

COUNTERFACTUAL HYPOTHESIS TESTING OF TUMOR MICROENVIRONMENT SCENARIOS THROUGH SEMANTIC IMAGE SYNTHESIS

Daniel Li^{1,2}, Qiang Ma^{1,*}, Jennifer Chen^{1,*}, Andrew Liu³, Justin Cheung², Yubin Xie², Herman Gudjonson², Tal Nawy², Dana Pe'er², Itsik Pe'er^{1,*}

1. Columbia University

2. Memorial Sloan Kettering Cancer Center

3. Google

Corresponding Author: ip2169@columbia.edu

ABSTRACT

Recent multiplexed protein imaging technologies make it possible to characterize cells, their spatial organization, and interactions within microenvironments at unprecedented resolution. Although observational data can reveal spatial associations, it does not allow users to infer biologically causative relationships and interactions between cells. To address this challenge, we develop a generative model that allows users to test hypotheses about the effect of cell-cell interactions on protein expression through in silico perturbation. Our Cell-Cell Interaction GAN (CCIGAN) model employs a generative adversarial network (GAN) architecture to generate biologically realistic multiplexed cell images from semantic cell segmentations. Our approach is unique in considering all imaging channels simultaneously, and we show that it successfully captures known tumor-immune cell interactions missed by other state-of-the-art GAN models, and yields biological insights without requiring in vivo manipulation. CCIGAN accepts data from multiple imaging technologies and can infer interactions from single images in any health or disease context.

Keywords— Semantic Image Synthesis, Generative Adversarial Networks, Multiplexed Imaging, Tumor Immunology

1 INTRODUCTION

The spatial patterns of cells in a neighborhood dictates signaling relationships and interactions with the cellular environment that are critical to development and disease. Tumor-immune cell interactions within the tumor microenvironment are a prime example, implicated in many facets of cancer pathogenesis and treatment (1; 2).

Recent advances in multiplexed tissue imaging, from fluorescence microscopy to ion mass spectrometry, are now enabling the acquisition of cell and protein localizations at a high dimensional subcellular resolution, providing unmatched insight into cellular interactions (3; 4). While these emerging technologies improve the resolution of tissue imaging, methods to effectively interpret and analyze the data are lacking.

A model that can predict changes in protein localization—not just abundance—as a result of changes in neighboring cellular environments would provide unique insight into important mechanistic questions, such as how the expression of inhibitory immune checkpoint receptors on a T cell change when it migrates into a particular tumor environment. Here, we present Cell-Cell Interaction Generative Adversarial Network (CCIGAN), a tool that addresses the complexity and limited nature of multiplexed imaging data, and allows for testing counterfactual cell interaction scenarios to elucidate interactive processes within tissues.

*Equal Contribution

Technology	Cell of Interest	Marker of interest	Neighboring cell type	Neighbor Protein Channel
MIBI	CD8 T, Endo.	PD-1	tumor	PD-L1
	CD3 T, Endo., Tumor	β -catenin	tumor, dendritic	dsDNA (11)(12)
	CD8 T, Tumor	pan-keratin	CD8 T	dsDNA (11)(12)
	Macrophage	PD-L1	tumor	PD-L1
t-CyCIF	CD8 T	PD-1	tumor	PD-L1
	CD8 T	LAG3	tumor	PD-L1

Table 1: Rediscovery of known cell-cell interactions as positive controls for CCIGAN

The high dimensionality and scarcity of multiplexed cell datasets pose unique challenges for image modeling and interpretation. Multiplexed imaging quantifies many proteins, each in a separate channel, resulting in high-dimensional data fundamentally different from typical red green blue (RGB) images and rendering conventional image analysis methods inadequate. Moreover, the time intensive nature of multiplexed image generation limits the data availability and scale. This restricts the number of observed instances containing specific cell neighborhoods observed and thus impedes modeling, exploratory analysis, and interpretation.

Existing methods for analyzing imaging data can be broadly grouped into cell classification and cell population analysis, neither of which take advantage of multiplexed imaging to investigate protein localization. Methods for cell classification (5) focus on clustering cell-types by protein characteristics but do not predict protein expression and localization. Bayesian methods for cell-cell interaction analysis (6) only predict scalar magnitudes of protein expressions and lack spatial information. Traditional binning and regression-type methods can be used to quantify neighboring cell contributions, but are fundamentally limited by issues such as bin size selection, high dimensionality, feature extraction and nonlinear effects, motivating the use of alternative neural-network-based methods. Some recent GAN-based methods augment cell imaging data for training purposes (7) or synthesize cell images into multiple channels (8), but none can be used to study cell cell interactions. The few existing conditional cell image synthesis models do not consider a cell’s spatial context or incorporate data multiplexity (9; 10).

We have developed a framework consisting of a deep generative network, CCIGAN, and a companion pipeline for image testing analysis. CCIGAN introduces a novel end-to-end deep architecture that is able to predict multiple channels of pixel-level protein magnitude and spatial information influenced by cell-cell interactions. CCIGAN can conditionally capture realistic but potentially unobserved cell scenarios for rapid counterfactual hypothesis testing. The accompanying data analysis pipeline interprets CCIGAN’s outputs at the neighborhood, single cell, and subcellular levels in a biologically meaningful manner. Lastly, we include an algorithm to search for specific cell scenarios that result in significant protein expression changes. CCIGAN successfully recapitulates well-studied patterns in tumor immuno-biology not captured by other state of the art GAN models and yields novel biological insights in immune cell-tumor interactions that are not plausible to deduce using solely ‘wet lab’ experimental perturbations. Our experiments provide insights into cellular interactions in disease pathogenesis that cannot be achieved by direct observation of multiplexed data alone.

2 RESULTS

We first report results that validate the CCIGAN methodology from a technical perspective. We then walk through a series of positive controls, where CCIGAN rediscovers known biological effects for various cell types (denoted as the cell of interest) due to neighboring cell influences (presented in different microenvironments) in the context of simultaneous different protein channels. We present 6 such controls across two different technologies (Table 1), and then leverage CCIGAN in discovery mode, to identify new potential effects.

The CCIGAN Model and Image Analysis Pipeline. We first considered multiplexed images of triple negative breast cancer (TNBC) biopsies (3) and primary lung cancer (13) that were input into CCIGAN. As a deep conditional image synthesis model, CCIGAN takes as input cell segmentation patches (a subsection of the full cell segmentation, as a cell-typed neighborhood) and outputs

a prediction of the proteins expressed on the cells within the segmentation patch, including their localization (Figure 1; see Online Methods 4.2 for data format and processing).

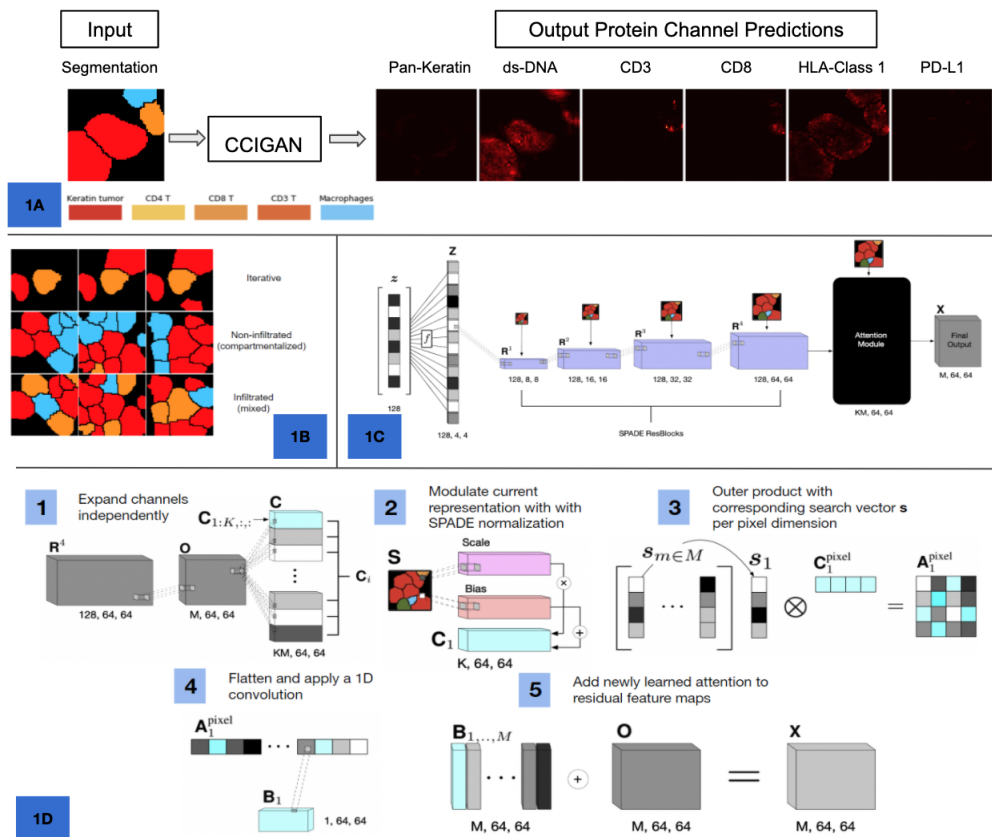


Figure 1: CCIGAN model overview and components (A) CCIGAN takes a labeled segmentation patch as input and returns predicted expression values (between 0 and 1) in each channel. (B) CCIGAN segmentation maps can be synthetically altered to pose different counterfactual biological scenarios for in silico hypothesis testing. Example scenario in the first row tests the effect on a CD8 T cell (orange) of adding tumor cells (red) around it. Second and third row pose different immune infiltration scenarios. (C) CCIGAN convolutional architecture (see Online Methods 4.1.1 for equations). A low dimensional noise vector (z) is convolutionally upsampled and normalized by an input segmentation through a series of spatially adaptive normalizing (SPADE) res-net blocks. The intermediate representation passes through a multiplexed attention module to reason biological relationships. (D) Attention module description (see Online Methods 4.1.2 for equations). The module first disentangles the fully connected representation into individual channels through grouped convolutions, then uses an outer product attention layer to model relationships and interactions between specific markers and cell types. Using an outer product forces attention at a pairwise pixel-level comparison for all combinations of a learned prior vector over markers and different cell types.

In doing so, CCIGAN learns a many-to-many mapping between different cell types and different protein markers. Due to its conditional nature, CCIGAN is able to take any segmentation patch as input, including those with user-defined cell types, cell shapes, and neighbors. Our model predicts protein expression changes in a cell of interest resulting from the influence of neighboring cells. We can ask questions such as “how would a tumor cell be affected by adding a CD8 T cell next to it?” (Fig. 1B) by artificially inserting an adjacent cell in the segmentation patch and observing the inferred protein expression change exerted on the cell of interest. Without CCIGAN, isolating and manipulating cells to investigate their independent protein expressions would be problematic because a simple deletion of adjacent cells around the cell of interest would fail to alter the cell

of interest's protein localization as it will still assume a neighborhood interaction (see S2.3 for an illustration).

Specifically, CCIGAN's convolutional architecture (Fig. 1C, Online Methods 4.1.1) is built on top of spatially adaptive normalization (SPADE) (14). CCIGAN associates the types, shapes, and relative locations of a cell of interest and its adjacent cells to predict spatial maps of protein expression. The ability of SPADE to modulate the neural network's activations based on the context of adjacent cells allows the network to effectively model the behaviors and interactions for cells of interest. A technical exploration into model interpretability is given in S1.4.

Addressing multiplexed data is critical. This is because channels dedicated to biological markers are relatively independent, compared to the typical RGB information. Without addressing this independence, a conventional GAN would equally attend to every location of the current latent representation, even if it is irrelevant to the current protein, making it difficult for a standard RGB multihead decoder to output multiple independent channels. CCIGAN is unique in its ability to tackle such complexity via a novel attention architecture. Since each protein channel is unique, the attention module individually models and conditions the final output of each protein channel on the cell types within the input segmentation patch (Fig. 1D). For example, for the protein marker pan-keratin (typically expressed only in tumor cells), the attention module focuses on generating the subcellular pattern of the tumor cell's pan-keratin while ignoring irrelevant cells. This is achieved through grouped convolutions, learned weights, and Einstein summations (S1.4).

In order to assess the outputs of CCIGAN and extend its utility, we use a series of biologically motivated image analysis techniques and a search algorithm, that we refer to collectively as the image analysis pipeline. Each of these techniques, cumulative expression, weighted centroid, and directional mass shift, quantify how one particular cell's protein expression at a pixel wise level reacts to newly introduced adjacent cells (Fig. 2).

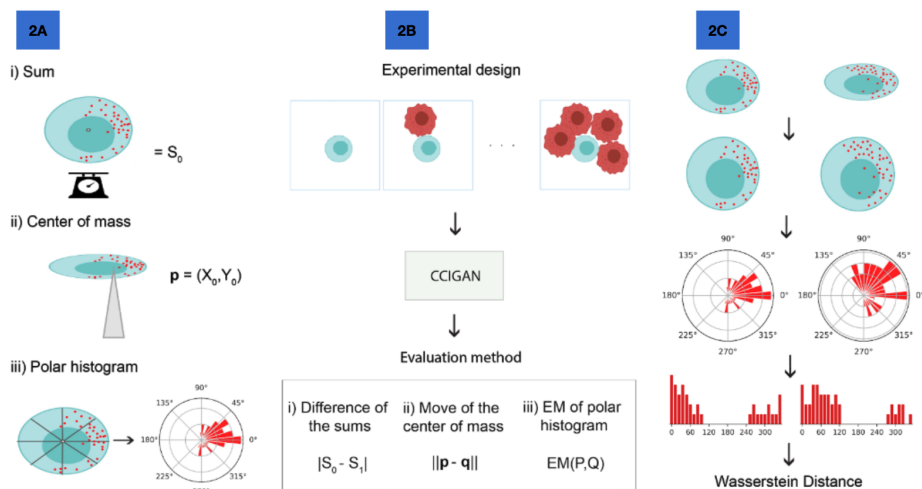


Figure 2: Image analysis pipeline. (A) We designed metrics to evaluate changes in protein expression (red dots) within a cell of interest (blue circle) due to changes in cell neighbors: (i) total protein expression in cell of interest, (ii) center of mass of a protein expression within cell, and (iii) spatial shift in expression. (B) To explore unknown interactions, we provide an automated search algorithm that scans CCIGAN predictions for a wide range of cell types and segmentation arrangements, and assesses changes in a specified metric. Red circles indicate a different cell type. (C) **EM (Earth Mover's) Score.** To evaluate differences in spatial-directional protein expression (directional mass movement), the cell is first warped into a standardized format with interpolated expression values, values are then binned, and a histogram is computed along its polar axis; finally, center of mass geometries are used as guiding directions to positively or negatively weight the earth mover's distance in computing the EM Score. Positive EM Score indicates a cell's protein shift towards another cell whereas a negative EM Score indicates a protein shift in the opposite direction.

Model evaluation and results. CCIGAN is unique among biological tools because current cell image analysis methods either do not perform multichannel or multicellular image synthesis for spatial protein map prediction, or consider protein expression counts only and do not predict a protein expression map.

We first benchmark our model (with and without the attention module) against other current state-of-the-art image synthesis techniques using our image analysis pipeline (Fig. 1B) to establish biological plausibility and using standard image evaluation criteria on the multiplexed ion beam imaging (MIBI) triple negative breast cancer dataset (3). Only CCIGAN is able to recapitulate established patterns of cell-cell interactions linked to known biological phenomena while other models fail to learn correct cell-cell interaction associations. We also report that CCIGAN's image generation outperforms or matches current methods for standard image evaluation criteria (Online Methods 4.3.1).

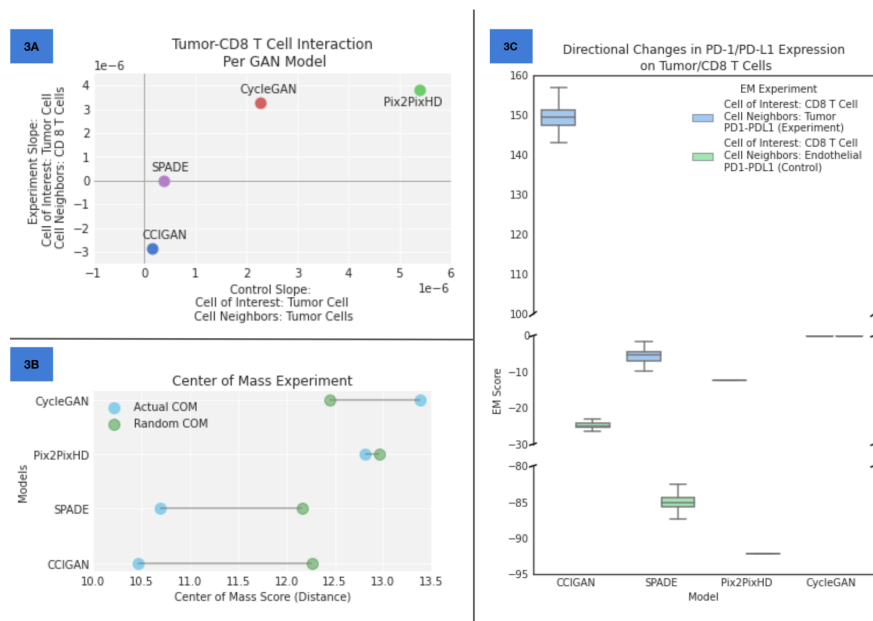


Figure 3: Evaluation of CCIGAN compared to existing GANs. Inference of cell-cell interactions from MIBI images of triple negative breast tumors. For nondeterministic models (CCIGAN and SPADE), we ran each experiment 50 times and use the mean in 3A and 3B, and provide error margins in 3C. **(A)** Predicted effects of neighboring CD8 T or endothelial control cells on pan-keratin levels in a tumor cell of interest. When CD8 T cell's surround a tumor cell, CCIGAN accurately predicts a decrease in total pan-keratin expression. In a control where a tumor cell is surrounded by tumor cells, CCIGAN accurately predicts no significant trend in pan keratin. Slopes represent rates of change in pan-keratin as a function of the number of cell neighbors. **(B)** Differences in predicted center of mass (COM) between T-cell-expressed PD-1 and tumor-cell-expressed PD-L1. Larger differences between the model (blue) and control (green) scores indicate the model is better at capturing the true response of PD-1 in a real scenario. **(C)** Differences in predicted spatial shift of protein expression. Higher Earth Mover's (EM) Scores (15) reflect stronger shifts. Predicted PD-1 expression shift in a T cell of interest resulting from PD-L1-expressing tumor cell neighbors or control endothelial cells. CCIGAN supports the tumor PD-L1 and CD8 T cell PD-1 interaction with large positive EM Score.

First, we examine trends in pan-keratin levels when CD8 T cells mediate tumor killing. Previous literature states that CD8 T cells release granzymes or induce extrinsic apoptosis which results in the release of enzymes that cleave the tumor cell's pan-keratin (15), disrupting cell structure (16) (17) (18) (19). We explored how a drop in pan-keratin levels in a tumor cell of interest could be used as a proxy for the likelihood of tumor cell death (prior to observed cell destruction) in the presence of adjacent CD8 T cells (Fig. 3A, Online Methods 4.3.4, S3.4). CCIGAN predicted decreases in tumor pan-keratin expression that became more dramatic when surrounded by more CD8 T cells, whereas

introducing adjacent tumor cells (control) produced no change. Other methods failed to detect a negative relationship and incorrectly reported changes under control conditions.

Next, we examined how PD-1 expression in a CD8 T cell of interest is affected by adjacent tumor cells (Fig. 2A.ii, Fig. 3B, Online Methods 4.3.2, S3.1). T cells located within the tumor microenvironment upregulate PD-1 expression as a result of influences from the tumor milieu (20; 21; 22). The T cell's PD-1 center of mass (COM) was used as a proxy for its localization, and adjacent tumor cells' PD-L1 center of mass was used as a proxy for cell contact. The score represents the distance between the PD-1 and PD-L1 centers of mass in pixels, with greater predictive accuracy represented by lower scores, since shorter distances are expected for this direct protein-protein interaction (13). Larger differences in score between model and a randomized control indicate the model is better at capturing the true PD-1 response. Using these metrics, CCIGAN predicts better localized PD-L1 expression than other methods such as Pix2PixHD (23) and CycleGAN (24), which do not yield any meaningful biological insight. CCIGAN also outperforms SPADE (14) in achieving the most accurate PD-1 center of mass prediction.

As a final model evaluation, we tested whether CCIGAN can identify cell-cell interaction effects on subcellular protein localization, by hypothesizing that iterative introduction of adjacent PD-L1-expressing tumor cells will induce directional shifts in PD-1 expression in a T cell of interest (Fig. 2A.iii, Fig. 3C, Online Methods 4.3.3). An Earth Mover's (EM) Score with a higher magnitude indicates a stronger cell-cell interaction; positive scores indicate the direction of protein localization moving towards cell neighbors and negative values otherwise. CCIGAN confirmed biological expectation by generating a large positive EM Score, indicating a shift in T-cell PD-1 expression towards the PD-L1 center of mass in adjacent tumor cells (Fig. 3C). In contrast, the addition of control endothelial cells only resulted in a small shift in T-cell PD-1 expression. PD-L1 is found on a variety of endothelial cell types and serves as an important immune checkpoint that protects normal cells from T-cell-driven autoimmune reactions (25; 26); thus, CCIGAN results for the endothelial control also conform to biological expectation. The difference in magnitude and directional effect of T-cell interaction with tumor and endothelial cells highlights the greater degree to which the PD-1/PD-L1 immune checkpoint is exploited by malignant cells to escape immune detection. Off-the-shelf methods SPADE, Pix2PixHD, and CycleGAN fail to capture this interaction.

Applications to MIBI data. To corroborate previously established biology, we design various cell microenvironments and assess these interactions utilizing the image analysis pipeline.

We applied an analysis (similar to Fig. 3C) on β -catenin expression in dendritic cells (DCs), antigen-presenting cells important in the adaptive immune response, upon interacting with tumors. Wnt signaling from tumors have been found to induce β -catenin expression in dendritic cells (DCs), suppressing DC activation (27). Increasing tumor cell presence adjacent to DCs triggers downstream immunosuppressive effects of the Wnt pathway, increasing directional β -catenin expression in the DC cell. To quantify changes of DC cell β -catenin expression due to neighboring cells, we measure mass movement between dsDNA expression of neighbors (indicating magnitude of cell presence) and a DC cell's β -catenin expression (Fig. 4A). Using CCIGAN to add tumor cells next to a DC cell of interest resulted in a directional increase in β -catenin (EM Score 333.9) that is greater than that observed upon adding endothelial cells (EM Score 233.6) or CD3 T cells (EM Score 89.3). These findings confirm that CCIGAN is capable of detecting expected directional patterns of the Wnt pathway activation in DC cells as a result of cell interactions in the TME.

In another study across TNBC patient groups within the MIBI dataset (Fig. 4B, Online Methods 4.3.5, S3.5), CCIGAN corroborated Keren et al.'s predictions on the PD-1/PD-L1 pathway and the variability of PD-L1 expression in macrophage and tumor cells in different TMEs. In compartmentalized TMEs, tumor and immune cells are spatially segregated, whereas in mixed TMEs, tumor and immune cells are mixed together. We trained 4 individual models on 2 exemplar patients chosen per TME type¹ from the TNBC cohort to demonstrate CCIGAN's ability to learn cell-cell interactions dependent on specific TMEs. In the mixed environment (Fig. 4B Upper), CCIGAN first predicted increases in PD-L1 expression in tumors and macrophages in infiltrated neighborhoods when compared to reference neighborhoods containing endothelial (control) cells. Reiterating Keren et al.'s finding in mixed TMEs, we observe tumors expressing the highest average PD-L1 expression, sub-

¹These 4 patients were chosen because they best demonstrate clearly delineated compartmentalized and mixed TMEs.

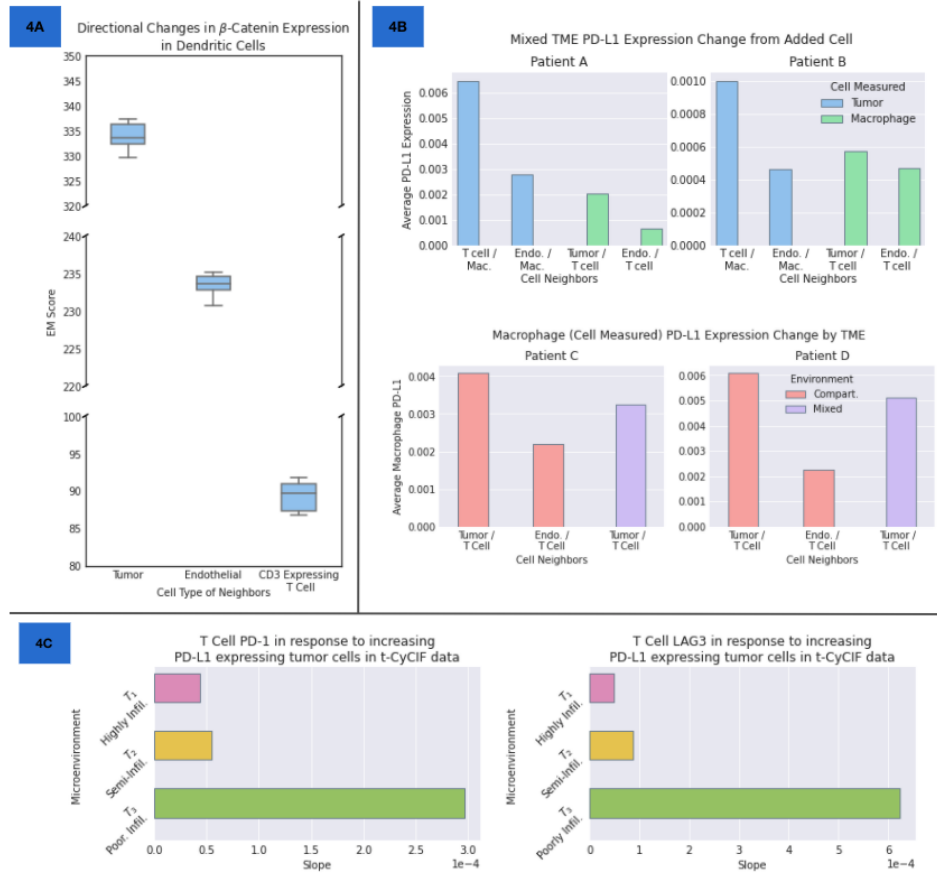


Figure 4: **Recapitulating discovered biological subcellular interactions** (A) Predicted shifts in directional β -catenin expression in dendritic cells (DCs) on adjacent tumor vs non-tumor cells. CCIGAN predicts high EM Score of DC cell β -catenin expression with a tumor cell neighbor, suggesting strong Wnt activation in the DC cell and confirming biological findings on the tumor's immunosuppressive effects on the DC. (B) Effects of tumor compartmentalization in tumor-immune cell interactions (1B Row 2), evaluated on MIBI data from TNBC patients in Keren et al. 200 test neighborhood segmentations were generated to simulate different environments. Control/reference: endothelial (endo.) cell. (Upper) Within the mixed tumor microenvironments experimental set up, PD-L1 expression in tumor and macrophage cells is induced by T cells & macrophages or tumors & T cells respectively, representing infiltrated neighborhoods. The PD-L1 expression in tumor or macrophage cells (legend) was measured next to their appropriate cell neighbor combinations (x -axis). (Lower) Macrophage PD-L1 expression increases induced by neighboring tumor cells within compartmentalized and mixed cell environments. Greater PD-L1 expression increase in macrophages in the compartmentalized environment than mixed environment recapitulates previous literature. (C) CCIGAN inference of cell-cell interactions from t-CyCIF images of primary lung cancer (13). (Left) PD-1 expression in CD8 T cells as a function of number of PD-L1 expressing tumor cell neighbors and type of tumor environment. Slope is calculated similarly to 3A. In an infiltrated environment (high lymphocyte presence within the tumor microenvironment, indicating strong anti-tumor immune response), reduced PD-1 expression on infiltrating T cells signals a more robust anti-tumor immune response. (Right) LAG3 expression in CD8 T cells as a function of PD-L1 expressing tumor cell neighbors and type of tumor environment. In an infiltrated environment, a decrease in T cell LAG3 expression indicates reduced T cell exhaustion in situations where there is a robust anti-tumor immune response.

stantially higher than macrophage PD-L1 expression, in their infiltrated neighborhoods. We also observed Keren et al.'s second claim of macrophages exhibiting the highest PD-L1 expression in compartmentalized TMEs (Fig. 4B Lower). Models trained on compartmentalized patient data were used to evaluate macrophage PD-L1 expression on both compartmentalized and mixed neighborhoods and predicted greater PD-L1 expression in macrophages that reside in tumor compartments as opposed to mixed environments (Fig. 4B Lower). This phenomenon further reveals CCIGAN's capability to learn general biological interactions even if trained on different TMEs. Ultimately, CCIGAN is robust when trained on individual patient samples and supports biological findings.

Applications to Tissue Cyclic Immunofluorescence (t-CyCIF). t-CyCIF data is obtained through a novel iterative immunofluorescence staining procedure that works with existing fluorescence microscopes (28). By applying CCIGAN to a t-CyCIF primary lung cancer image dataset (13), we demonstrate CCIGAN's ability to generalize across datasets and learn microenvironment-specific cell-cell interactions.

Utilizing a patient sample that exhibited progressive tumor cell growth, we split the image data according to three different stages of tumor-infiltration by immune cells – poorly-infiltrated (T_1), semi-infiltrated (T_2), and heavily-infiltrated (T_3) (S2.2.2). We trained one CCIGAN model per microenvironment to summarize its cell-cell interactions and allow for differential analysis of protein-level interactions.

Similar to the CD8/pan keratin trend experiment (Fig. 3A), we examined the correlation between PD-1 and PD-L1 in tumor and T-cell patches in each microenvironment-specific model (Fig. 4C Left, S3.6). In the non-infiltrated microenvironment model, we observe a substantially higher trend of T-cell PD-1 expression as a function of tumor area than in the semi-infiltrated which, in turn, was also slightly higher than the heavily-infiltrated microenvironments. These findings are consistent with the idea that tumors with poor lymphocyte-infiltration have had greater success in suppressing anti-tumor immune responses and thus, an increase in PD-1 would be expected in the T cell (cell of interest).

In a similar fashion (Fig. 4C Right, S3.6), we examined the association between exhaustion marker LAG3 (29) (30) expression on T cells, and PD-L1 expression on neighboring tumor cells. CCIGAN predicted an increase in LAG3 expression on the T cell (cell of interest) as tumor neighbors were added, with greater levels of induction from infiltrated to semi-infiltrated to non-infiltrated tumor scenarios. These results agree with the biological expectation that increased tumor PD-L1 expression would suppress T cell function and induce T cell exhaustion, and that T cells within tumor environments of higher lymphocyte-infiltration would have lower LAG3 expression, indicating less T cell exhaustion.

The t-CyCIF findings highlight the utility of CCIGAN in reliably reproducing biologically expected patterns of immune cell-tumor cell interaction, even from limited patient tumor samples. Moreover, CCIGAN robustly quantifies protein expression changes as a result of cell-cell interactions.

Search algorithm for detecting interactions. We developed a search algorithm to identify additional cell-cell interactions (Online Methods 4.4). Evaluating MIBI data (31), we find that while most changes in protein expressions are plausible and/or significant (Fig. 24), there are some erroneous correlations (e.g. a CD8 T cell expressing PD-L1 while in close contact to tumor cells) that are attributed to poor segmentations, imaging faults, and other noisy effects. In other experiments, we not only confirm our previous findings but also posit further interactions of biological significance for later exploration. A further discussion and investigation to these cell-cell interactions, and in particular, vimentin, are given in S4.2.

3 DISCUSSION

While multiplexed imaging technologies continue to rapidly increase in resolution and availability, tools to analyze and interpret their outputs remain primitive or nonexistent. To address these limitations, we introduced the idea of applying image synthesis in the form of a novel generative architecture, CCIGAN, and developed a computational image analysis pipeline to interpret CCIGAN's output. By simultaneously and generatively associating cell shapes, type, and sizes to their protein expression patterns, CCIGAN allows users to specify these parameters in a test environment to pose counterfactual cell-cell scenarios and observe their resulting protein expressions. Subse-

quently, users can employ the image analysis pipeline to interpret and quantify protein expression at the subcellular level as a function of their defined variables.

CCIGAN's capacity for generating subcellular protein predictions represents a step forward in our ability to understand cellular relationships within a microenvironment. Rather than assessing in vivo incidence of cell interaction phenomena, CCIGAN allows for hypothetical biological situations to be generated and analyzed. In other words, any individual cellular responses to a neighborhood of cells of any identity can be assessed without having to seek this specific occurrence within the available biological tissue sample. This addresses the issue of limited multiplexed data and greatly expands the power of individual datasets for causal inference.

On tumor data, we have shown that CCIGAN outperforms state-of-the-art image synthesis methods, and is able to capture, recapitulate, and quantify clinically established cell-cell interactions. CCIGAN performed well on both mass-spectroscopy and immunofluorescence based images, and can interrogate cell-cell interactions in any biological or disease contexts beyond cancer.

Despite recent advances in imaging technology, current datasets are of limited utility for studying cell-cell interactions, as protein markers are typically chosen for cell typing rather than functional protein localization. However, as these technologies become pervasive, more protein markers will also become available, furthering the need for CCIGAN. Although CCIGAN represents a promising rapid hypothesis testing tool for in silico prediction of cell-cell interactions, guided in vivo experiments are needed to corroborate CCIGAN's quantitative predictions. The greatest limitations in the current model are in the input data's noisy cell segmentations and cell type classifications. While the model performs well in learning and associating clear trends and could potentially correct for erroneous classifications, exceedingly rare cell-cell protein interactions may be interpreted as noise or with low confidence. Nonetheless, given the strong corroboration with previous clinically established biology and the utility of giving users freedom to posit new interactions to guide further experiments, we expect this work to be useful in the multiplexed imaging domain.

4 ONLINE METHODS

4.1 METHODS

4.1.1 MODEL OVERVIEW

For CCIGAN, we use SPADE residual blocks (14) as our generative backbone and DCGAN’s discriminator’s architecture (32). (14) have shown SPADE to be an effective way to inject conditioning into a generative model. The SPADE normalization layer serves as a replacement for previous layer normalization techniques. Instead of learning a universally shared per channel affine transformation, like in Batch or Instance Normalization, SPADE learns to predict affine transformations based on segmentation maps; each feature is uniquely transformed based on its cell type, size, and neighboring cells. The ability for SPADE to modulate activations based on the context of adjacent cell segmentations allows the network to effectively model the behaviors and interactions of cells. The input of CCIGAN is a noise vector $\mathbf{z} \in \mathbb{R}^{128}$ and a segmentation map \mathbf{S} . f denotes a linear layer $\mathbb{R}^{128} \mapsto \mathbb{R}^{2048}$. \mathbf{R}^i are feature map representations from SPADE resblocks and \mathbf{X} denotes the final output of M cell expressions. Below, each layer’s output dimensions are given next to their respective equations. Further details such as kernel size, activation functions, training regimen, and model interpretability are given in S1.1.

$$\mathbf{Z} \in \mathbb{R}^{(128,4,4)} = f(\mathbf{z}) \quad (1) \quad \mathbf{R}^3 \in \mathbb{R}^{(128,32,32)} = \text{SPADE_RESBLK}(\mathbf{R}^2, \mathbf{S}) \quad (4)$$

$$\mathbf{R}^1 \in \mathbb{R}^{(128,8,8)} = \text{SPADE_RESBLK}(\mathbf{Z}, \mathbf{S}) \quad (2) \quad \mathbf{R}^4 \in \mathbb{R}^{(128,64,64)} = \text{SPADE_RESBLK}(\mathbf{R}^3, \mathbf{S}) \quad (5)$$

$$\mathbf{R}^2 \in \mathbb{R}^{(128,16,16)} = \text{SPADE_RESBLK}(\mathbf{R}^1, \mathbf{S}) \quad (3) \quad \mathbf{X} \in \mathbb{R}^{(M,64,64)} = \text{ATTENTION}(\mathbf{R}^4, \mathbf{S}) \quad (6)$$

4.1.2 ATTENTION MODULE

Our architectural contribution is a protein marker dependent attention module in the final output layer. The goal of the attention module is to condition the final output of a channel on a protein marker m and \mathbf{S} ’s cell types. For example the protein marker, pan-keratin m_{pk} , is expressed exclusively in tumor cells but not in other cells. Appropriately, an attention mechanism should attend to tumor cells and ignore irrelevant cells in \mathbf{S} for m_{pk} . To replicate a marker searching for specific cell types that express it, we define a learned persistent vector for each marker denoted by $\mathbf{s}_{m \in M} \in \mathbb{R}^8$ that undergo a series of operations with the final feature map representation attending to m ’s specific cell types. It is also worthwhile to note that these persistent vectors \mathbf{s}_m offer a degree of model interpretability that mimic real world markers. The current input dimensions to the attention module are $\mathbb{R}^{(128,64,64)}$ following the last resblock \mathbf{R}^4 and m indexes from $1, \dots, M$.

$$\mathbf{O} \in \mathbb{R}^{(M,64,64)} = \text{CONV2D}(\mathbf{R}^4) \quad (7) \quad \mathbf{A}_i \in \mathbb{R}^{(|S| \times K, 64, 64)} = \mathbf{C}_i \otimes \mathbf{s}_m \quad (10)$$

$$\mathbf{C} \in \mathbb{R}^{(KM, 64, 64)} = \text{CONV2D}(\mathbf{O}) \quad (8) \quad \mathbf{B}_i \in \mathbb{R}^{(1, 64, 64)} = \sigma(\text{CONV2D}(\mathbf{A}_i)) \quad (11)$$

$$\mathbf{C}_i \in \mathbb{R}^{(K, 64, 64)} = \text{SPADE}(\mathbf{C}_{K(i-1):K i, :, :}, \mathbf{S}) \quad (9) \quad \mathbf{X} \in \mathbb{R}^{(M, 64, 64)} = \mathbf{O} + \mathbf{B}_{1, \dots, M} \quad (12)$$

After \mathbf{R}^4 , a bottleneck convolution is applied to match the original data’s dimension as \mathbf{O} (step 1), which is used in a residual manner with the final output. Intuitively at this stage, \mathbf{O} ’s feature maps resemble the target \mathbf{Y} , but we wish to further refine the output channels. We convolve \mathbf{O} into MK channeled features for each protein marker where $K = 8$. Considering each \mathbf{C}_i where $i \in \{1, \dots, M\}$ as a group of K channels, the model spatially adaptive normalizes each \mathbf{C}_i and computes an outer product with the corresponding persistent vector \mathbf{s}_i and \mathbf{C}_i . The resulting matrix is flattened and convolved (with a kernel size of 1 on the pixel level) from $\mathbf{A}_i \in \mathbb{R}^{(|S| \times K, 64, 64)} \mapsto \mathbb{R}^{(1, 64, 64)}$ followed by a sigmoid $\sigma(\cdot)$ activation. Lastly, the attentions $\mathbf{B}_{1, \dots, M}$ are added to \mathbf{O} to obtain the output \mathbf{X} .

Initially, the model has no priors over the interaction of protein markers and cell types. The proposed outer product attention layer (outer product and 1×1 convolution) excels at modeling these relationships and interactions between specific markers and cell types. By using an outer product, the model forces attention at a pairwise pixel level comparison for all combinations of elements between \mathbf{s}_m and \mathbf{A}_i . As training progresses, both the learned features over segmentation patches and the learned persistent vectors \mathbf{s}_m improve, in turn allowing the learned 1×1 convolution to reason about positive or negative relationships from the pairwise pixel combinations.

4.2 DATA AND DATA PROCESSING

We trained CCIGAN on two types of cell data, MIBI-TOF (multiplexed ion beam imaging by time-of-flight) and t-CyCIF (tissue cyclicimmunofluorescence) data. While these two types of multiplexed image data were obtained through different procedures, they share structural similarities.

Multiplexed cell images display multiple protein marker expression levels. These images are represented as high dimensional tensors $\mathbf{T} \in \mathbb{R}^{(M,H,W)}$, M being the number of markers, H being the height of the image, and W being the width. Each of these markers $m \in \{1, \dots, M\}$, are given as a channel taking on real values continuous in $[0, 1]$ at each (x, y) coordinate, indicating the expression level at a given protein. Protein markers' particular expression levels (either separately or in conjunction with other protein markers) demarcate different cellular subtypes and furthermore, are indicative of the functional properties of a cell. By simultaneously imaging over multiple protein markers, multiplexed images are able to identify cell type as well as provide detailed information of sub-cellular structure, cell neighbors, and interactions in the tumor microenvironment across these different marker settings.

4.2.1 MIBI-TOF

MIBI-TOF images are represented in $\mathbf{T} \in \mathbb{R}^{(M,2048,2048)}$. These images are then further processed at a cell by cell basis into $\mathbf{Y} \in \mathbb{R}^{(M,64,64)}$ patches, where a cell is at the center of the patch along with its neighbors. Next, we construct semantic segmentation maps $\mathbf{S} \in \mathbb{R}^{(C+1,64,64)}$, where a vector $\mathbf{S}_{:,i,j}$ is one-hot encoded based on a cell type $C = 17$, and the $C + 1$ -th channel denotes empty segmentation space. The data is train-test split at a 9:1 ratio at the MIBI-TOF image level to avoid cell neighborhood bias.

Data obtained through MIBI-TOF characterized tissue samples were collected from triple-negative breast cancer (TNBC) patients. MIBI-TOF images over 36 protein markers, but $M = 24$ markers were used in our training. A description of the technology and full list of these markers is given in S2.1.

4.2.2 T-CYCIF

T-CyCIF (tissue cyclicimmunofluorescence) images of primary lung squamous cell carcinoma are represented in $\mathbf{T} \in \mathbb{R}^{(M,H \approx 12000, W \approx 14000)}$. Segmentation patches are constructed in an similar fashion. The main salient difference from MIBI is the processed patch size $\mathbf{Y} \in \mathbb{R}^{(M,128,128)}$, and the segmentation patch size $\mathbf{Y} \in \mathbb{R}^{(C+1,128,128)}$. This was intentionally done to demonstrate the scalability of CCIGAN. 44 markers were imaged in t-CyCIF, however we excluded background cell markers to yield $M = 37$ markers. A description of the technology and full list of these markers is given in the S2.2.

Unlike MIBI, a significant amount of data processing was done in order to analyze the data. Full treatment of data is given in S2.2.

4.3 EVALUATION

To conduct fair experiments, all models were optimized, tuned, and set with similar parameters. They were also taken from their official online implementations and trained for 120 epochs or until convergence (max 150). CCIGAN is identical to our designed SPADE comparison baseline with the exception of the attention module.

4.3.1 IMAGE EVALUATION AND RECONSTRUCTION

First, we use the following evaluation metrics in order to compare with baseline results: adjusted L_1 and MSE score, L_1 and MSE score, structural similarity (SSIM) index (33) and cell based mutual information (MI) shown in Table 2. Bolded scores indicate the best scores. Equations and motivation are given in S3.

Three evaluation metrics were then used to conduct experiments and validate the trained model's utility in generating biologically meaningful cellular proteins in the tumor microenvironment and

Metrics	CCIGAN	SPADE	Pix2PixHD	CycleGAN
Adjusted L_1 Score	0.613	0.618	0.875	4.745
L_1 Score	0.594	0.602	0.745	3.959
Adjusted MSE Score	0.026	0.031	0.061	1.841
MSE Score	0.026	0.031	0.055	1.523
SSIM	0.810	0.802	0.709	0.394
Cell Mutual Information	10.46	10.25	9.26	7.96

Table 2: Comparison of conventional reconstruction metrics between different models.

ability to recapitulate and *quantify* previously established biological phenomena. Each subsection provides additional relevant information.

4.3.2 CENTER OF MASS (COM)

For a generated cell image, its weighted centroid, or center of mass, is the mean position of all the points in the cell weighted by a particular channel expression. Given a cell image $\mathbf{X} \in \mathbb{R}^{(H,W)}$, with indices of the segmented cell $V \subseteq \{1, \dots, H\} \times \{1, \dots, W\}$, the COM $\bar{\mathbf{p}} = (\bar{x}, \bar{y})$ is defined as $\bar{x} = \frac{\sum_{(x,y) \in V} x \mathbf{X}_{x,y}}{\sum_{(x,y) \in V} \mathbf{X}_{x,y}}$ and $\bar{y} = \frac{\sum_{(x,y) \in V} y \mathbf{X}_{x,y}}{\sum_{(x,y) \in V} \mathbf{X}_{x,y}}$.

For example, in the PD-1/PD-L1 experiment, we compute the COM of the CD8 T cell (cell of interest) weighted by PD-1 expression, given as $\bar{\mathbf{p}}^{\text{CD8}}$, and the COM of all tumor cells weighted by PD-L1 expression, given as $\bar{\mathbf{p}}^{\text{Tumor}}$. Since T cells located within the tumor microenvironment often have upregulated expression of PD-1, we assume that $\bar{\mathbf{p}}^{\text{CD8}}$ should have the roughly the same PD-L1 COM of all its surrounding tumor cells $\bar{\mathbf{p}}^{\text{Tumor}}$. The center of mass score is defined below as the relative distance between $\bar{\mathbf{p}}^{\text{CD8}}$ and $\bar{\mathbf{p}}^{\text{Tumor}}$, where N is defined as the number of patches:

$$\text{COM}_{\text{projection}} = \frac{1}{N} \sum_{i=1}^N \|\bar{\mathbf{p}}_i^{\text{CD8}} - \text{Proj}_{\text{CD8}}(\bar{\mathbf{p}}_i^{\text{Tumor}})\|_2 \quad (13)$$

The projection function $\text{Proj}(\cdot)$ is used to project $\bar{\mathbf{p}}^{\text{Tumor}}$ onto the CD8 T cell to ensure the expected COM of the tumor cells is inside of the CD8 T cell. As a reference we choose a random position $\bar{\mathbf{p}}^{\text{Random}}$ in the CD8 T cell (PD-1) which replaces $\bar{\mathbf{p}}^{\text{CD8}}$ in Eq. 13 and compute the random COM score to show the effectiveness of the result. An example illustration is given in the figure in S3.1.

4.3.3 EARTH MOVER'S SCORE

We develop Earth Mover's (EM) Score to measure the dissimilarity between two distributions and the direction of distributional shift. While Earth Mover's Distance is a nonnegative measurement of shift between two distributions, we incorporate the direction of shift in protein localization (positive indicating cell's protein shifting towards another cell, negative otherwise) into our EM Score. For our Pan Keratin/CD8 experiment, we used EM to evaluate the shift in protein localization between CD8 T cells and tumor cells. EM can be generalized to experiments involving other cell-cell interactions, but for the purposes of clarity, below we describe our EM approach in the context of the Pan Keratin/CD8 experiment.

For a segmentation map, we add T tumor cells around one CD8 T cell. The COM for the t -th tumor ($t \in \{1, \dots, T_i\}$) is defined as $\bar{\mathbf{p}}_t^{\text{Tumor}}$. Similarly, the PD-1 COMs of the CD8 T cell by adding the t -th tumor is defined by $\bar{\mathbf{p}}_t^{\text{CD8}}$. Initially when there are no tumor cells, $\bar{\mathbf{p}}_0^{\text{CD8}}$ is the centroid of the CD8 T cell.

We proceed to define vector a vector \mathbf{v}_t which points from $\bar{\mathbf{p}}_0^{\text{CD8}}$, to the COM of the t -th tumor cell $\bar{\mathbf{p}}_t^{\text{Tumor}}$. We also define vector \mathbf{u}_t which points from the previous COM $\bar{\mathbf{p}}_{t-1}^{\text{CD8}}$ to the current COM $\bar{\mathbf{p}}_t^{\text{CD8}}$ of the CD8 T cell. We define θ_t as the angle between $\mathbf{v}_t, \mathbf{u}_t$.

If $\cos \theta_t > 0$, that is to say if the cosine similarity is positive, the COM of a CD8 T cell $\bar{\mathbf{p}}_t^{\text{CD8}}$, moves correctly towards the COM of the added tumor cell $\bar{\mathbf{p}}_t^{\text{Tumor}}$. An illustration of the points and vectors is given in the figure in S3.2.

Formally:

$$\mathbf{v}_t = \bar{\mathbf{p}}_t^{\text{Tumor}} - \bar{\mathbf{p}}_0^{\text{CD8}}, \mathbf{u}_t = \bar{\mathbf{p}}_t^{\text{CD8}} - \bar{\mathbf{p}}_{t-1}^{\text{CD8}}, \cos \theta_t = \frac{\mathbf{u}_t \cdot \mathbf{v}_t}{\|\mathbf{u}_t\| \cdot \|\mathbf{v}_t\|} \quad (14)$$

After obtaining the directional information, we use EM Score (34) to evaluate the changes in PD-1 expression of the CD8 T cell. Earth Mover's Score, which measures the dissimilarity of two distributions, is used in this context to measure the protein localization shifts in PD-1 before and after adding a tumor cell. We consider each cell \mathbf{X} in polar coordinates (r, θ) with respect to its centroid, integrate its expression along the radius coordinates, and evaluate the resulting histogram $\text{hist}(\mathbf{X})$ along the angle coordinate. The 2nd figure in S3.2 shows an example histogram of cells by coordinate location.

This allows for the definition of distance for moving one histogram to another, i.e. $\text{em}(\mathbf{X}_i^t, \mathbf{X}_i^{t-1}) = d_{\text{EM}}(\text{hist}(\mathbf{X}_i^t), \text{hist}(\mathbf{X}_i^{t-1}))$, for the generated PD-1 expression of the CD8 T cell \mathbf{X}_i^t when adding the t -th tumor cell.

The final EM Score is defined as:

$$\text{EM Score} = \frac{1}{\sum_{i=1}^N T_i} \sum_{i=1}^N \sum_{t=1}^{T_i} \mathbb{1}(\|\mathbf{X}_i^t\| > \|\mathbf{X}_i^{t-1}\|) \cdot \text{em}(\mathbf{X}_i^t, \mathbf{X}_i^{t-1}) \cdot \cos \theta_{t,i} \quad (15)$$

where the indicator function $\mathbb{1}(\cdot) = 1$ if and only if $\|\mathbf{X}_i^t\| > \|\mathbf{X}_i^{t-1}\|$, otherwise $\mathbb{1}(\cdot) = 0$. This ensures that the biological constraint of PD-1 expression increasing as a response to added tumor cells is met. Recall, if $\cos \theta_t > 0$, $\bar{\mathbf{p}}_t^{\text{CD8}}$ has moved in the direction of $\bar{\mathbf{p}}_t^{\text{Tumor}}$, implying the shift in PD-1 expression is correct, and in turn increases the EM Score. By contrast, the EM Score decreases when $\bar{\mathbf{p}}_t^{\text{CD8}}$ moves in the opposite direction.

This can be adapted and used for any two channel protein interactions (such as in the β -catenin and dsDNA experiment for analyzing Wnt pathways).

4.3.4 PROTEIN EXPRESSION AND CELL SURFACE AREA TREND EXPERIMENTS

In this experiment, we used a Student's t -test as the statistical hypothesis test to evaluate the correlations of the protein expression of a specified cell as a function of the area/number of surrounding cells in the specified cell's microenvironment.

Given a generated protein channel $\mathbf{X}_i \in \mathbb{R}^{(H,W)}$ and the segmentation map channel for the surrounding cells $\mathbf{S}_i \in \mathbb{R}^{(H,W)}$, we compute the total area of the cells $a_i = \sum_{h=1}^H \sum_{w=1}^W \mathbf{S}_i$, and the total expression level of the specified cell $e_i = \sum_{h=1}^H \sum_{w=1}^W \mathbf{X}_i$. We then regress $\{e_i\}_{i=1}^N$ on $\{a_i\}_{i=1}^N$ and assess significance of the slope using a t -test against the null of no change in expression as a function of surrounding cells.

Additional figures of other models and explanations for pan-keratin and CD8 are given in S3.4. Additional figures for t-CyCIF experiments are given in S3.6 and S3.7.

4.3.5 TUMOR INFILTRATED AND COMPARTMENTALIZED MICROENVIRONMENTS

CCIGAN was used to compare protein localization in tumor infiltrated versus compartmentalized microenvironments. We used CCIGAN to predict on 200 directly manipulated mixed and non-mixed tumor environment segmentation patches. For each experiment, we challenge the cell in question with an opposing cell in a microenvironment with macrophages (for example T Cell is challenged with Tumor cell to result in increased PD-1 expression in the T Cell) and use endothelial cells as control cells to show our result has biological significance. Similar to Online Methods 4.3.4's experimental settings, we compute the average expression of a specific marker for the cells of interest for all patches.

Data resulting from the experiment is located in S3.5. The increase in PD-L1 for the above tumor and macrophage scenarios (S3.5 table 7) indicate that CCIGAN has appropriately captured previously reported biological outcomes and is capable of quantifying these phenomena at single cell levels. Furthermore, the model is adaptable to various different types of tumor architecture depending on its training set to produce different hypothesis testing environments.

4.4 SEARCH ALGORITHM

Here we provide a modular search algorithm framework to try to discover further cell-cell interactions in other channels. As a high level overview, the algorithm uses CCIGAN to automate and measure a specific cell's change in a specified protein's expression level due to user specified microenvironment changes. The algorithm allows a user to change and specify such changes as cell type, shape, size, and quantity. If the change in expression is greater than a user specified input, then the particular instance is logged. It is important to note that this tool is meant to guide and search for particular interesting interactions and still susceptible to issues such as noisy segmentations.

Algorithm 1: Search Algorithm

Input: Cell segmentation list $\{S_i\}_{i=1}^n$, the cell type c_f for the cell of interest, the protein index m that is measured, potential neighbors and their specified types $c_i \in C_{types}$ fixed noise δ , threshold β , a chosen metric d_y , and the generator G .

Randomly chose initial cell index $i_0 \in \{1, \dots, n\}$;

Input segmentation $S^{INPUT} = S_{i_0}$;

Assign cell type S^{INPUT} to c_f ;

Mask for the initial cell $U = \sum_{j=1}^n S_{i_0}[j, :, :]$;

Generated image $X_0 = G(\delta, S^{INPUT})$;

Specified channel $M_0 = U * X_0[m]$;

Show S^{INPUT} and M_0 ;

for $k = 1$ **to** $n - 1$ **do**

 Random index $i_k \in \{1, \dots, n\} \setminus \cup_{j=0}^{k-1} \{i_j\}$;

 Random index $c_i \in C_{types}$;

 Assign S_{i_k} to cell type c_i ;

$S^{INPUT} = S^{INPUT} + S_{i_k}$;

$X_k = G(\delta, S^{INPUT})$;

$M_k = U * X_k[m]$;

 Show S^{INPUT} and M_k ;

$E_k = d_y(M_{k-1}, M_k)$;

if $E_k > \beta \sum_{i,j} U_{i,j}$ **then**

 Log Significance;

if $\sum_{i,j} M_{k-1,i,j} < \sum_{i,j} M_{k,i,j}$ **then**

 Log Increase;

else

 Log Decrease;

REFERENCES

- [1] Gregory L. Beatty and Whitney L. Gladney. Immune escape mechanisms as a guide for cancer immunotherapy. *Clinical Cancer Research*, 21(4):687–692, 2015. ISSN 1078-0432. doi: 10.1158/1078-0432.CCR-14-1860. URL <https://clincancerres.aacrjournals.org/content/21/4/687>.
- [2] Els M. E. Verdegaal, Noel F. C. C. de Miranda, Marten Visser, Tom Harryvan, Marit M. van Buuren, Rikke S. Andersen, Sine R. Hadrup, Caroline E. van der Minne, Remko Schotte, Hergen Spits, John B. A. G. Haanen, Ellen H. W. Kapiteijn, Ton N. Schumacher, and Sjoerd H. van der Burg. Neoantigen landscape dynamics during human melanoma-t cell interactions. *Nature*, 536(7614):91–95, 2016. doi: 10.1038/nature18945. URL <https://doi.org/10.1038/nature18945>.
- [3] Michael Angelo, Sean C Bendall, Rachel Finck, Matthew B Hale, Chuck Hitzman, Alexander D Borowsky, Richard M Levenson, John B Lowe, Scot D Liu, Shuchun Zhao, Yasodha Natkunam, and Garry P Nolan. Multiplexed ion beam imaging of human breast tumors. *Nature medicine*, 20(4):436–442, 04 2014. doi: 10.1038/nm.3488. URL <https://pubmed.ncbi.nlm.nih.gov/24584119>.
- [4] Leeat Keren, Marc Bossé, Diana Marquez, Roshan Angoshtari, Samir Jain, Sushama Varma, Soo-Ryum Yang, Allison Kurian, David Valen, Robert West, Sean Bendall, and Michael Angelo. A structured tumor-immune microenvironment in triple negative breast cancer revealed by multiplexed ion beam imaging. *Cell*, 174:1373–1387.e19, 09 2018. doi: 10.1016/j.cell.2018.08.039.
- [5] Hartland W. Jackson, Jana R. Fischer, Vito R. T. Zanotelli, H. Raza Ali, Robert Mechera, Savas D. Soysal, Holger Moch, Simone Muenst, Zsuzsanna Varga, Walter P. Weber, and Bernd Bodenmiller. The single-cell pathology landscape of breast cancer. *Nature*, 578(7796):615–620, 2020. doi: 10.1038/s41586-019-1876-x. URL <https://doi.org/10.1038/s41586-019-1876-x>.
- [6] Damien Arnol, Denis Schapiro, Bernd Bodenmiller, Julio Saez-Rodriguez, and Oliver Stegle. Modeling cell-cell interactions from spatial molecular data with spatial variance component analysis. *Cell Reports*, 29(1):202 – 211.e6, 2019. ISSN 2211-1247. doi: <https://doi.org/10.1016/j.celrep.2019.08.077>. URL <http://www.sciencedirect.com/science/article/pii/S2211124719311325>.
- [7] Piotr Baniukiewicz, E. Josiah Lutton, Sharon Collier, and Till Bretschneider. Generative adversarial networks for augmenting training data of microscopic cell images. *Frontiers in Computer Science*, 1:10, 2019. ISSN 2624-9898. doi: 10.3389/fcomp.2019.00010. URL <https://www.frontiersin.org/article/10.3389/fcomp.2019.00010>.
- [8] Anton Osokin, Anatole Chessel, Rafael Edgardo Carazo-Salas, and Federico Vaggi. Gans for biological image synthesis. *CoRR*, abs/1708.04692, 2017. URL <http://arxiv.org/abs/1708.04692>.
- [9] Gregory R. Johnson, Rory M. Donovan-Maiye, and Mary M. Maleckar. Generative modeling with conditional autoencoders: Building an integrated cell, 2017.
- [10] Ting Zhao and Robert F. Murphy. Automated learning of generative models for subcellular location: Building blocks for systems biology. *Cytometry Part A*, 71A(12):978–990, 2007. doi: 10.1002/cyto.a.20487. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/cyto.a.20487>.
- [11] Leeat Keren, Marc Bosse, Steve Thompson, Tyler Risom, Kausalia Vijayaragavan, Erin McCaffrey, Diana Marquez, Roshan Angoshtari, Noah F Greenwald, Harris Fienberg, et al. Mibitof: A multiplexed imaging platform relates cellular phenotypes and tissue structure. *Science advances*, 5(10):eaax5851, 2019.

-
- [12] Leeat Keren, Marc Bosse, Diana Marquez, Roshan Angoshtari, Samir Jain, Sushama Varma, Soo-Ryum Yang, Allison Kurian, David Van Valen, Robert West, et al. A structured tumor-immune microenvironment in triple negative breast cancer revealed by multiplexed ion beam imaging. *Cell*, 174(6):1373–1387, 2018.
- [13] Rumana Rashid, Giorgio Gaglia, Yu-An Chen, Jia-Ren Lin, Ziming Du, Zoltan Maliga, Denis Schapiro, Clarence Yapp, Jeremy Muhlich, Artem Sokolov, et al. Highly multiplexed immunofluorescence images and single-cell data of immune markers in tonsil and lung cancer. *Scientific Data*, 6(1):1–10, 2019.
- [14] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. *CoRR*, abs/1903.07291, 2019. URL <http://arxiv.org/abs/1903.07291>.
- [15] Norio Azumi and Hector Battifora. The distribution of vimentin and keratin in epithelial and nonepithelial neoplasms: a comprehensive immunohistochemical formalin-and alcohol-fixed tumors. *American journal of clinical pathology*, 88(3):286–296, 1987.
- [16] Luis Martínez-Lostao, Alberto Anel, and Julián Pardo. How do cytotoxic lymphocytes kill cancer cells?, 2015.
- [17] RG Oshima. Apoptosis and keratin intermediate filaments. *Cell Death & Differentiation*, 9(5):486–492, 2002.
- [18] V Badock, U Steinhusen, K Bommert, B Wittmann-Liebold, and A Otto. Apoptosis-induced cleavage of keratin 15 and keratin 17 in a human breast epithelial cell line. *Cell Death & Differentiation*, 8(3):308–315, 2001.
- [19] Nam-On Ku, Pavel Strnad, Heike Bantel, and M Bishr Omary. Keratins: Biomarkers and modulators of apoptotic and necrotic cell death in the liver. *Hepatology*, 64(3):966–976, 2016.
- [20] Mojgan Ahmadzadeh, Laura A. Johnson, Bianca Heemskerk, John R. Wunderlich, Mark E. Dudley, Donald E. White, and Steven A. Rosenberg. Tumor antigen-specific CD8 T cells infiltrating the tumor express high levels of PD-1 and are functionally impaired. *Blood*, 114(8):1537–1544, 08 2009. ISSN 0006-4971. doi: 10.1182/blood-2008-12-195792. URL <https://doi.org/10.1182/blood-2008-12-195792>.
- [21] Shuming Chen, George A. Crabill, Theresa S. Pritchard, Tracee L. McMiller, Ping Wei, Drew M. Pardoll, Fan Pan, and Suzanne L. Topalian. Mechanisms regulating pd-11 expression on tumor and immune cells. *Journal for ImmunoTherapy of Cancer*, 7(1):305, 2019.
- [22] David Escors, María Gato-Cañas, Miren Zuazo, Hugo Arasanz, María Jesus García-Granda, Ruth Vera, and Grazyna Kochan. The intracellular signalosome of pd-11 in cancer cells. *Signal transduction and targeted therapy*, 3:26–26, 09 2018.
- [23] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [24] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, 2017.
- [25] Michael J Eppihimer, Jason Gunn, Gordon J Freeman, Edward A Greenfield, Tetyana Chernova, Jamie Erickson, and John P Leonard. Expression and regulation of the pd-11 immunoinhibitory molecule on microvascular endothelial cells. *Microcirculation (New York, N.Y. : 1994)*, 9(2):133–145, 04 2002.
- [26] Nancy Rodig, Timothy Ryan, Jessica A. Allen, Hong Pang, Nir Grabie, Tatyana Chernova, Edward A. Greenfield, Spencer C. Liang, Arlene H. Sharpe, Andrew H. Lichtman, and Gordon J. Freeman. Endothelial expression of pd-11 and pd-12 down-regulates cd8+ t cell activation and cytotoxicity. *European Journal of Immunology*, 33(11):3117–3126, 2003. doi: 10.1002/eji.200324270. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/eji.200324270>.

- [27] Daniel Swafford and Santhakumar Manicassamy. Wnt signaling in dendritic cells: its role in regulation of immunity and tolerance. *Discovery medicine*, 19(105):303, 2015.
- [28] Jia-Ren Lin, Benjamin Izar, Shu Wang, Clarence Yapp, Shaolin Mei, Parin M Shah, Sandro Santagata, and Peter K Sorger. Highly multiplexed immunofluorescence imaging of human tissues and tumors using t-cycif and conventional optical microscopes. *Elife*, 7, 2018.
- [29] Long Long, Xue Zhang, Fuchun Chen, Qi Pan, Pronnaphat Phiphatwatchara, Yuyang Zeng, and Honglei Chen. The promising immune checkpoint lag-3: from tumor microenvironment to cancer immunotherapy. *Genes & cancer*, 9(5-6):176, 2018.
- [30] Lawrence P Andrews, Ariel E Marciscano, Charles G Drake, and Dario AA Vignali. Lag 3 (cd 223) as a cancer immunotherapy target. *Immunological reviews*, 276(1):80–96, 2017.
- [31] S. C. Angelo, M. and Bendall, R. Finck, M. B. Hale, C. Hitzman, A. D. Borowsky, R. M. Levenson, J. B. Lowe, S. D. Liu, S. Zhao, Y. Natkunam, and G. P. Nolan. Multiplexed ion beam imaging of human breast tumors. *Nature Medicine*, 20:436–442, 2014.
- [32] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *arXiv e-prints*, art. arXiv:1511.06434, Nov 2015.
- [33] Zhou Wang, Alan C Bovik, Hamid R Sheikh, Eero P Simoncelli, et al. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [34] Yossi Rubner, Carlo Tomasi, and Leonidas J Guibas. The earth mover’s distance as a metric for image retrieval. *International journal of computer vision*, 40(2):99–121, 2000.
- [35] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [36] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *CoRR*, abs/1802.05957, 2018. URL <http://arxiv.org/abs/1802.05957>.
- [37] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2794–2802, 2017.
- [38] Satelli Arun and Li Shulin. Vimentin as a potential molecular target in cancer therapy or vimentin, an overview and its potential as a molecular target for cancer therapy. *Cell Mol Life Sci*, 68(18):3033–3046, 2011.
- [39] Martha E Kidd, Dale K Shumaker, and Karen M Ridge. The role of vimentin intermediate filaments in the progression of lung cancer. *American journal of respiratory cell and molecular biology*, 50(1):1–6, 2014.

Supplementary Materials

1 MODEL

1.1 MODEL ARCHITECTURE

The detailed architecture of our generator is shown on Table 3.

Layers	Output Size	Generator
Linear	(128, 4, 4)	Linear 128 × 2048
Upsampling	(128, 8, 8)	Upsampling 2 × 2
SPADE ResBlk-1	(128, 8, 8)	SPADE 128, Leaky ReLU Convolution 3 × 3 SPADE 128, Leaky ReLU Convolution 3 × 3
Upsampling	(128, 16, 16)	Upsampling 2 × 2
SPADE ResBlk-2	(128, 16, 16)	SPADE 128, Leaky ReLU Convolution 3 × 3 SPADE 128, Leaky ReLU Convolution 3 × 3
Upsampling	(128, 32, 32)	Upsampling 2 × 2
SPADE ResBlk-3	(64, 32, 32)	SPADE 128, Leaky ReLU Convolution 3 × 3 SPADE 64, Leaky ReLU Convolution 3 × 3 SPADE 64, Leaky ReLU Shortcut Convolution 3 × 3
Upsampling	(64, 64, 64)	Upsampling 2 × 2
SPADE ResBlk-4	(64, 64, 64)	SPADE 64, Leaky ReLU Convolution 3 × 3 SPADE 64, Leaky ReLU Convolution 3 × 3
Convolution	(24, 64, 64)	Leaky ReLU, Convolution 5 × 5
Convolution	(24 * 8, 64, 64)	Leaky ReLU, Convolution 5 × 5
Group SPADE	(24 * 8, 64, 64)	[SPADE 8] * 24
Modulation	(24 * 64, 64, 64) (24, 64, 64)	[Outer Product 8 ⊗ 8] * 24 Convolution 1 × 1, Sigmoid
Output	(24, 64, 64)	Sum residual, Sigmoid

Table 3: Architecture details of CCIGAN’s generator

where ResBlk is the residual block with skip connection used in ResNet (35), and SPADE is the spatially-adaptive normalization layer. The detailed architecture of our discriminator is shown on Table 4.

Layers	Output Size	Discriminator
Conv-1	(32, 32, 32)	Convolution 4×4 , stride 2 Instance Norm, Leaky ReLU
Conv-2	(64, 16, 16)	Convolution 4×4 , stride 2 Instance Norm, Leaky ReLU
Conv-3	(128, 8, 8)	Convolution 4×4 , stride 2 Instance Norm, Leaky ReLU
Conv-4	(256, 4, 4)	Convolution 4×4 , stride 2 Instance Norm, Leaky ReLU
Conv-5	(512, 2, 2)	Convolution 4×4 , stride 2 Instance Norm, Leaky ReLU
Conv-6	(1, 1, 1)	Convolution 3×3 , stride 2 Sigmoid

Table 4: Architecture details of CCIGAN’s discriminator

1.2 IMPLEMENTATION DETAILS AND TRAINING REGIMEN

Our implementation of the generator applies Spectral Norm to all layers (36). The discriminator’s input is the output of the generator concatenated with the segmentation patch $[\mathbf{X}, \mathbf{S}]$ and $[\mathbf{Y}, \mathbf{S}]$ for the ground truth. Finally CCIGAN uses ADAM ($lr_G = 0.0004$, $lr_D = 0.0001$) with GAN loss and feature matching loss. Full training details and loss functions are as follows.

1.3 MODEL TRAINING

G is the generator and D is the discriminator for CCIGAN. Given segmentation map \mathbf{S} , ground truth \mathbf{Y} and noise δ , the generated image is $\mathbf{X} = G(\mathbf{S}, \delta)$. The input of the discriminator is the cell image conditioned on the segmentation map \mathbf{S} . We use LSGAN loss (37) in CCIGAN, which is defined as follows:

$$L_{GAN}(G, D) = \mathbb{E}_{\mathbf{Y}, \mathbf{S}} [\|D(\mathbf{Y}, \mathbf{S})\|_2] + \mathbb{E}_{\mathbf{S}} [\|1 - D(G(\mathbf{S}, \delta), \mathbf{S})\|_2] \quad (16)$$

In addition to GAN loss, we also use feature matching loss (23) during training expressed as:

$$L_{FM}(G, D) = \mathbb{E}_{\mathbf{Y}, \mathbf{S}} \sum_{j=1}^J \frac{1}{N_j} [\|D_j(\mathbf{Y}, \mathbf{S}) - D_j(G(\mathbf{S}, \delta), \mathbf{S})\|_1] \quad (17)$$

where D_j is j -th layer feature map of the discriminator for $j \in \{1, \dots, J\}$, and N_j is the number of elements in j -th layer. Consequently, the objective function for training is given as follows:

$$\min_G \left(\left(\max_D L_{GAN}(G, D) \right) + \lambda L_{FM}(G, D) \right) \quad (18)$$

where $\lambda = 10$. Due to the size of cell patch is (64, 64), we do not use multi-scale discriminators and perceptual loss in CCIGAN and other baseline models e.g. SPADE and pix2pixHD.

In training, we use ADAM as the optimizer. The generator learning rate is $lr_G = 0.0004$ and the discriminator learning rate is $lr_D = 0.0001$. We train CCIGAN 120 epochs with a training set of 5648 cell patches. We train other baseline models for 120 epochs or until they converge (max 150). The full details of training of CCIGAN and baselines are shown as Table 5. The hyperparameters of each model are fine-tuned to get better performance. The training time was roughly equal for all models. In particular, CCIGAN was around 1.2 times slower than the SPADE baseline on a single Tesla V100 GPU.

Metrics	Ours	SPADE	Pix2PixHD	CycleGAN
lr_G	0.0004	0.0008	0.0002	0.0002
lr_D	0.0001	0.0001	0.0002	0.0002

Table 5: Hyperparameters of models

1.4 MODEL INTERPRETABILITY AND GENERATIVENESS

Examining the model’s persistent vectors s_m , we can try to understand if there is a match between real world protein markers and the representations of s_m . For example, the vector s_{pk} for pan-keratin attends to tumor cells and s_{CD8} attends to CD8 T cells at pixel pairwise levels. It follows that in a simple experiment where corresponding $s_{CD8} \leftrightarrow s_{pk}$ vectors are exchanged internally in the attention module (Eq. 10, Fig. 1D (Step 3) , outer product) we may observe a lower expression for tumor cells in channel m_{pk} and a lower expression for CD8 T cells in channel m_{CD8} since tumor cells do not express CD8 and CD8 T cells do not express pan-keratin. As a control, we also switch surface membrane markers HLA Class 1 and dsDNA markers as they are present in all cells and have very similar average expression values ($s_{HLAcl} \leftrightarrow s_{dsDNA}$). Accordingly, for our control, we expect to see negligible changes. We define the expression ratio as $\frac{\text{after}}{\text{before}} - 1$.

Protein Markers	CD8	pan keratin	HLA Class 1	dsDNA
Expression Ratios	-0.373	-0.145	-0.054	-0.0012

Table 6: s_m persistent vector interpretability experiments.

In Table 12, we can see a larger magnitude decrease of the expression ratios in the $s_{CD8} \leftrightarrow s_{pk}$ experiment and a minute difference in the $s_{HLAcl} \leftrightarrow s_{dsDNA}$. Further visualizations (Fig. 5) and discussion (model generativeness, Fig. 6) are as follows.

Fig. 5 shows the persistent vectors s_i for all proteins. Note the similarity between CD3 and CD8 T cell protein markers and the similarity between dsDNA and HLA Class 1 surface membrane proteins (expressed in all cells). It is also important to make the distinction that sparse markers (while different) are similar in state. This is due to the lack of training data for rare cell types, making it difficult for the model to reason on such a small sample size.




Figure 5: Persistent vectors s for various channels.

Fig. 6 shows the generativeness of CCIGAN through an uncertainty map over 100 instances (random noise). An uncertainty map shows the differences per pixel (x, y) location. The higher intensities indicate a higher probability of changing at the specified (x, y) location.




Figure 6: Uncertainty maps illustrating model generativeness.

2 DATA AND DATA PROCESSING

2.1 MIBI-TOF

Multiplexed ion beam imaging by time-of-flight (MIBI-TOF) represents a novel technology that can accurately quantify and spatially resolve cellular protein expressions at the single cell level within tissue samples. Given a tissue sample that is first stained with protein-specific antibodies tethered to elemental metals, MIBI-TOF bombards the sample with atomic ions (i.e. O_2^+) from a primary ion beam. This causes the release of elemental isotopes and tissue-specific material which can be quantified in a mass spectrometer (31).

For MIBI, the markers used (total 24) in our experiments are: Pan-Keratin, EGFR, Beta catenin, dsDNA, Ki67, CD3, CD8, CD4, FoxP3, MPO, HLA-DR, HLA-Class-1, CD209, CD11b, CD11c, CD68, CD63, Lag3, PD1, PD-L1, IDO, Vimentin, SMA, CD31. The markers we didn't use (total 12) in our experiments are: CD16, B7H3, CD45, CD45RO, Keratin17, CD20, CD163, CD56, Keratin6, CSF-1R, p53, CD138. These were not used primarily due to blank expressions.

2.2 t-CyCIF

The t-CyCIF images are obtained through iterative cycles of incubating the sample with antibodies (markers) conjugated directly with fluorophores that will bind to a specific protein of interest, imaging the samples with a microscope to record the light emitted at each location of the protein, then deactivating the fluorescence signal by soaking the sample in a compound (28). At each iteration, different antibodies that bind to different proteins of interests are applied. After obtaining multiple images with few channels at each iteration, the images are stitched together to produce a high-dimension image.

For t-CyCIF, the markers used (total 37) are: DAPI1, DAPI2, DAPI3, LAG3, ARL13B, DAPI4, KI67, KERATIN, PD1, DAPI5, CD45RB, CD3D, PDL1, DAPI6, CD4, CD45, CD8A, DAPI7, CD163, CD68, CD14, DAPI8, CD11B, FOXP3, CD21, DAPI9, IBA1, ASMA, CD20, DAPI10, CD19, GFAP, GTUBULIN, DAPI11, LAMINAC, BANF1, LAMINB. The markers we didn't use (total 7) are: A488background1, A555background1, A647background1, A488background2, A555background2, A647background2, A488background3.




Figure 7: Examples generated from a segmentation for certain channels for different models. The segmentation patch is the one hot encoded patch collapsed and colored into 1 channel. The horizontal labels represent protein markers and the vertical labels are each of the generative models.

2.2.1 T-CYCIF DATA PREPROCESSING

For t-CyCIF data, each cell was first clustered to one of many cell types. Using cell features (log cell expression data), we clustered the cells using the 26 non-DAPI markers. We exponentiated the cell features data to restore the data to raw cell intensities, then quantile clipped each marker to retain only the data from 1% to 99.5%. After 0-1 min-max rescaling the data, we clustered the cells using the Phenograph tool, developed in the Dana Pe'er Lab.

Let N denote the nuclear probability matrix of the image where $N_{i,j} = \Pr[\text{nucleus at index } (i, j)]$ and let N' be the 0-1 nuclear segmentation mask, $N'_{i,j} = \mathbb{1}[\text{cell nucleus at index } (i, j)]$. Cytoplasm probabilities were also provided, the cytoplasm probability mask is denoted as C , where $C_{i,j} = \Pr[\text{cytoplasm at index } (i, j)]$. N, N', C were provided to us from (?). We modified the cytoplasm probability mask to a 0-1 cytoplasm mask C' where $C'_{i,j} = \mathbb{1}[C_{i,j} > 0.4 \wedge N_{i,j} < 0.4]$ by choosing these thresholds ourselves. Then we overlaid the nuclear and cytoplasm masks together to form a mask T of the total cell, where $T_{i,j} = \mathbb{1}[\text{a cell occurs at index } (i, j)]$. Because there are small holes in this mask, we filled the holes using SK-image's morphology tools. To yield the final segmentation mask F , we dilated the nuclear segmentation mask N' with 2 iterations and perform an element-wise mask with mask T , where $F_{i,j} = \mathbb{1}[N'_{i,j} = 1 \wedge T_{i,j} = 1]$.

2.2.2 T-CYCIF DATA SPLIT

We split the t-CyCIF primary lung cancer dataset into four segments to test our hypothesis that CCIGAN can learn biological relationships specific to tumor microenvironments with varying characteristics. Three segments (highly, semi, and poorly infiltrated) were each manually selected based on the visible proportion of lymphocyte presence within the tumor microenvironment. Figures 8 and 9 show the segments in relation to the lymphocyte and keratin expression. More detailed descriptions of each individual dataset segment are as follows.

1. T_1 Infiltrated - High proportion of lymphocyte presence within the tumor microenvironment. Indicative of a strong inflammatory anti-tumor immune response.
2. T_2 Semi-infiltrated - Medium levels of lymphocyte infiltration within the tumor microenvironment.
3. T_3 Poorly-infiltrated - Few lymphocytes present within the tumor microenvironment indicative of a poor anti-tumor immune response.

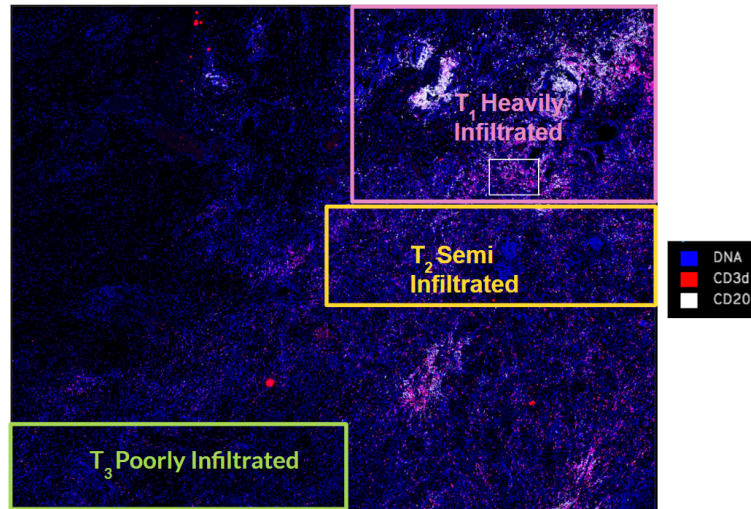


Figure 8: Splits of the tissue with lymphocyte cell protein markers highlighted




Figure 9: Splits of the tissue with keratin positive tumor cell type highlighted

2.3 DATA LIMITATIONS

Examples are given in Fig. 10.

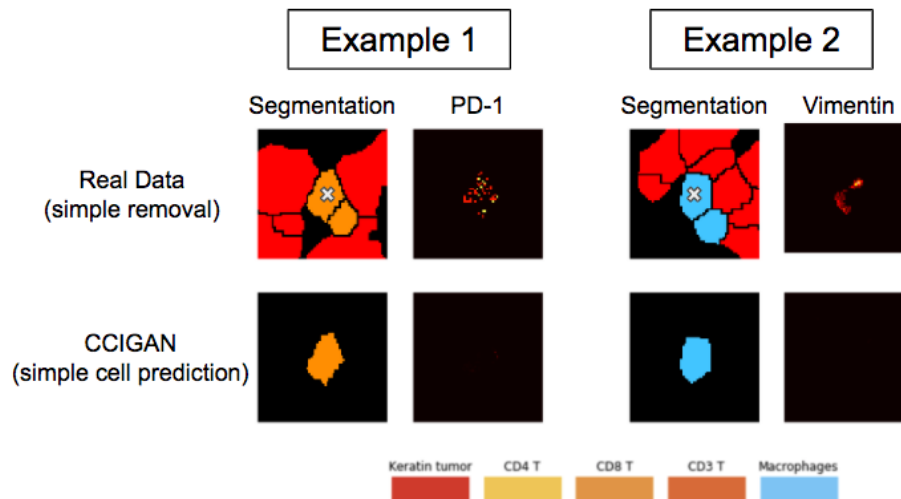


Figure 10: Data limitations regarding manually isolating cells and their protein localizations. Simply deleting cells surrounding a target cell does not reveal the true protein localization of the target cell (with the X marker), as the protein localization will still assume a neighboring interactions. In Example 1, the target CD8 T cell should only express PD-1 when it is surrounded by PD-L1 expressing tumor cells which upregulate the PD-1 expression. Note in the real data how deleting the surrounding cells do not change its PD-1 protein localization. CCIGAN learns that for an isolated CD8 T cell, PD-1 is not upregulated and therefore not expressed.

3 EVALUATION

3.1 CENTER OF MASS (COM)

Fig. 11 shows an illustration of the center of mass and projected center of mass of a patch of tumor and CD8 T cells.




Figure 11: An example illustration of the center of mass (COM) nomenclature. Note the projection onto the CD8 T cell. This provides a more consistent measurement across different patches by projecting p^{Tumor} onto the CD8 T cell.

3.2 EARTH MOVER'S SCORE

Figures 12 and 13 demonstrate the directional and histogram shifts as a function of adding more tumor cells.




Figure 12: An example illustration of the points and vector nomenclature from Online Methods 4.3.3. The blue dots are the expression of PD-1 and PD-L1 proteins. The cyan arrows show the vectors v_t and u_t . Note the shift in expression of the PD-1 as a response to the added tumor's PD-L1 expression.

3.3 RECONSTRUCTION METRICS

Given the generated image set $\mathcal{X} = \{\mathbf{X}_i\}_{i=1}^N$ and the ground truth set $\mathcal{Y} = \{\mathbf{Y}_i\}_{i=1}^N$ with $\mathbf{X}_i, \mathbf{Y}_i \in \mathbb{R}^{(M,H,W)}$, the L_1 /MSE score is defined as follows,

$$L(\mathcal{X}, \mathcal{Y}) = \sum_{i=1}^N \sum_{m=1}^M \|\text{sort}(U_i \odot \mathbf{X}_{i,m}) - \text{sort}(U_i \odot \mathbf{Y}_{i,m})\|_* \quad (19)$$

where $\|\cdot\|_*$ can be either L_1 or L_2 norm, \odot is the element-wise product, $\mathbf{X}_{i,m}$ and $\mathbf{Y}_{i,m}$ are the m -th channel of the i -th cell patch, $U_i \in \{0, 1\}^{(H,W)}$ is the mask matrix which masks all the cells in i -th

patch. For any matrix \mathbf{A} , $\text{sort}(\mathbf{A})$ is the sort function that sorts all entries of \mathbf{A} . The sorting function ensures our metrics are position independent and only measures the intensity of the generated image and ground truth. The score function $L(\mathcal{X}, \mathcal{Y})$ only computes the loss of sorted expression inside of the cells. Then we add penalization for expression outside of cells. The adjusted L_1/MSE score is introduced as follows,

$$L_{\text{adj}}(\mathcal{X}, \mathcal{Y}) = \sum_{m=1}^M \left(\|\text{sort}(\mathbf{U}_i \odot \mathbf{X}_{i,m}) - \text{sort}(\mathbf{U}_i \odot \mathbf{Y}_{i,m})\| - \|\text{sort}((\mathbf{1}_d - \mathbf{U}_i) \odot \mathbf{X}_{i,m}) - \text{sort}((\mathbf{1}_d - \mathbf{U}_i) \odot \mathbf{Y}_{i,m})\| \right) \quad (20)$$

where $\mathbf{1}_d$ is the matrix with all entries equal to 1. A smaller score indicates a better result.

For any two images $\mathbf{X}, \mathbf{Y} \in [0, 1]^{(H,W)}$, the SSIM and MI are defined as:

$$\text{SSIM}(\mathbf{X}, \mathbf{Y}) = \frac{(2\mu_X\mu_Y + c_1)(2\sigma_{XY} + c_2)}{(\mu_X^2 + \mu_Y^2 + c_1)(\sigma_X^2 + \sigma_Y^2 + c_2)} \quad (21) \quad I(\mathbf{X}; \mathbf{Y}) = H(\mathbf{X}) + H(\mathbf{Y}) - H(\mathbf{X}, \mathbf{Y}) \quad (22)$$

where $H(\cdot)$ is entropy, μ_X and σ_X are the mean and standard deviation of \mathbf{X} , c_1, c_2 are constants. Then the SSIM between \mathcal{X}, \mathcal{Y} is

$$\text{SSIM}(\mathcal{X}, \mathcal{Y}) = \frac{1}{N} \frac{1}{M} \sum_{i=1}^N \sum_{m=1}^M \text{SSIM}(\mathbf{X}_{i,m}, \mathbf{Y}_{i,m}) \quad (23)$$

In cell based MI, test patches are processed at a cell-cell basis where their mutual information is computed with the corresponding cell in the ground truth. For the generated image \mathbf{X}_i of the i -th patch, we assume there are T_i cells in the i -th patch. Then for each cell t , the pixels of m -th channel of the t -th cell in the i -th patch can be expressed as a vector $\mathbf{x}_{i,m}^t$. Hence, the cell based MI is formulated as:

$$I(\mathcal{X}; \mathcal{Y}) = \frac{1}{\sum_{i=1}^N T_i} \frac{1}{M} \sum_{i=1}^N \sum_{m=1}^M \left(\sum_{t=1}^{T_i} I(\mathbf{x}_{i,m}^t; \mathbf{y}_{i,m}^t) \right) \quad (24)$$

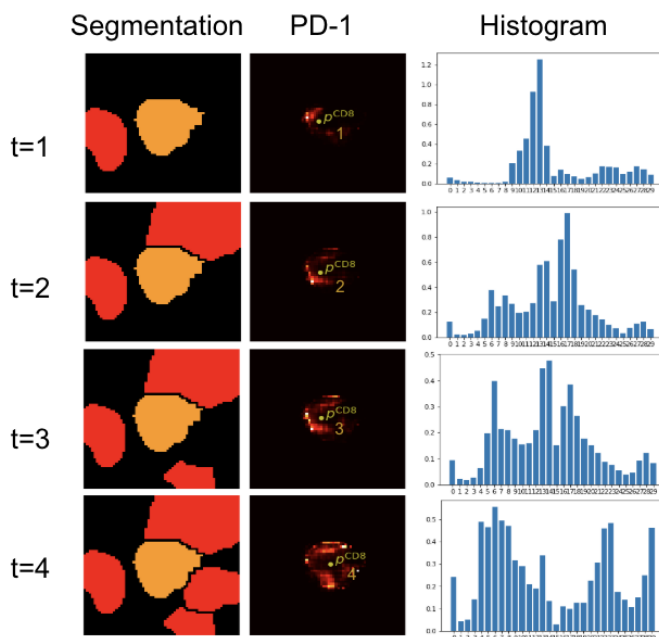


Figure 13: Example illustration of how a CD8 T cell's (orange) PD-1 histogram changes as a function of iteratively added tumor cells.

The SSIM measures the similarity between the generated image and the ground truth. For SSIM, we use HLA Class I and dsDNA due to their expressions in all cells. If all channels were considered, the SSIM would be uninformative due to the majority of the channels being blank or sparse. The MI measures the information shared between generated image and ground truth at a cell by cell basis where we consider all channels. Consider the example where a model generates no expression in marker m but the real data has expression in m , the MI would be 0 and vice versa. Higher SSIM and MI values indicate better results.

3.4 PAN-KERATIN AND CD8 EXPRESSION

The pan-keratin/CD8 experiment is similar to Fig. 22's orientation except the center cell (cell of interest) is a tumor cell (red) and the adjacent neighboring cells are CD8 T cells (orange). CCIGAN predicted a decrease in tumor cell pan-keratin expression with respect to increasing CD8 T cell area/number (Fig. 14). This is juxtaposed to the tumor cell control where there is no change in the pan-keratin level as the number of neighboring tumor cells is increased.

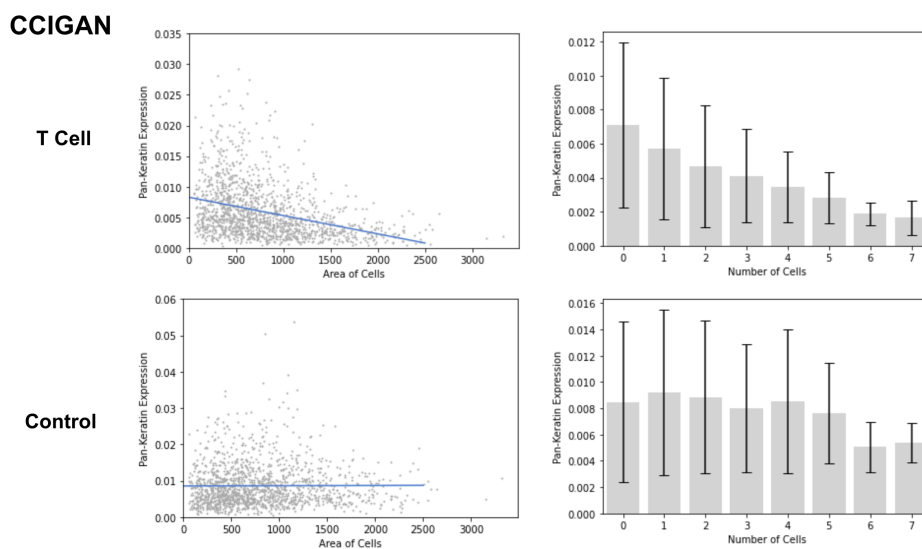


Figure 14: CCIGAN experiment for adding CD8 T cells and tumor cells (control) around a tumor cell.

SPADE does not predict a decrease in tumor cell pan-keratin expression with respect to increasing CD8 T cell area/number and shows no difference in pan-keratin expression trends between the T cell and control groups (Fig. 15).

pix2pixHD erroneously predicts an increase in tumor cell pan-keratin expression with respect to increasing CD8 T cell area/number and shows no difference in pan-keratin expression trends between the T cell and control groups (Fig. 16).

CycleGAN fails to predict a decrease in tumor cell pan-keratin expression with respect to increasing CD8 T cell area/number and shows no difference in pan-keratin expression trends between the T cell and control groups (Fig. 17).

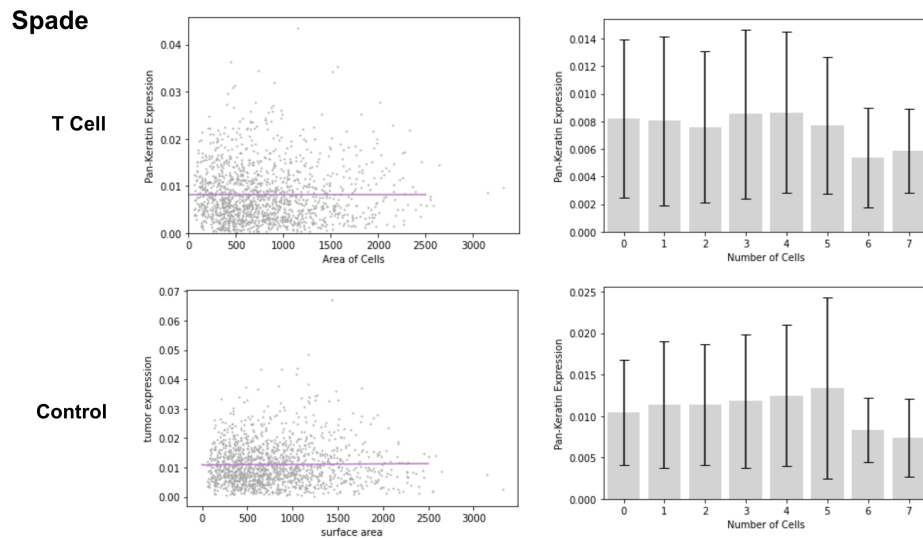


Figure 15: SPADE experiment for adding CD8 T cells and tumor cells (control) around a tumor cell.

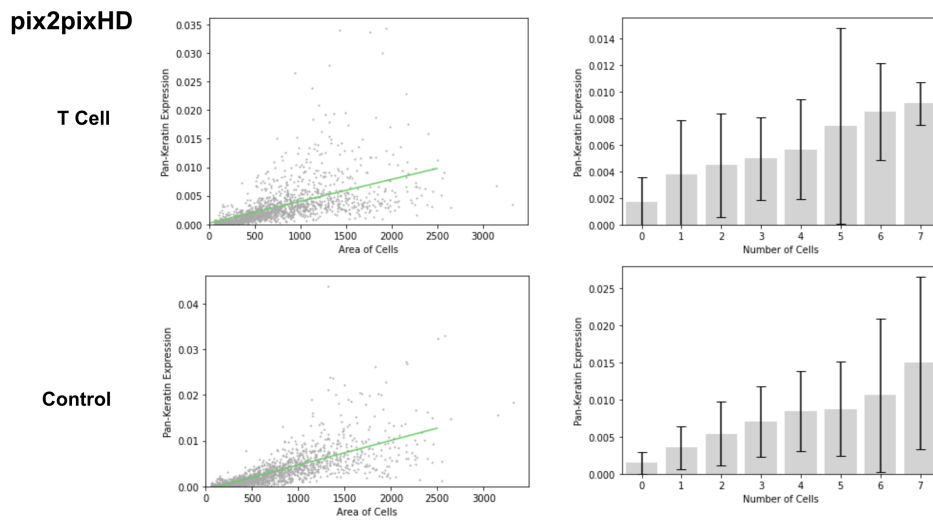


Figure 16: pix2pixHD experiment for adding CD8 T cells and tumor cells (control) around a tumor cell.

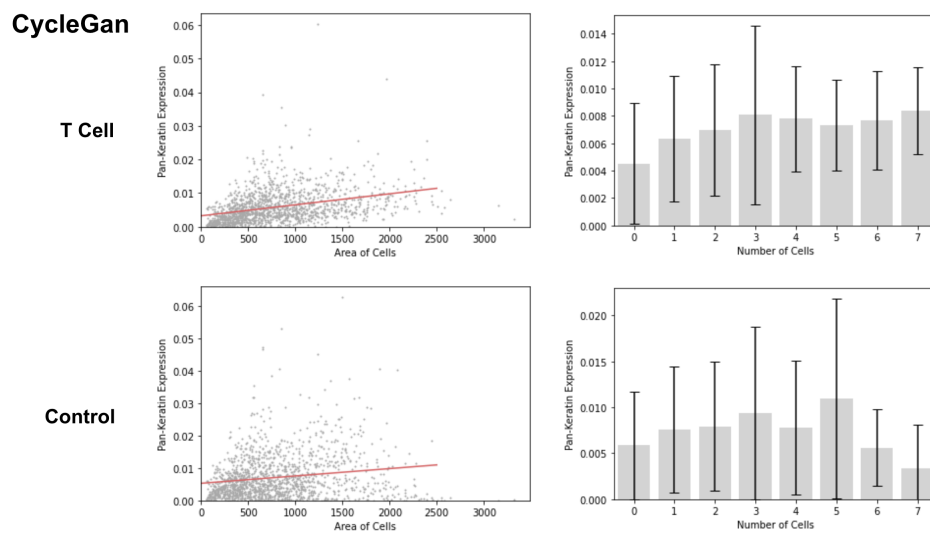


Figure 17: CycleGAN experiment for adding CD8 T cells and tumor cells (control) around a tumor cell.

3.5 TUMOR INFILTRATED AND COMPARTMENTALIZED MICROENVIRONMENTS

Keren et al. (2018) determined that in situations of mixed tumor-immune environments, where immune cells freely infiltrated the tumor, the tumor cells predominantly expressed PD-L1. Conversely, in situations of compartmentalized tumors, where there is a greater degree of physical separation between immune and tumor cells, macrophages were the predominant source of expressed PD-L1, particularly at the tumor boundary.

These findings were recapitulated by CCIGAN. For a patient with a mixed tumor environment, when trained with mixed patient samples, CCIGAN reported increased PD-L1 expression on tumor cells. Furthermore, CCIGAN was able to quantify this difference in expression at the single cell level, reporting a tumor to macrophage PD-L1 expression ratio of approximately 3.2 and 1.75 for patients A and B respectively.

Conversely, when trained with compartmentalized patient samples, CCIGAN reported increased PD-L1 expression on macrophages adjacent to tumor cells when compared to macrophages adjacent to normal endothelial (inert) cells. This difference was quantified as a ratio of PD-L1 expression of tumor-adjacent macrophages to endothelial-adjacent macrophages, approximately 1.85 and 2.7 for patient C and patient D respectively.

Below are tables of the data used to generate Fig. 4B of the main paper. Results from testing increased PD-1/PD-L1 expression from the bolded cell being challenged with another cell type in its microenvironment are located in table 7. The 3rd column shows summed pixel intensity of the specified protein expression.

In table 8, the second row shows that even when using the trained compartmentalized model to predict on mixed segmentation patches, CCIGAN still reports a 26% (patient C, 0.00408/0.00324) and 19% (patient D, 0.00608/0.00510) increase of macrophage PD-L1 expression in the compartmentalized compared to mixed microenvironment (Table 8). This observation of higher macrophage PD-L1 expression in the compartmentalized environment supports Keren et al.'s findings.

Experiment	Microenvironments	Patient A	Patient B
PD-1 (T cell)	T cell / Tumor / Macrophages	0.01886	0.00131
	T cell / Endothelial / Macrophages	0.00558	0.00107
PD-L1 (Tumor)	T cell / Tumor / Macrophages	0.00649	0.00100
	Endothelial / Tumor / Macrophages	0.00279	0.00046
PD-L1 (Macrophages)	T cell / Tumor / Macrophages	0.00204	0.00057
	T cell / Endothelial / Macrophages	0.00068	0.00047

Table 7: Average PD-1/PD-L1 expression on the mixed tumor environment. The bolded cells indicate which cells are being measured.

Experiment (Macrophages)	Microenvironments	Patient C	Patient D
PD-L1 (Compartmentalized)	T cell / Tumor / Macrophages	0.00408	0.00608
	T cell / Endothelial / Macrophages	0.00220	0.00225
PD-L1 (Mixed)	T cell / Tumor / Macrophages	0.00324	0.00510

Table 8: Average PDL1 expression of macrophages/monocytes on the compartmentalized tumor environment.

3.6 T-CYCIF PD-1/PD-L1 TREND EXPERIMENT (FIG. 4C LEFT) EXPLANATIONS

By training CCIGAN models on different segments of the t-CyCIF dataset (S2.2.2), we investigate the relationship between PD-1 and PD-L1. Using the same experimental setting as S3.4, our cell of interest (center cell) is a CD8 T cell and we iteratively add tumor cells as adjacent neighboring cells. Fig. 23 shows an example patch of a CD8 T cell (orange) in the center and iteratively adding tumor cells (yellow). As surrounding tumor cell surface area increases and surrounding PD-L1 expression increases, we expect PD-1 in the CD8 T cell to be upregulated, as it is an indicator of T cell exhaustion. PD-1 expression trend differs depending on the level of tumor-infiltration in the tumor microenvironment. In a poorly infiltrated microenvironment, the PD-1 expression

Environment	T_1	T_2	T_3
Slope	2.629×10^{-3}	2.953×10^{-3}	3.134×10^{-3}
t -test	36.715	14.341	23.803
p -value	3.079×10^{-225}	1.945×10^{-44}	2.848×10^{-110}

Table 9: Slope and statistical values for t-CyCIF PD-1/PD-L1 Trend Experiment (S3.7) with respect to surrounding tumor cell surface area.

Environment	T_1	T_2	T_3
Slope	4.339×10^{-5}	5.450×10^{-5}	2.970×10^{-4}
t -test	35.248	14.477	25.608
p -value	2.098×10^{-211}	3.285×10^{-45}	4.520×10^{-125}

Table 10: Slope and statistical values for t-CyCIF PD-1/PD-L1 Trend Experiment with respect to total surrounding PD-L1 expression.

trend should be greater than in a highly infiltrated microenvironment, since low infiltration indicates greater immunosuppression and a higher rate of T cell exhaustion. Our results as shown in main paper Fig. 4C Left are fully displayed in Fig. 19, which illustrates sample runs, and table 9, which shows the full trend and statistical values from Fig. 4C Left. Additionally, we plotted the slope trend of CD8 T cell PD-1 expression against surrounding total tumor PD-L1 expression in table 10

T Cell PD-1 in response to increasing PD-L1 expressing tumor cells in t-CyCIF data

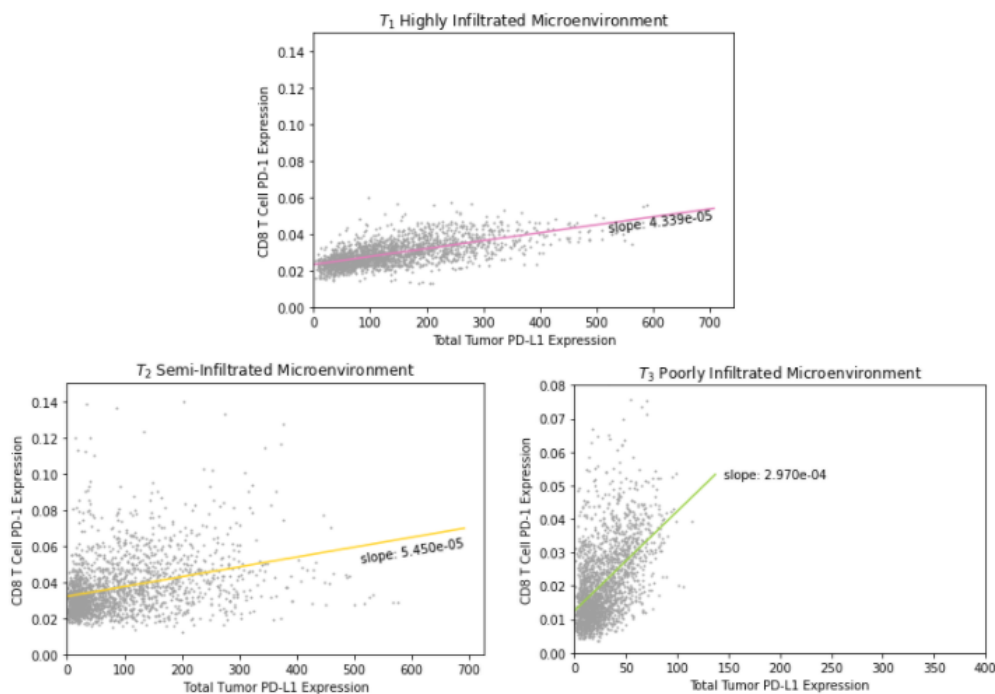


Figure 18: PD-1/PD-L1 experiment on tumor microenvironments with varying levels of tumor-infiltration in t-CyCIF (S3.6) with respect to total surrounding PD-L1 expression.

T Cell PD-1 in response to increasing area of surrounding tumor cells in t-CyCIF data

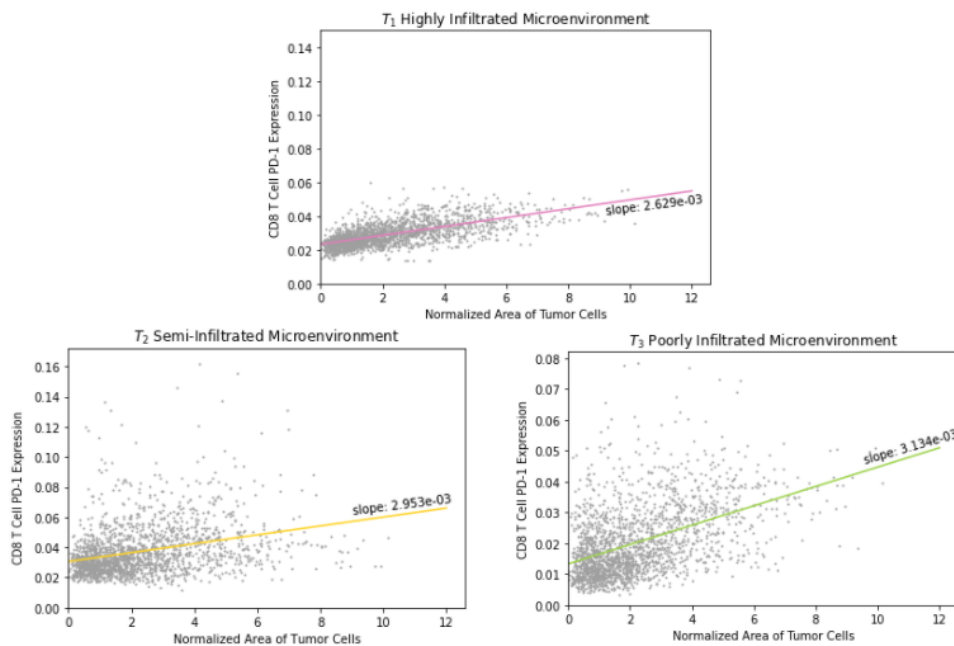


Figure 19: PD-1/PD-L1 experiment on tumor microenvironments with varying levels of tumor-infiltration in t-CyCIF (S3.6) with respect to surrounding tumor cell surface area.

T Cell PD-1 in response to increasing PD-L1 expressing tumor cells in t-CyCIF data

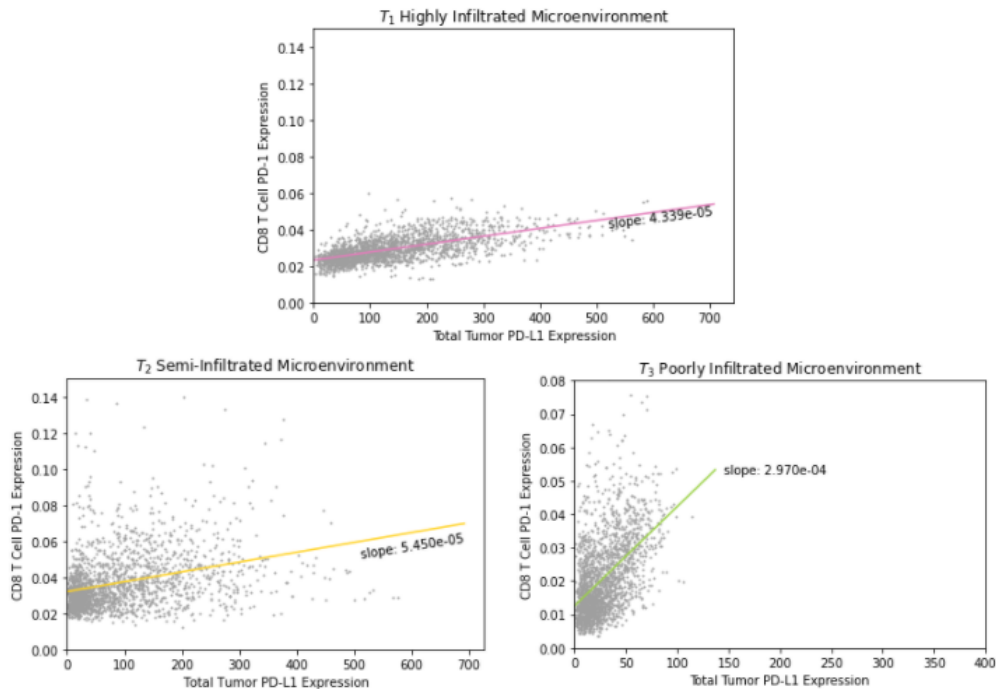


Figure 20: PD-1/PD-L1 experiment on tumor microenvironments with varying levels of tumor-infiltration in t-CyCIF (S3.6) with respect to total surrounding PD-L1 expression.

Environment	T_1	T_2	T_3
Slope	4.843×10^{-5}	8.838×10^{-5}	6.242×10^{-4}
t -test	35.778	12.711	21.693
p -value	2.172×10^{-216}	1.238×10^{-35}	9.723×10^{-94}

Table 11: Slope and statistical values for t-CyCIF LAG3/PD-L1 Trend Experiment.

3.7 T-CYCIF LAG3/PD-L1 TREND EXPERIMENT (FIG. 4C RIGHT) EXPLANATIONS

Furthermore, we investigate the relationship between LAG3 and PD-L1. Using the same experimental setting as S3.6, our cell of interest (center cell) is a CD8 T cell and the adjacent neighboring cells are tumor cells. As tumor cells are added and the surrounding PD-L1 expression increases, we expect LAG3 in the CD8 T cell to be upregulated, as it is another indicator of T cell exhaustion. Similarly to S3.6, we expect a higher trend of LAG3 upregulation with respect to PD-L1 in microenvironments with lower tumor-infiltration. Our results as shown in main paper Fig. 4C Right are fully displayed in Figure 21, which illustrates sample runs, and table 11, which shows the full trend and statistical values from Fig. 4C Right.

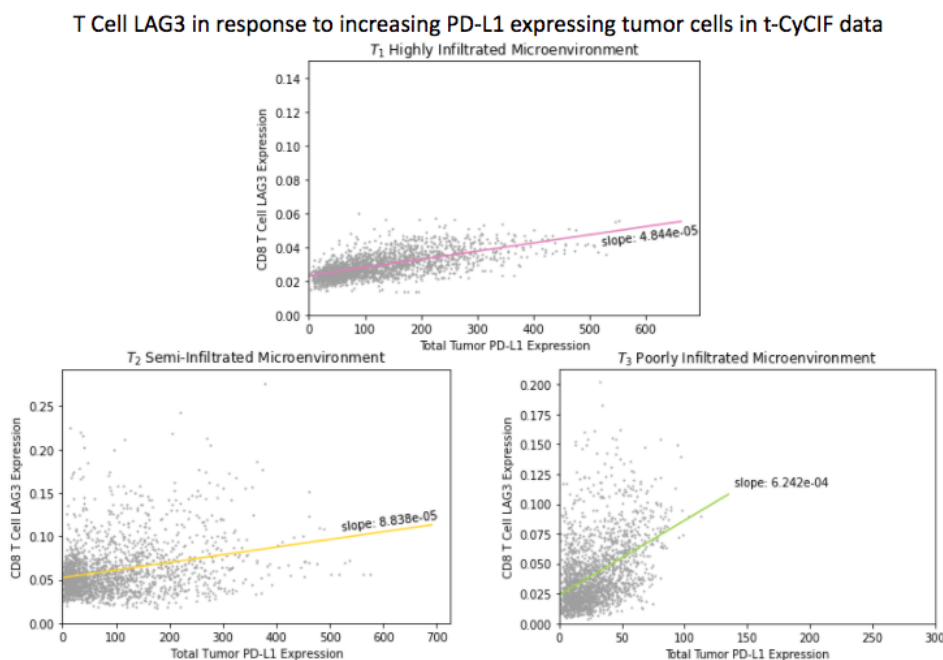


Figure 21: LAG3/PD-1 experiment on tumor microenvironments with varying levels of tumor-infiltration in t-CyCIF

4 FURTHER EXPERIMENTS AND EXPERIMENT DETAILS




Figure 22: Process of iteratively adding tumor cells in MIBI-TOF. The added red cells are tumor cells (PD-L1) and the center orange cell indicates a CD8 T cell (cell of interest, PD-1). For this process, we focus on each instance of an added tumor.




Figure 23: Iteratively Adding Cells for t-CyCIF data. A random dsDNA protein channel is shown.

4.1 EXPERIMENTAL SETUP

In our general experiments that involved iteratively manipulating cell patches, we created an experimental dataset of approximately 1000 patches for each type of data (MIBI-TOF or t-CyCIF) and manipulated the cell types to the necessary cell types for each experiment. For a patch s in the experimental dataset with n cells in the patch, we expanded the patch into n individual patches $P_s = \{p_1^{(s)}, p_2^{(s)}, \dots, p_n^{(s)}\}$ where for $0 \leq i \leq n$, $p_i^{(s)}$ = patch of cells from 1 to i . An example can be seen in Fig. 22 for MIBI data where PD-1 expression in the CD 8 T cell of interest reacts to

newly introduced PD-L1 expressing tumor cells. Another example is shown for t-CyCIF data in Fig. 23. Following this, changes in protein expressions in the cell of interest due to newly introduced cells were analyzed using a variety of techniques (center of mass, summation, mass shift, regression trend).

4.2 SEARCH ALGORITHM DISCUSSION




Figure 24: **Search algorithm results.** Specific cell-cell interactions and the associated magnitude of change in protein expression. The cell types in boxes on the y-axis represent the cell of interest in each experiment; cell names adjacent to these boxes represent neighbors. The circle diameter represents the relative frequency of this specific cell - cell interaction inducing a non-trivial change in the expression of a given protein. The hue of each circle denotes the order magnitude of change in expression level. For example, a CD8 T Cell (boxed, y-axis) with tumor cell neighbors (y-axis) experiences frequent and significant interactions that induce a change in the T cell's PD-1 expression (x-axis). These interactions are predicted to induce an average of 3 orders of magnitude increase in PD-1 expression.

In this section we discuss potential biological trends, erroneous correlations, and algorithm settings. It is important to note that some of the spurious correlations are easily explained by poor segmentations (known issue).

Fig. 24 displays the results of our search algorithm, which shows which particular cell-cell interactions are significantly captured in CCIGAN. For four primary cells— tumor, macrophage, CD8, and endothelial, we measure their expression level change in a subset of protein markers and calculate the relative frequency of this cell-cell interaction. Using a subset of 1268 test cell patches, we counted the number of significant logged interactions that produced an increase or decrease in expression level greater than a chosen threshold. Expression level change was measured by simply calculating the difference in summed pixel intensity (protein expression) in a particular channel.

For the relative frequency of a particular cell-cell interaction, the normalization schema is as follows. Using the logged counts of significant changes in expression level, we take the max of either the increase counts or the decrease counts. Then, to quantify the ratio of relative importance within a primary cell group, we divide the max counts by the total number of logged counts (for either only increases or decreases). For example, say we are measuring tumor as the primary cell. If 50

decreases and 200 increases in expression level are recorded, with 1000 total increases logged for the tumor primary cell, $200/1000 = 0.2$ is the relative frequency after normalization.

Algorithm settings were set on the same sensitivity threshold for experiments. Additionally, figure generation normalization was done across protein markers (within different cell groups). Finally, interactions comprising of 5% or less of total logged interactions were disregarded.

Individual marker discussions are given below:

1. Vimentin

It was observed that nearly all scenarios of cell-cell interaction accounted for by CCIGAN resulted in an varying increases of vimentin expression in the cell of interest. While not necessarily biologically explained, the observed changes are plausible as vimentin is a structural protein found in all cells and changes in its levels may not be attributable to specific cell interactions. A continued investigation is given in S4.3 regarding tumor and CD 8 T cell vimentin expressions.

2. PD-1

A slight increase in the PD-1 expression is seen on tumor cells when surrounded by CD8 T cells. However, this increase was negligible and can be attributed to noisy segmentations where PD-1 expressing CD8 T cells are located near tumor cells, suggesting in the training data that a tumor cell expresses PD-1.

3. PD-L1

The model suggests PD-L1 expression on CD8 T cells. This is due to noisy segmentations similar to the situation in PD-1.

4. CD8

The model predicted an increase in CD8 expression on the tumor cell in several scenarios despite this not being biologically expected. This also due to the same noisy segmentations.

5. Pan Keratin

The decreased levels in Pan Keratin expression within the tumor cell line with neighboring non-tumor cells was biologically expected. However, the slight increase in pan keratin expression when the tumor cell of interest had tumor cell neighbors may be due to the same noisy segmentations as mentioned before.

4.3 VIMENTIN TABLE EXPRESSIONS

We measure trends of vimentin expression in tumor cells due to varying amounts of surrounding CD 8 T cells. The cell of interest is a tumor cell and we iteratively change surrounding adjacent cells from tumor cells to CD8 T cells with different probabilities. Then we measure the number of instances a change of vimentin is detected, and what that change is.

Vimentin plays an important role in tumorigenesis both as a structural protein and in cellular signalling. Numerous studies have noted vimentin's regulatory importance in a variety of cell signalling pathways which promote tumor survival and resistance to cellular stress (38) (39). CCIGAN experiments wherein a tumor cell of interest was surrounded by increasing numbers of CD8 T cells caused up to a 91.6% increase in the total tumor vimentin expression (S4.2.1). Increasing CD8 T cell presence around a tumor cell is likely to be associated with increasing amounts of pro-apoptotic signalling on the tumor cell, causing a large degree of cellular stress. Thus the increase in the tumor's vimentin would be a biologically expected outcome given the protein's role in promoting cellular survival pathways.

Cell of Interest	Adjacent Cell Neighbor %	% of Patches Indicating Increase in Vimentin
Tumor	100% CD8 0% Tumor	91.6%
Tumor	75% CD8 25% Tumor	81.6%
Tumor	50% CD8 50% Tumor	60.4%
Tumor	25% CD8 75% Tumor	26.7%

Table 12: Table indicating net changes in Vimentin expression in a tumor cells due to increasing CD8 T cell presence.