

EPiC Series in Computing

Volume 114, 2026, Pages 125–134

Proceedings of the Workshop on Mathematical Modeling and Scientific Computing



Distributional and Information-Theoretic Decomposition Technique for Characterisation of 3D Cell-Protein Interactions

Iuliia Kurnaeva and Anna Maslovskaya

Laboratory for Simulation and Intelligent Analysis of Complex Systems, Innopolis University, Innopolis, Russia
a.maslovskaya@innopolis.ru

Abstract

Three-dimensional immunofluorescent images of the brain are the subject of interdisciplinary research at the intersection of image analysis, statistics, and modelling. By examining the interaction between microglia and protein aggregates in such images, it is possible to assess the brain's immune response and draw conclusions that are important for the diagnosis of Parkinson's disease. Typically, this interaction is described by a single number per field of view, the average or the overlap fraction – implicitly assuming that the relationship between an aggregate's volume and its coverage is linear and uniform across all aggregates. We propose an approach that, instead of a single number, describes the interaction using a distribution and evaluates it using methods from information theory. We distinguish between overlap occurrence and overlap depth, using a population-averaged model for occurrence, quantile regression for depth, and mutual information for both. This allows us to see how aggregate volume affects each variable and to detect disease-related shifts that are lost when averaging. We applied this approach to 16.7 million three-dimensional aggregates from postmortem human brain tissue of donors with Parkinson's disease and age-matched HC. Aggregate volume is a strong predictor of overlap occurrence, but we found no evidence that this relationship differs between groups. The only disease-related signal is in the median overlap depth: in the putamen, overlap is shallower in Parkinson's disease, and the median registers a shift several times stronger than the mean, although a cohort of 27 donors does not establish this at conventional significance. Because aggregates within a donor are not independent, all inference is donor-level or donor-clustered; aggregate-level counts are descriptive, not independent observations. The volume of the aggregate is approximately an order of magnitude more informative about the depth of overlap than about the mere fact of overlap, and large aggregates overlap less completely in all donors from both groups.

1 Introduction

Parkinson's disease is characterised by the progressive accumulation of aggregates of phosphorylated α-synuclein (pSyn) [5], which triggers a neuroinflammatory response. Microglia – the brain's resident immune cells, engulf these aggregates and degrade them intracellularly; this is

V. Turova, A. Kovtanyuk, A. Maslovskaya and J. Zimmer (eds.), MMSC 2026 (EPiC Series in Computing, vol. 114), pp. 125–134

called phagocytic clearance [11]. In Parkinson's disease, this process is disrupted, and the aggregates accumulate instead of being degraded. This has been demonstrated in experiments on cells and animals, as well as in people with a genetic predisposition to Parkinson's disease [9,10].

Phagocytosis is a physical process, and the volume of the target aggregate constrains it. Yet tools for extracting quantitative information about the microglial state from images remain scarce [4], and the interaction between microglia and protein aggregates is still routinely reduced to global averages or a proportion of scalar overlap per field of view, which assumes that the relationship between aggregate volume and capture is linear and homogeneous. Such a mean-based approach completely masks underlying distributional shifts; for instance, microglia might capture aggregates with the same overall frequency, but fundamentally alter the depth of engulfment, or lose their size-selective filtering capability entirely. Here and throughout, we use *volume* for the measured quantity V and *size* only in the general sense of "size dependency".

This work continues our previous study [7], in which the interaction between microglia and α-synuclein aggregates in three-dimensional immunofluorescence data was profiled with spatial statistics and explainable deep learning. Here we go beyond that description and, instead of a single summary per field of view, decompose the interaction at the level of individual aggregates into two variables: the occurrence of contact and its depth, assessing each with a population-averaged model, quantile regression and mutual information. Focusing on a per-aggregate analysis, we investigate how the physical volume of a pSyn aggregate dictates both the probability of its capture and the subsequent depth of microglial engulfment, aiming to uncover structural dysfunctions that standard mean-based models fail to detect.

2 Data and target variables: occurrence and depth of contact

We analyse the interaction between microglia and pSyn aggregates at the level of individual aggregates. The overall workflow is summarised in Fig. 1: segmentation converts the image stacks into a table, where each row corresponds to a single aggregate. This table is then split into two target variables: the occurrence of contact with microglia and its depth. Each of these is analysed using a population-averaged model, a distributional model, and an information-theoretic measure.

For our analysis, we used publicly available immunofluorescence images of postmortem human brains from donors with Parkinson's disease at Braak stage 3/4 and age-matched healthy controls (hereafter HC) [1,2]. The source study comprises four imaging components: IBA1/pSyn and MAP2/pSyn imaging for each diagnostic group. We used only the IBA1/pSyn components from the putamen and substantia nigra, comprising 1,402 fields of view (729 putamen and 673 substantia nigra), from which 16.7 million aggregates were segmented. The MAP2/pSyn components were not used.

The dataset covers the nigrostriatal pathway: the substantia nigra (SN) and the putamen. The images consist of high-resolution three-dimensional z-stacks acquired using spinning-disk confocal microscopy with a 100× immersion objective. Binary masks for microglia (IBA1-positive) and pSyn aggregates were generated and processed independently using 3D connected-component labelling with full 26-connectivity, the most permissive criterion that avoids artifactual fragmentation of cellular processes.

Segmentation turns the image stacks into one row per aggregate; Table 1 gives the resulting counts. The dataset contains postmortem tissue from 27 donors (14 with Parkinson's disease, 13 age-matched HC), imaged at an anisotropic voxel size of $0.5 \times 0.11 \times 0.11$ μm, which gives

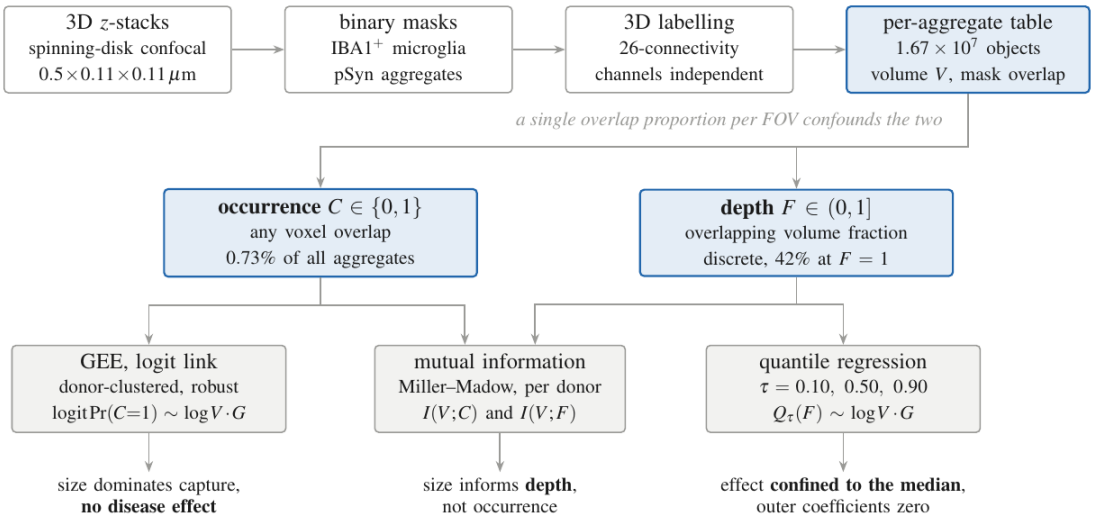


Figure 1: Analysis pipeline.

a voxel volume of 0.00605 μm^3. For each aggregate we recorded its physical volume V and its voxel-level overlap with the microglial mask, and from these we defined two target variables.

Capture C is a binary flag: it is 1 if any voxel of the aggregate lies inside the microglial mask, and 0 otherwise. It records the *occurrence* of phagocytic contact. *Overlap fraction* $F \in (0, 1]$ is defined for captured aggregates as the proportion of the aggregate volume that lies inside the mask, and records the *depth* of engulfment. Figure 2 shows what the two mean physically: an aggregate barely grazed and one entirely engulfed are identical in C and as far apart as they can be in F. A single overlap proportion per field of view, the conventional readout, cannot tell them apart. Capture is also rare, fewer than one aggregate in a hundred, which is what makes the depth of the rare events worth examining separately.

Since the analysis is based on the spatial overlap between the two masks rather than on a lysosomal marker, C and F are used here as proxy measures of microglia-aggregate interaction; we therefore refer to them as overlap occurrence and overlap depth.

Table 1: The segmented dataset.

Quantity	**Putamen**	**Substantia nigra**	**Total**
Fields of view	729	673	1,402
Segmented aggregates	8.29×10^6	8.42×10^6	1.67×10^7
Captured aggregates	38,774	82,837	121,611
Capture rate	0.47%	0.98%	0.73%

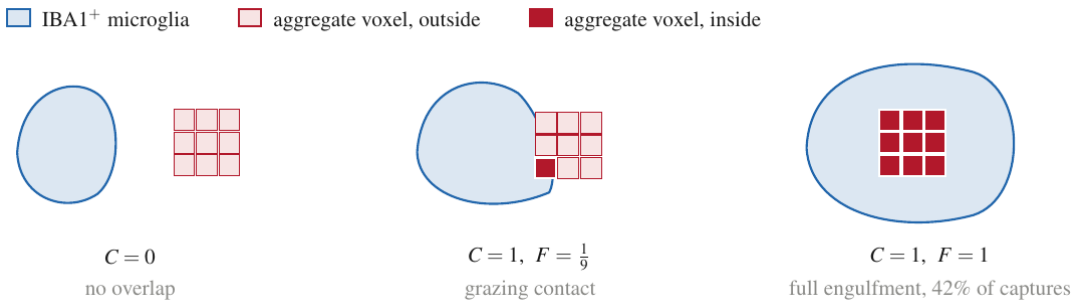


Figure 2: Two target variables: occurrence and depth.

3 Statistical framework

3.1 Population-averaged baseline

Because multiple fields of view originate from the same donor, the observations are not strictly independent, and treating them as such in a standard logistic regression would deflate standard errors and produce falsely optimistic p-values. We therefore modelled the binary probability of capture with a generalized estimating equation (GEE): binomial family, logit link, exchangeable working correlation clustered by donor, and robust covariance [8]:

$$\operatorname{logit}\Pr(C=1) = \beta_0 + \beta_1 \log V + \beta_2 G + \beta_3 (\log V \cdot G), \tag{1}$$

where G is the diagnosis indicator, coded as $G = 0$ for HC (reference group) and $G = 1$ for Parkinson's disease donors, and β_3 tests whether the size dependency changes with the disease state.Here and throughout, log denotes the natural logarithm. With this coding, β_1 is the log-odds slope with respect to $\log V$ in the reference group, and the slope in the PD group is $\beta_1 + \beta_3$. Model (1) was fitted separately within each region and in a pooled form with region as an additional factor; both appear in Table 2.

3.2 Distributional decomposition via quantile regression

To overcome the averaging fallacy of the GEE model [8] and investigate structural shifts in the overlap fraction, we transition from modelling the mathematical expectation to modelling conditional quantiles. The variable F is a poor candidate for a mean-based model, as Fig. 3 shows. It is bounded, discrete rather than continuous, and left-skewed with a large atom on the upper boundary: 42% of captured aggregates are engulfed completely. A mean over F is therefore largely determined by the size of that boundary mass, and the upper quantile is pinned to the boundary and cannot move whatever the covariates do. Quantile regression instead minimises the asymmetric absolute loss $\rho_\tau(u) = u(\tau - \mathbf{1}\{u < 0\})$ and estimates the conditional τ-quantile directly [6]. When fitted on captured aggregates only, the model reads

$$Q_\tau(F \mid V, G) = \beta_0(\tau) + \beta_1(\tau) \log V + \beta_2(\tau) G + \beta_3(\tau) (\log V \cdot G), \tag{2}$$

for $\tau \in \{0.10, 0.50, 0.90\}$, corresponding to superficial contact, typical engulfment, and deep engulfment. Because $\beta_3(\tau)$ is estimated separately at each τ, a shift confined to one regime cannot be diluted by the others.

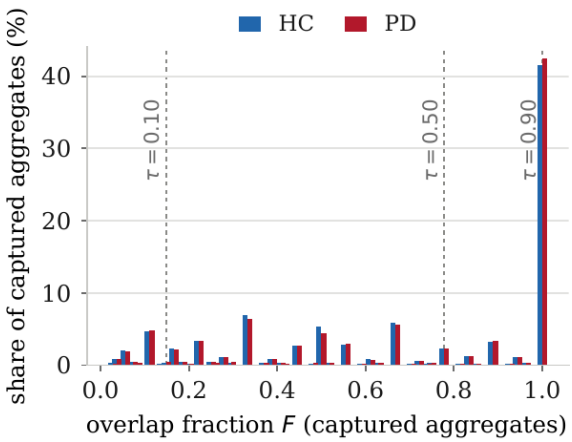


Figure 3: Distribution of engulfment depth F across the 121,611 captured aggregates. F is discrete with 42% of its mass at $F = 1$; dashed lines mark the estimated quantiles (0.148, 0.778, 1.000).

Model (2) is fitted on the pooled aggregates, which are not independent within a donor. We therefore attach confidence intervals to $\beta_3(\tau)$ by resampling donors with replacement and refitting; we use 10,000 bootstrap replicates and report the 2.5th and 97.5th percentiles as the 95% confidence interval. These results are reported in Table 3. This analysis tests whether the size dependence of overlap depth differs between groups at each quantile, but it does not provide a marginal group contrast. We therefore complement it with a separate donor-level analysis in Table 4, in which each donor contributes one value of a chosen summary statistic and the two groups are compared directly. The two analyses address different questions: Table 3 tests for interaction, while Table 4 tests for a marginal shift in the level of overlap depth.

3.3 Information-theoretic assessment

The two regression models each impose a link function and a linear predictor. Shannon's mutual information imposes neither, and is therefore a useful independent measure of how much the size of an aggregate determines what happens to it [3]. For a discretised target Y,

$$I(V;Y) = \sum_{v} \sum_{y} p(v,y) \log_2 \frac{p(v,y)}{p(v)\,p(y)}. \tag{3}$$

We evaluate (3) for two targets. With $Y = C$, over all aggregates, it measures how much size determines whether contact happened; with $Y = F$, over captured aggregates only, how much it determines the depth of that contact. Both are estimated within each donor on quantile-binned contingency tables, using twenty volume bins for $I(V;C)$ and eight bins per axis for $I(V;F)$, with the Miller–Madow correction for the upward bias of the plug-in estimator at finite sample size. Neither bin count follows from a rule, so we repeated the estimation over counts from four to twenty-five per axis; both conclusions – the gap between the two coefficients and the absence of a group difference – hold at every binning. Mutual information cannot exceed the entropy of its target, and those two entropies are nothing like each other here: a binary event occurring in

Table 2: GEE estimates from model (1). Reference group: HC ($G = 0$). The OR in the $\log V$ row is the odds ratio per unit of $\log V$ in the reference group; the corresponding OR in the PD group is obtained by multiplying it by the OR in the $\log V \cdot G$ row. For example, in the putamen the PD slope is $\exp(0.322 - 0.081) = 1.273$. The last row comes from the pooled fit, in which region R enters as an additional factor; it tests whether that change itself differs between the two regions.

Region	**Term**	β **(log-odds)**	**OR (95% CI)**	p**-value**
Putamen	$\log V$	0.322	1.380 (1.280–1.489)	$5.4 \cdot 10^{-17}$
Putamen	$\log V \cdot G$	−0.081	0.922 (0.779–1.091)	0.345
Substantia nigra	$\log V$	0.253	1.287 (1.234–1.343)	$2.0 \cdot 10^{-31}$
Substantia nigra	$\log V \cdot G$	−0.062	0.940 (0.865–1.020)	0.137
Pooled	$\log V \cdot G \cdot R$	+0.028	1.028 (0.822–1.287)	0.809

under one aggregate per hundred leaves under a tenth of a bit of uncertainty to resolve, while a depth spread over the unit interval leaves around two. Raw bits are therefore not comparable between the two targets, and we report the uncertainty coefficient $I(V;Y)/H(Y)$, the fraction of the available uncertainty that aggregate size actually removes. Because overlap is rare, occurrence was analysed using all aggregates with donor-clustered GEE and robust covariance, whereas depth was analysed conditionally among the 121,611 overlapping aggregates. The information-theoretic estimates used the Miller-Madow correction, and sensitivity to binning was assessed over four to twenty-five bins per axis. The precision and generalisability of the findings are limited by the number of donors rather than by the aggregate count.

4 Results

4.1 Size dependence of overlap occurrence and its relation to diagnosis

Aggregate volume is the dominant predictor of whether an aggregate overlaps the microglial mask, in both regions and by a wide margin (Table 2). In the reference group (HC), a 2.72-fold increase in physical volume, one unit on the log scale, raises the odds of overlap by about a third; the corresponding effect in the PD group is smaller but of the same sign, and the difference does not reach conventional significance.

The diagnosis-by-volume interaction is a different matter: it is not significant in either region, and the pooled fit gives no evidence that it differs between them. Taken at face value, the population-averaged analysis provides no evidence that Parkinson's disease modifies the size dependency of overlap occurrence. That is as much as a model of contact *occurrence* can establish, and the rest of this paper asks what happens to the depth of the contacts that do occur.

Table 3: Interaction coefficient $\beta_3(\tau)$ of model (2), fitted on all overlapping aggregates. Confidence intervals come from resampling donors with replacement; at the outer quantiles the coefficients are numerically zero and no interval is informative.

Region	$\beta_3(0.10)$	$\beta_3(0.50)$	**95% CI, donor bootstrap**	$\beta_3(0.90)$
Putamen	$+2.1 \cdot 10^{-6}$	-0.0172	$[-0.051,\ +0.056]$	$-7.1 \cdot 10^{-9}$
Substantia nigra	$-7.7 \cdot 10^{-7}$	$+0.0281$	$[-0.030,\ +0.080]$	$+1.8 \cdot 10^{-8}$

4.2 Median shift in putaminal engulfment depth

Depth behaves differently from occurrence, and how differently depends on where the distribution is read.

Fitting model (2) to the pooled aggregates puts the whole interaction at the median: the coefficients at the two outer quantiles are numerically zero, which is what the shape of F dictates, since the upper quantile sits on the boundary atom and cannot move whatever the covariates do. Once the confidence intervals are widened to account for the nesting of aggregates within donors, however, neither median coefficient excludes zero, and this route does not separate the two regions (Table 3).

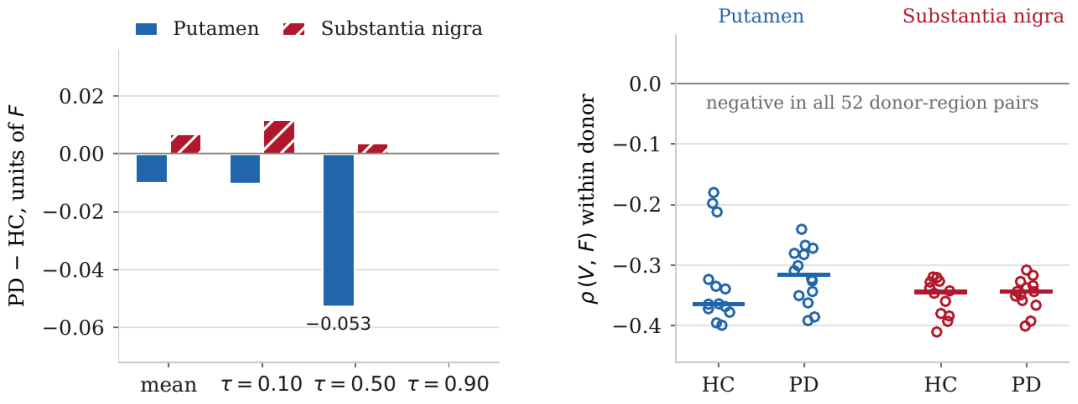
The interaction test in Table 3 asks whether the volume dependence of overlap depth changes with diagnosis. It does not test whether the level of overlap depth itself is lower in PD, and a nonzero interaction is not required for such a shift to exist. We therefore complement the pooled analysis with a donor-level comparison in Table 4, in which each donor contributes one value of a chosen summary statistic of F, and the two groups are compared directly.

Note that $\beta_3(\tau)$ measures how the volume dependency of F differs between groups, not the group-level shift in the τ-quantile. The latter is $\beta_2(\tau) + \beta_3(\tau) \log V$ and can be nonzero even when $\beta_3(\tau) = 0$. We therefore report the group-level shift in the median separately, in the donor-level analysis of Table 4.

The donor-level comparison shows the same asymmetry between summaries, and there the two regions do come apart.

In the putamen, every summary that can move is lower in Parkinson's disease, and the median moves five times as far as the mean of the same variable (Table 4, Fig. 4a). This marginal shift is not captured by the interaction coefficient $\beta_3(0.5)$ of Table 3, which tests whether the size dependency of F changes with diagnosis; the donor-level analysis of Table 4 is the appropriate source for the group-level contrast. That advantage is one of magnitude rather than of discrimination: the rank correlations for mean and median are almost identical, because the two summaries order the donors in nearly the same way. The median is the summary that registers the size of the shift, not the one that separates the groups more sharply. The substantia nigra shows no shift under any summary.

Neither route reaches conventional significance with this cohort, and we report the putaminal shift as an effect size rather than a result. It is at least internally consistent across statistics: the share of aggregates overlapping completely moves in the same direction in the same donors, at a comparable magnitude. The pattern is one of readjusted central tendency rather than abolished phagocytosis, and it is concentrated in the middle of the distribution: the lower decile moves about as far as the mean, and the upper one cannot move at all.



(a) (b)

Figure 4: (a) Donor-level shift in each summary of the overlap depth, in units of F itself. In the putamen, the median moves several times further than the mean, while the upper quantile, pinned to the atom at $F = 1$, cannot move at all. The substantia nigra shows no consistent shift. None of these differences reaches conventional significance; see Table 4. (b) Within-donor rank correlation between aggregate volume and overlap depth, one point per donor and region, with group medians as bars. Larger aggregates overlap proportionally less completely in every donor of both cohorts, and we found no evidence that the disease changes the strength of that relation. Panel (a) is an absolute difference in units of F; panel (b) is a dimensionless correlation coefficient.

Table 4: Donor-level comparison. Each donor contributes one value of the statistic; entries are group means of that per-donor value, the rank-biserial correlation r, and a two-sided Mann–Whitney p (asymptotic, with tie and continuity corrections). The first four rows of each block summarise the depth F; the last two give the uncertainty coefficient, the fraction of the target's entropy that aggregate size removes. Putamen: 13 HC against 14 PD; substantia nigra: 12 against 13. At $\tau = 0.90$ every donor returns exactly 1, so no comparison is possible.

Region	**Statistic**	**HC**	**PD**	**difference**	r	p
Putamen	mean of F	0.677	0.667	−0.010	−0.38	0.094
Putamen	$\tau = 0.10$	0.150	0.140	−0.010	−0.33	0.122
Putamen	$\tau = 0.50$ (median)	0.778	0.725	**−0.053**	−0.37	0.079
Putamen	$\tau = 0.90$	1.000	1.000	0.000	—	—
Putamen	$I(V;C)/H(C)$	0.0036	0.0029	−0.0007	+0.10	0.680
Putamen	$I(V;F)/H(F)$	0.0224	0.0204	−0.0021	−0.15	0.512
Substantia nigra	mean of F	0.690	0.697	+0.007	+0.04	0.892
Substantia nigra	$\tau = 0.10$	0.151	0.163	+0.012	−0.01	1.000
Substantia nigra	$\tau = 0.50$ (median)	0.827	0.831	+0.004	+0.05	0.846
Substantia nigra	$\tau = 0.90$	1.000	1.000	0.000	—	—
Substantia nigra	$I(V;C)/H(C)$	0.0018	0.0017	−0.0001	−0.10	0.683
Substantia nigra	$I(V;F)/H(F)$	0.0303	0.0297	−0.0006	−0.09	0.724

4.3 Information-theoretic asymmetry between occurrence and depth

Asking how much an aggregate's size tells us about its fate splits into two different questions, and once each answer is expressed as a fraction of the uncertainty available to it, they come out roughly an order of magnitude apart (Table 4). Size removes only a few thousandths of the uncertainty about whether microglia will make contact. It removes several times more of the uncertainty about how deeply the aggregate overlaps the microglial mask once contact has happened, about seven times more in the putamen and seventeen times more in the substantia nigra.

The relation behind that asymmetry is the most reproducible feature of the dataset. Within every donor the association between volume and overlap depth is negative and of similar strength (Fig. 4b): larger aggregates overlap proportionally less completely, with a within-donor rank correlation near -0.34 and no exception in any of the 52 donor-by-region combinations. The volume dependency that matters therefore resides in the depth of overlap rather than in its occurrence.

Neither measure separates the two groups at conventional significance. Both uncertainty coefficients and the rank correlation are statistically indistinguishable between PD and HC in both regions, and the data therefore provide no evidence that the disease modifies the dependence of depth on volume.

5 Discussion and conclusion

This study proposes an approach for analysing the interaction between microglia and protein aggregates in three-dimensional immunofluorescence images. Rather than describing the interaction with a single numerical value per field of view, we divided it into two distinct characteristics: whether contact occurred and what its depth was. This approach reveals details that are lost when averaging. To analyse each characteristic, we used three complementary tools: a generalized estimating equation model as a statistical benchmark, quantile regression as a distributional analysis, and mutual information as an independent information-theoretic measure. The main contribution of this work is methodological: a reproducible pipeline for analysing three-dimensional cell-protein interactions that can be applied to other problems.

The two aspects of the interaction behave differently. Aggregate volume is a strong predictor of whether contact occurs: the odds ratio is high. But it explains only a small fraction of the uncertainty about that event, because contact is rare. For depth, the picture is different: the relative information is several-fold larger than for occurrence, although the absolute fraction of uncertainty explained remains modest. In relative terms, aggregate volume is approximately an order of magnitude more informative for depth than for the occurrence of contact. The two characteristics therefore cannot be reduced to a single number, and the choice of summary-mean, median, or quantile-determines whether a shift will be detected at all.

Only one signal is associated with the disease: in the putamen, the median overlap depth is lower in Parkinson's disease, and the median reflects this shift several times more strongly than the mean. With 27 donors, we report this as an effect size rather than a confirmed clinical result; no comparable shift is present in the substantia nigra. In all donors from both groups, larger aggregates overlap less completely, and we found no evidence that the disease alters this relationship. The data are consistent with a shift in the level of overlap rather than a change in its dependence on volume. This illustrates a general problem in spatial biomedical data: for rare events, models based on means and linear relationships mask distributional shifts, whereas quantile regression combined with mutual information provides a more complete picture without

requiring complex neural-network models.

Two groups of limitations apply. The first concerns the data and the design. The data are cross-sectional and observational; all donors are at Braak stage 3/4; and donor-level covariates are not available in the public release, so age, sex, and time since death cannot be accounted for. Depth is defined only where contact occurred, which makes the second stage of the analysis conditional on the first; a formal two-part treatment of that dependence is left to future work. The second group concerns measurement and segmentation. Segmentation relies on 26-connectivity on a strongly anisotropic grid, where axial neighbours are several times farther apart than lateral ones, and this affects where object boundaries fall. Because both masks are three-dimensional, the overlap fraction separates enclosure from surface contact; without a lysosomal marker, however, it cannot resolve whether an enclosed aggregate has been internalised into a phagosome or merely lies within the optical footprint of a process. Finally, the discreteness of F confines the analysis to the interior of its distribution.

A further caveat concerns statistical power. The effective sample size is the number of donors (27), not the number of aggregates: aggregates within a donor are not independent, and aggregate-level p-values would be anticonservative. With 27 donors, the analysis is underpowered to detect small group differences, and the confidence intervals reported in Tables 2 and 3 are consistent with both small positive and small negative effects. Absence of a statistically significant interaction should therefore not be read as evidence that no interaction exists.

Future work will proceed in three directions. First, two-part models that jointly describe the fact of contact and its depth. Second, external validation on independent datasets and larger cohorts, together with additional markers that allow internalisation to be distinguished from surface contact. Third, adaptation of the proposed pipeline to other tasks involving pairwise interactions between objects in three-dimensional images.

References

[1] J.S. Beckwith, J.C. Breiter, and S.F. Lee. α-synuclein oligomer-driven cell susceptibility in early Parkinson's disease. BioStudies, accession S-BIAD2216, 2025. [Accessed: 01.07.2026].

[2] J.S. Beckwith, J.C. Breiter, and S.F. Lee. α-synuclein oligomer-driven cell susceptibility in early Parkinson's disease. bioRxiv, 2025.

[3] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley-Interscience, Hoboken, NJ, USA, 2nd edition, 2006.

[4] B.M. Davis et al. Characterizing microglia activation: a spatial statistics approach to maximize information extraction. *Sci. Rep.*, 7(1):1576, 2017.

[5] I. Frigerio et al. Synaptic enrichment of pSer129 alpha-synuclein correlates with dopaminergic denervation in early-stage Parkinson's disease. *Nat. Commun.*, 16(1):6630, 2025.

[6] R. Koenker and G. Bassett. Regression quantiles. *Econometrica*, 46(1):33–50, 1978.

[7] I. Kurnaeva and A. Maslovskaya. Spatial statistics and explainable deep learning for 3D cell–protein interaction profiling. *Big Data Cogn. Comput.*, 10(10):333, 2026.

[8] K.-Y. Liang and S.L. Zeger. Longitudinal data analysis using generalized linear models. *Biometrika*, 73(1):13–22, 1986.

[9] C.-S. Pei et al. α-synuclein disrupts microglial autophagy through STAT1-dependent suppression of Ulk1 transcription. *J. Neuroinflammation*, 21(1):275, 2024.

[10] H. Scheiblich et al. Microglia rescue neurons from aggregate-induced neuronal dysfunction and death through tunneling nanotubes. *Neuron*, 112(18):3106–3125.e8, 2024.

[11] N. Stefanova. Microglia in Parkinson's disease. *J. Parkinson's Dis.*, 12(s1):S105–S112, 2022.