GPU-Based Collocation Visualization of Immune Niches in 3D Cyclic Immunofluorescence Images

Uroš Šmajdek
uros.smajdek@fri.uni-lj.si

Eva Boneš
eva.bones@fri.uni-lj.si

Ciril Bohak
ciril.bohak@fri.uni-lj.si

Faculty for Computer and Information Science
University of Ljubljana

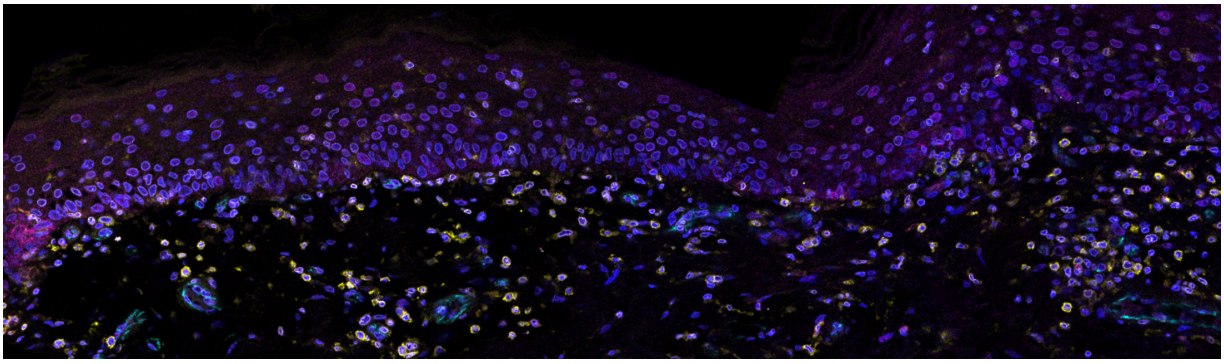

Figure 1: A selection of biomarker channels in the Melanoma 3D CycIF dataset visualized with Vitessce [1].

Abstract

This work presents our contributions to the Bio+MedVis Challenge 2025, specifically the 3D Microscopy Imaging Challenge. We integrated a GPU-based method into the Vitessce visualization framework to compute and visualize voxel-wise collocation scores for selected biomarker channels in large volumetric CycIF datasets. Our approach enables interactive exploration of immune niches by highlighting regions of interest (ROIs) where collocation exceeds a user-adjustable threshold. This report details the method, its technical implementation, discusses its limitations, and shows examples of its use with biologically relevant markers from the challenge dataset. We also discuss the limitations of this approach and outline directions for future work.

1 Introduction

Highly multiplexed tissue imaging methods like Cyclic Immunofluorescence (CycIF) are essential for investigating the subcellular complexities of cancer [2]. CycIF (an example of the data seen in Figure 1) has been instrumental in studying immune-tumor interactions in diseases like melanoma at single-cell precision [3]. Recent extensions of these techniques into three dimensions offer an even more comprehensive analysis of the tumor microenvironment, allowing for a deeper understanding of spatial cell-cell interactions [4].

However, identifying and visualizing these intricate interactions within large, complex 3D volumes remains a significant computational and visual challenge. To address this, we developed a GPU-based method within the Vitessce [1] visualization framework. Our method computes per-voxel collocation of selected markers and uses user-driven thresholding to detect regions of interest with high collocation. This approach balances computational efficiency with visual interpretability to highlight biologically meaningful spatial co-occurrences.

2 Method Description and Implementation

Our method is integrated into the volume rendering pipeline of Vitessce. The volumetric CycIF data is loaded onto the GPU as a multi-channel 3D texture, and we developed a custom GLSL fragment shader that executes for every rendered voxel.

The shader performs the following real-time steps:

1. **Data Sampling:** For a given voxel coordinate, the shader samples the intensity values from the user-selected biomarker channels in the 3D texture.

2. **Normalization:** Intensity values are normalized to a floating-point range of [0, 1] based on user-defined sliders for each channel. This step is crucial for mitigating bias from markers with different intensity scales.

3. **Collocation Calculation:** The normalized intensities of the selected channels $(C_1, C_2, ..., C_n)$ are multiplied ($Score = \prod C_i$) to compute a final collocation score. This operation emphasizes regions where all selected markers are strongly expressed. Our implementation natively supports combining from two to six channels.

4. **Color Mapping and Filtering:** The final collocation score is mapped to a user-defined color gradient. A threshold, passed to the shader, filters out voxels with low scores, isolating the most relevant ROIs.

Because all computations occur on the GPU, any change in parameters (channel selection, thresholds, colors) results in an instantaneous update, maintaining interactive framerates on large datasets. To improve usability, we also extended the Vitessce interface with independent opacity sliders for each channel, allowing users to visually separate the raw channel data from the collocation overlay (Figure 2).

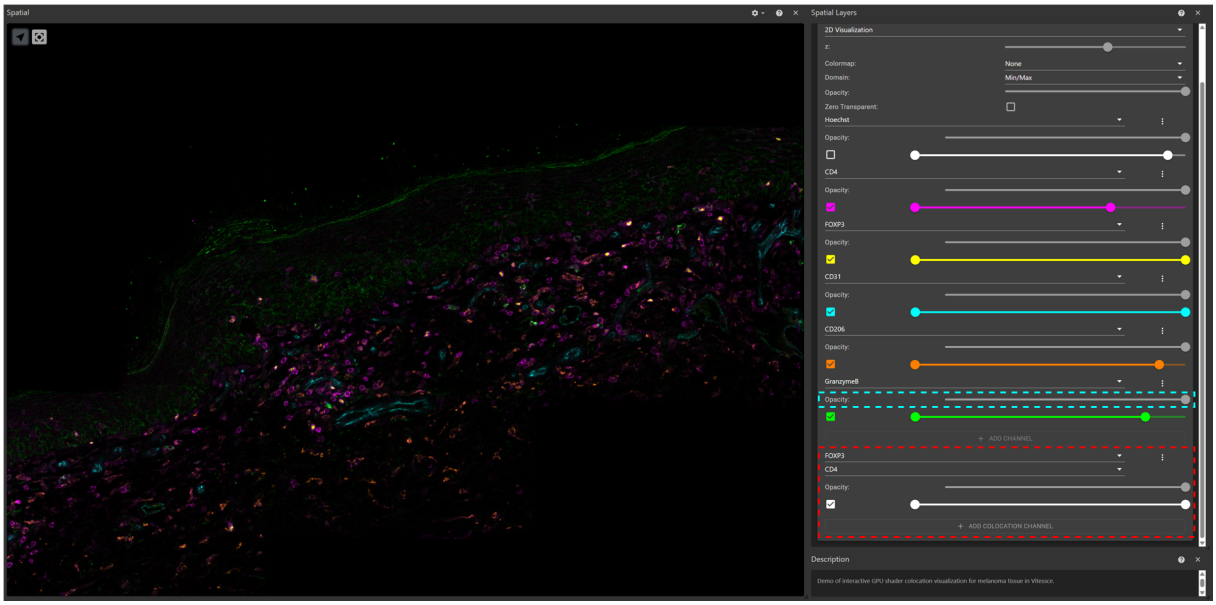

Figure 2: User interface of our implementation integrated within Vitessce. The volumetric tissue visualization is displayed on the left, while interactive controls for channel selection, opacity, collocation threshold, and color adjustments are on the right. The added collocation controls are highlighted in red, while the opacity controls are highlighted in blue.

3 Visualization Examples

Our method effectively demonstrates biomarker relationships at both global and local scales, as shown in the supplementary material. The first example (Figure 1 in supplementary material) shows a global-scale visualization of the collocation between CD4 and FOXP3 markers. While a traditional overlay makes it difficult to distinguish co-expression from simple visual overlap, our method highlights only the regions where both markers are strongly collocated, making the immune niche distribution clearly visible across the tissue. The second example (Figure 3 in supplementary material) demonstrates a cellular-level analysis of CD31 and CD206 markers, revealing small, localized overlaps that are not apparent in a standard overlay.

4 Limitations and Discussion

The primary goal of our method is to provide a fast and intuitive way to highlight regions of co-expression. The current approach for this analysis often involves toggling the visibility of channels and visually searching for color-blended regions. Our approach reduces the cognitive load of this task by explicitly calculating and displaying a new *collocation channel*.

A consideration of our simple multiplicative approach is that it can disproportionately emphasize markers with a high overall intensity, even after normalization. While effective for identifying regions where multiple markers are highly expressed (e.g., A+B+C+), it is not designed for more complex spatial relationships, such as mutual exclusion (e.g., A+B-). This could be fairly easily addressed in the future by allowing users to define custom logic in the shader, such as $C_A \times C_B \times (1 - C_C)$ to find where markers A and B are present but C is absent.

5 Conclusion and Future Work

Our GPU-based collocation visualization is a simple, fast, and effective method for highlighting overlapping biomarkers in 3D microscopy data. With its independent opacity controls and interactive threshold tuning, it supports detailed exploration of immune niches in large volumetric tissue images.

There are still plenty of options for future improvement. An immediate plan is to package the implementation as a publicly available Vitessce plugin to make it accessible to the broader research community. We also want to integrate some post-processing controls for brightness, contrast, and exposure. While these adjustments would not alter the underlying ROI detection, they would give scientists more flexibility in creating clear figures for publication and presentation.

Future work could also focus on expanding the tool's analytical capabilities. This could include adding other combination operations, like summation or user-defined logic, to complement the current multiplicative approach. Additionally, a system to quantitatively score identified ROIs and guide users to them with a camera navigation tool could be a helpful feature to explore.

Acknowledgment

This research was conducted as part of the basic research project *Cell visualization of unified microscopic data and procedurally generated sub-cellular structures* [project number J2-50221], funded by the Slovenian Research and Innovation Agency (Javna agencija za znanstvenoraziskovalno in inovacijsko dejavnost RS) from the state budget.

References

[1] M. S. Keller, I. Gold, C. McCallum, et al. Vitessce: integrative visualization of multimodal and spatially resolved single-cell data. *Nat Methods*, 22(1):63–67, Jan. 2025. doi: 10.1038/s41592-024-02436-x 1

[2] J. R. Lin, M. Fallahi-Sichani, J. Y. Chen, et al. Cyclic Immunofluorescence (CycIF), A Highly Multiplexed Method for Single-cell Imaging. *Current protocols in chemical biology*, 8(4):251–264, 2016. doi: 10.1002/cpch.14 1

[3] A. J. Nirmal, Z. Maliga, T. Vallius, et al. The Spatial Landscape of Progression and Immunoediting in Primary Melanoma at Single-Cell Resolution. *Cancer Discovery*, 12(6):1518–1541, 2022. doi: 10.1158/2159-8290.CD-21-1357 1

[4] C. Yapp, A. J. Nirmal, F. Zhou, et al. Multiplexed 3D Analysis of Immune States and Niches in Human Tissue. *bioRxiv [Preprint]*, pp. 2023–11, 2024. doi: 10.1101/2023.11.10.566670 1

GPU-Based Collocation Visualization of Immune Niches in 3D Cyclic Immunofluorescence Images

— Supplementary Material —

Uroš Šmajdek
uros.smajdek@fri.uni-lj.si

Eva Boneš
eva.bones@fri.uni-lj.si

Ciril Bohak
ciril.bohak@fri.uni-lj.si

Faculty for Computer and Information Science
University of Ljubljana

Figures 1 and 2 show a global-scale example of visualizing the collocation between CD4 and FOXP3 markers. In the traditional overlay (top), the high signal intensity of the CD4 marker visually saturates the image, making it difficult to discern the co-expression of the FOXP3 marker. Our method (bottom) highlights only the regions where both markers are strongly collocated, making the immune niche distribution clearly visible over a larger tissue region.

Figures 3 and 4 illustrate a cellular-level example of our method using the CD31 and CD206 markers. The left image shows a traditional overlay of separate channels. The right image includes our collocation visualization, revealing small, localized overlaps that are not apparent in the left image alone.

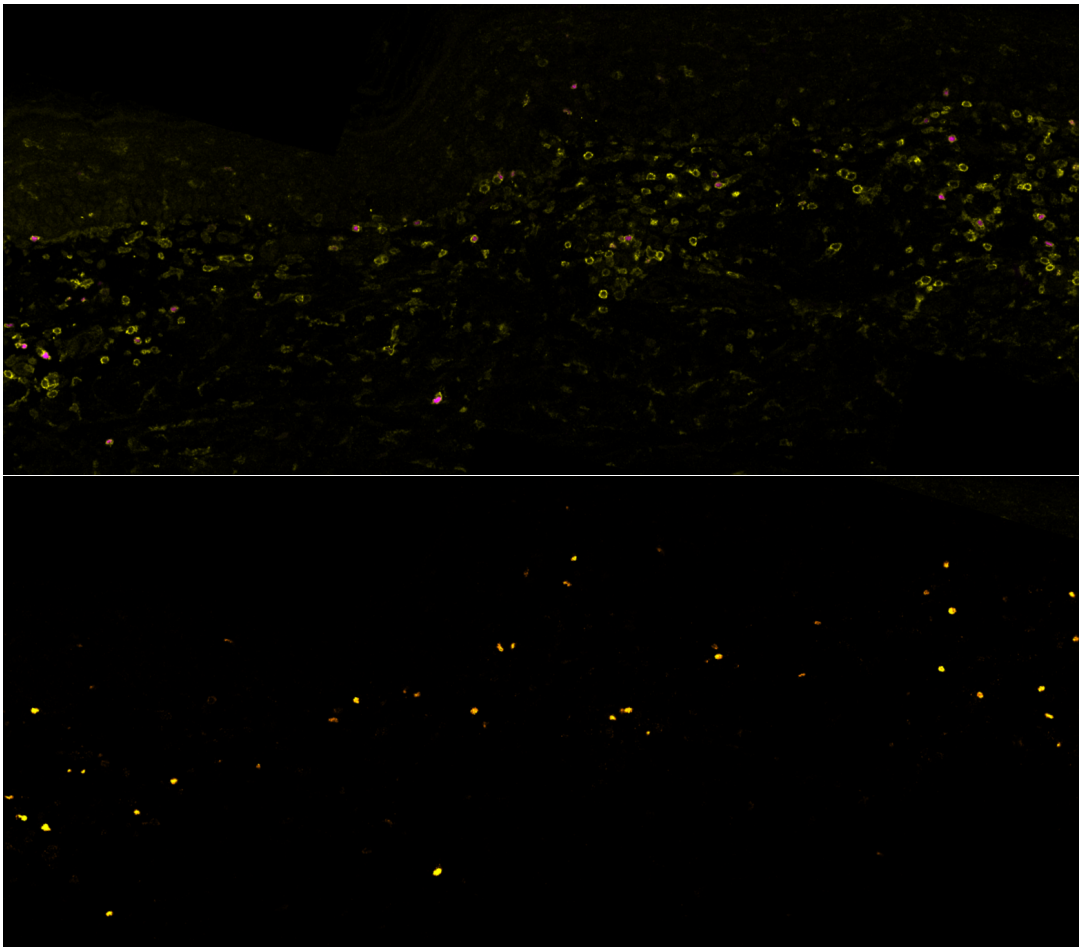

Figure 1: A global-scale example. **Top:** Traditional overlay of CD4 (yellow) and FOXP3 (pink) markers. **Bottom:** Our visualization highlighting regions where both markers co-occur (orange).

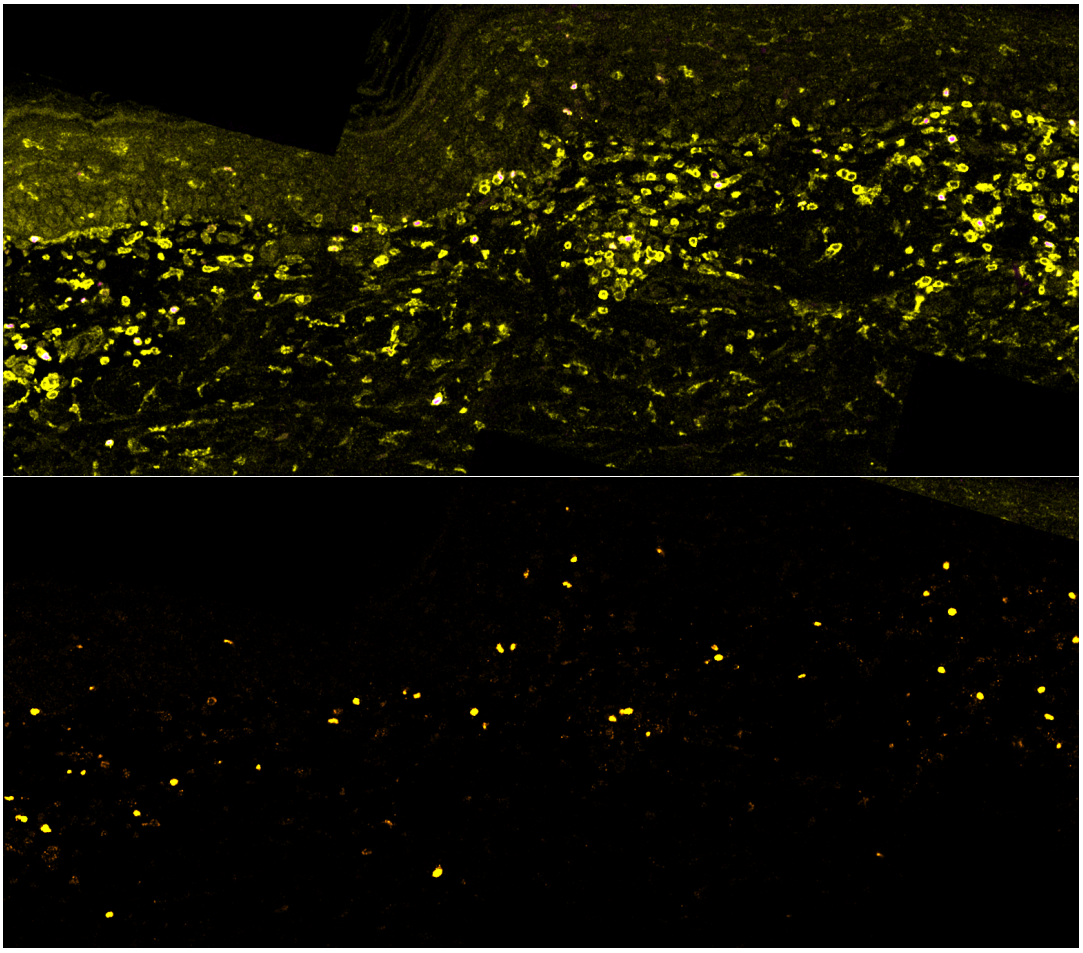

Figure 2: Screenshots from Figure 1 with manually increased brightness and exposure for improved visibility in print.

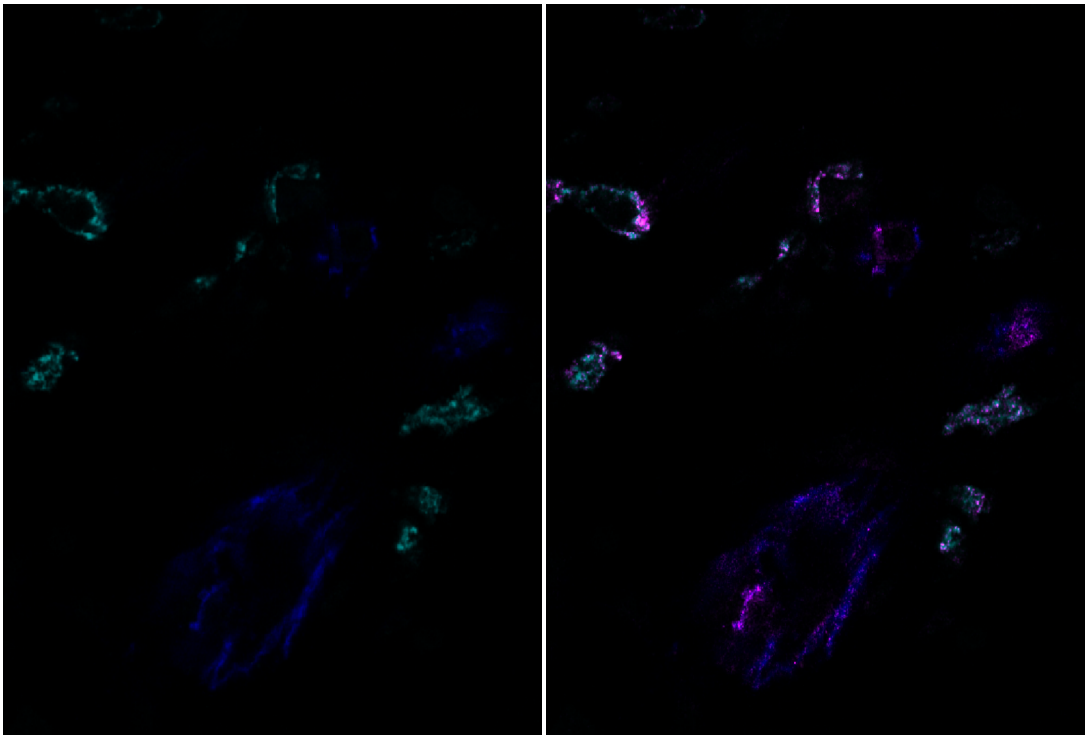

Figure 3: A local, cellular-level example. **Left:** Overlay of CD31 (dark blue) and CD206 (light blue). **Right:** Visualization of their collocation in pink overlayed over both channels.

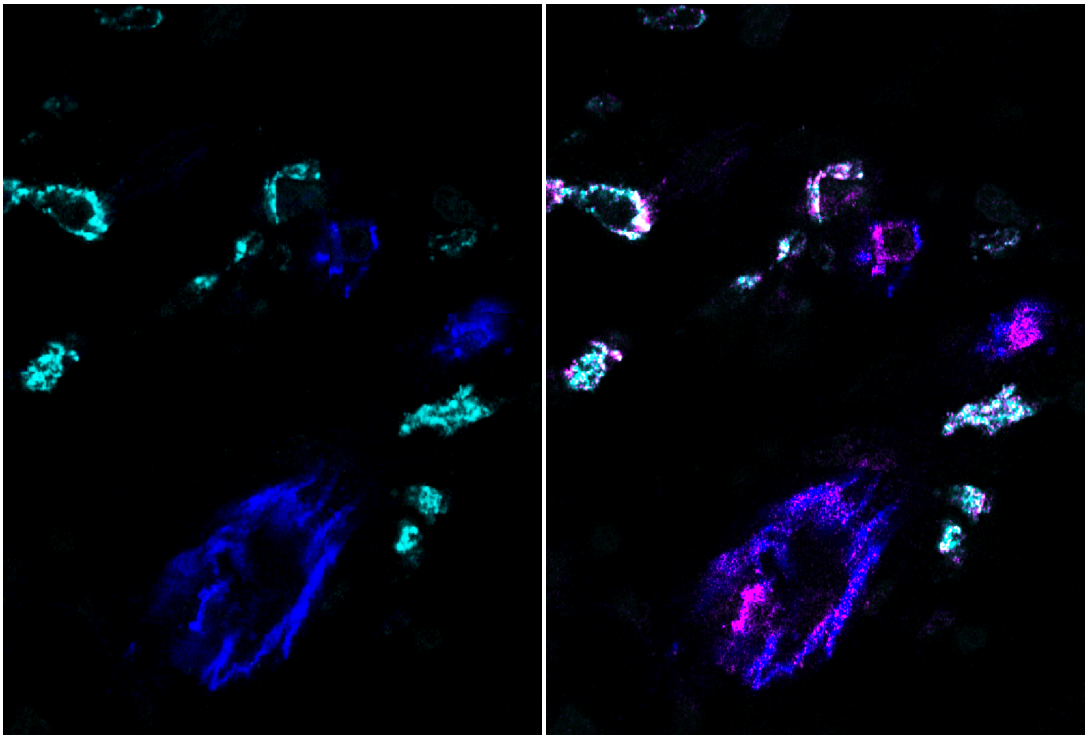

Figure 4: Screenshots from Figure 3 with manually increased brightness and exposure for improved visibility in print.

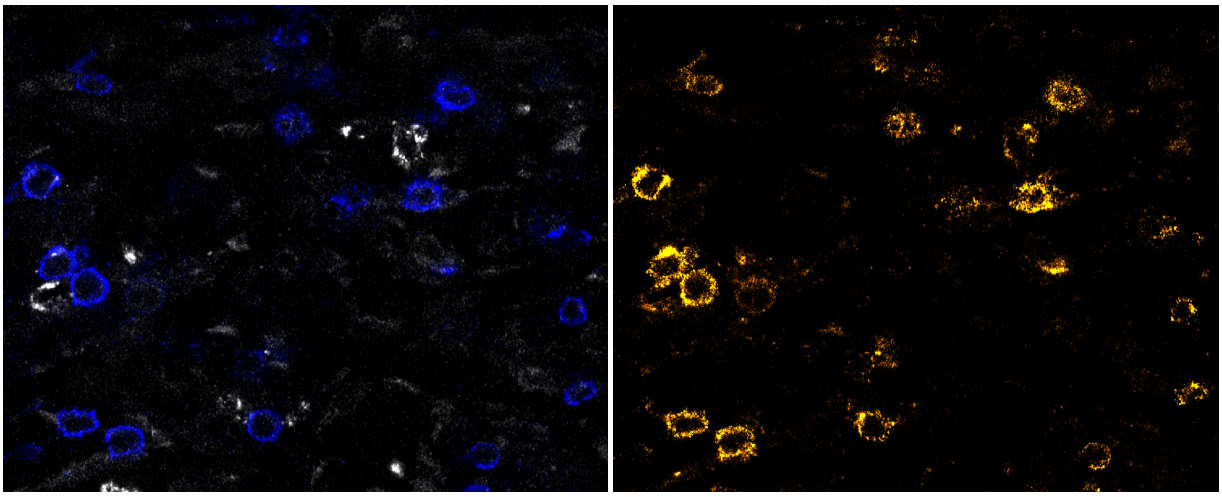

Figure 5: Another example demonstrating the method's clarity. **Left:** In the traditional overlay, we can't really see where the markers are overlapping and where only one is present. **Right:** Our visualization makes the collocation immediately apparent.