3 and CCL4 actively recruit CD8+ T cells. Nat Rev Immunol. 2006 Jun;6(6):427–
863 427.

864 40. Mukaida N, Sasaki SI, Baba T. CCL4 Signaling in the Tumor Microenvironment. Adv Exp Med
865 Biol. 2020;1231:23–32.

866 41. Jiang P, Gu S, Pan D, Fu J, Sahu A, Hu X, et al. Signatures of T cell dysfunction and exclusion
867 predict cancer immunotherapy response. Nat Med. 2018 Oct;24(10):1550–8.

868 42. Chen M, Liu H, Li Z, Ming AL, Chen H. Mechanism of PKM2 affecting cancer immunity and
869 metabolism in Tumor Microenvironment. J Cancer. 2021;12(12):3566.

870 43. Zheng Y, Wang Y, Lu Z, Wan J, Jiang L, Song D, et al. PGAM1 Inhibition Promotes HCC
871 Ferroptosis and Synergizes with Anti-PD-1 Immunotherapy. Adv Sci. 2023;10(29):2301928.

872 44. Gu Y, Tang S, Wang Z, Cai L, Lian H, Shen Y, et al. A pan-cancer analysis of the prognostic and
873 immunological role of β-actin (ACTB) in human cancers. Bioengineered. 12(1):6166–85.

874 45. Mudd TW, Lu C, Klement JD, Liu K. MS4A1 expression and function in T cells in the colorectal
875 cancer tumor microenvironment. Cell Immunol. 2021 Feb;360:104260.

876 46. Zhang Y, Chen H, Mo H, Hu X, Gao R, Zhao Y, et al. Single-cell analyses reveal key immune
877 cell subsets associated with response to PD-L1 blockade in triple-negative breast cancer.
878 Cancer Cell. 2021 Dec 13;39(12):1578-1593.e8.

879 47. Caushi JX, Zhang J, Ji Z, Vaghasia A, Zhang B, Hsiue EHC, et al. Transcriptional programs of
880 neoantigen-specific TIL in anti-PD-1-treated lung cancers. Nature. 2021 Aug;596(7870):126–
881 32.

882 48. Yost KE, Satpathy AT, Wells DK, Qi Y, Wang C, Kageyama R, et al. Clonal replacement of
883 tumor-specific T cells following PD-1 blockade. Nat Med. 2019 Aug;25(8):1251–9.

884 49. Mei Y, Wang X, Zhang J, Liu D, He J, Huang C, et al. Siglec-9 acts as an immune-checkpoint
885 molecule on macrophages in glioblastoma, restricting T-cell priming and immunotherapy
886 response. Nat Cancer. 2023 Sep;4(9):1273–91.

887 50. Bassez A, Vos H, Van Dyck L, Floris G, Arijs I, Desmedt C, et al. A single-cell map of
888 intratumoral changes during anti-PD1 treatment of patients with breast cancer. Nat Med.
889 2021 May;27(5):820–32.

890 51. Phenotype, specificity and avidity of antitumour CD8+ T cells in melanoma | Nature. [cited
891 2024 Aug 29]. Available from: https://www.nature.com/articles/s41586-021-03704-y

892 52. Tietscher S, Wagner J, Anzeneder T, Langwieder C, Rees M, Sobottka B, et al. A
893 comprehensive single-cell map of T cell exhaustion-associated immune environments in
894 human breast cancer. Nat Commun. 2023 Jan 6;14(1):98.

895 53. Tang Z, Kang B, Li C, Chen T, Zhang Z. GEPIA2: an enhanced web server for large-scale
896 expression profiling and interactive analysis. Nucleic Acids Res. 2019 Jul 2;47(W1):W556–60.

897 54. Cui H, Wang C, Maan H, Pang K, Luo F, Duan N, et al. scGPT: toward building a foundation
898 model for single-cell multi-omics using generative AI. Nat Methods. 2024 Aug;21(8):1470–
899 80.

900 55. Szałata A, Hrovatin K, Becker S, Tejada-Lapuerta A, Cui H, Wang B, et al. Transformers in
901 single-cell omics: a review and new perspectives. Nat Methods. 2024 Aug;21(8):1430–43.

902 56. Picelli S, Faridani OR, Björklund ÅK, Winberg G, Sagasser S, Sandberg R. Full-length RNA-seq
903 from single cells using Smart-seq2. Nat Protoc. 2014 Jan;9(1):171–81.

904 57. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, et al. Scikit-learn:
905 Machine Learning in Python. arXiv; 2018 [cited 2024 Aug 29]. Available from:
906 http://arxiv.org/abs/1201.0490

907