

Introduction: Neuroscience has demonstrated that the complexity of thought, behavior, and general nervous coordination in humans is mirrored in the architecture and physiology of the brain and nervous system. The human brain has approximately 86 billion neurons and about as many non-neurons and accounts for 20% of the metabolic activity of the body despite being 2% of its mass[1]. While current research in the field attempts to elucidate the specific mechanisms by which the organ processes input, responds, and generates consciousness, systems-level approaches provide a more general framework within which to understand the organization of the brain and make biological hypotheses without full understanding of the biology[2]. Characterization in this framework will provide guidance in applications, for example in determining an upper limit to expected decoder complexity or the locations in the brain providing the maximum information for utilization in the development of brain-machine interfaces. The theory of complex systems implies that a system will evolve to match the complexity profile of its environment across the scales on which it operates, enabling adaptation to the full range of possible stimuli[3]. A common representation of this phenomenon in biological systems is hierarchical modularity, a structure composed of nested, complete structures. In the brain, this manifests as the division of tasks between lobes, broad and fine regions within them, the cell-type neighborhoods that constitute them[4]. Dimensionality reduction of data on the brain provides an avenue to understand the complexity of the system. Accurate, low dimensional embeddings of high dimensional data demonstrate simplistic structure, and similar approximate dimensionalities imply similar complexity otherwise obscured by differences in measurement. Here, I will ask two questions. First, what does the performance of dimensionality reduction on single unit action potential measurements say about the complexity of electrical activity in the brain, and can some of this be accounted for by spatial gradients? Second, how does the dimensionality of single unit recordings compare to that of cell-type neighborhoods? Is spiking more complex because it inherently reflects the behavior of these neighborhoods, or are they equivalently complex? And do they vary in the same ways? To answer these questions, sorted single unit spike data from the primary motor cortex of mice measured using Neuropixel high density silicon probes during a visual discrimination task completed in Steinmetz et al. (2019)[5] will be analyzed alongside predicted cell-type proportions at $200\ \mu\text{m}^3$ resolution from the same region, generated in my own line of research (Sodicoff, unpublished). Utilizing principal component analysis, in its well-known form as well as its nonlinear and kernel-based derivatives, I will characterize its intrinsic dimensionality of both datasets and relate the generated dimensions to spatial gradients.

Methods: Code for this section was generated in R and can be found in **Appendix 2**

Spiking Data: Single unit action potential measurements come from Steinmetz et al. (2019). To extract the spikes used, the raw data for all experiments completed for the study were downloaded and loaded into R. For each experiment, probe channels were associated with a location in the Allen Common Coordinate Framework, allowing for localization of spikes measured on the channel to regions. Only the channel of highest amplitude was used from the provided two-dimensional spike waveforms, and these waveforms were then filtered such that they only originated in channels within the MOp. Three of the provided experiments included spikes from the MOp, requiring normalization between them by scaling first by row, corresponding to all measurements at a timestep, and then by column, corresponding to the sample, to fully integrate the datasets.

Cell-Type Proportion Data: Cell-type proportion data was generated in the development of a distribution-free method for deconvolution of cell type from spatial transcriptomic data using labeled single cell transcriptomic data (Sodicoff, unpublished). scRNA-seq data was provided by members of the NIH Brain Initiative Cell Census Network, and the spatial data derived from the Allen Institute's quantified in-situ hybridization (ISH) data from their Mouse Brain Atlas [6], both for the MOp. The later dataset provides average staining intensity for individual genes over $200\ \mu\text{m}^3$ "voxels", corresponding to approximately 900 voxels for the region. To

select genes for utilization in the method, a Kruskal Wallis test was used to find differentially expressed genes between annotated clusters from the single cell measurements. A novel, nonnegative derivative of partial least squares discriminant analysis (Sodicoff, unpublished) was utilized to determine cell-type signatures, defined by expression on the selected genes for all provided annotations. NMF regression was applied to the scaled, quantified ISH data to determine cell type loadings by voxel, followed by scaling of each voxel to sum to one.

Principal Component Analysis: Both processed datasets were analyzed with three forms of principal component analysis. Such techniques were selected for their direct relationship to the underlying dimensionality of the space, as opposed to techniques like UMAP and local linear embedding, as well as their quantification of variance explained by component. Data was centered and scaled before application of the selected pipelines. Canonical principal component analysis, implemented by the R *stats* package[7] function *prcomp*, finds a low dimensional representation of the data by finding orthogonal axes through the data accounting for the highest variance. Nonlinear PCA (NLPCA), as implemented in the *pcaMethods*[8], uses an auto-associative neural network to find a set of components that can best be used to regenerate the original data. Kernel PCA (KPCA) calculates the covariance matrix utilized for calculation of principal components not from the original data, instead in a feature space defined by transformation of the original data by a kernel function. Here, the *compute.kernels* and *kernel.pca* functions from the *mixKernel* package[9] were utilized to calculate the decomposition here. *compute.kernels* was used to find the K matrix using a Gaussian kernel with $\sigma = 0.025$ and $= 0.15$ for the spiking and cell type data, respectively, selected via manual iteration. For the NLPCA and KPCA, a prior number of components to find is required, and 10 was selected based on preliminary analysis with the PCA, assuming that nonlinear components could only better fit the data. These three methods were selected to search for three separate low-dimensional structures, with linear and nonlinear functions in Euclidean space and a manifold with structure described by similarity between samples.

Visualization and Statistical Analysis: For all PCAs performed, plots of the samples defined by their coordinates in the first two component spaces were generated. For the spiking data, these plots were colored by the experiment from which they were measured, as well as the registered X, Y, and Z position within the brain. For the cell type data, the plots were only colored by the X, Y, and Z indexes of the voxels. For all PCAs, the cumulative variance for the first 10 principal components was calculated and plotted. For the PCA of the spiking data, the 8 eigenvectors accounting for the highest variance were plotted, to determine sources of variation in that dataset. For all PCAs conducted for both datasets, linear regressions were generated attempting to predict the X, Y, and Z values of the samples from the first two principal components, individually, to determine if the principal components explain spatial variability in the data.

Results: All referenced p -values from the described regression analyses can be found in **Table 1**.

Analysis of PCA on Spiking Data: In **Figure 1**, the cumulative variance accounted for by the first 10 principal components is shown. The first three components account for approximately 70% of the variance, with PCs 2 and 3 accounting for similar fractions of the total. The first 6 account for 90% of the variance, demonstrating that the data can fairly accurately be reconstructed from those 6 features. In **Figure 2**, we plot the samples by their scores on PCs 1 and 2 and color by the experiment from which they were derived, showing strong overlap of the provided sources after preprocessing, minimizing the impact of technical effects on the decomposition. **Figures 3, 4, and 5** show this same plot colored by the X, Y, and Z coordinates, corresponding to anterior-posterior, superior-inferior, and left-right position, of the channels from which the spikes were recorded, respectively. Visual observation did not yield a clear relationship between the principal components and the spatial variables, though linear models attempting to fit the spatial variables as a function of the first two principal components found significance on the coefficients for PC 2 for X and Z. In **Figure 6**, the first eight eigenvectors are shown. Upon inspection, the first three are primarily composed of peaks and troughs seemingly

aligned with the de- and repolarization of the neurons associated with passage of the action potential. The others are characterized by oscillatory behavior after initiation. Given the expected congruence between known potential traces and the first three components, it is sensible they account for the majority of the variance

Analysis of NLPCA on Spiking Data: In **Figure 7**, the cumulative variance accounted for by the first 10 principal components is shown. The first three components account for approximately 45% of the variance, with PCs 2 and 3 accounting for similar fractions of the total. The first 5 account for 75% of the variance, with the final five accounting for about even amounts. This differs from the PCA, likely because orthogonality is more difficult to constrain with the method, meaning there is more covariance between the components. In **Figure 8**, we plot the samples by their scores on PCs 1 and 2 and color by the experiment, again showing a high level of integration between the datasets, more meaningful because

NLPCA accounts for possible discriminating nonlinear effects. **Figures 9, 10, and 11** show this same plot colored by the X, Y, and Z coordinates of the channels from which the spikes were recorded, respectively. Visual observation again did not yield a clear relationship between the principal components and the spatial variables. However, the fit linear models demonstrated a highly significant relationship ($p < 0.0005$) between X and PC 1, and significant relationships between PC 2 and all spatial variables.

Analysis of KPCA on Spiking Data: In **Figure 12**, the cumulative variance accounted for by the first 10 principal components is shown. The first component accounts for nearly half of the variance, the first two for more than 75%, the first three for approximately 88% of the variance, with additional components. This provides the most compact representation of the data. In **Figure 13**, we again plot the samples by their scores, colored by the experiment, again showing a high level of integration between the datasets. We note that the distribution is approximately oblong, with maximum and minimum scores on PCs 1 and 2 of 1 and 0, respectively. **Figures 14, 15, and 16** show this plot colored by the X, Y, and Z coordinates of the channels from which the spikes were recorded, again not yielding insight on the relationship. Regression analysis demonstrated significant relationships between the visualized principal components and the Y and Z values.

Comparison to Analysis of Cell Type Proportion Data: In **Figure 17**, the cumulative variance accounted for by the first 10 principal components found with PCA on the second dataset is displayed. The first component accounts for half of the variance, the first 3 for approximately 75% of the variance, and the first 6 compose approximately 90%. This closely matches the results from PCA on the spiking data, suggesting [[][]]. However, there was greater variation found in the structure of the data and its relationship to spatial variables. In **Figure 18**, we see the structure of the data plotted on the first two components, with most points on a curve from (-1,-0.15) to (0.25,0.45) to (0.35,-0.45). This plot shows a clear relationship between X coordinate and the PCs, with smaller X values associated with negative PC 1 scores and high magnitude PC 2 scores. **Figure 19** depicts the relationship between Y coordinate and components, with larger Y values associated with extreme PC 2 values and small values near zero. **Figure 20**, colored by Z coordinate, generally has more negative PC 2 values on the positive half of PC 2. These relationships are well supported by the regression analysis, with PC 1 having a significant relationship with all spatial features, and PC 2 with all but X.

Table 1: p -values for coefficients from regression analysis, where orange is $p < 0.05$, yellow is $p < 0.005$, and orange is $p < 0.0005$

Spiking Data	X	Y	Z
PC_1_PCA	0.424	0.109	0.851
PC_2_PCA	1.91E-12	0.172	0.0136
PC_1_NLPCA	2.05E-07	0.152	0.157
PC_2_NLPCA	0.00104	0.0388	0.01
PC_1_KPCA	0.189	0.0102	0.0175
PC_2_KPCA	0.0993	0.021	0.0464
Cell Type Data	X	Y	Z
PC_1_PCA	2.61E-12	<2E-16	5.5E-09
PC_2_PCA	0.98	1.94E-05	0.0419
PC_1_NLPCA	1.67E-07	<2E-16	0.000411
PC_2_NLPCA	3.08E-10	0.0249	8.13E-16
PC_1_KPCA	<2E-16	0.0308	<2E-16
PC_2_KPCA	4.8E-16	0.000383	<2E-16

Application of the NLPCA here yields divergent results from the spiking data. In **Figure 21**, we note that the first component derived accounts for only 20% of the variance, with no subsequent component providing much less information than the previous, and the first eight components cumulatively accounting for 90% of the variance. Considering the plots of sample scores, colored by X, Y, and Z in **Figures 22, 23, and 24**, respectively, we note, overall, a similar continuous distribution of samples, with a large number of samples clustered near (-0.2, -0.1). This cluster is relatively homogeneous, with small X, Y, and Z values, while the distribution throughout the rest of the plot is difficult to infer. Regression analysis yielded that both of the first two components are predictors of all spatial features. The results presented here differ strikingly from the analysis of the spiking data, for which fewer components were required to accurately reconstruct the data (though both representations require more dimensions than that generated by the PCA) and for which those components were less closely associated with the spatial distribution.

Finally, we consider the results of KPCA on the cell type data. In **Figure 25**, we see that the first component accounts for 70% of the variance, and the first 4 for about 90%. As KPCA with a Gaussian kernel essentially searches for patterns of similarity, it is consistent with previous results that a majority of the variance may fall along one continuum. In **Figures 26, 27, and 28**, we compare scores on the first two components for all samples to their spatial features. While the compression associated with KPCA, demonstrated in the previous section, makes visual interpretation difficult, regression analysis again finds a significance for prediction of all spatial variables from both principal components. This result is highly reasonable given that the samples represent discrete portions of a continuous gradient of cell-type neighborhoods found within the MOp, and thus association between them follows spatial gradients.

Discussion: Results of investigation into both proposed questions yielded insights on the spiking and cell type data as well as the PCA methods used. Spiking behavior and cell type distribution can be described with a similar linear dimensionality of approximately 6, though nonlinear methods might allow for more compact representations of the cell type data as the data conforms to a continuous, high dimensional curve. While the components generated for both sets generally relate to the provided spatial features as per the results of regression analysis, spatial variability in the cell type data is better explained by those components. This is important because it, first, demonstrates that while spiking phenomena and cell type composition might have similar complexity, they are less linked than was originally expected and, second, because spiking behavior is thus only loosely associated with the architecture of the region.

To more fully answer this question, a joint representation of spiking and cell type behavior would need to be completed by integration within a common framework. This was originally proposed for this project, but on account of difficulty in alignment of the provided coordinate systems as well as concerns about the sparsity of the spiking data, localized along a small number of probes with channels arranged linearly, in the space of the region, resulted in dropping this aim. To better complete this task, a different spatial transcriptomic modality allowing for resolution at the single cell level, like Drop-seq, would need to be utilized to study the cells found directly along the probe. A study like this would allow for full characterization of neuronal cell types by their associated spiking patterns and allow for inference on causes of variation in spiking based on the measured cell type neighborhood. If this project is continued by my laboratory, we will also attempt to link data on connectivity to morpho-electrical properties, as understanding the source of measured signals could provide insight into similarities not accounted for by local features.

Acknowledgements: Dr. Joshua Welch and April Kriebel in the Department of Computational Medicine and Bioinformatics at the Medical School were essential resources on the unpublished cell type deconvolution project and troubleshooting related to downloading and processing the data, and provided critical feedback on concept development for this project.

References:

- [1] Herculano-Houzel, S. (2012). The remarkable, yet not extraordinary, human brain as a scaled-up primate brain and its associated cost. *Proceedings of the National Academy of Sciences*, 109(Supplement 1), 10661-10668.
- [2] Fallani, F. D. V., & Bassett, D. S. (2019). Network neuroscience for optimizing brain–computer interfaces. *Physics of life reviews*, 31, 304-309.
- [3] Siegenfeld, A. F., & Bar-Yam, Y. (2020). An introduction to complex systems science and its applications. *Complexity*, 2020.
- [4] Meunier, D., Lambiotte, R., Fornito, A., Ersche, K., & Bullmore, E. T. (2009). Hierarchical modularity in human brain functional networks. *Frontiers in neuroinformatics*, 3, 37.
- [5] Steinmetz, N. A., Zátka-Haas, P., Carandini, M., & Harris, K. D. (2019). Distributed coding of choice, action and engagement across the mouse brain. *Nature*, 576(7786), 266-273.
- [6] Sunkin, S. M., Ng, L., Lau, C., Dolbeare, T., Gilbert, T. L., Thompson, C. L., ... & Dang, C. (2012). Allen Brain Atlas: an integrated spatio-temporal portal for exploring the central nervous system. *Nucleic acids research*, 41(D1), D996-D1008.
- [7] R Core Team (2012). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, URL <http://www.R-project.org/>
- [8] Stacklies W, Redestig H, Scholz M, Walther D, Selbig J (2007). “pcaMethods – a Bioconductor package providing PCA methods for incomplete data.” *Bioinformatics*, 23, 1164–1167.
- [9] Mariette J, Brouard C, Flamary R, Vialaneix N (2022). *mixKernel: Omics Data Integration Using Kernel Methods*. R package version 0.8.

Appendix 1: Figures

Figure 1: Cumulative variance explained for PCA principal components found from spiking data

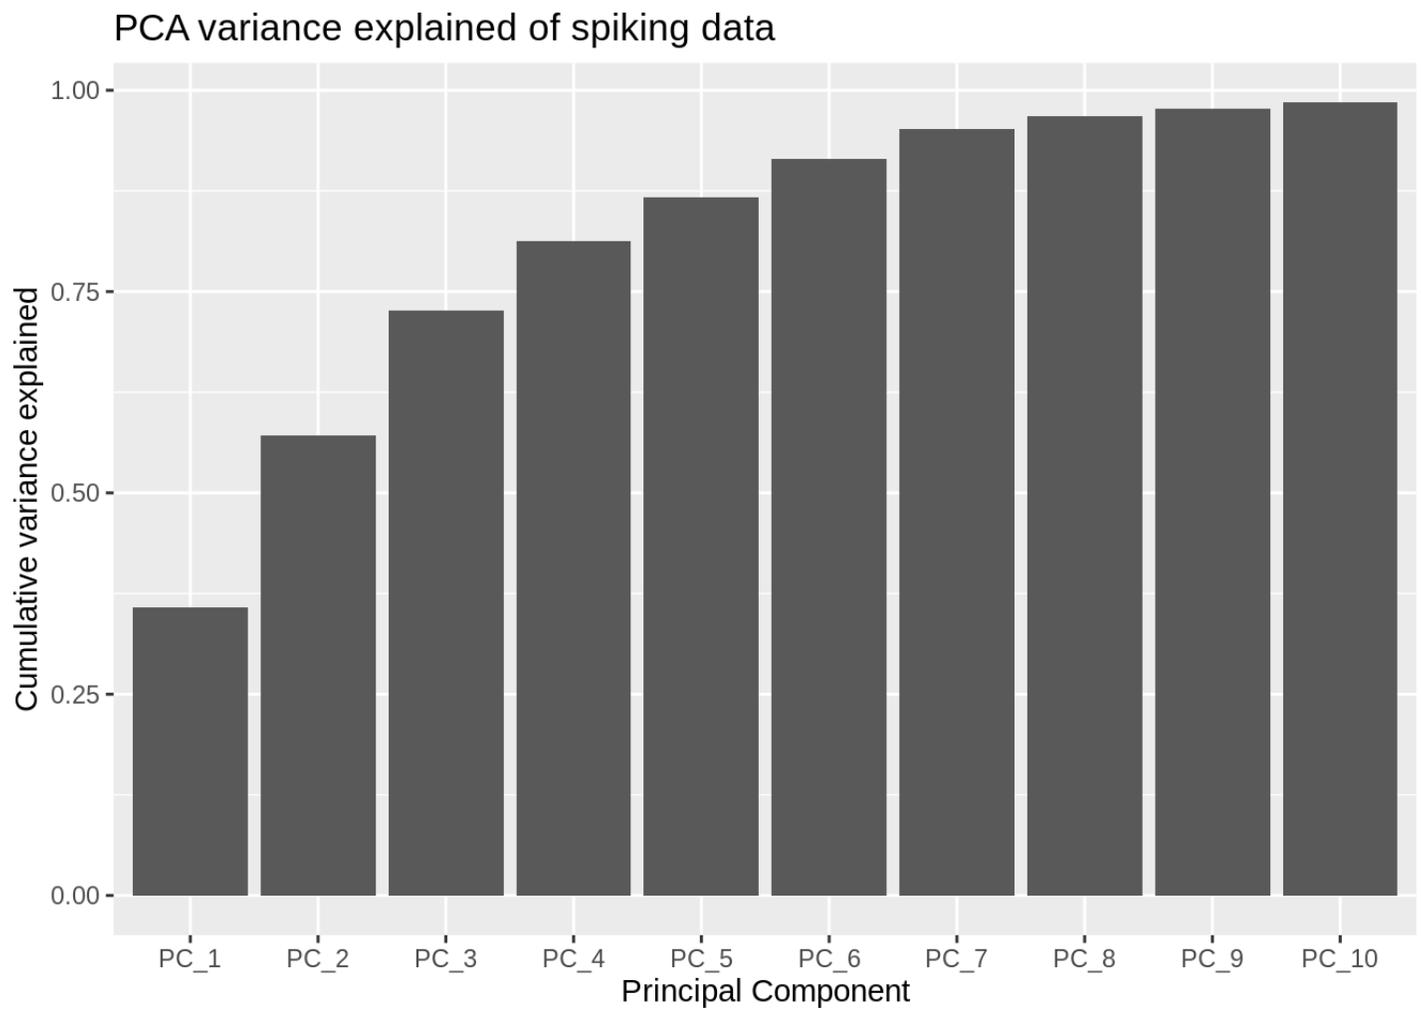


Figure 2: Plot of scores for PC 1 and PC 2 generated from PCA of spiking data, colored by experiment of origin.

PCA plot, colored by experiment

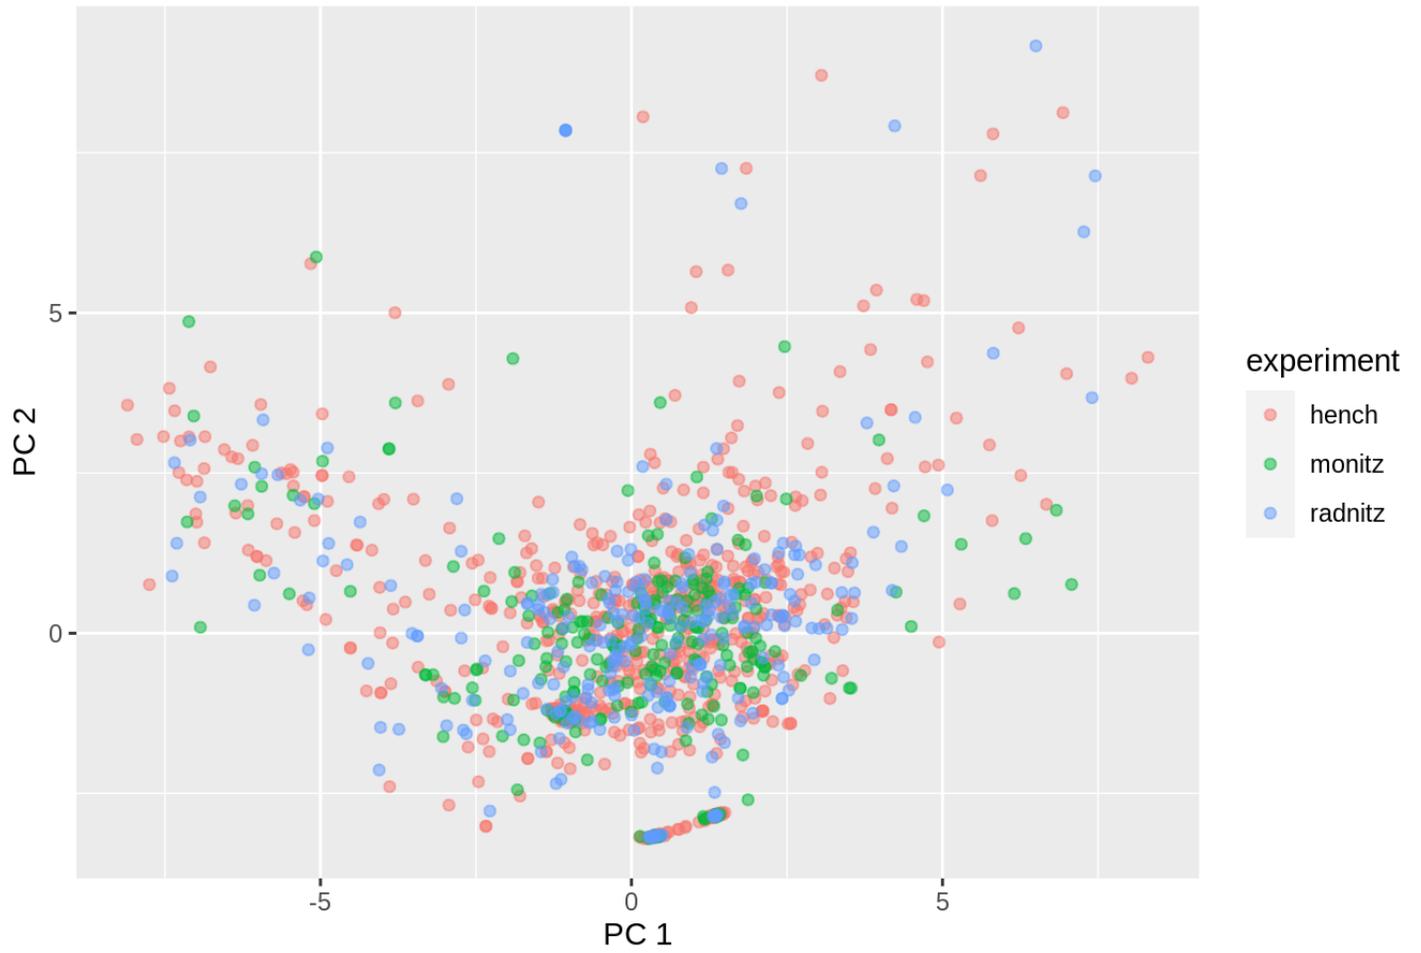


Figure 3: Plot of scores for PC 1 and PC 2 generated from PCA of spiking data, colored by X position.

PCA plot, colored by x

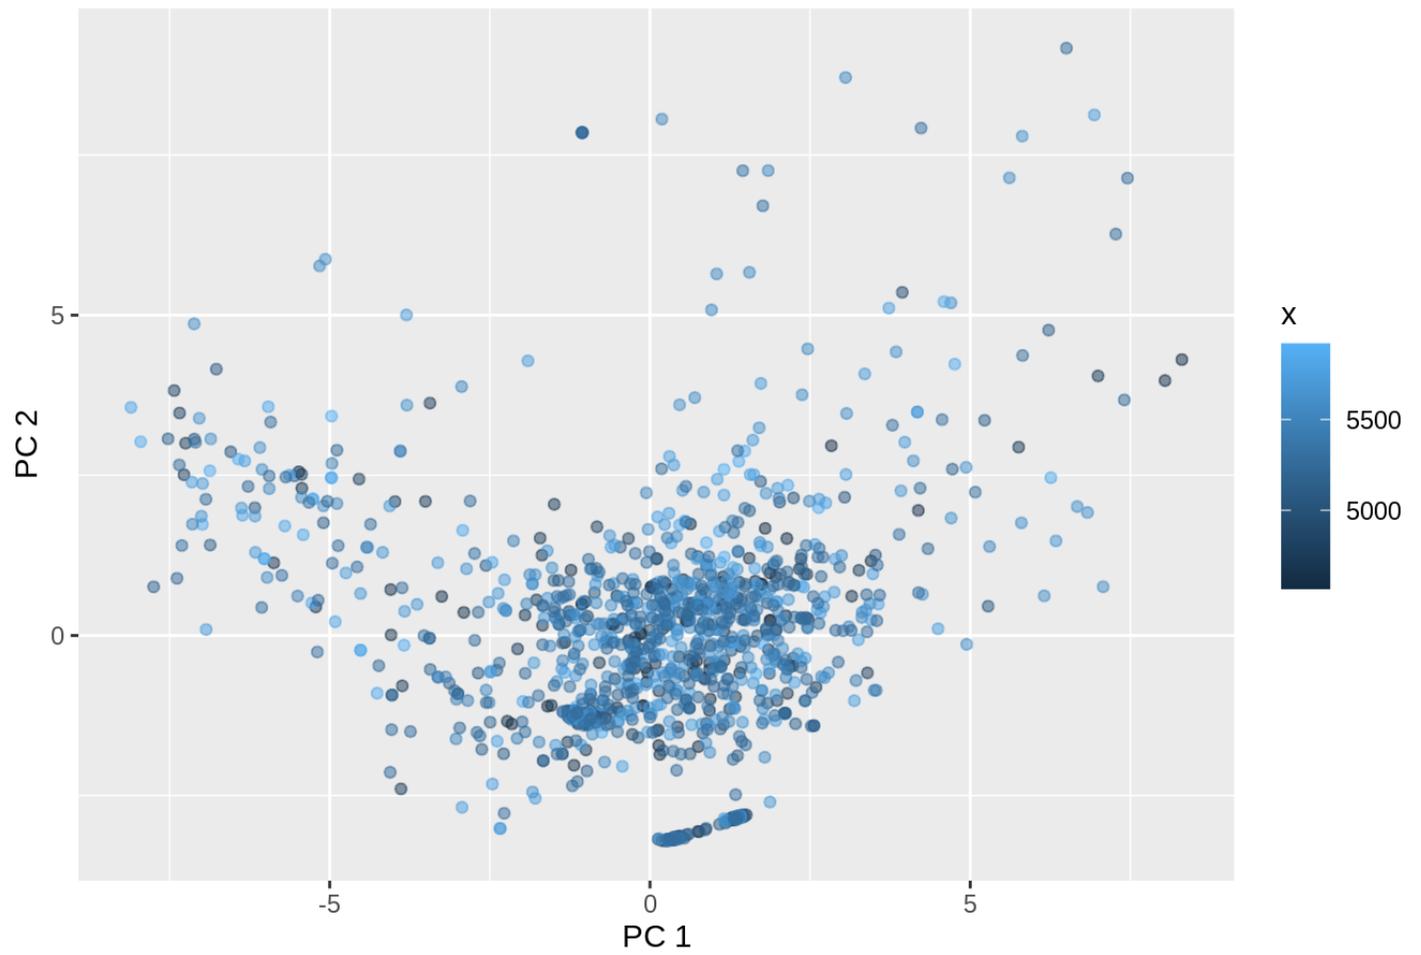


Figure 4: Plot of scores for PC 1 and PC 2 generated from PCA of spiking data, colored by Y position.

PCA plot, colored by y

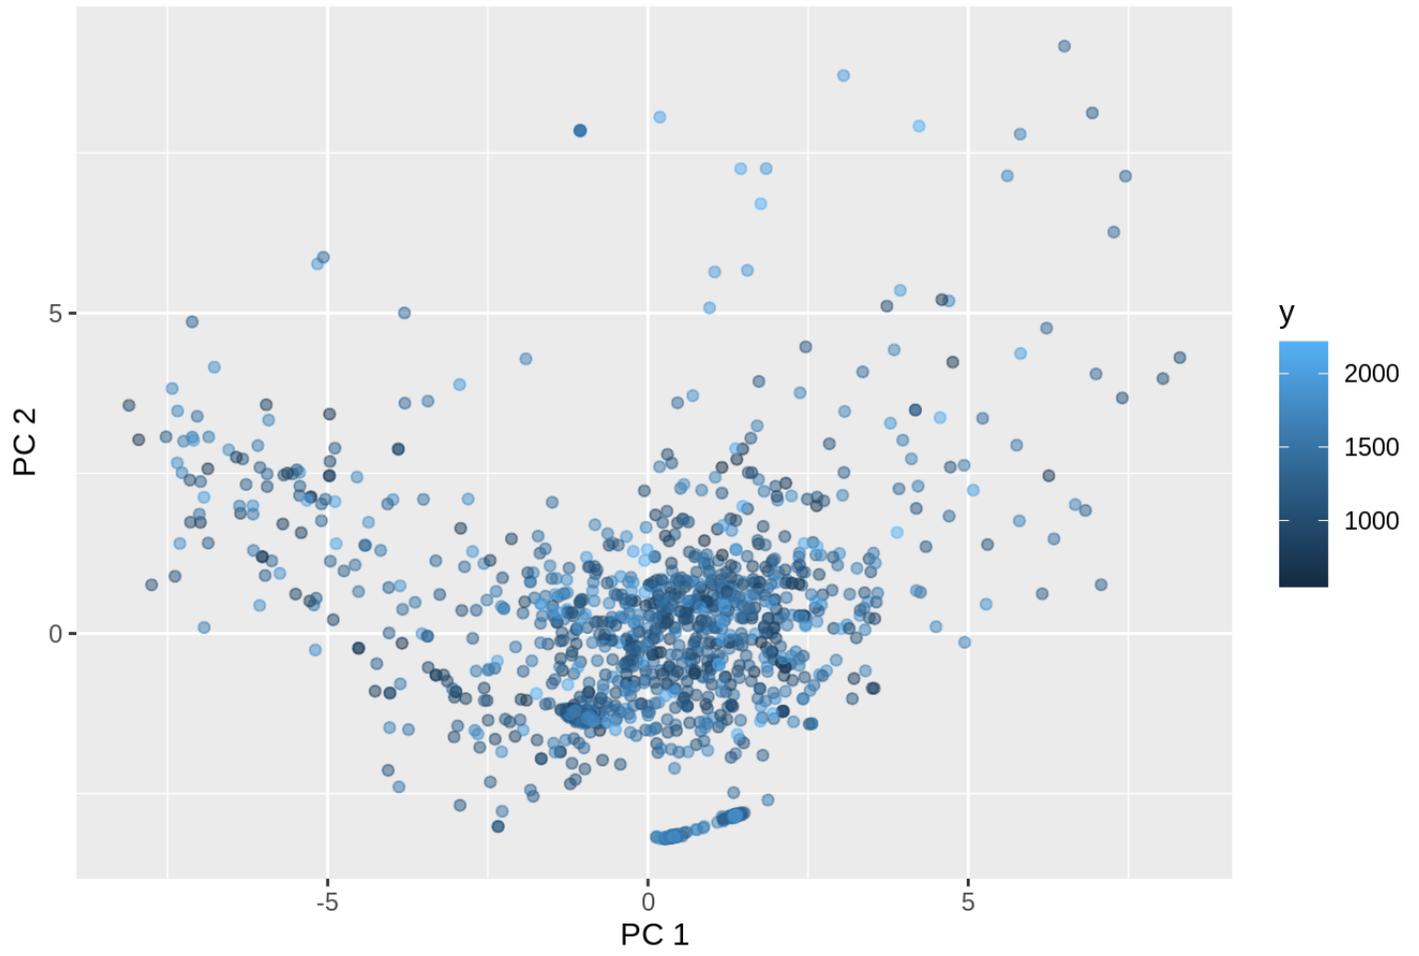


Figure 5: Plot of scores for PC 1 and PC 2 generated from PCA of spiking data, colored by Z position.

PCA plot, colored by z

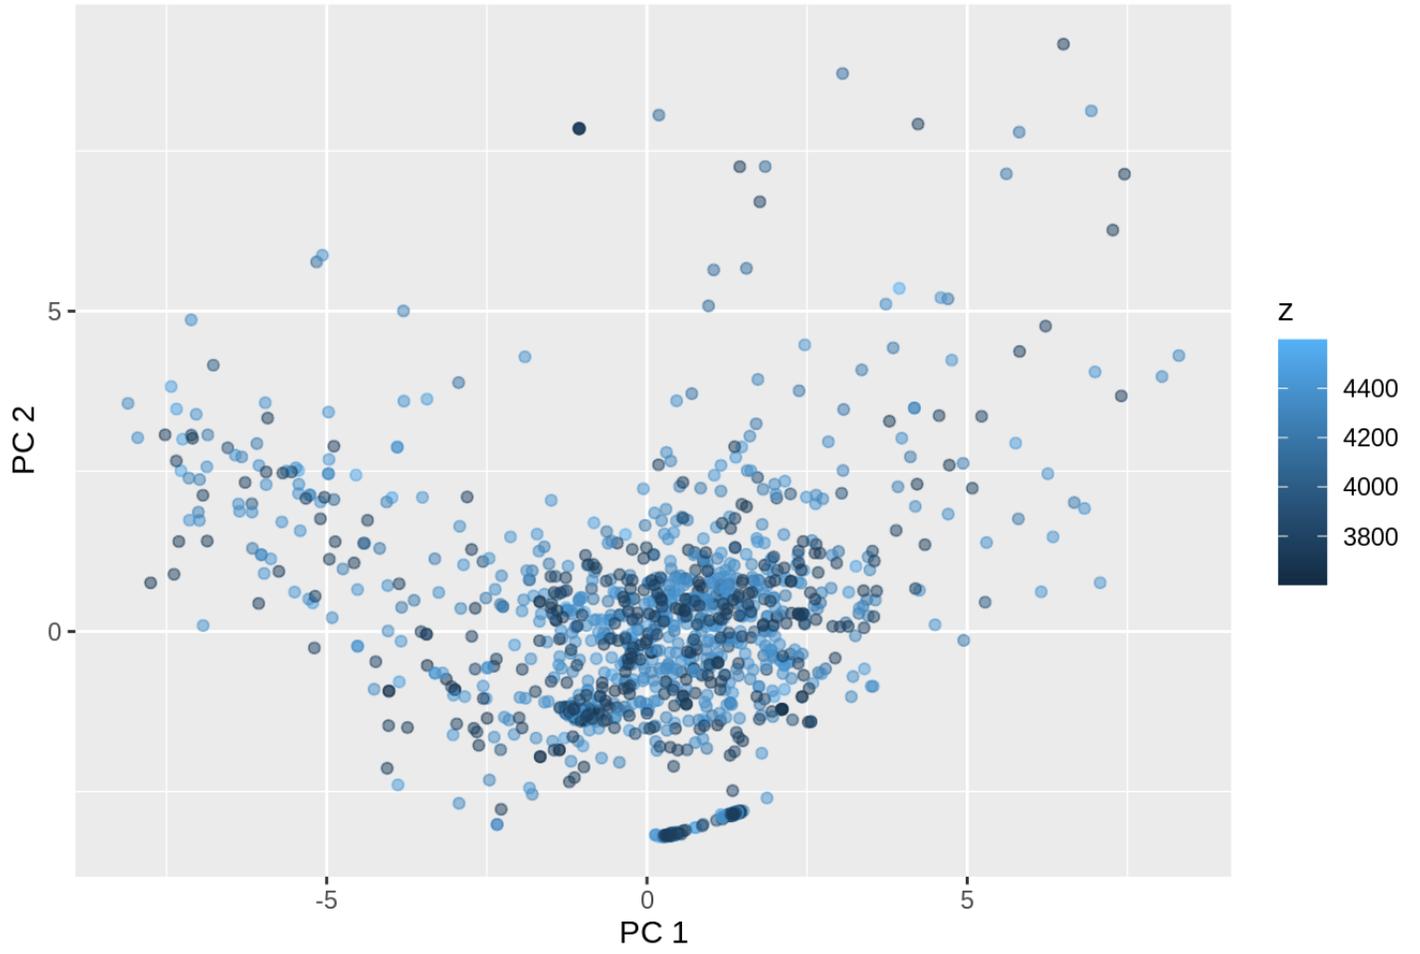


Figure 6: Plots of the eight eigenvectors accounting for the highest variance obtained from application of PCA to spiking data.

Traces of eigenvectors for first 8 PCs

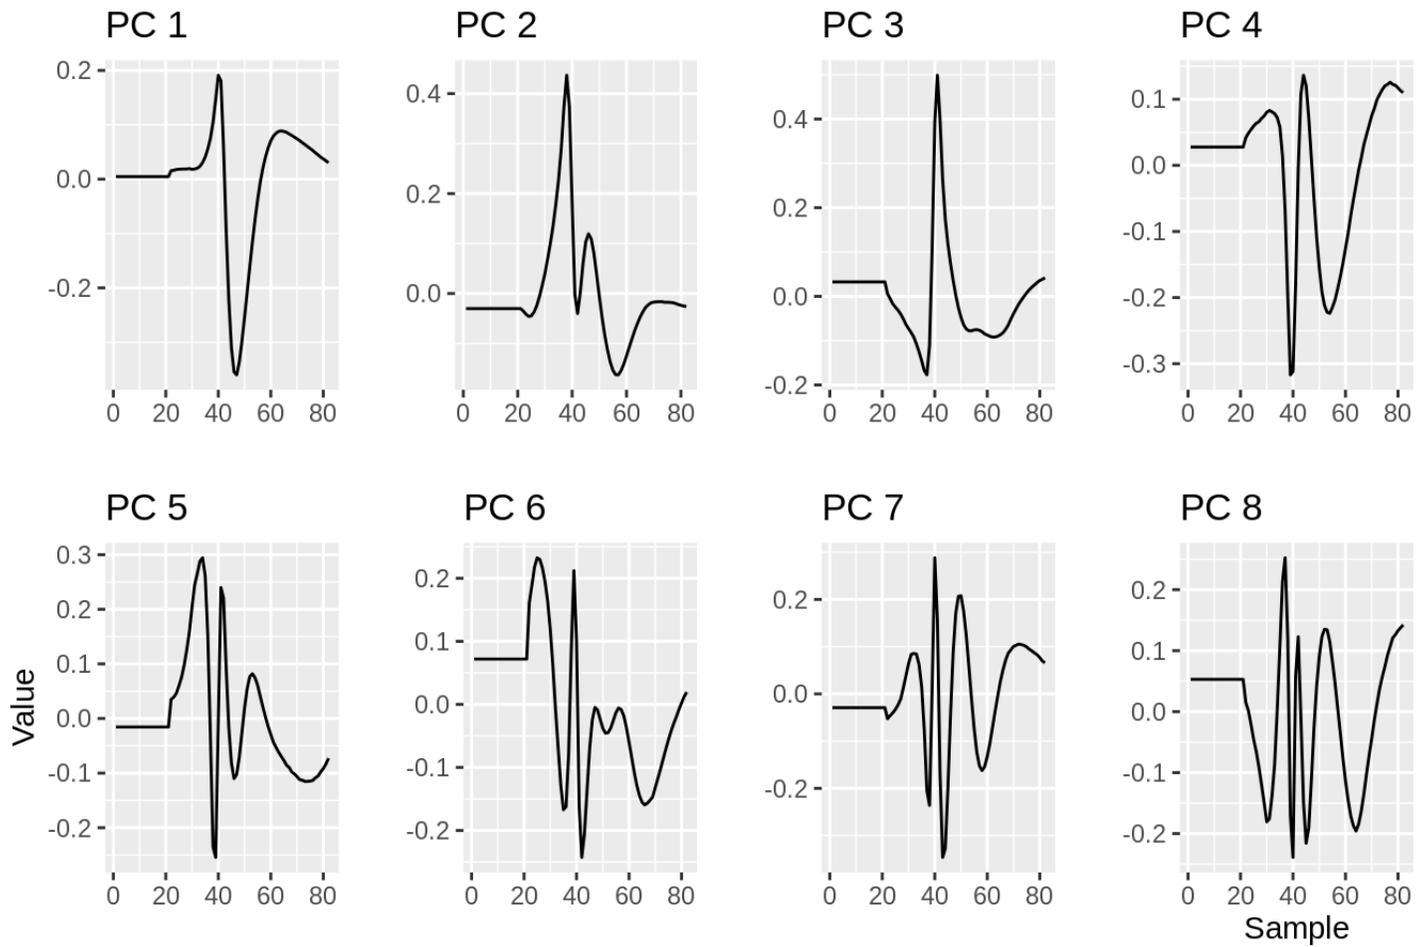


Figure 7: Cumulative variance explained for NLPCA principal components found from spiking data

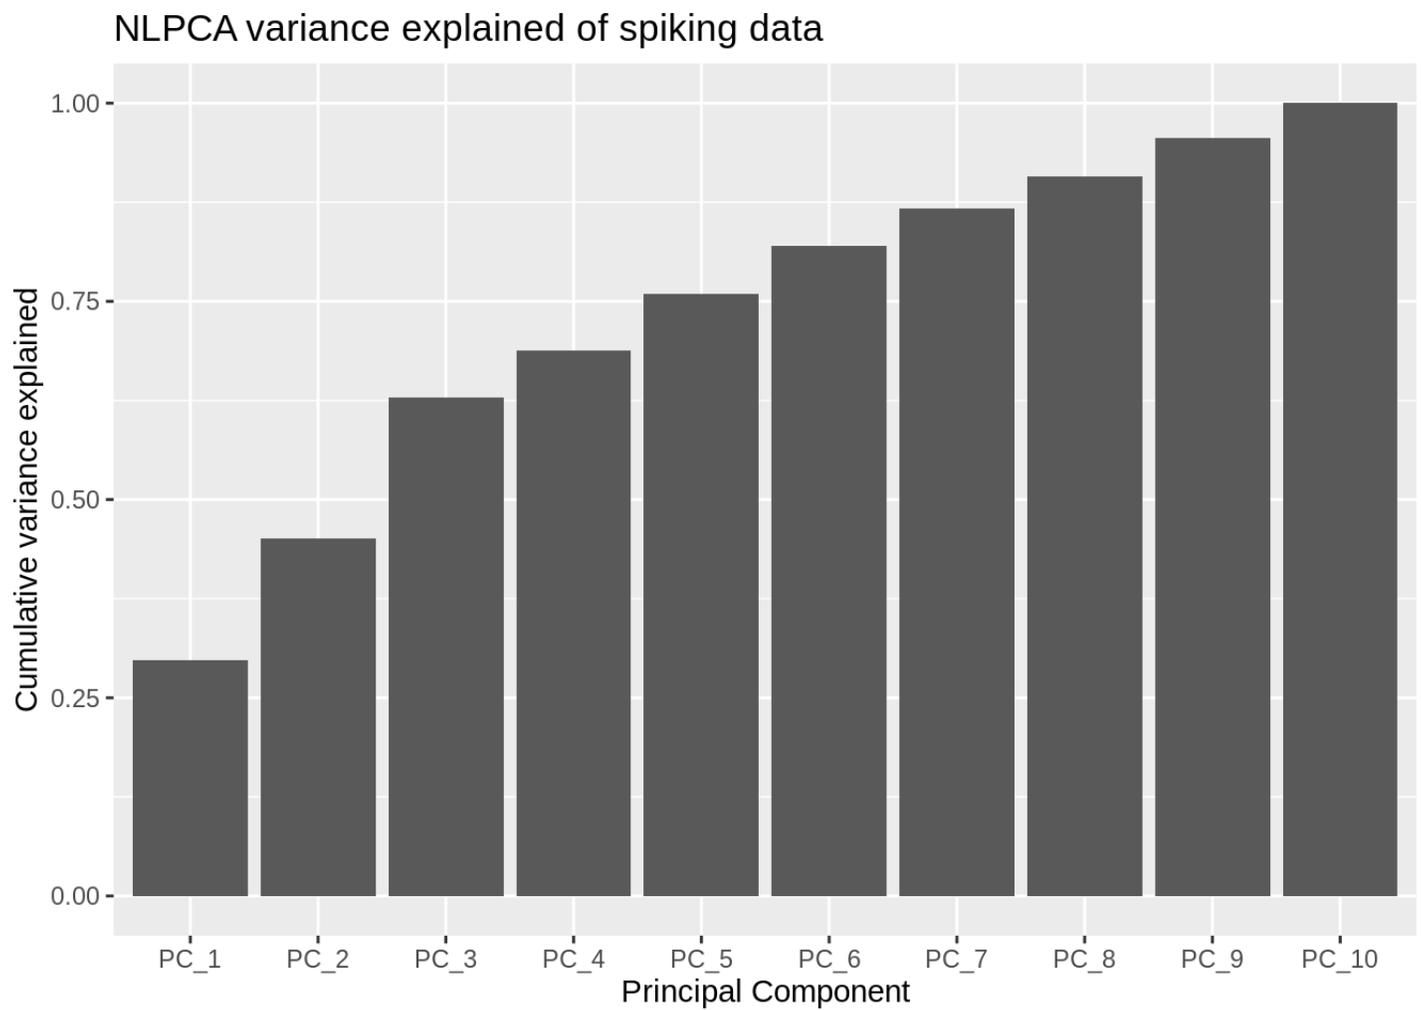


Figure 8: Plot of scores for PC 1 and PC 2 generated from NLPCA of spiking data, colored by experiment.

NLPCA plot, colored by experiment

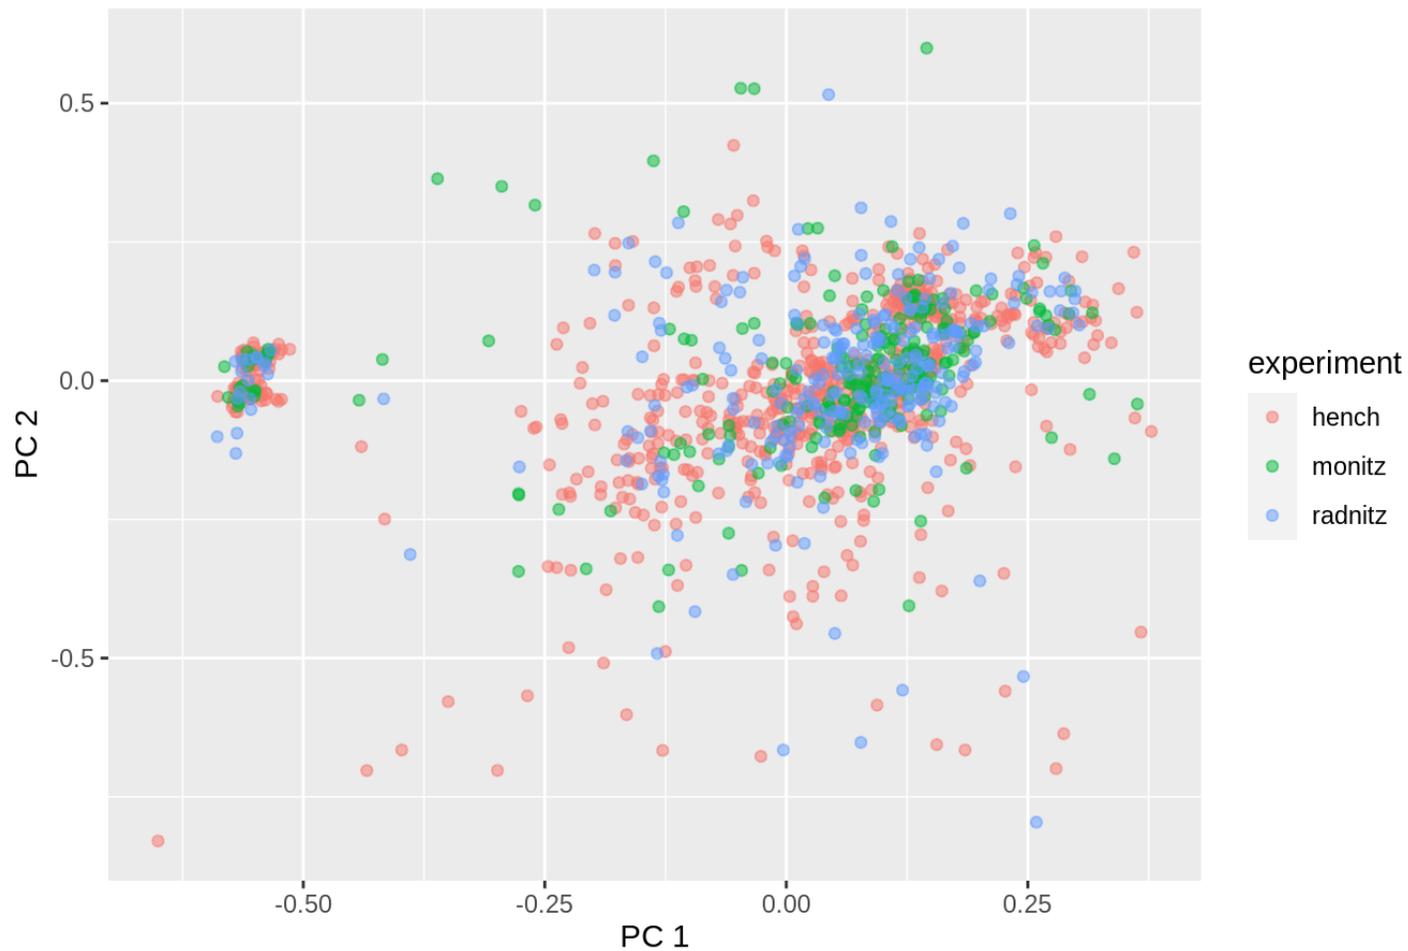


Figure 9: Plot of scores for PC 1 and PC 2 generated from NLPCA of spiking data, colored by X position.

NLPCA plot, colored by x

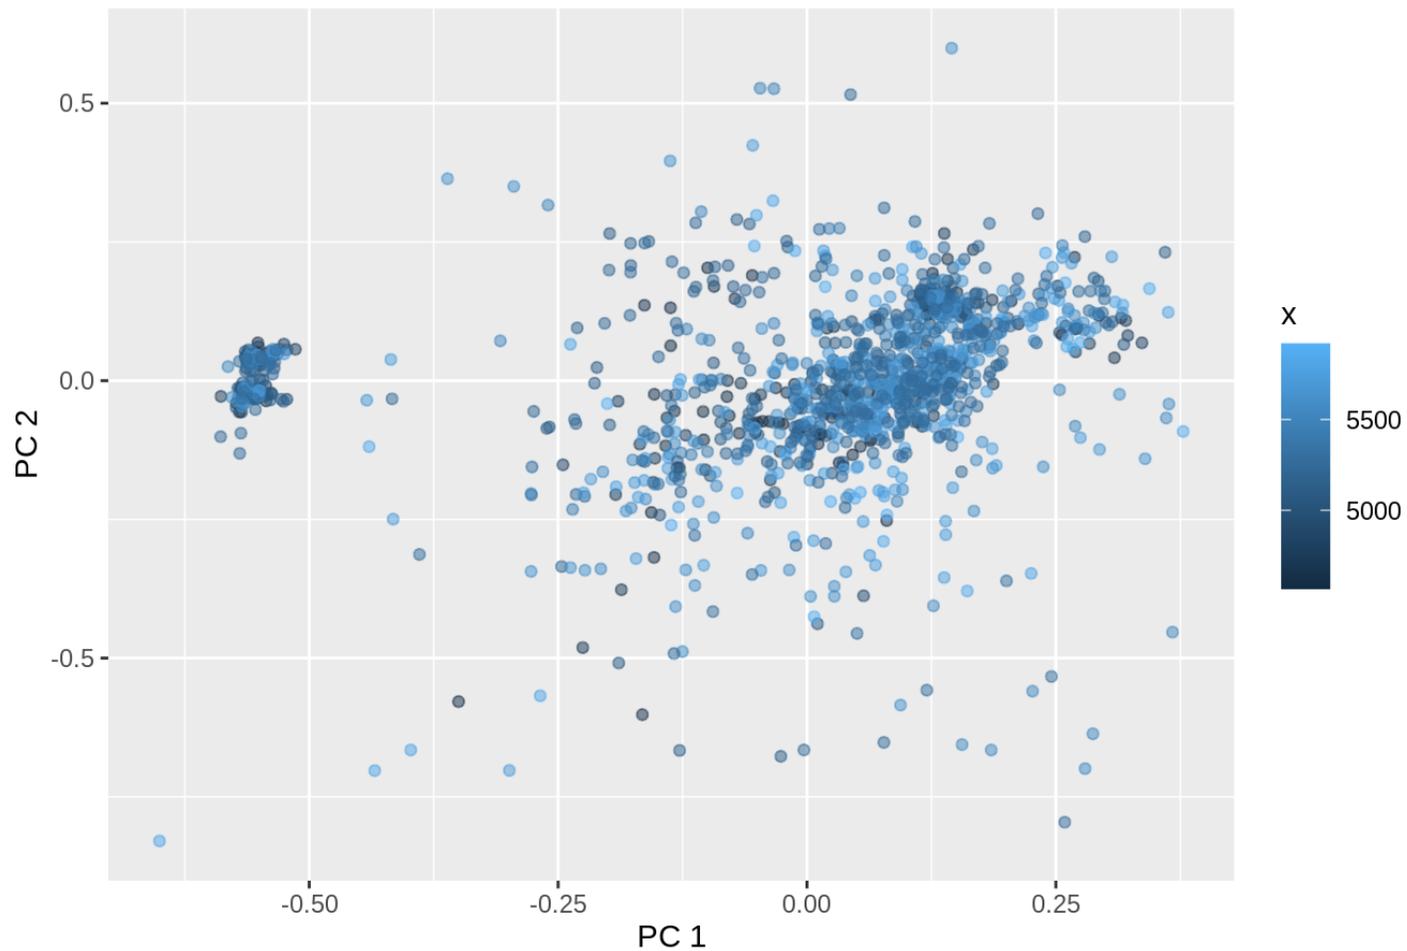


Figure 10: Plot of scores for PC 1 and PC 2 generated from NLPCA of spiking data, colored by Y position.

NLPCA plot, colored by y

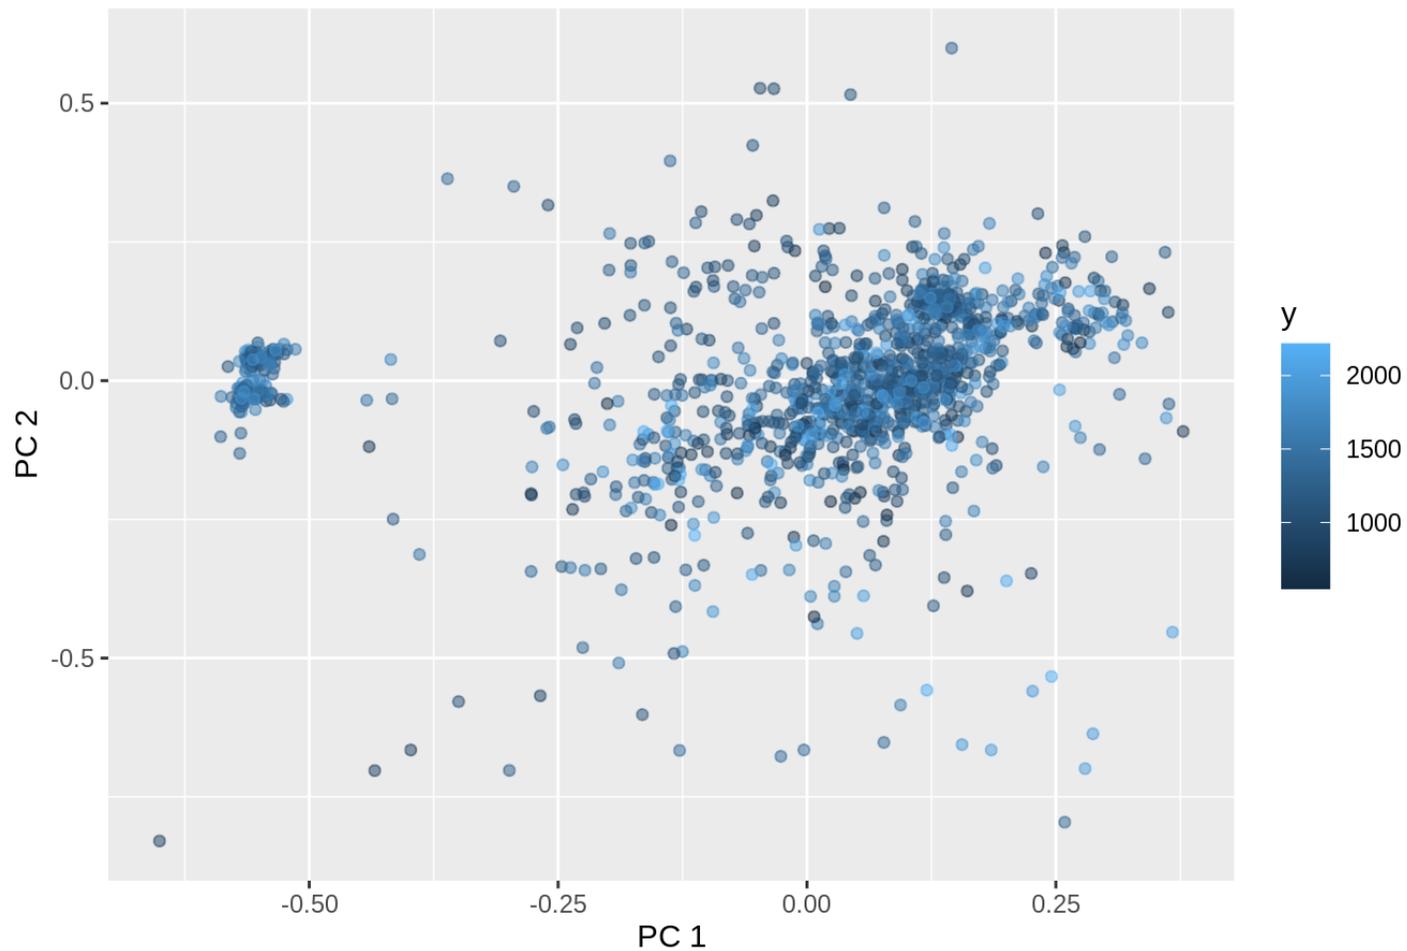


Figure 11: Plot of scores for PC 1 and PC 2 generated from NLPCA of spiking data, colored by Z positon.

NLPCA plot, colored by z

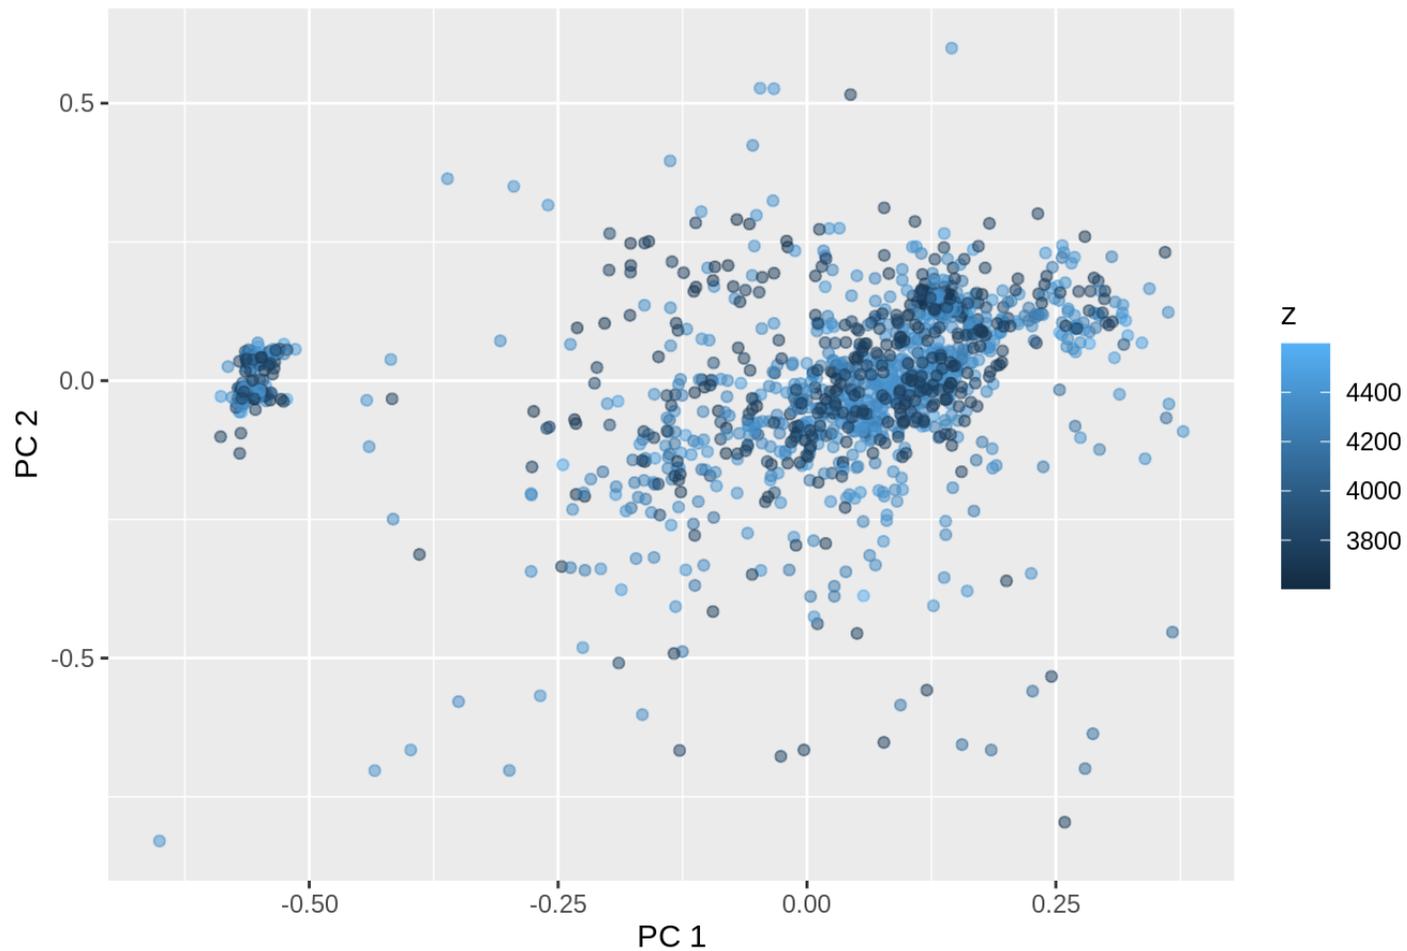


Figure 12: Cumulative variance explained for KPCA principal components found from spiking data

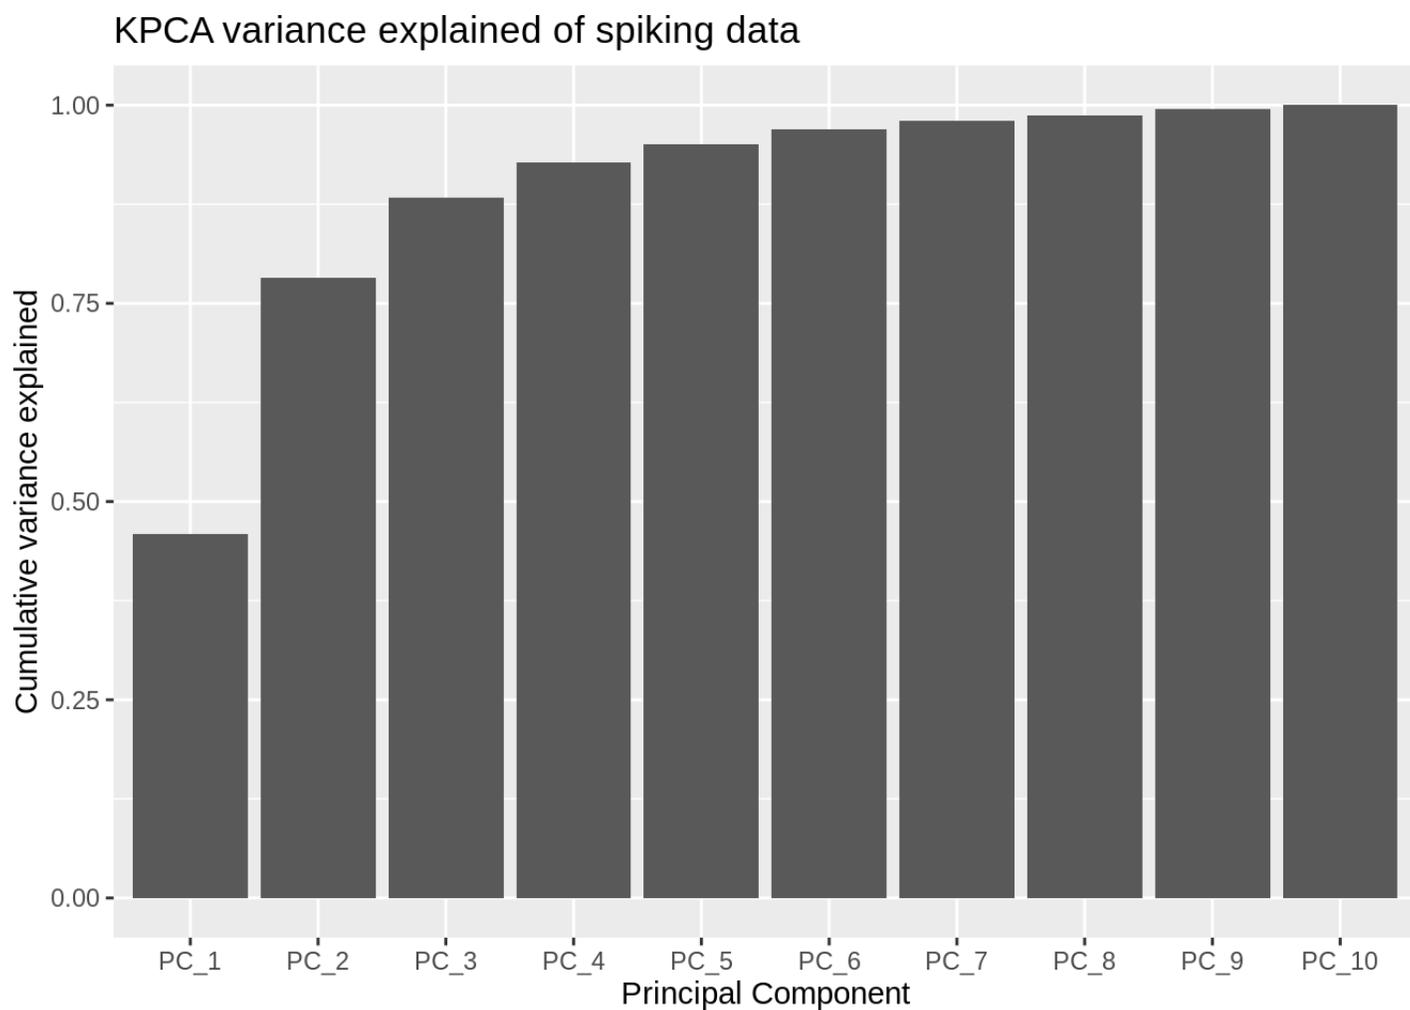


Figure 13: Plot of scores for PC 1 and PC 2 generated from KPCA of spiking data, colored by experiment.

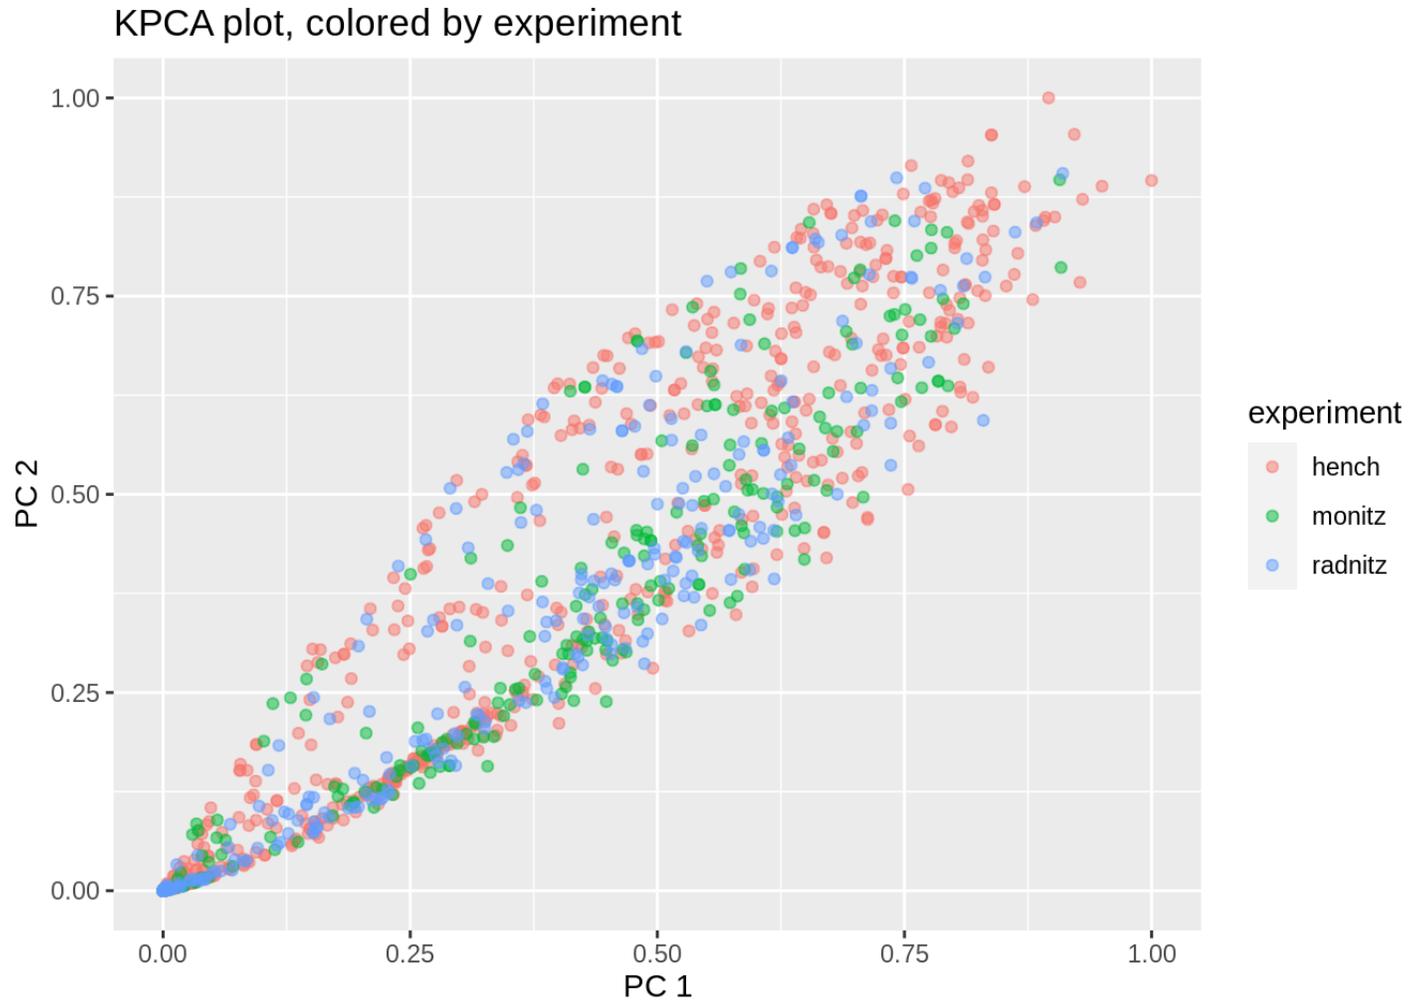


Figure 14: Plot of scores for PC 1 and PC 2 generated from KPCA of spiking data, colored by X position.

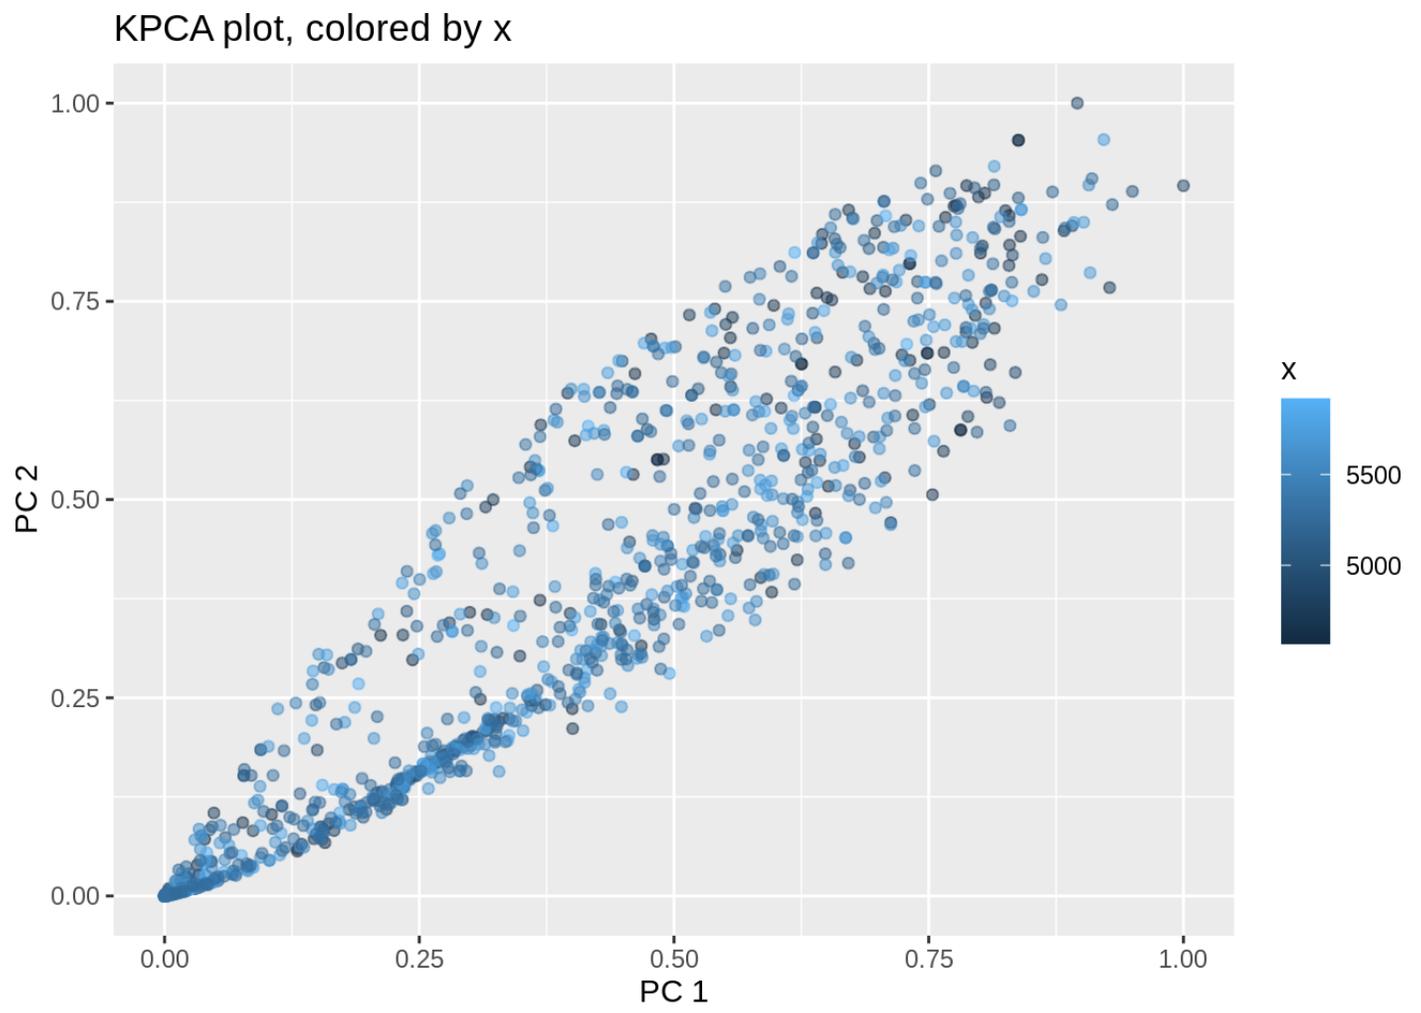


Figure 15: Plot of scores for PC 1 and PC 2 generated from KPCA of spiking data, colored by Y position.

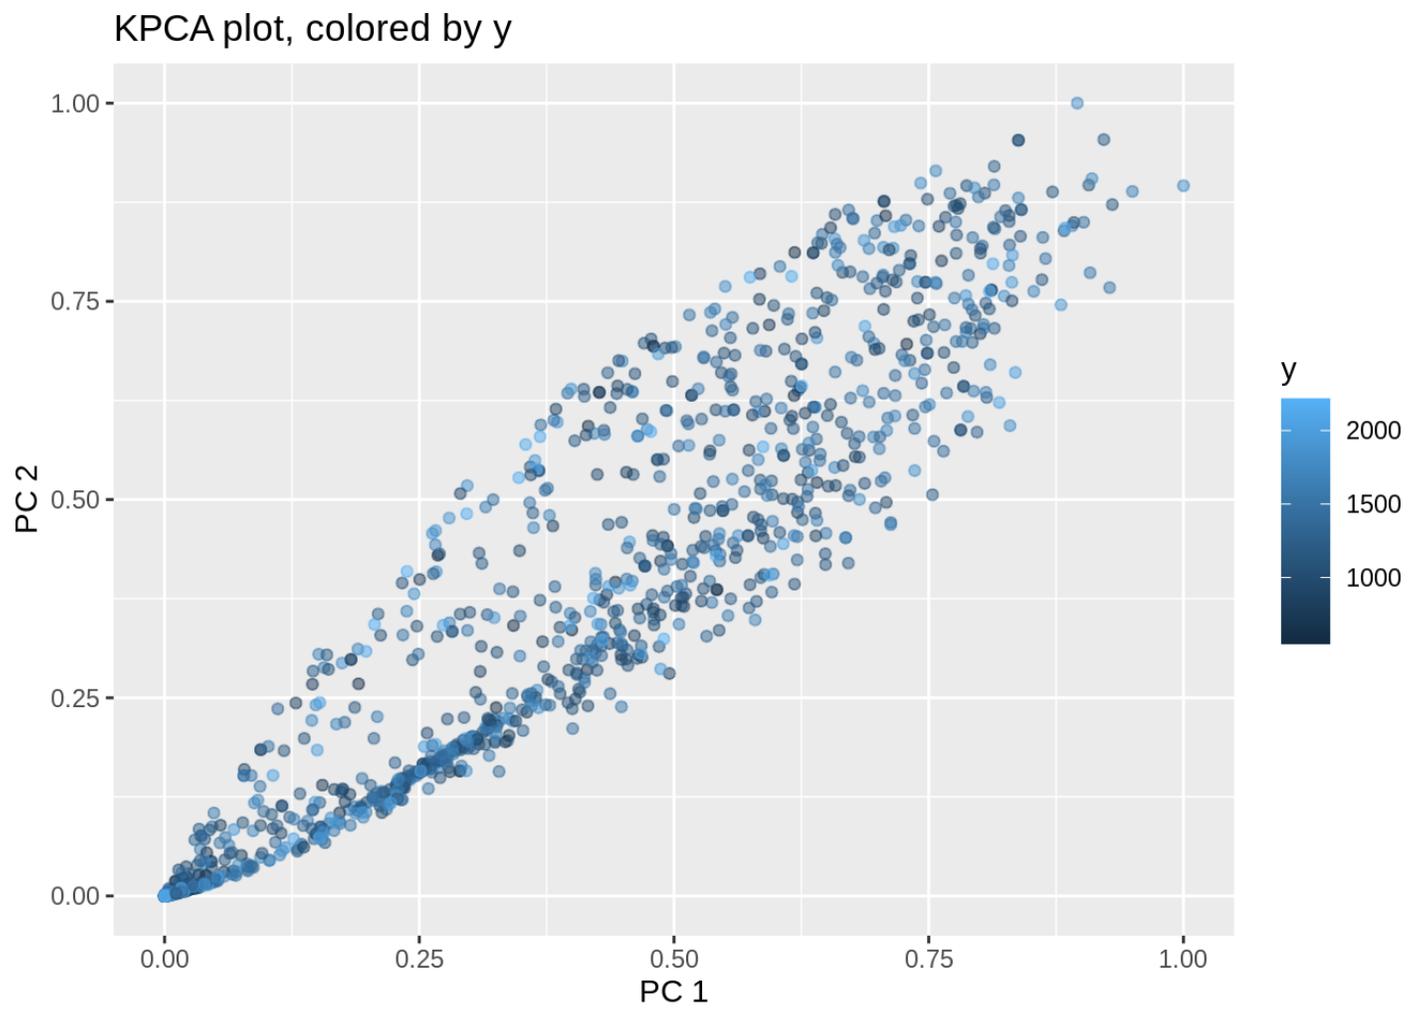


Figure 16: Plot of scores for PC 1 and PC 2 generated from KPCA of spiking data, colored by Z position.



Figure 17: Cumulative variance explained for PCA principal components found from cell type data

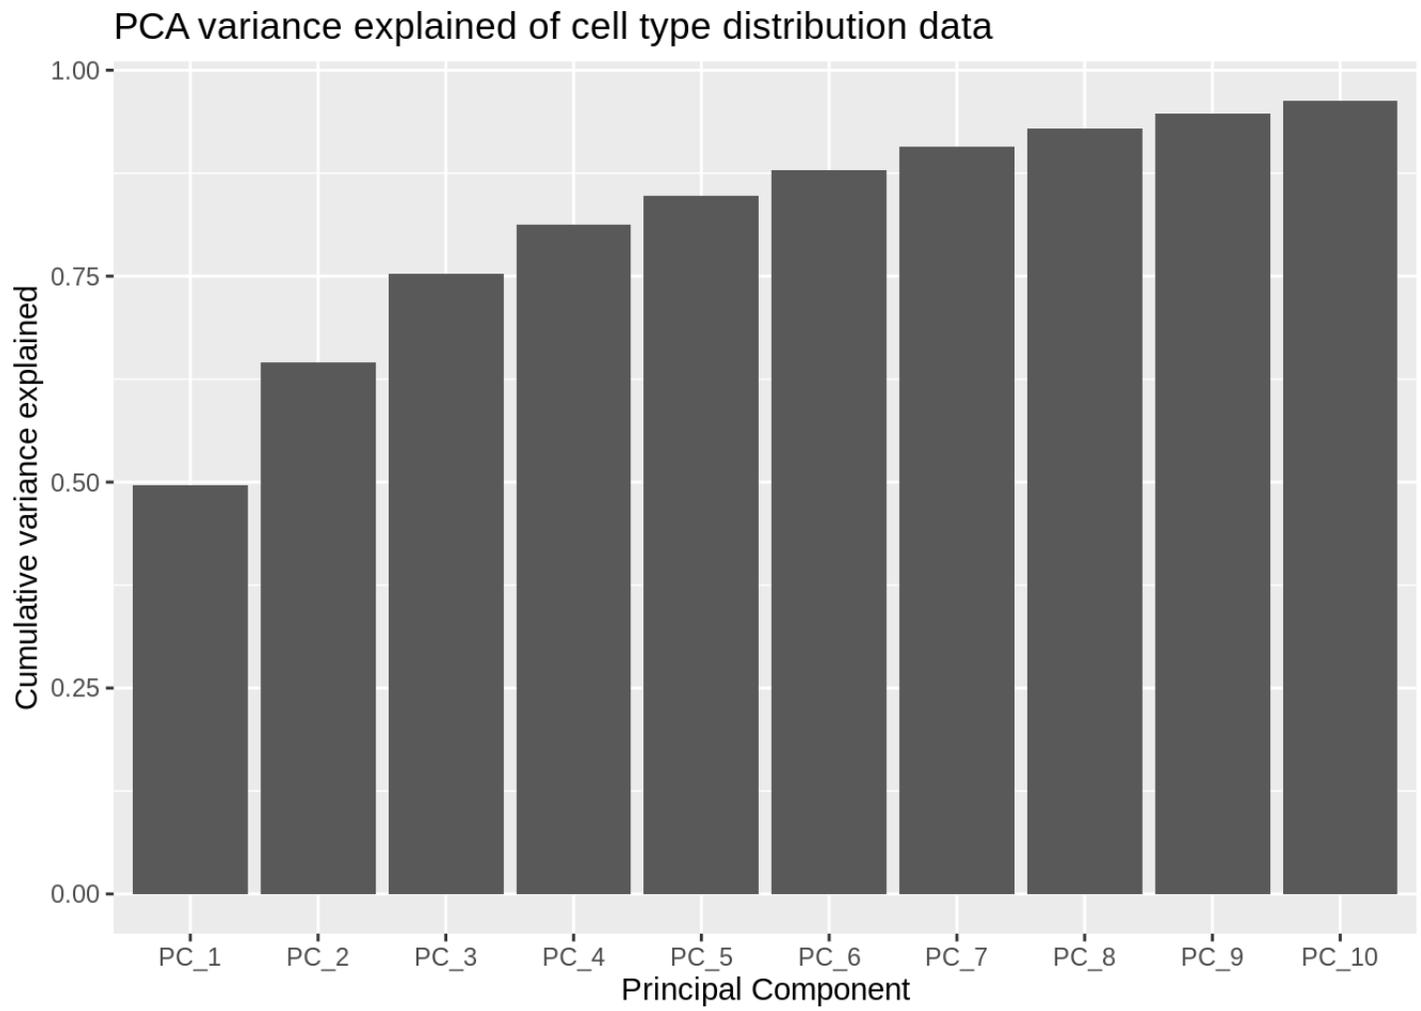


Figure 18: Plot of scores for PC 1 and PC 2 generated from PCA of cell type data, colored by X position.

PCA plot of cell type distribution data, colored by x

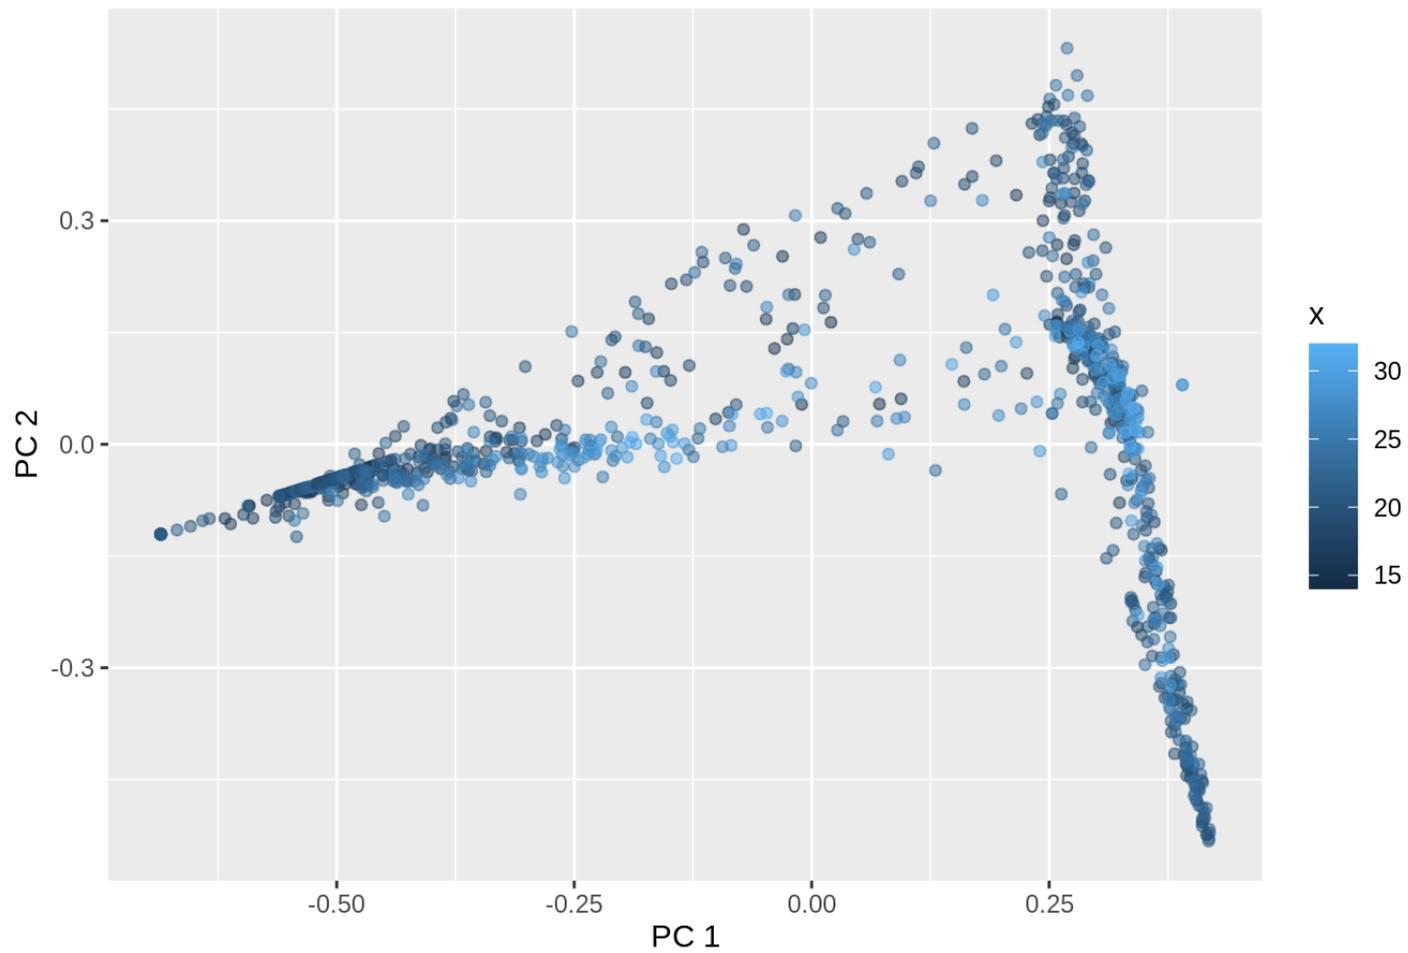


Figure 19: Plot of scores for PC 1 and PC 2 generated from PCA of cell type data, colored by Y position.

PCA plot of cell type distribution data, colored by y

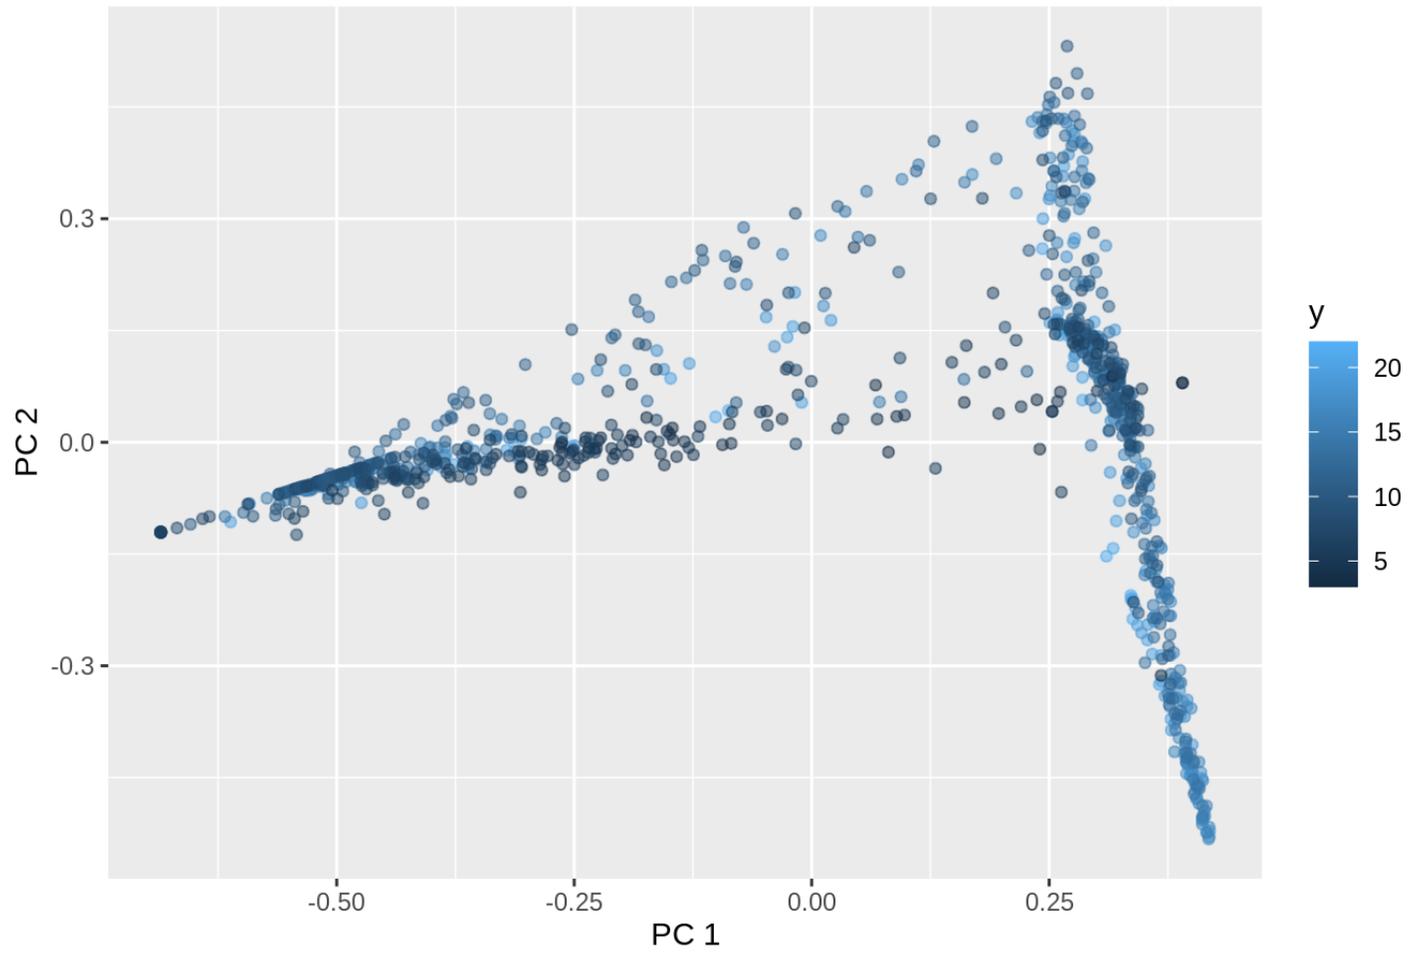


Figure 20: Plot of scores for PC 1 and PC 2 generated from PCA of cell type data, colored by Z position.

PCA plot of cell type distribution data, colored by z

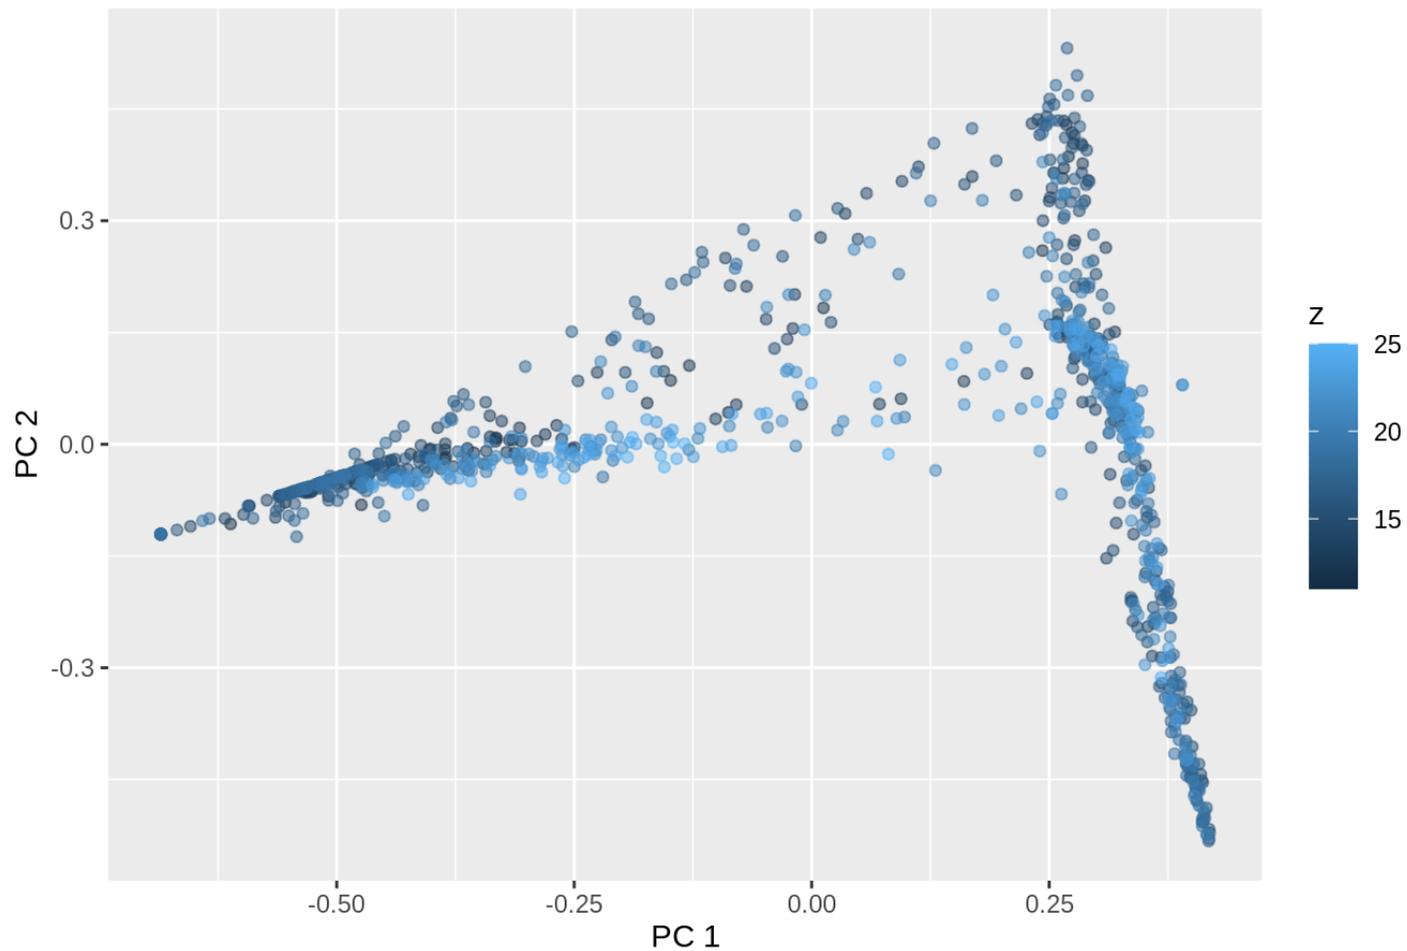


Figure 21: Cumulative variance explained for NLPCA principal components found from cell type data

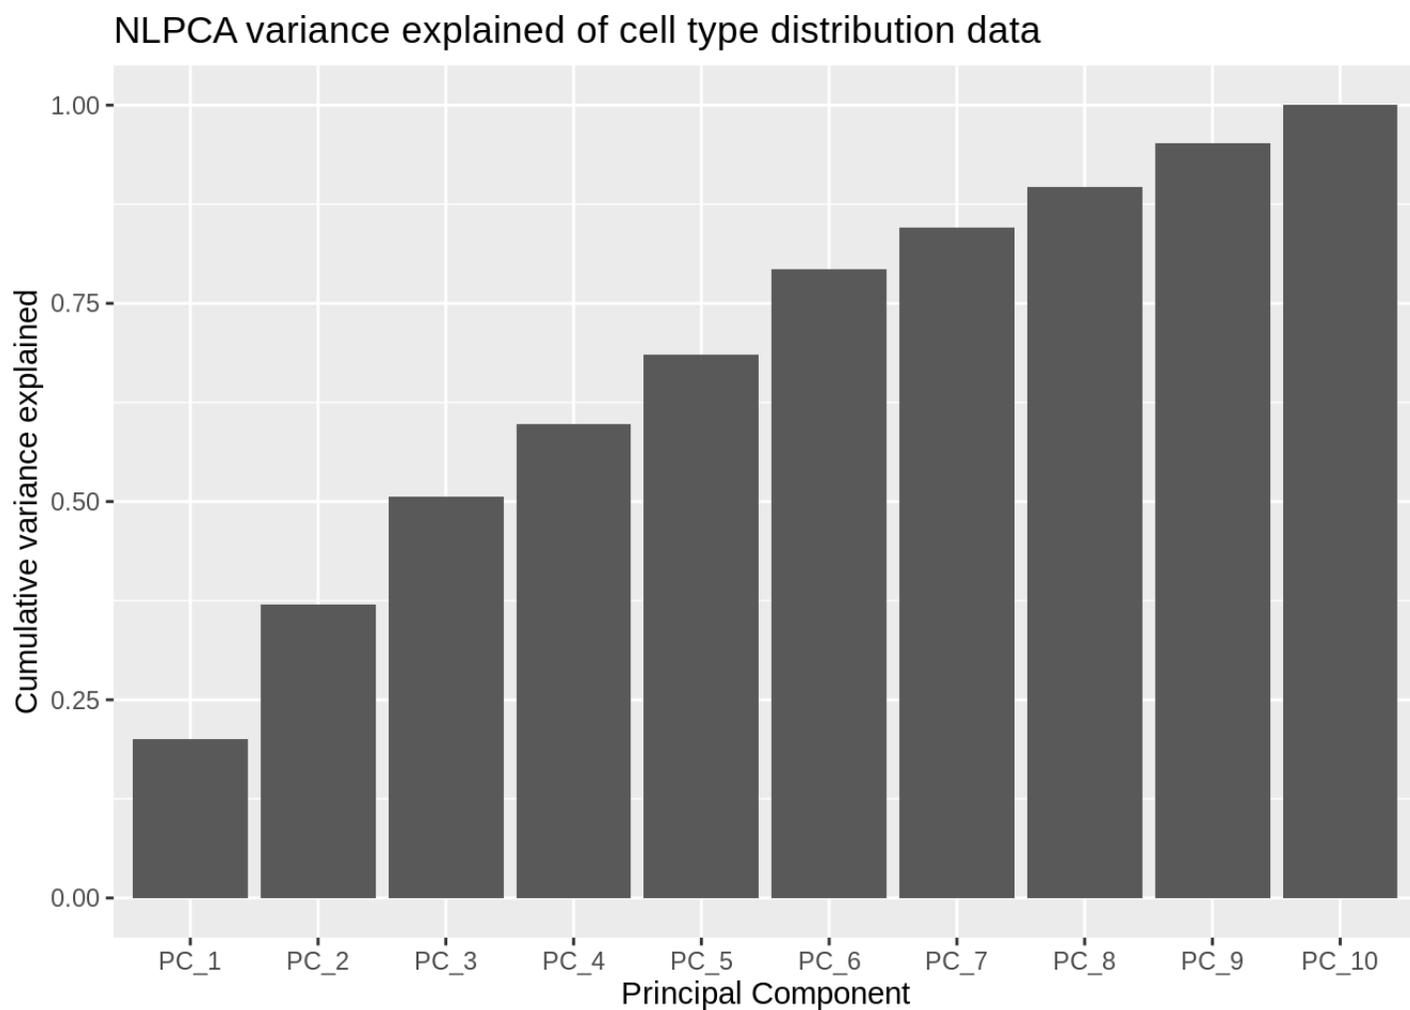


Figure 22: Plot of scores for PC 1 and PC 2 generated from NLPCA of cell type data, colored by X position.

NLPCA plot, colored by x

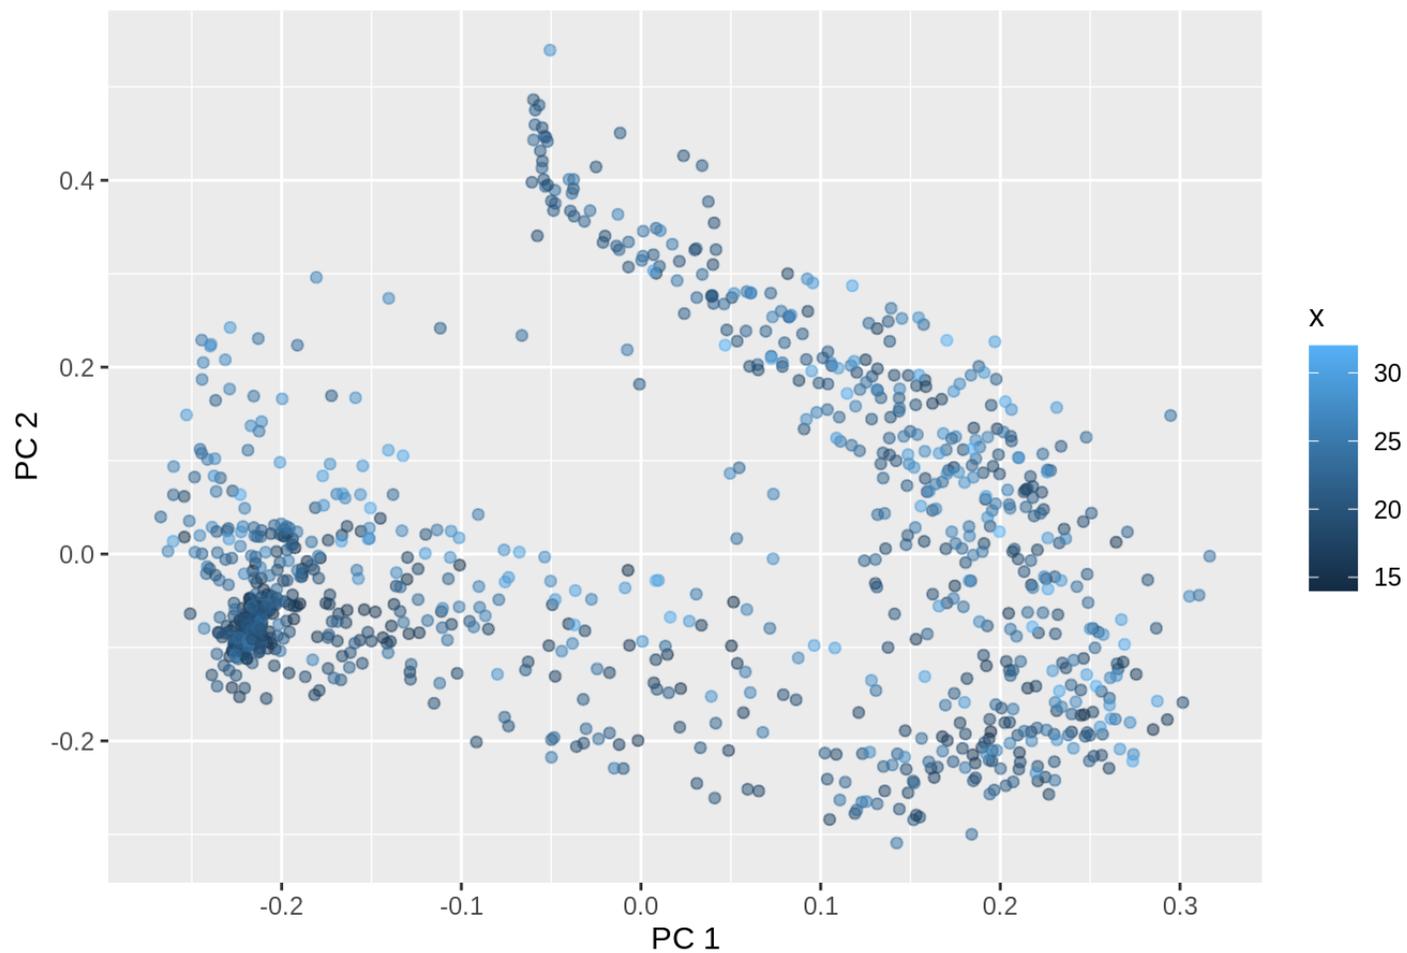


Figure 23: Plot of scores for PC 1 and PC 2 generated from NLPCA of cell type data, colored by Y position.

NLPCA plot, colored by y

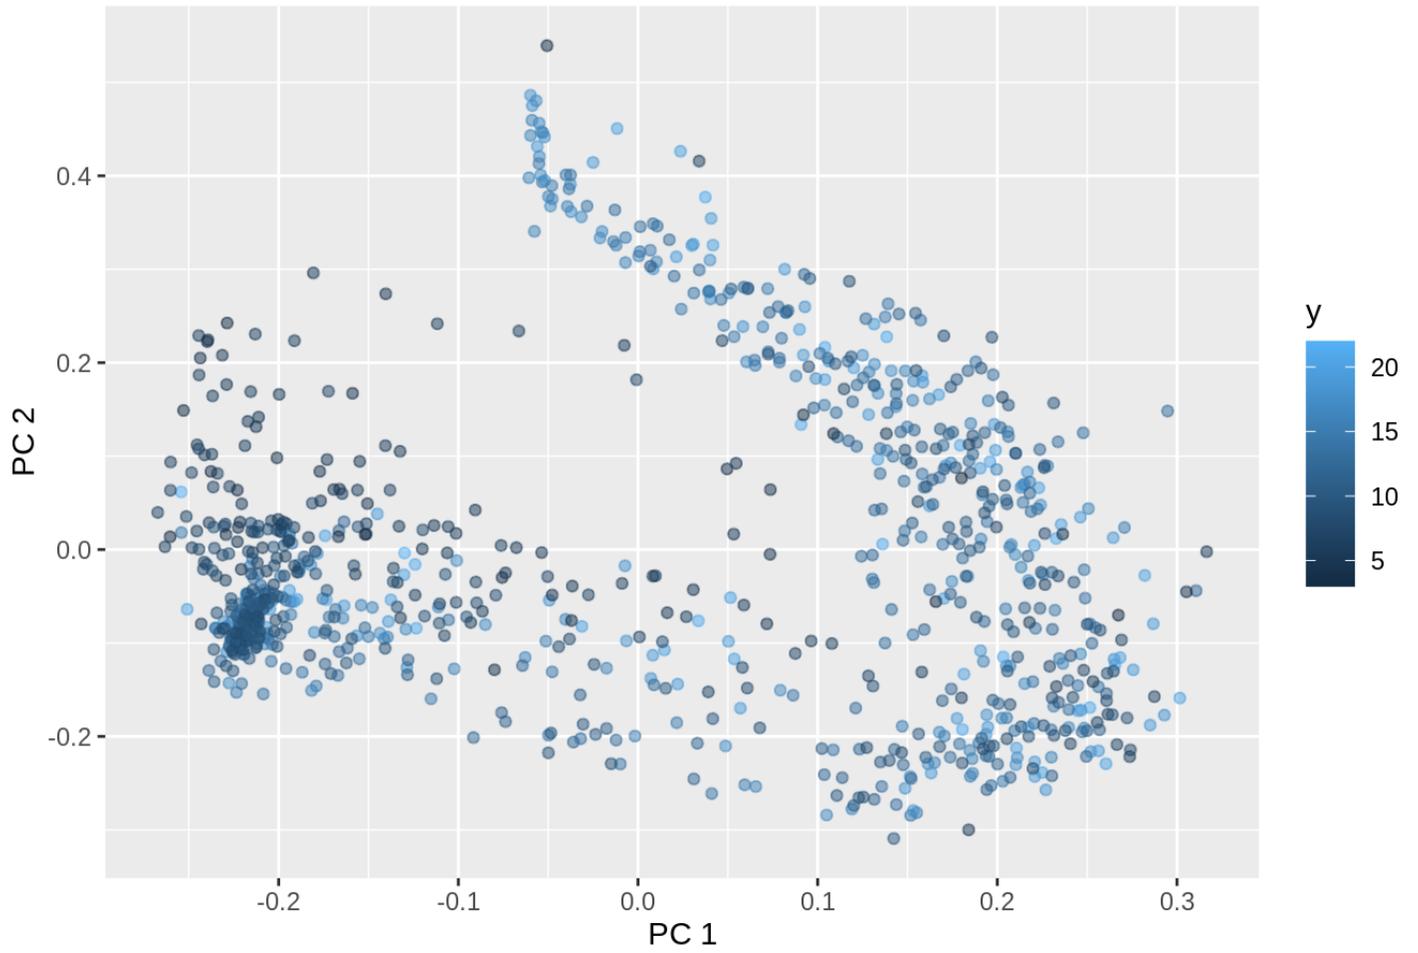


Figure 24: Plot of scores for PC 1 and PC 2 generated from NLPCA of cell type data, colored by Z position.

NLPCA plot, colored by z

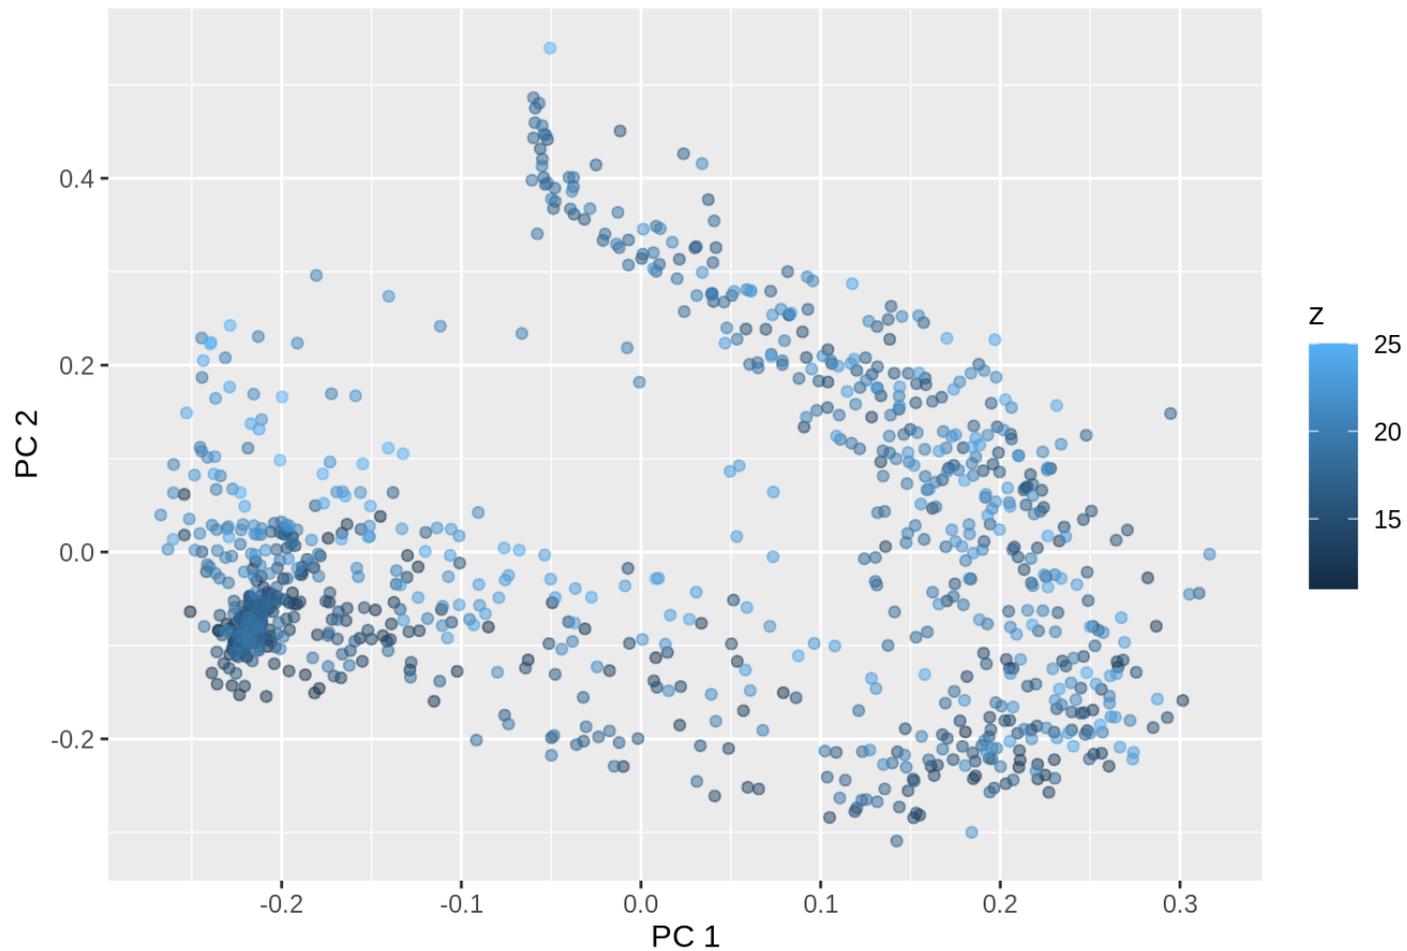


Figure 25: Cumulative variance explained for KPCA principal components found from cell type data

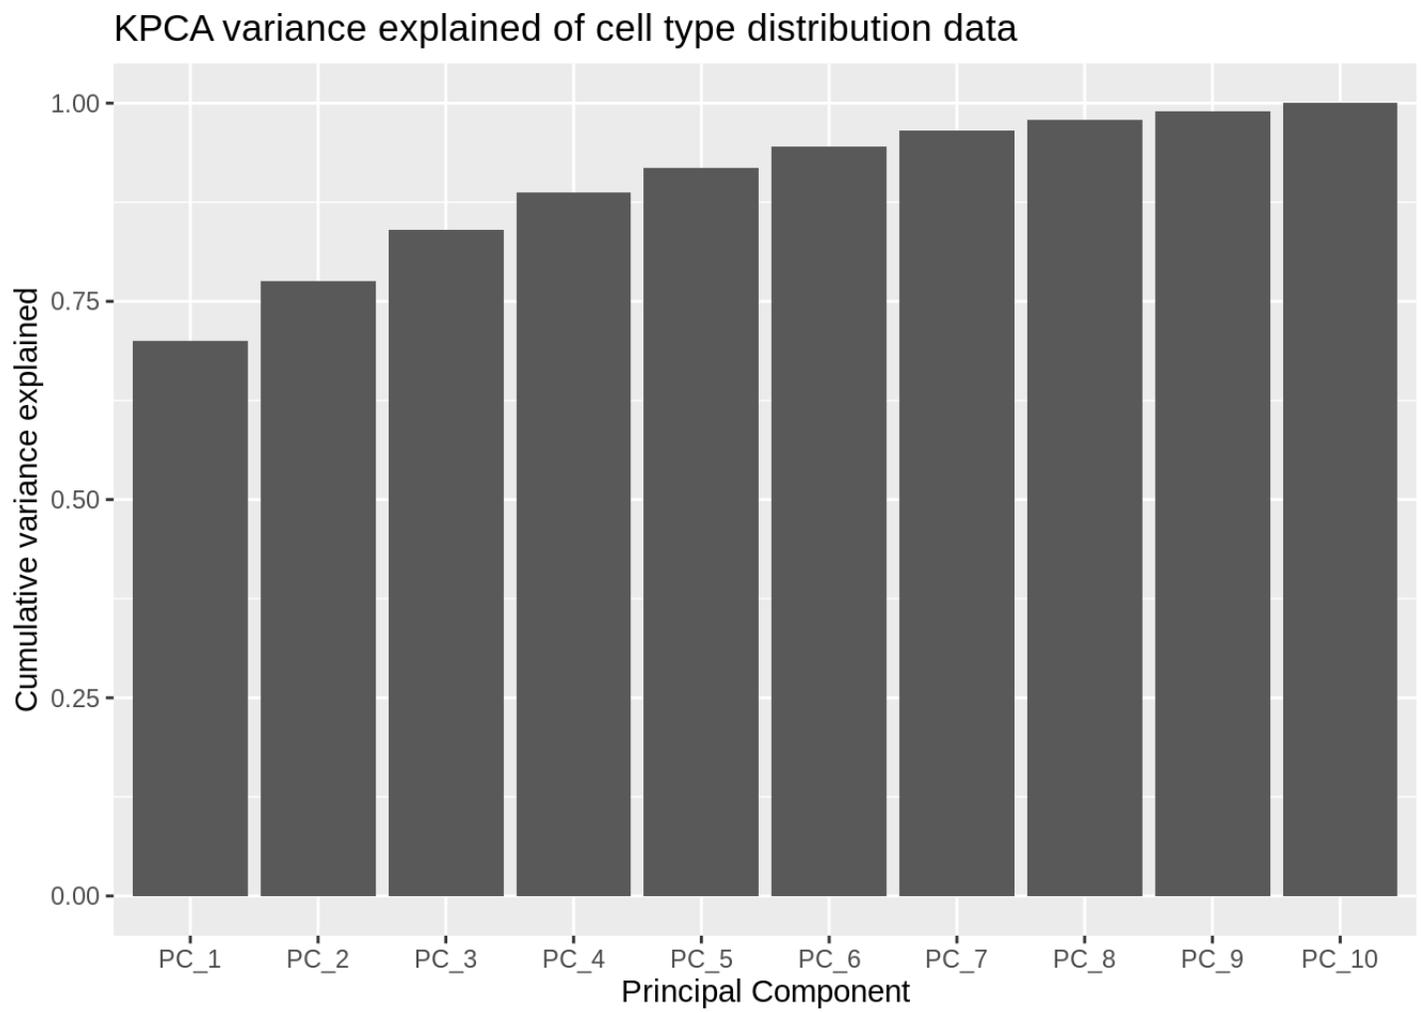


Figure 26: Plot of scores for PC 1 and PC 2 generated from KPCA of cell type data, colored by X position.

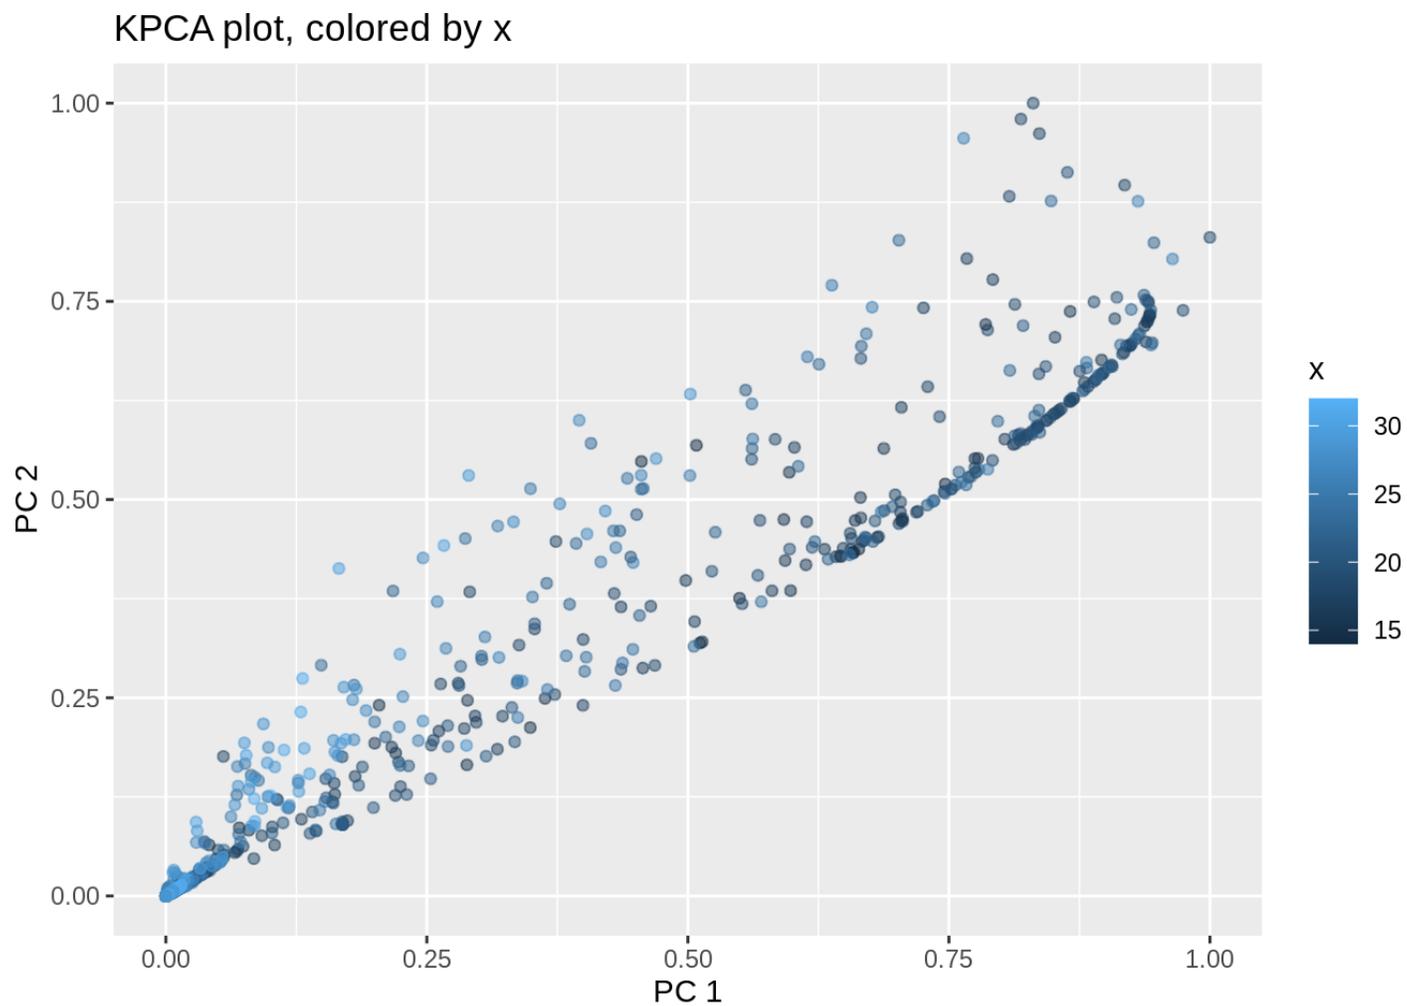


Figure 27: Plot of scores for PC 1 and PC 2 generated from KPCA of cell type data, colored by Y position.

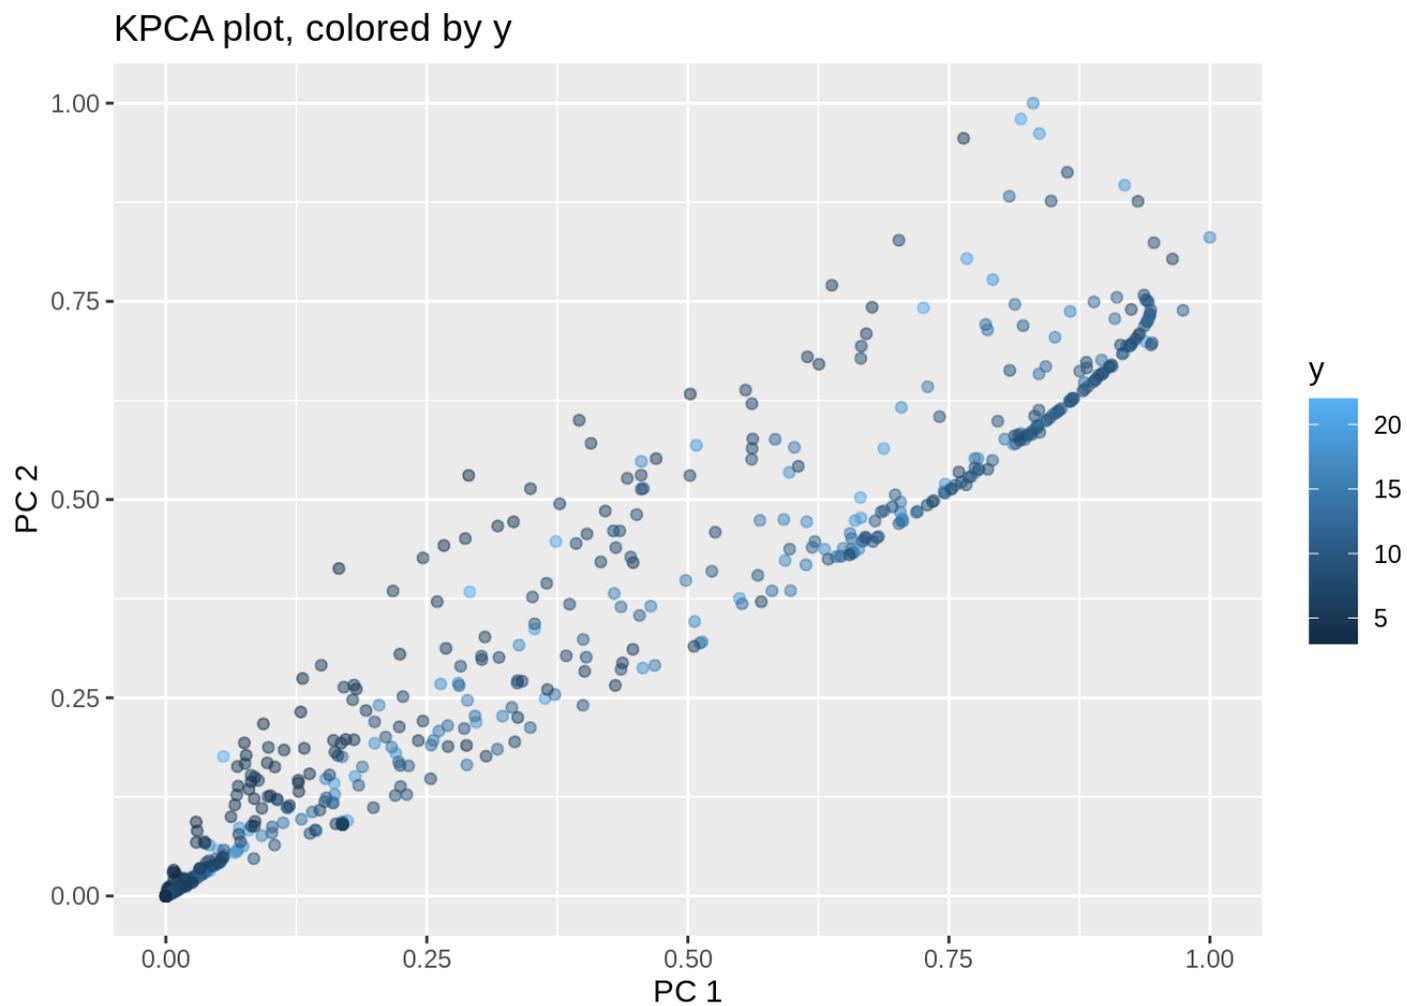


Figure 28: Plot of scores for PC 1 and PC 2 generated from PCA of cell type data, colored by Z position.

