

StormGraph: An automated graph-based algorithm for quantitative clustering analysis of single-molecule localization microscopy data.

Joshua M. Scurll^{1,2}, Libin Abraham^{2,3}, Da Wei Zheng¹, Reza Tafteh⁴, Keng C. Chou⁴,
Michael R. Gold^{2,3}, Daniel Coombs¹

¹ Department of Mathematics and Institute of Applied Mathematics, 1984 Mathematics Road, University of British Columbia, Vancouver, British Columbia V6T 1Z2 Canada,

² Life Sciences Institute I³ and Cell Research Groups, University of British Columbia, 2350 Health Sciences Mall, Vancouver, British Columbia V6T 1Z3 Canada,

³ Department of Microbiology and Immunology, University of British Columbia, 2350 Health Sciences Mall, Vancouver, British Columbia V6T 1Z3 Canada,

⁴ Department of Chemistry, University of British Columbia, Vancouver, British Columbia V6T 1Z1 Canada

Corresponding authors

Joshua M. Scurll and Daniel Coombs

Department of Mathematics and Institute of Applied Mathematics

1984 Mathematics Road, University of British Columbia

Vancouver, British Columbia V6T 1Z2 Canada

Phone: +1-604-446-0450 (JMS); +1-604-822-2859 (DC)

E-mail: jscurll.ubc@gmail.com; coombs@math.ubc.ca

Keywords: single-molecule localization microscopy, dSTORM, heterogeneous clustering, graph community detection, hierarchical clustering, B cell receptor, co-localization.

1 **Abstract**

2 Clustering of proteins is crucial for many cellular processes and can be imaged at nanoscale resolution using
3 single-molecule localization microscopy (SMLM). Existing cluster analysis methods for SMLM data suffer
4 from major limitations, such as unsuitability for heterogeneous datasets, failure to account for uncertainties
5 in localization data, excessive computation time, or inability to analyze three-dimensional data. To address
6 these shortcomings, we developed StormGraph, an algorithm using graph theory and community detection
7 to identify and quantify clusters in heterogeneous 2D and 3D SMLM datasets. StormGraph accounts for
8 localization uncertainties and, by determining thresholds adaptively, it allows many heterogeneous samples
9 to be analyzed using identical parameters. Consequently, StormGraph improves the potential accuracy,
10 objectivity, and throughput of cluster analysis. Furthermore, StormGraph generates a hierarchical clustering,
11 and it quantifies cluster colocalization for two-color SMLM data. We use simulated data to show that
12 StormGraph is superior to existing algorithms. Finally, we demonstrate its application to two-dimensional
13 B-cell antigen receptor clustering and three-dimensional intracellular LAMP-1 clustering.

14 Introduction

15 Single-molecule localization microscopy (SMLM) techniques, such as direct stochastic optical reconstruc-
16 tion microscopy (dSTORM) (1; 2) and photoactivated localization microscopy (PALM) (3), overcome the
17 diffraction limit of conventional microscopy by acquiring many sequential images, each containing very few
18 fluorescing labels. Individual labels can then be computationally super-resolved and precisely localized to
19 generate a list of localization coordinates, often with estimated positional uncertainties (4; 5; 6). This is
20 possible in both two and three dimensions (7; 8; 9; 10).

21 SMLM is commonly used to investigate nanoscale clustering of cell-membrane and intracellular proteins
22 (11; 12; 13; 14; 15; 16; 17; 18; 19; 20), which usually exhibits both cell-to-cell and within-cell heterogene-
23 ity. Notwithstanding, clustering is frequently analyzed using spatial summary statistics that fail to capture
24 the heterogeneity of clusters within a sample, such as Ripley’s functions (21; 22). Instead, clusters can be
25 individually quantified by using a clustering algorithm to assign localizations to specific clusters. However,
26 using existing clustering algorithms, it is difficult to accurately and objectively analyze multiple heteroge-
27 neous samples. Subjective bias can result from algorithm parameter selection, or from selection of a small
28 number of “representative samples” to analyze using slow or cumbersome algorithms. Failure to account for
29 localization uncertainties can also make conclusions unreliable.

30 The most widely used clustering algorithms in SMLM literature, including Density-Based Spatial Cluster-
31 ing of Applications with Noise (DBSCAN) (23), identify clusters based on a user-specified minimum number
32 of points within a user-specified radius. However, these parameters are difficult to select and, for hetero-
33 geneous samples, should be sample-specific. Recently, algorithms based on Voronoi diagrams, for example
34 ClusterViSu, have been developed for 2D (24; 25) and 3D SMLM data (26). Importantly, however, none of
35 the above clustering algorithms account for localization uncertainties. A pixelated variant of DBSCAN (27)
36 partially addresses this deficiency but is limited to 2D datasets. Simulation-aided DBSCAN (28) offers a more
37 complete strategy, but it remains limited by user-determined DBSCAN parameter selection. A Bayesian,
38 model-based cluster identification method (29) uses localization uncertainties and has been extended to 3D
39 (30), but it assumes that clusters are circular or spherical and it is prohibitively slow.

40 To address all of the limitations of existing cluster analysis methods described above, we developed
41 StormGraph, a comprehensive graph-based clustering algorithm inspired by PhenoGraph (31), which was
42 developed for single-cell cytometry data. StormGraph converts SMLM data into a graph using localization
43 coordinates and their uncertainties to specify nodes and weighted edges. It then utilizes graph theory and
44 community detection algorithms (32) to assign nodes to specific clusters. Crucially, StormGraph determines

45 key thresholds from the data for each region of interest (ROI) adaptively, using at most three user-specified
46 parameters that can remain fixed across experiments, enabling unbiased comparison of results. Unlike the
47 Bayesian method, StormGraph makes no assumptions about the shapes of clusters and it is at least an
48 order of magnitude faster. Moreover, StormGraph has both 2D and 3D implementations and it can quantify
49 cluster overlap for two-color SMLM data.

50 Furthermore, SMLM data often exhibits hierarchical clustering at multiple spatial scales. For example,
51 SMLM revealed multiscale organization of RNA polymerase in *Escherichia coli* (33). To enable multiscale
52 cluster analysis, StormGraph generates a hierarchical clustering. This is an advantage over existing meth-
53 ods, for which users must repeatedly change parameters to perform multiscale analysis. Notwithstanding,
54 StormGraph also generates an appropriate single-level clustering.

55 Here, we describe StormGraph and its capabilities, and we use simulated data to compare its accuracy
56 to that of DBSCAN and ClusterViSu. We then apply StormGraph to characterize nanoscale clustering of
57 B-cell antigen receptors (BCRs) from heterogeneous 2D SMLM data. We also demonstrate StormGraph's
58 ability to quantify 3D clusters of the lysosomal protein LAMP-1 and to quantify cluster overlap for two-
59 color SMLM data. Because of cell-to-cell heterogeneity as well as non-uniformity of cellular compartments
60 such as the plasma membrane, it is essential to compile SMLM data from ROIs from multiple cells for each
61 experimental condition. To make this practical, we developed software to crop ROIs and batch process
62 StormGraph analysis for multiple samples. We will make the software and source code freely available upon
63 publication of this manuscript.

64 Results

65 The StormGraph algorithm

66 To identify clusters in SMLM data, dense localization neighborhoods must be identified. To this end,
67 StormGraph first determines an ROI-specific length scale r_0 from the data using either of two methods (see
68 Methods and Figure S1). One method, which seeks a balance between inter-localization and inter-cluster
69 distances, is fully automatic but heuristic. The other is semi-automatic, using k -nearest neighbor (kNN)
70 distances with user-defined k . The automatic method reduces user input but is designed primarily for data
71 with very few dispersed localizations between clusters, whereas the kNN method is universally applicable.

72 Next, using the localizations as nodes (Figure 1a), StormGraph essentially constructs a weighted r_0 -
73 neighborhood graph (Figure 1b) as follows. Define

$$s_{ij} = \begin{cases} 1 - r_{ij}/r_0 & \text{if } r_{ij} \leq r_0, \\ 0 & \text{if } r_{ij} > r_0, \end{cases}$$

74 where r_{ij} is the Euclidean distance between nodes i and j . If localization coordinate uncertainties are
75 unknown, StormGraph assigns to each node pair $\{i, j\}$ an edge of weight $W_{ij} = s_{ij}$. Otherwise, StormGraph
76 uses the uncertainties to estimate $\langle s_{ij} \rangle$, the expectation of s_{ij} , from Monte Carlo simulations (Methods) and
77 assigns $W_{ij} = \langle s_{ij} \rangle$.

78 At this stage, unclustered localizations are identified and removed by applying a threshold to the weighted
79 node degree,

$$\text{deg}(i) = \sum_{j \neq i} W_{ij},$$

80 a proxy for local density. In principle, nodes are classified as unclustered and removed if their degree
81 falls below a data-dependent threshold (Figure 1c). StormGraph automatically determines this threshold
82 from random point clouds using the user-defined significance parameter α (Methods). The value of α is
83 the maximum probability of a completely randomly distributed localization being classified as clustered, a
84 subjective choice for the user but typically 0.05. Figure S2 shows the effects of varying the parameters α
85 and k . Removed localizations are reported as an algorithm output.

86 The graph is then regenerated using a new r_0 value determined automatically from only the retained
87 nodes. If localization coordinate uncertainties are available, edges are subsequently pruned from the graph

88 to ensure that any pair of edges have at least an estimated 50% probability of co-occurring in the r_0 -
89 neighborhood graph for the unknown true localization positions (Methods). StormGraph then finds a hi-
90 erarchy of node clusters (Figure 1e) using the multi-level Infomap algorithm (34), followed by additional
91 cluster merging when warranted (Methods).

92 To obtain a single-level clustering from the hierarchy, we developed a novel, fast method motivated
93 by the idea of consensus clustering (35; 36). Briefly, clusters are recursively divided into their coarsest
94 constituent subclusters if they closely match the connected components of an alternative neighborhood
95 graph (Methods). Optionally, the user may specify the minimum number of points per reported cluster
96 (MinCluSize). As output, StormGraph provides the single-level and hierarchical cluster assignments of every
97 localization. Combined with localization coordinates, this provides the information necessary to quantify
98 individual cluster properties, such as area (Figure 1d). Our software automatically quantifies the single-level
99 and coarsest-level clusterings.

100 Lastly, a common caveat of SMLM is multiple counting of single molecules, often causing single molecules
101 to spuriously appear as clusters. This can be due to multiple labeling of single molecules or individual
102 fluorophores yielding multiple localizations. Therefore, StormGraph includes optional functionality that
103 uses localization uncertainties to reclassify as unclustered localizations any putative clusters that cannot be
104 confidently distinguished from multiply counted single molecules (Methods).

105 **Validation using simulated data and comparison to other algorithms**

106 To compare StormGraph with DBSCAN and ClusterViSu, we simulated a wide variety of $2\ \mu\text{m} \times 2\ \mu\text{m}$ ROIs
107 containing isolated and heterogeneously aggregated circular nanoclusters (e.g. Figure 2a; Methods). Outside
108 the clusters we added randomly distributed molecules. Individual simulated molecules were allowed to yield
109 multiple localizations, each with a positional uncertainty sampled from a real dSTORM experiment. We
110 tested both the automatic and kNN ($k = 10, 15$ or 20) methods for determining r_0 while maintaining $\alpha = 0.05$.
111 We set a minimum cluster size of 5 localizations in both StormGraph and ClusterViSu. For DBSCAN, we
112 tested 16 different parameter choices based on the underlying parameters used for data simulation, although
113 such knowledge is generally unavailable for real data. To assess cluster assignments from each algorithm, we
114 used normalized mutual information (NMI) (37) and mean F-measure (38). Higher values indicate superior
115 performance.

116 StormGraph consistently outperformed ClusterViSu regardless of whether localization uncertainties were
117 used and regardless of the method used to determine r_0 (Figures 2b–c and S3). DBSCAN’s performance

118 was very sensitive to the choice of parameters and no single choice was suitable for all of the data (Figure
119 S4), demonstrating its unsuitability for batch processing analysis of heterogeneous samples. StormGraph was
120 generally superior to DBSCAN regardless of parameter choices among those tested. Moreover, StormGraph’s
121 performance was not very sensitive to varying k from 10 to 20, particularly when localization uncertainties
122 were used (Figures 2b and S3).

123 For simulated data with nanoclusters of 50 nm radius, we were able to manually identify a level of
124 clustering from StormGraph’s cluster hierarchy that accurately recovered the ground-truth nanoclusters that
125 composed larger aggregations (Figure S5). This demonstrates that StormGraph is able to identify meaningful
126 clusters at multiple scales. Additionally, we performed tests using simulated data without multiple counting
127 of single molecules and found that StormGraph still outperformed ClusterViSu and DBSCAN (Figure S6).

128 Finally, we compared StormGraph to Ripley’s H-function (22) using simulated circular clusters. Ripley’s
129 H-function was biased towards the clusters containing the most points, as mathematically expected, and
130 it did not provide an accurate measure of cluster radius (Figure S7). Conversely, StormGraph provided
131 excellent estimates (Figure S7).

132 **StormGraph quantifies heterogeneous B-cell receptor clustering from dSTORM** 133 **data in 2D**

134 To test StormGraph on real SMLM data, we used it to analyze the clustering of immunoglobulin M (IgM)-
135 isotype B-cell antigen receptors (BCRs) on the cell membranes of B lymphocytes. IgM-BCRs are thought
136 to exist in nanoclusters on resting B cells (12; 39) and form larger “microclusters” during B-cell activation
137 induced by antigen engagement (40; 41). The exact changes in IgM-BCR arrangement are controversial,
138 however (12; 39; 41).

139 Using dSTORM, we imaged fluorescently labeled IgM-BCRs on *ex vivo* murine splenic B cells that
140 were either resting or treated with bivalent antibodies against the BCR’s Ig κ light chain, used as antigen
141 surrogates. Localization coordinates and their associated uncertainties were computationally determined
142 from the fluorescence data. We then used StormGraph ($\alpha = 0.05$, MinCluSize = 5 localizations) to batch
143 process the analysis of IgM-BCR clustering in, respectively, 28 and 24 rectangular ROIs that were $> 1 \mu\text{m}^2$,
144 from separate cells, and entirely within cell boundaries (Figure 3a). We applied StormGraph’s functionality
145 to remove from the results any clusters of localizations that could not be confidently distinguished from
146 overcounted single molecules (Methods).

147 Using StormGraph’s single-level clustering results, we compared cluster areas between conditions. Using

148 $k = 15$, we found that the mean area of IgM-BCR clusters was significantly larger on anti-Ig κ -treated cells
149 than on resting cells (Figure 3b(i), $p < 10^{-5}$), as expected. This difference was mainly due to an increase
150 in the size and frequency of clusters $> 6000 \text{ nm}^2$, rather than a uniform increase in cluster areas (Figure
151 3b(ii)). In fact, the majority of clusters present on anti-Ig κ -treated cells were small multimers that were
152 comparable to, or even smaller than, the IgM-BCR clusters on untreated cells. The automatic (no k value)
153 implementation of StormGraph yielded consistent conclusions (Figure S8).

154 Large BCR clusters have been associated with chronic BCR signaling in a subset of activated B-cell like
155 (ABC) diffuse large B-cell lymphomas (DLBCLs). Diffraction-limited microscopy revealed large IgM-BCR
156 microclusters in the absence of any stimulus on the ABC DLBCL cell lines HBL-1 and TMD8 but not on
157 the Burkitt's lymphoma cell line BJAB (42). To further investigate this observation, we used dSTORM
158 to image IgM-BCRs on HBL-1, TMD8 and BJAB cells. We then batch processed StormGraph analysis
159 of, respectively, 39, 33 and 81 ROIs $> 1 \mu\text{m}^2$ (Figure 3c), which contained between 5×10^3 and 3×10^5
160 localizations.

161 Using $k = 15$, StormGraph revealed that the mean areas of IgM-BCR clusters on HBL-1 and TMD8
162 cells were significantly larger than on BJAB cells ($p < 10^{-4}$ and $p < 10^{-14}$ respectively; Figure 3d(i)).
163 Interestingly, the relative distributions of IgM-BCR cluster areas on BJAB and HBL-1 cells resembled those
164 for resting and anti-Ig κ -treated B cells, respectively. Both BJAB and HBL-1 cells had many small IgM-BCR
165 clusters, but HBL-1 displayed a notable increase in the size and frequency of large clusters exceeding 10^4 nm^2 .
166 In contrast, TMD8 cells displayed an overall increase in cluster areas compared to BJAB (Figure 3d(ii)).
167 Again, the automatic implementation of StormGraph yielded similar results (Figure S8). Our observations
168 reveal that IgM-BCR organization can differ substantially between DLBCL cell lines of the same ABC
169 subtype. Furthermore, assuming that BCR signaling is mostly due to large clusters, the presence of many
170 small IgM-BCR clusters suggests that only a fraction of IgM-BCRs participate in chronic BCR signaling in
171 HBL-1 cells, unless the clusters are highly dynamic.

172 Finally, we repeated all StormGraph analyses ignoring localization uncertainties. Although the magnitude
173 and statistical significance of our results were altered, the overall conclusions were unchanged (Figure S8).
174 Importantly, this demonstrates that StormGraph can still detect differences in clustering when localization
175 uncertainties are unavailable.

176 **StormGraph is robust to changes in global density of SMLM localizations**

177 Because the average density of SMLM localizations can vary between ROIs, a clustering algorithm must not
178 depend on the average localization density if batch processing is to be implemented and clustering results
179 are to be compared across samples with different localization densities. We verified that StormGraph is
180 robust in this regard by repeating StormGraph analysis after randomly removing 0%, 25%, 50% or 75%
181 of the localizations from a dSTORM ROI containing heterogeneous clusters (Figure 4a). Although small,
182 low-density clusters were eventually lost, the identification and area quantification of large, unambiguous
183 clusters was robust, and the overall distribution of cluster areas was not significantly impacted ($p > 0.05$;
184 Figure 4b).

185 We also tested StormGraph's sensitivity to random noise by artificially adding random localizations
186 (with uncertainties) to the same ROI (Figure S9). StormGraph's ability to detect all but small, low-density
187 clusters was again robust, and its overall sensitivity to random noise was minimized by including localization
188 uncertainties and using the kNN method to determine r_0 . This implementation with $k = 15$ resulted in
189 no statistically significant ($p < 0.05$) change in the distribution of cluster areas until the ratio of true
190 localizations to artificial localizations was < 2 .

191 **Two-color analysis of cluster overlap**

192 To quantify colocalization of differently colored (e.g. red and blue) clusters in two-color SMLM data, our
193 software quantifies the total area of overlap divided by each of the following: (1) total red cluster area; (2)
194 total blue cluster area; and (3) total area covered by clusters of either color, yielding the Jaccard index
195 (43) (Figure 5d). Our software also reports analogous quantities using numbers of localizations instead
196 of areas (not shown). To estimate the maximal experimentally observable colocalization, colocalization
197 analysis should first be applied to the same molecular species labeled with two different probes. This rarely
198 yields 100% colocalization for several reasons, including differing affinities of antibody-fluorophore conjugates,
199 differing photophysical properties of fluorophores, and the inability of two probes to occupy the same binding
200 site.

201 To demonstrate cluster overlap analysis by StormGraph, we performed such a positive control experiment
202 by simultaneously labeling cell-surface IgG-BCRs on murine A20 B cells with anti-IgG antibodies conjugated
203 to either Alexa Fluor 647 (AF647) or Cy3B fluorophores. These antibodies were bivalent, thus inducing
204 formation of large clusters prior to cell fixation. Both color channels were imaged using dSTORM and
205 aligned using custom MATLAB code to correct for chromatic aberrations. We then analyzed multiple ROIs

206 using StormGraph (Figure 5). On average, we found 79% overlap of the IgG-AF647 clusters with the IgG-
207 Cy3B clusters and 66% overlap of the IgG-Cy3B clusters with the IgG-AF647 clusters (Figure 5e). This
208 difference is likely due to differing qualities of the AF647- and Cy3B-conjugated antibodies. The Jaccard
209 index cannot exceed either one-sided overlap score, and we obtained an average Jaccard index of 0.5. In a
210 similar experiment but for tubulin, ClusterViSu obtained $\sim 40\%$ overlap of each probe with the other. This
211 shows that StormGraph performs well as part of a pipeline for analyzing cluster colocalization by SMLM.

212 Clustering in three dimensions

213 To extend StormGraph to 3D, we considered some particular features of 3D SMLM. StormGraph implicitly
214 assumes that all dimensions should be weighted equally during graph construction, but most 3D SMLM
215 techniques achieve lower axial resolution than lateral resolution. Therefore, StormGraph pre-processes the
216 data, for cluster identification but not subsequent quantification, by rescaling the axial (z) dimension so
217 that average axial and lateral positional uncertainties, when known, become equal. Furthermore, 3D SMLM
218 localizations are often concentrated around a focal plane, causing their axial distribution to be nonuniform.
219 Accordingly, StormGraph uses the parameter α to obtain a z -dependent node-degree threshold from random
220 point clouds with normally distributed z -coordinates (Methods). This provides a clear advantage over
221 DBSCAN, which is unable to adapt to axial variation in localization density. For situations with localizations
222 distributed uniformly in z , StormGraph still retains the option to use a constant threshold instead.

223 We compared the performances of StormGraph and DBSCAN in 3D using simulated 3D data (Meth-
224 ods). As in 2D, we found that, overall, StormGraph was superior to DBSCAN regardless of parameter
225 choices (Figure S10). We also performed 2D clustering of the xy -projections of our simulated 3D data. To
226 achieve results with DBSCAN that were comparable between 2D and 3D, it was inevitably necessary to
227 use different parameter values. In contrast, StormGraph required no changes to parameters, thus making it
228 easy to switch between 2D and 3D analyses. Nonetheless, including the z -component of 3D data improved
229 clustering accuracy (Figure S10). This is because localizations and clusters that are separated only in z are
230 indistinguishable in the xy -projection.

231 To illustrate StormGraph's application to 3D SMLM data, we used dSTORM to image intracellular
232 lysosomal-associated membrane protein 1 (LAMP-1). We simultaneously immunostained LAMP-1 in B16
233 melanoma cells with two different labels, AF647 and Cy3B, and applied StormGraph ($k = 15$, $\alpha = 0.1$,
234 $\text{MinCluSize} = 5$ localizations) to a 3D ROI with axial variation in localization density and known localization
235 uncertainties (Figure 6a–b). StormGraph detected 363 LAMP-1 AF647 clusters and 129 LAMP-1 Cy3B

236 clusters (Figure 6c-d). The AF647 clusters had volumes ranging from $1.5 \times 10^3 \text{ nm}^3$ to $7.1 \times 10^7 \text{ nm}^3$ with a
237 median of $3.5 \times 10^5 \text{ nm}^3$, and Cy3B clusters had volumes ranging from $3.1 \times 10^3 \text{ nm}^3$ to $3.7 \times 10^7 \text{ nm}^3$ with
238 a median of $9.0 \times 10^5 \text{ nm}^3$ (Figure 6e). The discrepancy in cluster volumes was likely caused by variance in
239 labeling or probe detection. Indeed, we detected almost four times as many AF647 localizations as Cy3B
240 localizations (9.0×10^4 versus 2.5×10^4). Hence, when performing one-color SMLM, choosing the optimal
241 fluorescent label can improve cluster detection and quantification.

242 Additionally, we computed volumetric overlap between AF647 clusters and Cy3B clusters (Figure 6f-g).
243 To our knowledge, our software is the first to offer this functionality for two-color, 3D SMLM data. We
244 found that 31% of the total AF647 cluster volume overlapped with Cy3B clusters, whereas 50% of the total
245 Cy3B cluster volume overlapped with AF647 clusters. The comparatively low overlap of AF647 with Cy3B
246 was likely due to weaker labeling or detection with Cy3B than AF647. The Jaccard index was 0.24. In sum,
247 our results for LAMP-1 clearly show that StormGraph can identify and quantify clusters of localizations in
248 3D SMLM ROIs and, furthermore, that it can detect overlap between 3D clusters in two-color data despite
249 experimental limitations.

250 Discussion

251 By converting 2D or 3D SMLM localization data into a neighborhood graph, StormGraph leverages concepts
252 from graph theory, especially community detection, to assign localizations to individual clusters that can be
253 quantified. It enables analysis of clustering at multiple scales within datasets by generating a hierarchical
254 clustering, but it also provides a single-level clustering to simplify interpretation of results. The StormGraph
255 can be run in MATLAB using either a script or a simple graphical user interface. The software automatically
256 quantifies clusters and it includes MATLAB functions for data visualization in 2D or 3D. *The software and*
257 *a user manual are available from the authors on request.*

258 StormGraph has three optional, user-definable parameters. If by visual inspection the vast majority
259 (> approximately 90%) of localizations are organized into clear, well-separated clusters, then all three user-
260 definable parameters could be discarded from StormGraph. Otherwise, we provide guidelines for their
261 selection in the Methods. Each parameter specifies either a number of localizations (k and MinCluSize) or
262 a probability (α) and can therefore be set without any knowledge of the scale or density of the localization
263 data. StormGraph then adaptively determines scale- and density-dependent thresholds from the data using
264 k , or automatically without k , and using α respectively. This allows disparate datasets to be analyzed using
265 identical parameters, which increases both the objectivity and, by means of batch processing, the potential
266 throughput of cluster analysis. This is in sharp contrast to DBSCAN, where two user-defined parameters
267 explicitly define a threshold density.

268 Another distinguishing feature of StormGraph is its full utilization of individual localizations' positional
269 uncertainties, when available. This is important because these uncertainties locally determine the minimum
270 scale at which clusters can be reliably resolved. Only StormGraph and Bayesian methods (29; 30) have this
271 feature in addition to being suitable for batch processing of heterogeneous ROIs. However, StormGraph
272 has significant advantages over the Bayesian methods. First and foremost, it has superior computational
273 efficiency. On a standard desktop computer, StormGraph analyzed a 2D ROI containing 5,349 localizations
274 in less than 40 seconds and a 2D ROI containing 26,941 localizations in less than 3 minutes. The largest
275 ROIs that we analyzed using StormGraph contained more than 10^5 localizations. On the other hand, the
276 2D Bayesian method took ~ 50 minutes to analyze the 5,349-localization ROI and it failed to analyze the
277 26,941-localization ROI due to memory limitations. The second notable advantage of StormGraph is that,
278 unlike the Bayesian methods, it makes no assumptions about the shapes of clusters. These advantages make
279 StormGraph widely applicable and make it the standout choice for cluster analysis of SMLM data.

280 Using simulated data, we demonstrated that StormGraph is superior to both DBSCAN and the most

281 recent algorithm based on Voronoi diagrams, ClusterViSu, at assigning localizations to clusters. We also
282 applied StormGraph to actual dSTORM data. This revealed the presence of many small BCR clusters, in
283 addition to the expected large clusters, on the cell membranes of activated B cells. Our discoveries here
284 highlight the importance of being able to objectively analyze nanoscale protein clustering in heterogeneous
285 samples. By providing improved, high-throughput quantitative characterization of nanoscale receptor clus-
286 tering, StormGraph should enable new insights into the relationship between receptor clustering and receptor
287 signaling.

288 It should be noted that measurement errors generally cause clusters to appear slightly larger than the
289 true underlying molecular clusters, and they can also cause over- or under-estimation of cluster overlap for
290 two-color data. StormGraph does not correct for this during cluster quantification. However, for data with
291 approximately Gaussian clusters, mathematical correction methods (44) could be applied to StormGraph's
292 clusters prior to quantification. Nonetheless, we believe that StormGraph will advance cluster analysis in the
293 SMLM field thanks to its generality, its ability to utilize information about localization uncertainties, and
294 its potential to increase the throughput of single-molecule localization cluster analysis via batch processing
295 of heterogeneous datasets.

296 Methods

297 Calculation of the length scale r_0

298 (1) The fully automatic, heuristic method

299 To automatically determine a length scale r_0 without user input, we implement a variation of the elbow
300 method heuristic. For values of ε ranging from 0 to a sufficiently large value based on the optimal affinity
301 scale stated by Arias-Castro (45), we construct the ε -neighborhood graph for the data. We then plot the
302 number of connected components (including singletons) against ε . This must be monotonically decreasing
303 and typically bears resemblance to a decaying exponential or logistic function. As ε increases, an “elbow”
304 region occurs as rapid linking of nodes within clusters at small values of ε transitions to slower linking of
305 distinct clusters and dispersed nodes at larger values of ε . Eventually all nodes would belong to a single
306 connected component.

307 Sometimes, a natural number of clusters will be evident as a horizontal (i.e. constant) plateau occurring
308 at > 1 connected component in this plot. In such cases, we find the plateau corresponding to the largest fold
309 increase in the area or volume of the ε -neighborhood. Let ε_1 be the value of ε at the start of this plateau,
310 and let $\varepsilon_2 = 2^{1/d}\varepsilon_1$, where d is the dimensionality of the data, be chosen such that the ε_2 -neighborhood
311 is twice the area or volume of the ε_1 -neighborhood. If the ε_1 - and ε_2 -neighborhood graphs have the same
312 number of connected components, then we set $r_0 = \varepsilon_2$ (Figure S1).

313 Otherwise, we fit a curve $f(\varepsilon)$ to the number of connected components versus ε (Figure S1). We choose
314 $f(\varepsilon)$ to be the sum of a constant b and either one or two generalized logistic functions of the form

$$L(\varepsilon) = \frac{a}{(1 + \exp(s(\varepsilon - \varepsilon_0)))^{1/\nu}},$$

315 where $b \geq 0$, $a \geq 0$, $s \geq 0$, $\nu > 0$, and ε_0 are coefficients to be fit. To avoid overfitting, we only include
316 the second logistic function if it yields a substantial improvement in the goodness of fit and we restrict its
317 allowable values of ν . The elbow of this curve is not mathematically well defined, but intuitively it is related
318 to the concavity: the curve achieves maximum (positive) concavity as it approaches the elbow region, and
319 then its concavity decreases as it traverses the elbow region. StormGraph chooses the length scale r_0 to be
320 towards the end of the elbow region as follows. Let ε_{\max} be the value of ε at which $f''(\varepsilon)$, the concavity of
321 $f(\varepsilon)$, is maximized. StormGraph sets r_0 to be the value of $\varepsilon > \varepsilon_{\max}$ where $f''(\varepsilon)$ first falls below 2% of its
322 maximum value (Figure S1).

323 When localization uncertainties are available in the data, they are initially excluded when utilizing the
324 elbow method to set the initial length scale r_0 , which is used for classifying localizations as either clustered
325 or unclustered. The uncertainties are subsequently taken into account during the final use of the elbow
326 method, which sets the value of r_0 that is used for construction of the final graph following elimination of
327 unclustered localizations. Specifically, the graph in which we count the number of connected components for
328 a given ε is constructed from Monte Carlo simulated realizations of the data with two nodes connected to
329 each other by an edge if and only if they are within a distance ε of each other in at least 75% of the Monte
330 Carlo simulations. Note that edge weights are not relevant here because they do not affect the number of
331 connected components.

332 **(2) The kNN method**

333 To determine the length scale r_0 for a selected ROI using a k -nearest neighbors (kNN) approach, StormGraph
334 first finds the distance of every point in the ROI to its k^{th} nearest neighbor. If localization uncertainties are
335 available in the data, this is performed for 100 Monte Carlo simulated realizations of the data, and the 95%
336 confidence level for the k^{th} nearest neighbor distance is obtained for every localization. The distribution
337 of k^{th} nearest neighbor distances is also obtained for Monte Carlo simulations of random data with the
338 same global average point density as the ROI. A histogram of k^{th} nearest neighbor distances should initially
339 increase more rapidly for clustered data than for random data, but the histograms for clustered and random
340 data will eventually intersect each other (Figure S1). Inspired by the automated version of ClusterViSu
341 (25), StormGraph defines r_0 as the distance at which these histograms of k^{th} nearest neighbor distances first
342 intersect. Points closer than r_0 to their k^{th} nearest neighbor are more likely to exist in clustered data, while
343 points farther than r_0 from their k^{th} nearest neighbor are more likely to exist in random data. Moreover,
344 points in clusters will tend to have more than k neighbors within a distance r_0 , while randomly distributed
345 points will tend to have fewer than k neighbors within a distance r_0 . However, if this first histogram
346 intersection occurs after the median of the random data's histogram, this indicates that, on average, the real
347 data is actually more dispersed than the random data, and in this case StormGraph defines r_0 simply as the
348 median of the random data's k^{th} nearest neighbor distances.

349 **Simulating multiple data realizations and calculation of graph edge weights**

350 StormGraph uses Monte Carlo simulations to simulate multiple realizations of the data by resampling each
351 localization's coordinates. The new x , y and, if applicable, z coordinates for a particular localization are

352 drawn independently from normal distributions centered at the original observed localization position. The
353 standard deviations are equal to the corresponding uncertainties recorded in the data. StormGraph then
354 determines the graph edge weights $W_{ij} = \langle s_{ij} \rangle$ from the Monte Carlo simulations by calculating $\langle s_{ij} \rangle$ to be
355 the mean of the simulated values of s_{ij} for each specific node pair $\{i, j\}$.

356 **Thresholding of node degrees to eliminate unclustered nodes**

357 Setting $\alpha = 1$ skips the thresholding step altogether, allowing all nodes to be considered for clustering.
358 Otherwise, to set the node-degree threshold, StormGraph first constructs r_0 -neighborhood graphs with edge
359 weights s_{ij} for simulated random point clouds with the same global average point density as the SMLM data.
360 For 2D data (and for 3D data with uniform axial acquisition), the random points are uniformly distributed
361 in x and y (and z). Then StormGraph sets the degree threshold as the $((1 - \alpha) \times 100)^{\text{th}}$ percentile of
362 the aggregated degree distribution of the random simulations. For 3D data with localizations concentrated
363 around a focal plane, StormGraph simulates random data with z -coordinates that are distributed normally
364 with the same interquartile range as the data. StormGraph then obtains a z -dependent node-degree threshold
365 by fitting a Gaussian curve to node degree versus z for the simulated random points and finding the $(1 -$
366 $\alpha) \times 100\%$ confidence upper bound curve. Thus, for both 2D and 3D data, an expected $\alpha \times 100\%$ of nodes
367 in any of the random simulations would have degrees exceeding the threshold.

368 For actual data, because the edge weights are calculated by averaging s_{ij} over Monte Carlo simulations,
369 the number of localizations that would be classified as clustered in random data would usually be less than
370 $\alpha \times 100\%$. Hence, this averaging using localization uncertainties reduces the detection of spurious, small
371 clusters arising from random spatial fluctuations in density.

372 If localization uncertainties are not known, then we take a different approach to reduce detection of
373 spurious clusters. Preliminary clusters are defined using a community detection algorithm. A node is then
374 classified as unclustered if it meets any of the following four criteria: (1) it belongs to a preliminary cluster
375 whose mean degree is below the threshold; (2) its own degree is below the threshold and is also a lower outlier
376 ($< \text{lower quartile (LQ)} - 1.5 \times \text{interquartile range (IQR)}$) for its preliminary cluster; (3) its own degree passes
377 the threshold but is a strong lower outlier ($< \text{LQ} - 3 \times \text{IQR}$) for its preliminary cluster; (4) its own degree
378 is less than half of the threshold. The first criterion provides robustness by spatially averaging node degrees
379 over small areas. This prevents the inclusion of spurious, small clusters. The other three criteria prevent the
380 inclusion of nodes that are visually separate from a cluster but still within a distance r_0 of one.

381 To avoid biases arising from the choice of algorithm used for the preliminary clustering, StormGraph

382 performs this twice, independently, using two different community detection algorithms, and it then classifies
383 nodes as unclustered if either method does. The two algorithms used are the two-level version of Infomap
384 (46) and the Louvain method (47), which are two of the top performing community detection algorithms (32).
385 Infomap is an information theoretic algorithm based on flow on the graph, while the Louvain method is one
386 of several algorithms that aims to maximize a property of the graph called “modularity”. See Supplementary
387 Note 1 for further technical details.

388 **Edge pruning**

389 When localization uncertainties are used in the StormGraph algorithm, we prune edges from the final graph
390 that is constructed from only the nodes that are retained after thresholding node degrees. To do this, we
391 delete every edge that has nonzero s_{ij} in less than 75% of the Monte Carlo simulations that were used to
392 calculate the edge weights. This guarantees that any pair of retained edges have at least an estimated 50%
393 probability of co-occurring in the r_0 -neighborhood graph for any realization of the data, and the unknown
394 true localization positions is one possible realization. In turn, this prevents the linking of clusters that are
395 disconnected in most realizations of the r_0 -neighborhood graph but connected in the average graph.

396 **Merging clusters at the top of the multi-level Infomap hierarchy**

397 To facilitate the identification and quantification of particularly large clusters, StormGraph creates an ad-
398 ditional level at the top of the multi-level Infomap cluster hierarchy, if possible, by merging sufficiently
399 interconnected clusters. It is natural to consider the connected components of a graph to be the clusters
400 at the coarsest level of a cluster hierarchy. We therefore use this concept to define the top level of Storm-
401 Graph’s cluster hierarchy by merging Infomap clusters that form connected components. However, due to the
402 uncertainties in SMLM data, StormGraph only merges clusters if they form stable connected components,
403 which we define as connected components that would remain connected following the random removal or
404 displacement of any one node. Oftentimes, this step results in no merging of clusters and so no additional
405 level of clustering is created.

406 **Algorithm to obtain single-level clustering from cluster hierarchy**

407 Although various methods exist to select one level from a cluster hierarchy, for example silhouette scores
408 (48) and the gap statistic (49), existing methods are either very computationally intensive or otherwise

409 incompatible with StormGraph. We therefore developed our own fast algorithm to obtain a single-level
410 clustering from the cluster hierarchy output by StormGraph, which we describe here.

411 The hierarchical clustering output by StormGraph is generated from an r_0 -neighborhood graph. An
412 alternative type of graph commonly used for clustering problems is the symmetric k -nearest neighbor (kNN)
413 graph, in which two nodes are connected by an edge if either of them is among the k nearest neighbors of
414 the other. A related graph is the mutual kNN graph, a subgraph of the symmetric kNN graph, in which two
415 nodes are connected by an edge if and only if each node is among the k nearest neighbors of the other. One
416 simple clustering algorithm would be to identify the connected components in a symmetric kNN graph or in
417 a mutual kNN graph, where k is an adjustable parameter.

418 In a symmetric kNN graph, it is guaranteed that every node has at least k edges. However, as k increases,
419 nodes in low-density regions between two distinct clusters quickly become connected to both clusters, while
420 the high-density regions inside the clusters may remain fragmented into multiple connected components
421 until higher values of k . A mutual kNN graph, in which every node is guaranteed to have at most k
422 edges, more faithfully represents such clusters by preventing nodes in low-density regions from making too
423 many connections. However, mutual kNN graphs often suffer from having singletons and small connected
424 components due to the weak connectivity. We therefore chose to combine the concepts of both the symmetric
425 kNN and mutual kNN graphs.

426 For a set of points V and positive integers M and $K > M$, we define $G_{M,K}(V)$ to be the union of the
427 symmetric M NN graph and the mutual K NN graph for vertices V . This is still a subgraph of the symmetric
428 K NN graph, but it has stronger connectivity than the mutual K NN graph by guaranteeing that every node
429 has at least M edges, which in turn ensures that $G_{M,K}(V)$ contains no connected components with fewer
430 than $(M + 1)$ nodes.

431 For each cluster at the top level of the cluster hierarchy, StormGraph decides whether to split the cluster
432 into its subclusters at the next level down in the hierarchy according to the algorithm described below. If
433 the split is rejected, then StormGraph keeps the current cluster and does not examine any of the finer levels
434 of the hierarchy within that cluster. If the split is accepted, then this process is repeated recursively for
435 each of the newly accepted subclusters. A split is automatically rejected if more than 1% of the points in
436 the cluster belong to subclusters with fewer than the minimum number of points, specified by the user, that
437 constitute a cluster.

438 Let V be the set of nodes in a cluster C , let $A = \{C_1, C_2, \dots, C_n\}$ be the set of n subclusters of C at
439 the next finest level of the cluster hierarchy, and let $B(M, K) = \{C'_1, C'_2, \dots, C'_{n'}\}$ be the set of n' connected

440 components of the graph $G_{M,K}(V)$. StormGraph decides whether to split cluster C into its constituent
441 subclusters A using the following algorithm:

- 442 1. Construct $G_{2,K}(V)$ for all integers $K \in \{6, \dots, K_1\}$, where K_1 is the smallest integer such that
443 $G_{2,K_1}(V)$ is connected. We empirically chose the minimum value of K to be 6 because this usually
444 results in randomly distributed points forming a single connected component.
- 445 2. Find the value of K for which $B(2, K)$ is most similar to A according to some measure of similarity.
446 Denote this value of K by K^* .
- 447 3. Split cluster C into subclusters A if the similarity between A and $B(2, K^*)$ is greater than both a
448 threshold similarity and the similarity between C and $B(2, K^*)$.

449 The most obvious choices for a similarity measure to score the similarity between two clusterings of the
450 nodes V are normalized mutual information (NMI) (37) and mean F-measure (38). We require a similarity
451 measure that is defined even if one of the clusterings being compared consists of only a single cluster. This
452 eliminates NMI as a suitable choice, so we use mean F-measure.

453 Let $F(A, B)$ denote the similarity of clustering A to clustering B as measured by the mean F-measure.
454 The F-measure or F_1 score for a binary classification problem in which a cluster C_i is compared to a reference
455 cluster C'_i (usually the ground-truth cluster that the cluster C_i , found by a clustering algorithm, is supposed
456 to recover) is defined as the harmonic mean of precision (P) and recall (R):

$$F_1(C_i, C'_i) = \frac{2 \cdot P(C_i, C'_i) \cdot R(C_i, C'_i)}{P(C_i, C'_i) + R(C_i, C'_i)}.$$

457 The precision $P(C_i, C'_i)$ is the fraction of C_i that belongs to C'_i , and the recall $R(C_i, C'_i)$ is the fraction of
458 C'_i that belongs to C_i . The mean F-measure $F(A, B)$ is then defined as the weighted arithmetic mean of the
459 maximum F-measures for each of the clusters C'_i in B :

$$F(A, B) = \frac{\sum_{i=1}^{n'} |C'_i| \max_{1 \leq j \leq n} \{F_1(C_j, C'_i)\}}{\sum_{i=1}^{n'} |C'_i|},$$

460 where $|C'_i|$ denotes the number of points in C'_i .

461 The mean F-measure is not symmetric, i.e. $F(A, B) \neq F(B, A)$, which is not desirable in our situation
462 where we wish to compare two clusterings, neither of which is necessarily ground-truth. To avoid having
463 to choose one of the clusterings A and B to be the reference, we define a symmetric similarity measure,

464 $\bar{F}(A, B)$, as the arithmetic mean of $F(A, B)$ and $F(B, A)$:

$$\bar{F}(A, B) = \frac{1}{2} (F(A, B) + F(B, A)).$$

465 This is the similarity measure that we use in our algorithm for obtaining a single-level clustering from the
466 hierarchy. It ranges from 0 to 1, and $\bar{F}(A, B) = 1$ if and only if A and B are identical. We impose a minimum
467 similarity score of $\bar{F}_{\min} = 0.8$ for a cluster split to be considered. Thus, we split cluster C into its highest
468 level of subclusters, A , if A is at least 80% similar to $B(M, K^*)$ and is also a closer match to $B(M, K^*)$ than
469 the single, unified cluster C is. The 80% similarity threshold prevents the fragmentation of a cluster if there
470 is not substantial consensus between the two independent subclusterings. This threshold could be tuned to
471 make it more or less difficult to split a cluster into finer levels of subclusters. In particular, a threshold of
472 $\bar{F}_{\min} = 1$ would demand perfect agreement between the subclusters of C and the alternative, independent
473 clustering $B(M, K^*)$ for the subclusters to be accepted as a better clustering of V than a single cluster. We
474 chose a threshold of 0.8 to allow some leniency.

475 **Identifying clusters that can be confidently distinguished from multiply counted** 476 **single molecules**

477 Localizations arising from multiply counted single molecules may be falsely identified as clusters. As an
478 optional step during StormGraph analysis, clusters of localizations that cannot be distinguished with high
479 confidence, due to their positional uncertainties, from multiply counted single molecules can be identified
480 and subsequently reclassified as unclustered (cluster label 0). To do this, StormGraph checks each cluster
481 systematically as follows.

482 First, for each pair of localizations, $\mathbf{X}^i = (x^i, y^i, z^i)$ and $\mathbf{X}^j = (x^j, y^j, z^j)$, in the cluster, let $\mathbf{Y}^{ij} = \mathbf{X}^i - \mathbf{X}^j$
483 be their vector difference, and let Σ^{ij} be the covariance matrix for the coordinates of \mathbf{Y}^{ij} . The off-diagonal
484 elements of Σ^{ij} are assumed to all be 0 (i.e. the uncertainty in each coordinate of a localization is assumed
485 to be independent of its other coordinates). Assuming each molecule to be approximated by a point particle
486 of zero size, the m^{th} diagonal element of Σ^{ij} is $V_m^{ij} = \sigma_m^i{}^2 + \sigma_m^j{}^2$, where σ_m^i denotes the standard deviation
487 for the uncertainty in the m^{th} coordinate of localization i , as given by the input data.

488 This assumes that the true position is identical for all localizations originating from the same molecule.
489 In practice, the fluorophore positions may be different from the actual molecule positions. For example,
490 when molecules are detected using antibodies, the fluorophore conjugated to the antibody may be located

491 as much as 10 nm away from the antibody’s binding site. In addition, if each molecule can be labeled by
492 more than one fluorophore, then the true positions of localizations originating from a single molecule will
493 not only be different from the actual molecule but also from each other. If the sizes of the molecule and
494 fluorescent label are not negligible, they can be approximately taken into account in the following way. For
495 mathematical simplicity, we approximate the uncertainty due to the molecule and label size as an isotropic
496 Gaussian distribution with variance $(r/3)^2$, where r is the effective radius of the molecule and fluorescent
497 label combined, which is specified by the user based on underlying biophysical knowledge. We then add this
498 variance term twice (once each for localizations \mathbf{X}^i and \mathbf{X}^j) to each of the diagonal elements in Σ^{ij} . For
499 our simulated data, this was not necessary as the true position of every localization was at the centre of a
500 simulated molecule. For our BCR dSTORM data, we used $r = 8$ nm.

501 Next, we construct the statistic $Z^{ij} = \sum_{m=1}^d Y_m^{ij2} / V_m^{ij}$ for each pair of localizations, where d is the number
502 of dimensions (2 or 3) and Y_m^{ij} denotes the m^{th} coordinate of the vector \mathbf{Y}^{ij} . If two localizations \mathbf{X}^i and \mathbf{X}^j
503 have the same true position, then Z^{ij} is chi-squared distributed with d degrees of freedom. We then look
504 for pairs of localizations for which Z^{ij} exceeds a desired quantile of the appropriate chi-squared distribution,
505 indicating confidence that they originated from different molecules. Because we are testing multiple pairs
506 of localizations for significance, we correct for multiple hypothesis testing using the Šidák correction. If we
507 desire a significance level of $1 - q$, then we look for pairs of localizations for which Z_{ij} exceeds the $(q^{1/N})^{\text{th}}$
508 quantile of the chi-squared distribution with d degrees of freedom. Here, N is the number of localizations
509 in the cluster. Even though there are $N(N - 1)$ pairs of localizations, the null hypotheses are that each
510 localization originated from the same molecule as all other localizations in the cluster, and so there are only
511 N hypotheses to test. By default, StormGraph uses a significance level of 0.05, so it uses the $(0.95^{1/N})^{\text{th}}$
512 quantile. Finally, since a cluster must always contain at least three localizations (we do not consider pairs
513 of localizations to be clusters), StormGraph increases confidence further by demanding that at least two
514 localizations are each, probabilistically, sufficiently far from at least two other localizations. This way, a
515 single outlying localization within a cluster is not sufficient on its own to qualify the cluster as containing
516 multiple molecules with high confidence.

517 Guidelines for StormGraph parameter selection

518 StormGraph has three user-controllable parameters. The parameter α controls the node-degree threshold
519 used to identify and remove unclustered nodes prior to clustering. For data that does not suffer from
520 overcounting of molecules, or for which overcounting has already been corrected, α is effectively the maximum

521 false positive rate (FPR) for classifying localizations as clustered if all localizations in a random distribution
522 should be classified as unclustered. When overcounting is present in the data, the FPR may be greater
523 than α . Nevertheless, for any given $\alpha < 1$, StormGraph takes steps to minimize the FPR as far as possible.
524 Hence, we suggest setting α as the maximum fraction of localizations that the user would accept as being
525 clustered if they were completely randomly distributed. For most applications, we recommend $\alpha = 0.05$, the
526 default value. Larger values of α may be suitable if the user is already confident that the localizations are
527 strongly clustered but there is large variation in the density of clusters. For example, $\alpha = 0.5$ would simply
528 demand that clusters are at least as dense as the average density of a random distribution, but this could
529 result in as many as 50% of localizations in a random distribution qualifying as clustered. Alternatively, the
530 user can choose to skip the thresholding step and instead allow all localizations to be possibly assigned to
531 clusters by setting $\alpha = 1$, which ultimately removes all use of α from the StormGraph algorithm.

532 The optional parameter k specifies the number of nearest neighbors to use when calculating the graph
533 neighborhood radius r_0 . The value of k , if set, is the minimum (respectively maximum) number of neighbors
534 that most clustered (respectively unclustered) localizations should have. It should be smaller than the
535 number of localizations in a typical cluster, but preferably larger than the estimated number of times that
536 a typical single molecule might blink. These values can be estimated by visual inspection of localization
537 clusters within cell boundaries and on the coverslip outside of cells. Increasing k , and consequently r_0 , can
538 influence the exact placement of cluster boundaries, and hence cluster quantification, by allowing more low-
539 density localizations on the periphery of clusters to be included in the clusters. This highlights the inherent
540 ambiguity in clustering problems, which results from the lack of a clear definition of a cluster. Nonetheless,
541 we found values of k between 10 and 20 to be generally appropriate. Alternatively, if k is not set, StormGraph
542 will determine r_0 heuristically without any user input.

543 Finally, the user can optionally set the minimum number of localizations that a cluster must contain,
544 MinCluSize. One possible strategy for setting its value is to investigate background regions outside of cells,
545 where most clusters of localizations are likely to be due to individual fluorescent labels stuck to the coverslip,
546 and assess how many localizations are typical of these apparent clusters. However, because StormGraph
547 provides an option to use localization uncertainties to identify and reclassify localization clusters that could
548 have arisen just from overcounting of single molecules, clusters that could be due to single molecules can
549 be automatically removed from analysis without the need for a minimum cluster size parameter. Note that
550 StormGraph requires all clusters to contain at least three localizations, even if MinCluSize is not set.

551 Computational approximations in StormGraph

552 In order to improve computational efficiency, StormGraph includes some computational approximations.
553 Firstly, neighborhood searches about each node are performed using the MATLAB function “rangesearch”,
554 which uses a k-d tree, as this is faster than computing distances between all pairs of nodes. Without
555 uncertainties in localization positions, rangesearch is implemented with a search radius of r_0 . However,
556 when Monte Carlo simulations are used to perturb localization positions using their uncertainties, it is
557 inefficient to perform rangesearch for every simulation. Instead, we perform rangesearch just once, using an
558 expanded search radius, to identify candidate edges for the graph. StormGraph then calculates expected
559 edge weights only for the candidate edges. Since the computational time for rangesearch increases as the
560 search radius increases, we chose $(r_0 + 6 \times \text{mean localization uncertainty})$ as the expanded search radius
561 because most pairs of nodes separated by distances greater than this would have only negligible or zero edge
562 weights anyway. Increasing the search radius further would not only make rangesearch slower, but it could
563 also add more edges to the graph and consequently increase the computational cost of community detection,
564 even though the additional edges would be mostly negligible.

565 Secondly, StormGraph limits nodes to having no more than 500 neighbors in the graph. This is to
566 prevent extremely dense, large clusters from dramatically slowing down community detection, since the
567 computational time required by Infomap scales with the number of edges in the graph. In practice, for
568 reasonably chosen values of k , e.g. in the range from 10 to 20, and for r_0 values determined heuristically,
569 very few nodes, if any, in most datasets should have this many neighbors.

570 Lastly, we note that StormGraph is not deterministic, meaning that it can give slightly different results
571 each time that it is run. This is for two reasons. The first reason is because StormGraph uses Infomap
572 or the Louvain method to perform community detection. Infomap seeks to optimize the map equation and
573 the Louvain method seeks to optimize modularity. In both cases, the full optimization problem is NP-
574 hard. Therefore, both methods take a greedy approach to the optimization, which generally finds a local,
575 but not necessarily global, optimum. They then select the best optimum from multiple iterations started
576 from random initiations. In StormGraph, the default number of iterations used for finding the final cluster
577 hierarchy is 50. Results can be improved at the expense of increasing computational time by increasing the
578 number of iterations. Conversely, computational time can be reduced at the expense of cluster accuracy by
579 decreasing the number of iterations. The second reason for slight variability in results is the use of Monte
580 Carlo simulations by StormGraph. This variability can be decreased, again at the expense of increasing
581 computational cost, by increasing the number of Monte Carlo simulations.

582 The non-deterministic nature of StormGraph is only a minor drawback, as variability in clustering results
583 for a single dataset is small. To demonstrate this, we repeatedly applied StormGraph using identical settings
584 to a heterogeneous dSTORM ROI containing visually ambiguous clusters. We did this in both 2D and 3D
585 and for both the automatic and kNN methods for determining r_0 , each time generating 11 StormGraph
586 repeats. We then assessed the similarity of cluster assignments from each of the last 10 repeats to the first
587 one using NMI, which can range from 0 to 1. We always achieved $\text{NMI} > 0.94$, indicating very high similarity
588 (Figure S11).

589 **Simulating SMLM data in 2D and 3D**

590 In both 2D and 3D, we distributed 3000 molecules into circular nanoclusters with a fixed radius, r , and
591 fixed molecular density, ρ . Each molecule was assigned uncertainties, which were sampled randomly from
592 a real dSTORM dataset, in its x -, y - and (for 3D) z -coordinates and a number of blinks, which was drawn
593 from a geometric distribution (50) supported on $\{1, 2, 3, \dots\}$ with success probability parameter λ . Within
594 each nanocluster, molecules were distributed uniformly at random, and for each molecule the observed
595 localizations (blinks) were drawn from a normal distribution with mean equal to the molecule's position
596 and standard deviations equal to the uncertainties assigned to the molecule. Every observed localization
597 was assigned the same uncertainties as its associated molecule. The total number of nanoclusters, N_{nano} ,
598 was determined by the total number of molecules in clusters (3000) and the density, ρ , of molecules within
599 clusters.

600 The nanoclusters were positioned inside a $2 \mu\text{m} \times 2 \mu\text{m}$ ROI in 2D or a $2 \mu\text{m} \times 2 \mu\text{m} \times 1 \mu\text{m}$ ROI in
601 3D such that some existed as isolated nanoclusters and others were randomly aggregated into larger clusters
602 according to the following process, which was adapted from a Dirichlet process: for i from 1 to N_{nano} , draw
603 a random number from the uniform distribution on $[0,1]$; if it is less than or equal to $((p + 10)/(p + i - 1))^q$
604 for positive integers p and q , then place the i^{th} nanocluster away from existing clusters; otherwise, add the
605 i^{th} nanocluster to a randomly selected existing cluster, excluding the first 10 nanoclusters that were placed.
606 If a nanocluster was added to an existing cluster, it was placed such that its centre was exactly a distance
607 $2r$ from the centre of another nanocluster in the same aggregate cluster, and without overlapping with any
608 other existing nanoclusters in the aggregate cluster.

609 This process ensures that there are at least 10 isolated nanoclusters and a variable number of larger
610 aggregate clusters of variable size, thus creating heterogeneous clusters. The heterogeneity is controlled by
611 the parameters p and q . In our simulations, we fixed $p = 5$ and varied q from 1 to 5, with larger values of q

612 resulting in larger (and fewer) cluster aggregates. Outside of the clusters, we added molecules uniformly at
613 random at a specified average density, and the number and positions of observed localizations corresponding
614 to each of these background molecules were drawn from geometric and normal distributions respectively, as
615 described for the in-cluster localizations.

616 If the simulations were performed in 3D, points were then randomly removed such that the probability
617 of a localization being observed in the final simulated data decayed according to a Gaussian profile as the
618 axial distance from a central focal plane increased. This was to imitate the realistic scenario for most 3D
619 SMLM techniques in which fluorescent blink events are more likely to be collected and localized the closer
620 they are to the focal plane.

621 We generated 64 2D datasets with multiple blinking of molecules (e.g. Figures 2a(i) and S5) by varying
622 the following parameters: (1) the radius of the nanoclusters (20 nm, 30 nm or 50 nm), (2) the density of
623 clustered molecules (0.01 nm^{-2} or 0.02 nm^{-2}), (3) the density of the random molecules (1%, 5%, 10%, 20%
624 or 40% of the within-cluster molecular density), (4) the average number of blinks per molecule (4/3, 2 or 4;
625 these values provide examples ranging from cases in which most molecules blink only once to cases where
626 the molecules could be bivalent and labeled by fluorophores that blink on average twice, which is typical
627 for the photoactivatable fluorophore mEos2 (4; 51)), and (5) the propensity for nanoclusters to coalesce into
628 larger aggregate clusters (parameter q).

629 We generated 130 3D datasets analogously but using within-cluster molecular densities of $1 \times 10^{-4} \text{ nm}^{-3}$
630 and $2 \times 10^{-4} \text{ nm}^{-3}$. In 3D, we used nanoclusters of radii 30 nm and 50 nm, and we used densities of random,
631 unclustered molecules equal to 1%, 5%, 10% or 20% of the within-cluster molecular density. At 20%, clusters
632 were barely visible in 2D projections of the simulated 3D data onto the xy -plane.

633 **Running ClusterViSu on simulated data**

634 The ClusterViSu algorithm consists of running a series of two functions provided as part of its source code,
635 specifically the functions “VoronoiMonteCarlo” and “VoronoiSegmentation”. However, the authors did not
636 provide a script for running ClusterViSu. Hence, for users with zero programming expertise, it can only
637 be run using a graphical user interface that requires each file to be loaded and analyzed separately. Also,
638 ClusterViSu outputs the bounding polygon for each detected cluster but not the actual cluster assignments of
639 the localizations, which we needed to compute NMI and mean F-measure scores for assessing the performance
640 of cluster assignment. Therefore, we wrote our own custom MATLAB script (available upon request) to run
641 and batch process ClusterViSu from its source code and subsequently determine the cluster assignments of

642 the localizations. In addition, we found that ClusterViSu prefers input ROIs to be at least $18 \mu\text{m} \times 18 \mu\text{m}$,
643 so we rescaled our $2 \mu\text{m} \times 2 \mu\text{m}$ simulated data by a factor of 9, which drastically improved ClusterViSu's
644 performance, at least in terms of computational time.

645 Furthermore, we only included ClusterViSu results for simulated datasets on which ClusterViSu analysis
646 completed in under 2 hours. This resulted in 15 out of 64 simulated datasets being excluded from our
647 summary of test results for ClusterViSu, but these 15 datasets were still included for assessing StormGraph
648 and DBSCAN. However, these 15 datasets were excluded in Figures 2c and S3b, where NMI or mean F-
649 measure results for StormGraph and DBSCAN are shown as a ratio to the NMI or mean F-measure results
650 for ClusterViSu.

651 **Functionalization of glass coverslips for cell adherence**

652 Glass coverslips were cleaned and functionalized as previously described (52). Briefly, acid-cleaned glass cov-
653 erslips (Marienfeld #1.5H, $18 \text{ mm} \times 18 \text{ mm}$; catalogue #0107032, Lauda-Königshofen, Germany) were incu-
654 bated with 0.01% poly-L-lysine (Sigma-Aldrich; catalogue #P4707) or $0.25 \mu\text{g}/\text{cm}^3$ of the non-stimulatory
655 M5/114 anti-MHCII monoclonal antibody (Millipore; catalogue #MABF33) or $2 \mu\text{g}/\text{cm}^2$ fibronectin (Sigma
656 Aldrich; catalogue #F4759) for at least 3 h at 37°C . The slides were then washed with phosphate-buffered
657 saline (PBS) prior to being used for experiments.

658 **Monovalent Fab fragments and antibodies**

659 The anti-mouse-Ig κ antibody for clustering BCRs was purchased from Southern Biotech (Birmingham,
660 AL; catalogue #1050-01). AF647-conjugated anti-mouse-IgM Fab fragments (catalogue #115-607-020) and
661 AF647-conjugated anti-human-IgM Fab fragments (catalogue #109-607-043) were from Jackson ImmunoRe-
662 search Laboratories (West Grove, PA). All Fab fragments were routinely tested for aggregation using dynamic
663 light scattering (Zetasizer Nano) and unimodal size distributions were observed. Anti-LAMP-1 antibody was
664 purchased from Abcam (catalogue #ab24170). AF647-conjugated goat anti-mouse-IgG (catalogue #A21236)
665 and AF647-conjugated goat anti-rabbit-IgG (catalogue #A21244) were purchased from ThermoFisher Sci-
666 entific. Goat anti-mouse-IgG (Jackson ImmunoResearch Laboratories; catalogue #115-005-008) and goat
667 anti-rabbit-IgG (Jackson ImmunoResearch Laboratories; catalogue #111-001-008) were conjugated to Cy3B
668 using a Pierce antibody conjugation kit (catalogue #44985).

669 Cell labeling for dSTORM

670 (1) Murine splenic B cells

671 Animal protocols were approved by the University of British Columbia and all animal experiments were
672 carried out in accordance with institutional regulations. Splenic B cells were obtained from 6- to 10-week
673 old C57BL/6 mice (Jackson Laboratory) of either sex using a B-cell isolation kit (Stemcell Technologies;
674 catalogue #19854) to deplete non-B cells. To induce IgM-BCR clustering, 5×10^6 *ex vivo* splenic B cells/mL
675 were stimulated with 20 μ g/mL anti-Ig κ in PBS for 10 min at 37 °C. A similar volume of PBS was added to
676 control samples (resting B cells). All subsequent procedures were performed at 4 °C. Cells were washed three
677 times with ice-cold PBS, and IgM-BCRs on the cell surface were labeled using AF647-conjugated, monovalent
678 anti-mouse-IgM Fab fragments for 15 min. These Fab fragments bind to the constant region of the μ heavy
679 chain of IgM-BCRs, which is distinct from sites on the IgM-BCR that the anti-Ig κ treatment antibody binds
680 to. Following multiple PBS washes, cells were settled onto pre-cooled anti-MHCII-functionalized coverslips
681 for 10 min and subsequently fixed with PBS containing 4% paraformaldehyde and 0.2% glutaraldehyde for 90
682 min. The coverslips were washed thoroughly with PBS and fiducial markers (100 nm diameter; ThermoFisher
683 Scientific, catalogue #F8799) were allowed to settle onto the coverslip overnight at 4 °C. Unbound fiducial
684 markers were removed by PBS washes and the stuck particles were used for real-time drift stabilization (53).

685 (2) Human and murine B-lymphoma cell lines

686 A20 and BJAB B-lymphoma cells were obtained from American Type Culture Collection (ATCC). HBL-1
687 cells were obtained from Dr. Izidore S. Lossos, Sylvester Comprehensive Cancer Center, University of Miami
688 (Miami, FL). TMD8 cells were a gift from Dr. Neetu Gupta, Lerner Research Institute, Cleveland Clinic
689 (Cleveland, OH). All B-cell lines were cultured in RPMI-1640 (Life Technologies; catalogue #21870-076),
690 supplemented with 10% heat-inactivated fetal bovine serum, 2 mM L-glutamine, 50 μ M β -mercaptoethanol,
691 1 mM sodium pyruvate, 50 U/mL penicillin, and 50 μ g/mL streptomycin (complete medium). All cell lines
692 were authenticated by STR DNA profile analysis.

693 All staining procedures were performed at 4 °C. Cell-surface IgM-BCRs on BJAB, HBL-1 and TMD8 cells
694 were labeled using AF647-conjugated anti-human-IgM Fab fragments for 15 min. Cell-surface IgG-BCRs on
695 A20 cells (ATCC) were labeled using both AF647-conjugated anti-mouse-IgG and Cy3B-conjugated anti-
696 mouse-IgG at 1:1 stoichiometry for 15 min. Fc receptors on A20 cells were blocked prior to staining using the
697 2.4G2 rat anti-Fc γ receptor monoclonal antibody. Cells were washed in PBS and subsequently fixed with

698 ice-cold PBS containing 4% paraformaldehyde and 0.2% glutaraldehyde for 60 min. Following multiple PBS
699 washes, the cells were settled onto pre-cooled poly-L-lysine-coated coverslips for 15 min and subsequently
700 fixed again for 30 min. The coverslips were washed thoroughly with PBS and fiducial markers were added
701 and incubated overnight at 4 °C.

702 **(3) B16 melanoma cell lines**

703 B16F1 melanoma cells (ATCC) were grown in RPMI-1640 complete medium. Approximately 3×10^4 cells
704 were seeded on fibronectin-coated coverslips for 1 h and fixed with PBS containing 4% paraformaldehyde
705 for 30 min. Cells were permeabilized with 0.1% Triton X-100 for 10 min, washed with PBS, and incubated
706 for 30 min at room temperature (RT) with Image-IT FX Signal Enhancer (Life Technologies, catalogue
707 #I36933) to neutralize surface charge. Cells were washed briefly in PBS and then incubated with BlockAid
708 blocking solution (Life Technologies; catalogue #B10710) for 1 h at RT. The cells were incubated with
709 anti-LAMP-1 antibody (diluted in BlockAid) for 4 h at RT. Following PBS washes, cells were incubated
710 with both AF647-conjugated anti-rabbit-IgG and Cy3B-conjugated anti-rabbit-IgG at 1:1 stoichiometry for
711 90 min. Cells were washed in PBS and subsequently fixed again with 4% paraformaldehyde for 10 min. The
712 coverslips were washed thoroughly with PBS and fiducial markers were added and incubated overnight at
713 4 °C.

714 **dSTORM**

715 Imaging was performed using a custom-built microscope with a sample drift-stabilization system that has
716 been described previously (53; 54). Briefly, three lasers were used in the excitation path. These were a 639 nm
717 laser (Genesis MX639, Coherent) for exciting the AF647, a 532 nm laser (Laser quantum, Opus) for exciting
718 the photo-switchable Cy3B, and a 405 nm laser (LRD 0405, Laserglow Technologies) for reactivating the
719 AF647 and Cy3B. All three lasers were coupled into an inverted microscope equipped with an apochromatic
720 TIRF oil-immersion objective lens (60x; NA 1.49; Nikon). The emission fluorescence was separated using
721 appropriate dichroic mirrors and filters (Semrock) (53; 54), and detected by EM-CCD cameras (Ixon, Andor).
722 A feedback loop was employed to lock the position of the sample during image acquisition using immobile
723 fiducial markers. Sample drift was controlled to be less than 1 nm laterally and 2.5 nm axially.

724 **dSTORM image acquisition and reconstruction**

725 Imaging was performed in an oxygen-scavenging GLOX-thiol buffer consisting of 50 mM Tris-HCl, pH
726 8.0, 10 mM NaCl, 0.5 mg/ml glucose oxidase, 40 μ g/ml catalase, 10% (w/v) glucose and 140 mM 2-
727 mercaptoethanol (55). The coverslip with attached cells was mounted onto a depression slide filled with
728 imaging buffer and sealed with Twinsil two-component silicone-glue (Picodent; catalogue #13001000).

729 For SMLM imaging, a laser power density of 1 kW/cm² for the 639 nm and 532 nm lasers was used
730 to activate the AF647 and Cy3B, respectively. For each sample, 4×10^4 images were acquired for each
731 color channel at 50 Hz. Localization coordinates and their associated uncertainties were computationally
732 determined simultaneously by fitting a function to the intensity profile of each fluorescence event using
733 MATLAB (Figure S12), as described previously (54). Expressed as standard deviations, lateral uncertainties
734 were typically < 10 nm while axial uncertainties were typically < 40 nm (Figure S12).

735 For two-color SMLM, image acquisition was performed sequentially for each color with AF647 imaged
736 first to prevent photobleaching by the Cy3B excitation laser. Two-color SMLM images were acquired using
737 a beam splitter with appropriate filters to direct each signal to one of two independent cameras. Alignment
738 of these two colors was carried out using $\sim 4 \times 10^4$ images of fluorescent beads simultaneously recorded at
739 various positions to find an optimal geometric transformation. The resulting color-alignment error is ~ 10 nm
740 root mean squared.

741 **Acknowledgements**

742 We thank Alejandra Herrera-Reyes for helpful discussions and code, Ki Woong Sung for preliminary compu-
743 tational work, Dr. Vivian Qian Liu for assistance with dSTORM data fitting, Dr. Neetu Gupta for TMD8
744 cells, Dr. Izidore S. Lossos for HBL-1 cells, and Dr. David R.L. Scriven for helpful discussion. This work
745 was supported by funding from the Natural Science and Engineering Research Council of Canada (Discovery
746 Grant RGPIN-2015-04611 to DC, an Undergraduate Student Research Award to DWZ, Discovery Grant
747 RGPIN-2014-03581 to KCC, and Discovery Grant RGPIN-2017-04862 to MRG), the Canadian Cancer So-
748 ciety Research Institute (Innovation Grant 704254 to DC and MRG), the Canadian Institutes of Health
749 Research (PJT-19426 to MRG), and Canada Foundation for Innovation (to KCC).

750 **Author contributions**

751 JMS conceived the project, developed and tested the StormGraph algorithm and software, proposed exper-
752 iments, performed data analysis, produced figures, and wrote the manuscript. LA proposed and performed
753 experiments, wrote experimental methods, produced figures, and provided essential feedback about the al-
754 gorithm and software. DWZ assisted with software development, simulation of data, and algorithm testing.
755 RT built the dSTORM microscope and assisted with dSTORM data fitting and processing. KCC provided
756 code for fitting dSTORM localizations and aligning two-color dSTORM data. MRG and DC supervised the
757 project and wrote the manuscript. All authors approved the final manuscript.

758 **Competing interests statement**

759 The authors declare no competing interests.

References

- [1] Heilemann, M. *et al.* Subdiffraction-resolution fluorescence imaging with conventional fluorescent probes. *Angew Chem Int Ed Engl* **47**, 6172–6176 (2008).
- [2] van de Linde, S. *et al.* Direct stochastic optical reconstruction microscopy with standard fluorescent probes. *Nat Protoc* **6**, 991–1009 (2011).
- [3] Betzig, E. *et al.* Imaging intracellular fluorescent proteins at nanometer resolution. *Science* **313**, 1642–1645 (2006).
- [4] Shivanandan, A., Deschout, H., Scarselli, M. & Radenovic, A. Challenges in quantitative single molecule localization microscopy. *FEBS Letters* **588**, 3595–3602 (2014).
- [5] Thompson, R., Larson, D. & Webb, W. Precise nanometer localization analysis for individual fluorescent probes. *Biophys. J.* **82**, 2775–2783 (2002).
- [6] Mortensen, K., Churchman, L., Spudich, J. & Flyvbjerg, H. Optimized localization analysis for single-molecule tracking and super-resolution microscopy. *Nat. Methods* **7**, 377–381 (2010).
- [7] Juette, M. *et al.* Three-dimensional sub-100 nm resolution fluorescence microscopy of thick samples. *Nat. Methods* **5**, 527–529 (2008).
- [8] Huang, B., Wang, W., Bates, M. & Zhuang, X. Three-dimensional super-resolution imaging by stochastic optical reconstruction microscopy. *Science* **319**, 810–813 (2008).
- [9] Pavani, S. *et al.* Three-dimensional, single-molecule fluorescence imaging beyond the diffraction limit by using a double-helix point spread function. *Proc. Natl. Acad. Sci.* **106**, 2995–2999 (2009).
- [10] Aquino, D. *et al.* Two-color nanoscopy of three-dimensional volumes by 4Pi detection of stochastically switched fluorophores. *Nat. Methods* **8**, 353–359 (2011).
- [11] Owen, D. *et al.* PALM imaging and cluster analysis of protein heterogeneity at the cell surface. *J. Biophotonics* **3**, 446–454 (2010).
- [12] Mattila, P. *et al.* The actin and tetraspanin networks organize receptor nanoclusters to regulate B cell receptor-mediated signaling. *Immunity* **38**, 461–474 (2013).
- [13] Grove, J. *et al.* Flat clathrin lattices: stable features of the plasma membrane. *Molecular Biology of the Cell* **25**, 3581–94 (2014).

- [14] Pollitt, A. *et al.* Syk and Src family kinases regulate C-type lectin receptor 2 (CLEC-2)-mediated clustering of podoplanin and platelet adhesion to lymphatic endothelial cells. *The Journal of Biological Chemistry* **289**, 35695–710 (2014).
- [15] Hartley, J. *et al.* Super-resolution imaging and quantitative analysis of membrane protein/lipid raft clustering mediated by cell-surface self-assembly of hybrid nanoconjugates. *Chembiochem* **16**, 1725–9 (2015).
- [16] Roh, K.-H., Lillemeier, B. F., Wang, F. & Davis, M. M. The coreceptor CD4 is expressed in distinct nanoclusters and does not colocalize with T-cell receptor and active protein tyrosine kinase p56lck. *Proceedings of the National Academy of Sciences* **112**, E1604–E1613 (2015).
- [17] Gao, L., Chen, J., Gao, J., Wang, H. & Xiong, W. Super-resolution microscopy reveals the insulin-resistance-regulated reorganization of GLUT4 on plasma membranes. *Journal of Cell Science* **130**, 396–405 (2017).
- [18] Rahbek-Clemmensen, T. *et al.* Super-resolution microscopy reveals functional organization of dopamine transporters into cholesterol and neuronal activity-dependent nanodomains. *Nature Communications* **8**, 740 (2017).
- [19] Eryilmaz, M. *et al.* Localization microscopy analyses of MRE11 clusters in 3D-conserved cell nuclei of different cell lines. *Cancers* **10**, 25 (2018).
- [20] Gao, J. *et al.* Revealing the cellular localization of STAT1 during the cell cycle by super-resolution imaging. *Scientific Reports* **5**, 9045 (2015).
- [21] Ripley, B. Modelling spatial patterns. *J. R. Stat. Soc. Series B Stat. Methodol.* **39**, 172–192 (1977).
- [22] Kiskowski, M., Hancock, J. & Kenworthy, A. On the use of Ripley’s K-function and its derivatives to analyze domain size. *Biophys J.* **97**, 1095–1103 (2009).
- [23] Ester, M., Kriegel, H., Sander, J. & Xu, X. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining.*, 226–231. KDD-96 (AAAI Press, 1996).
- [24] Levet, F. *et al.* SR-Tesseler: a method to segment and quantify localization-based super-resolution microscopy data. *Nature Methods* **12**, 1065–1071 (2015).

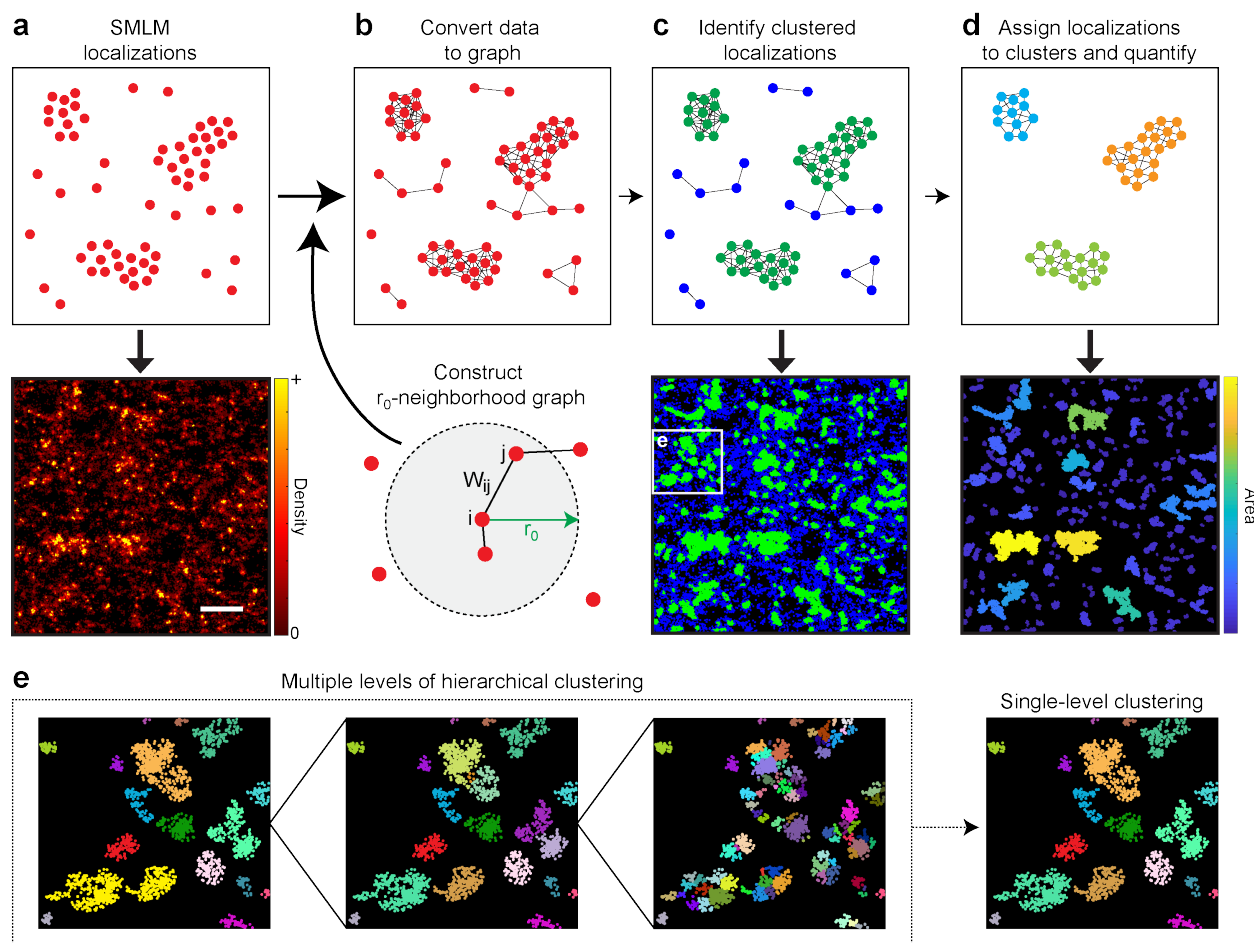
- [25] Andronov, L., Orlov, I., Lutz, Y., Vonesch, J. & Klaholz, B. ClusterViSu, a method for clustering of protein complexes by Voronoi tessellation in super-resolution microscopy. *Sci. Rep.* **6** (2016).
- [26] Andronov, L. *et al.* 3DClusterViSu: 3D clustering analysis of super-resolution microscopy data by 3D Voronoi tessellations. *Bioinformatics* **34**, 3004–3012 (2018).
- [27] Mazouchi, A. & Milstein, J. Fast Optimized Cluster Algorithm for Localizations (FOCAL): a spatial cluster analysis for super-resolved microscopy. *Bioinformatics* **32**, 747–754 (2016).
- [28] Nan, X. *et al.* Single-molecule superresolution imaging allows quantitative analysis of RAF multimer formation and signaling. *Proc Natl Acad Sci USA* **110**, 18519–18524 (2013).
- [29] Rubin-Delanchy, P. *et al.* Bayesian cluster identification in single-molecule localization microscopy data. *Nature Methods* **12**, 1072–1076 (2015).
- [30] Griffié, J. *et al.* 3D Bayesian cluster analysis of super-resolution data reveals LAT recruitment to the T cell synapse. *Sci. Rep.* **7** (2017).
- [31] Levine, J. *et al.* Data-driven phenotypic dissection of aml reveals progenitor-like cells that correlate with prognosis. *Cell* **162**, 184–197 (2015).
- [32] Fortunato, S. Community detection in graphs. *Physics Reports* **486**, 75–174 (2010).
- [33] Endesfelder, U. *et al.* Multiscale spatial organization of RNA polymerase in *Escherichia coli*. *Biophys J.* **105**, 172–181 (2013).
- [34] Rosvall, M. & Bergstrom, C. Multilevel compression of random walks on networks reveals hierarchical organization in large integrated systems. *PLoS ONE* **6**, e18209 (2011).
- [35] Strehl, A. & Ghosh, J. Cluster ensembles – a knowledge reuse framework for combining multiple partitions. *J. Mach. Learn. Res.* **3**, 583–617 (2002).
- [36] Lancichinetti, A. & Fortunato, S. Consensus clustering in complex networks. *Scientific Reports* **2** (2012).
- [37] Danon, L., Díaz-Guilera, A., Duch, J. & Arenas, A. Comparing community structure identification. *J. Stat. Mech.* 09008 (2005).
- [38] Aghaeepour, N. *et al.* Critical assessment of automated flow cytometry data analysis techniques. *Nat. Methods* **10**, 228–238 (2013).

- [39] Maity, P. *et al.* B cell antigen receptors of the IgM and IgD classes are clustered in different protein islands that are altered during B cell activation. *Sci. Signal.* **8**, ra93 (2015).
- [40] Depoil, D. *et al.* CD19 is essential for B cell activation by promoting B cell receptor-antigen microcluster formation in response to membrane-bound ligand. *Nature Immunology* **9**, 63–72 (2008).
- [41] Harwood, N. E. & Batista, F. D. Early events in B cell activation. *Annual Review of Immunology* **28**, 185–210 (2010).
- [42] Davis, R. *et al.* Chronic active B-cell-receptor signalling in diffuse large B-cell lymphoma. *Nature* **463**, 88–92 (2010).
- [43] Jaccard, P. The distribution of the flora of the alpine zone. *New Phytologist* **11**, 37–50 (1912).
- [44] Shivanandan, A., Unnikrishnan, J. & Radenovic, A. Accounting for limited detection efficiency and localization precision in cluster analysis in single molecule localization microscopy. *PLoS ONE* **10**, e0118767 (2015).
- [45] Arias-Castro, E. Clustering based on pairwise distances when the data is of mixed dimensions. *IEEE Transactions on Information Theory* **57**, 1692–1706 (2011).
- [46] Rosvall, M. & Bergstrom, C. Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences* **105**, 1118–1123 (2008).
- [47] Blondel, V., Guillaume, J., Lambiotte, R. & Lefebvre, E. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* **10**, 10008 (2008).
- [48] Rousseeuw, P. J. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics* **20**, 53–65 (1987).
- [49] Tibshirani, R., Walther, G. & Hastie, T. Estimating the number of clusters in a data set via the gap statistic. *J. R. Statist. Soc. B* **63**, 411–423 (2001).
- [50] Nino, D., Rafiei, N., Wang, Y., Zilman, A. & Milstein, J. Molecular counting with localization microscopy: A Bayesian estimate based on fluorophore statistics. *Biophysical Journal* **112**, 1777–1785 (2017).
- [51] Fricke, F., Beaudouin, J., Eils, R. & Heilemann, M. One, two or three? Probing the stoichiometry of membrane proteins by single-molecule localization microscopy. *Scientific Reports* **5** (2015).

- [52] Abraham, L. *et al.* Limitations of Qdot labelling compared to directly-conjugated probes for single particle tracking of B cell receptor mobility. *Scientific Reports* **7** (2017).
- [53] Tafteh, R. *et al.* Real-time 3D stabilization of a super-resolution microscope using an electrically tunable lens. *Opt. Express* **24**, 22959–22970 (2016).
- [54] Tafteh, R., Scriven, D. R. L., Moore, E. D. W. & Chou, K. C. Single molecule localization deep within thick cells; a novel super-resolution microscope. *J. Biophoton.* **9**, 155–160 (2016).
- [55] Dempsey, G. T., Vaughan, J. C., Chen, K. H., Bates, M. & Zhuang, X. Evaluation of fluorophores for optimal performance in localization-based super-resolution imaging. *Nature Methods* **8**, 1027–1036 (2011).

Figures and figure legends

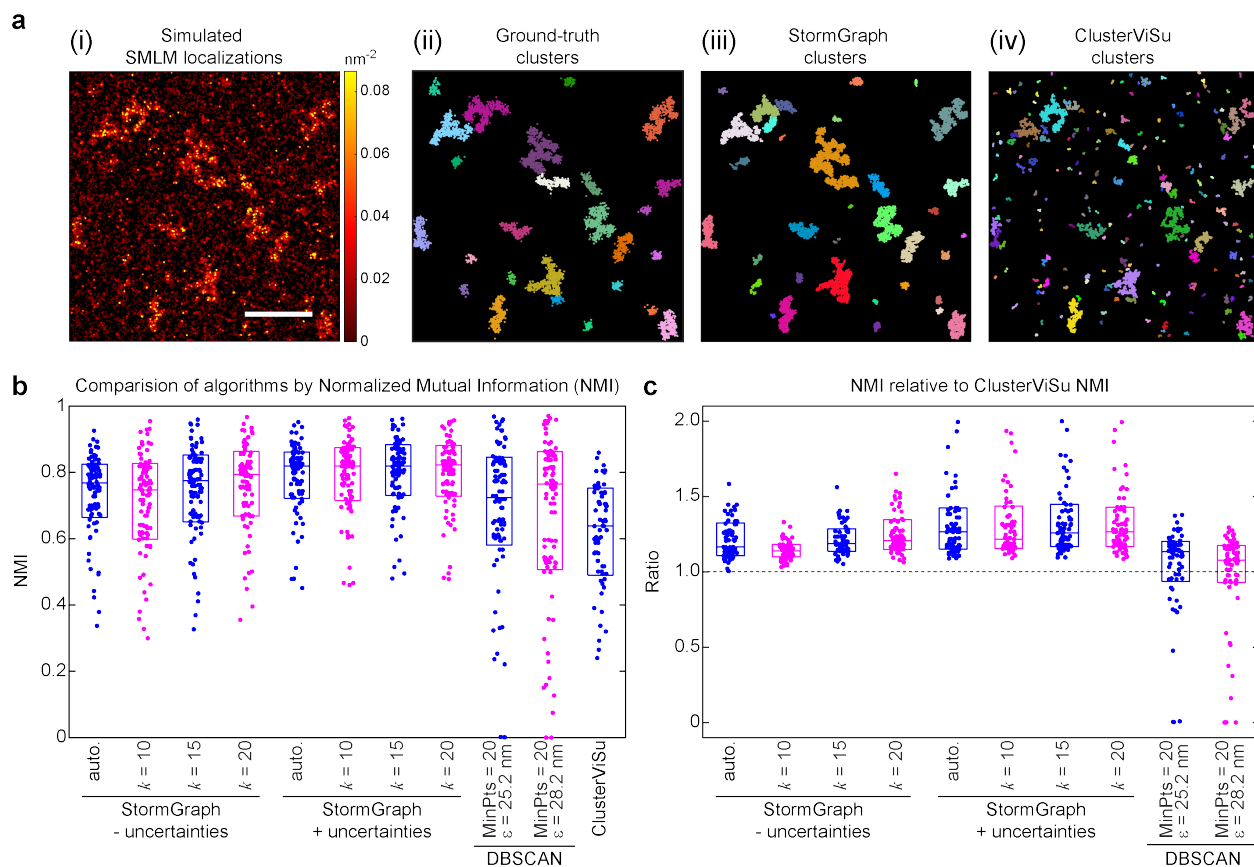
Figure 1.



Scurll *et al*, Figure 1

Figure 1: **Schematic illustration of StormGraph's workflow.** (a–d) SMLM localizations are used as nodes (a) and converted into a weighted graph (b). This graph is based on r_0 -neighborhood graphs, where edges connect each node to all other nodes within a distance r_0 . Two nodes i and j are connected by an edge of weight W_{ij} , which describes their similarity based on the distance between them and, if known, their positional uncertainties. Nodes are classified as either clustered (green) or unclustered (blue) based on their node degree, i.e. sum of adjacent edge weights, (c). A new graph is constructed from only the clustered nodes, which are then assigned to specific clusters using a community detection algorithm (d). Cluster properties (e.g. area) can then be quantified. The bottom panels in (a), (c), and (d) illustrate each step for an actual SMLM region of interest (scale bar = 500 nm). (e) StormGraph identifies a hierarchy of clusters at multiple scales and then generates an appropriate single-level clustering from the hierarchy. Shown are three different levels from the cluster hierarchy for the region in the white box in the lower panel of (c), along with the single-level clustering for this region. Colors distinguish different clusters.

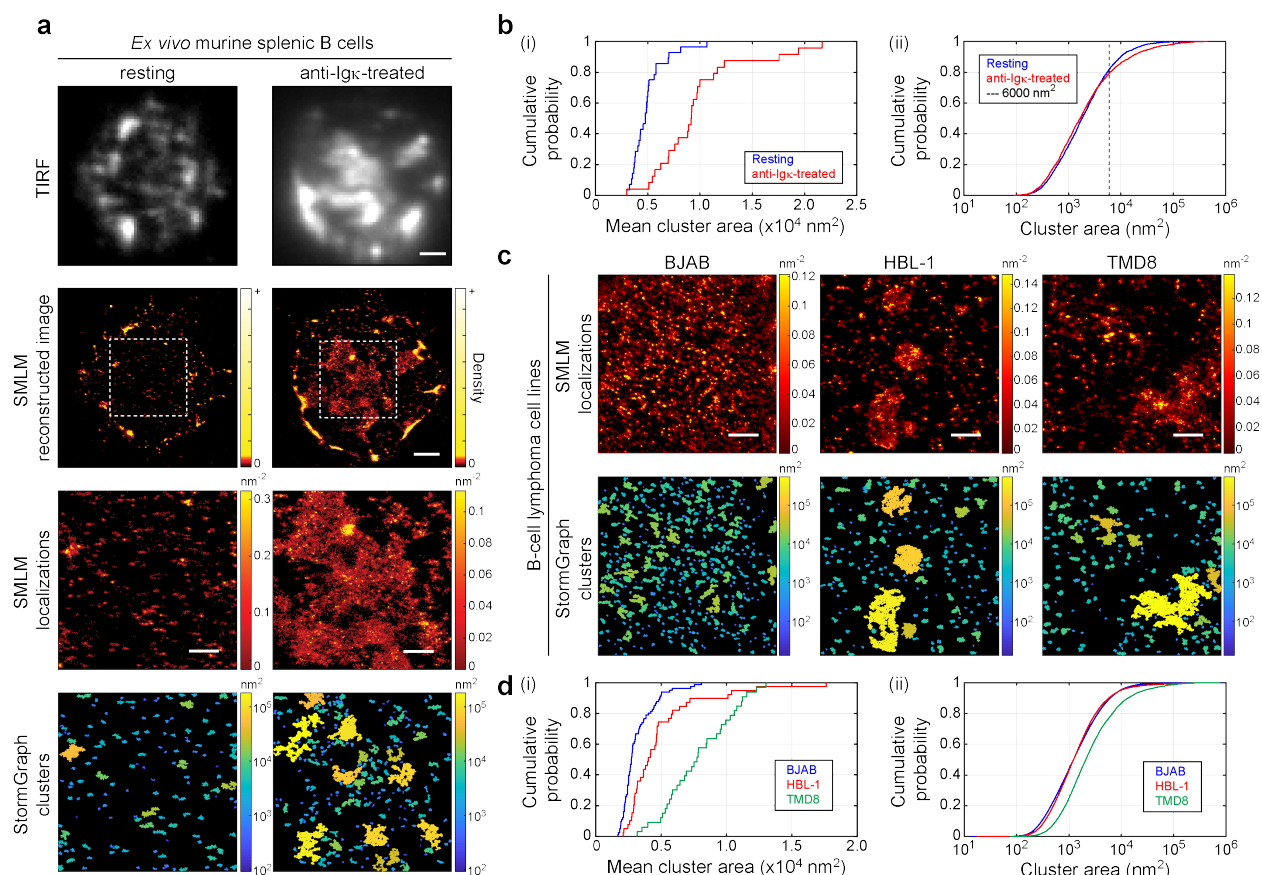
Figure 2.



Scurll et al, Figure 2

Figure 2: **StormGraph consistently outperforms ClusterViSu and DBSCAN on simulated data.** (a) (i) An example simulated dataset (color bar = density, scale bar = 500 nm), (ii) its ground-truth clusters, and cluster assignment results for (iii) StormGraph and (iv) ClusterViSu. Colors distinguish distinct clusters. (b) Accuracy of assigning data points to clusters as assessed by normalized mutual information (NMI; 1 = perfect). A total of 64 simulated datasets were analyzed using StormGraph, ClusterViSu, and DBSCAN. StormGraph was run either with (+) or without (-) localization uncertainties. The value of r_0 used by StormGraph was determined using the heuristic method (auto.) or the k -nearest neighbor method with $k = 10, 15$ or 20 . DBSCAN was implemented using 16 different selections of its two parameters, MinPts and ϵ , of which the two best-performing are shown here. ClusterViSu results are only shown for the 49 datasets on which the analysis was completed in under 2 h. Boxes show medians and interquartile ranges. (c) The same results as in panel (b) are displayed normalized to ClusterViSu's performance for each of the 49 simulated datasets for which analysis by ClusterViSu was completed in under 2 h. StormGraph was consistently more accurate than ClusterViSu at assigning points to clusters, indicated by ratios > 1 .

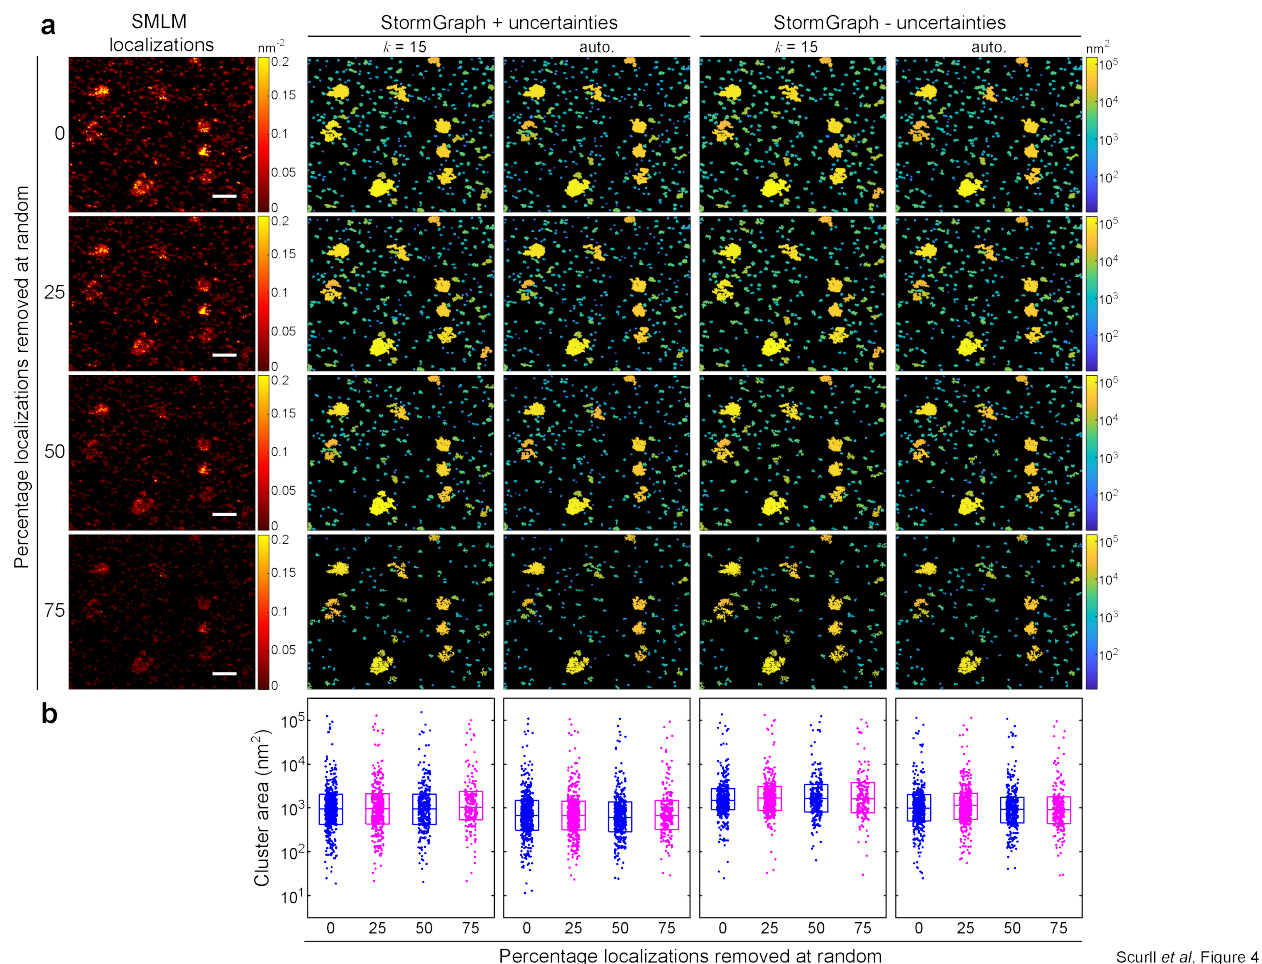
Figure 3.



Scurll *et al.*, Figure 3

Figure 3: StormGraph analysis of IgM-isotype B-cell antigen receptors (IgM-BCRs) on *ex vivo* murine splenic B cells (a–b) and human B-lymphoma cell lines (c–d) imaged using dSTORM. (a) StormGraph analysis of IgM-BCRs imaged by dSTORM on *ex vivo* murine splenic B cells that were either resting (left panels) or treated with bivalent anti-Ig κ antibodies (anti-Ig κ -treated; right panels). Top row: Total internal reflection fluorescence microscopy images taken prior to dSTORM imaging of the same cells. Scale bar = 1 μ m. Second row: Reconstructed images of IgM-BCRs from dSTORM localizations. Scale bar = 1 μ m. Third row: IgM-BCR dSTORM localizations in the ROIs shown in the second row. Scale bar = 500 nm, color bar = density (nm $^{-2}$). Bottom row: Clusters identified by StormGraph, colored by their areas (nm 2). **(b)** Cumulative distribution functions (CDFs) for cluster areas in ROIs from 28 **resting** (blue) and 24 **anti-Ig κ -treated** (red) *ex vivo* murine splenic B cells. Panel (i) shows the mean cluster area in each ROI. Panel (ii) shows all cluster areas from all ROIs. The increased mean cluster areas in anti-Ig κ -treated cells compared to resting cells is due to clusters larger than ~ 6000 nm 2 . **(c)** StormGraph analysis of IgM-BCRs imaged by dSTORM on resting BJAB, HBL-1 and TMD8 cells. Top: ROIs containing IgM-BCR dSTORM localizations from one Burkitt's lymphoma (BJAB) and two ABC DLBCL (HBL-1 and TMD8) cells. Scale bar = 500 nm, color bar = density (nm $^{-2}$). Bottom: Clusters identified by StormGraph, colored by their areas (nm 2). **(d)** CDFs for cluster areas in ROIs from 81 **BJAB** (blue), 39 **HBL-1** (red), and 33 **TMD8** (green) cells. Panel (i) shows the mean cluster area in each ROI. Panel (ii) shows the areas of all clusters in all ROIs. The larger mean area of clusters on HBL-1 cells than on BJAB cells is due to small numbers of very large clusters. All StormGraph results shown here were generated using localization uncertainties and $k = 15$.

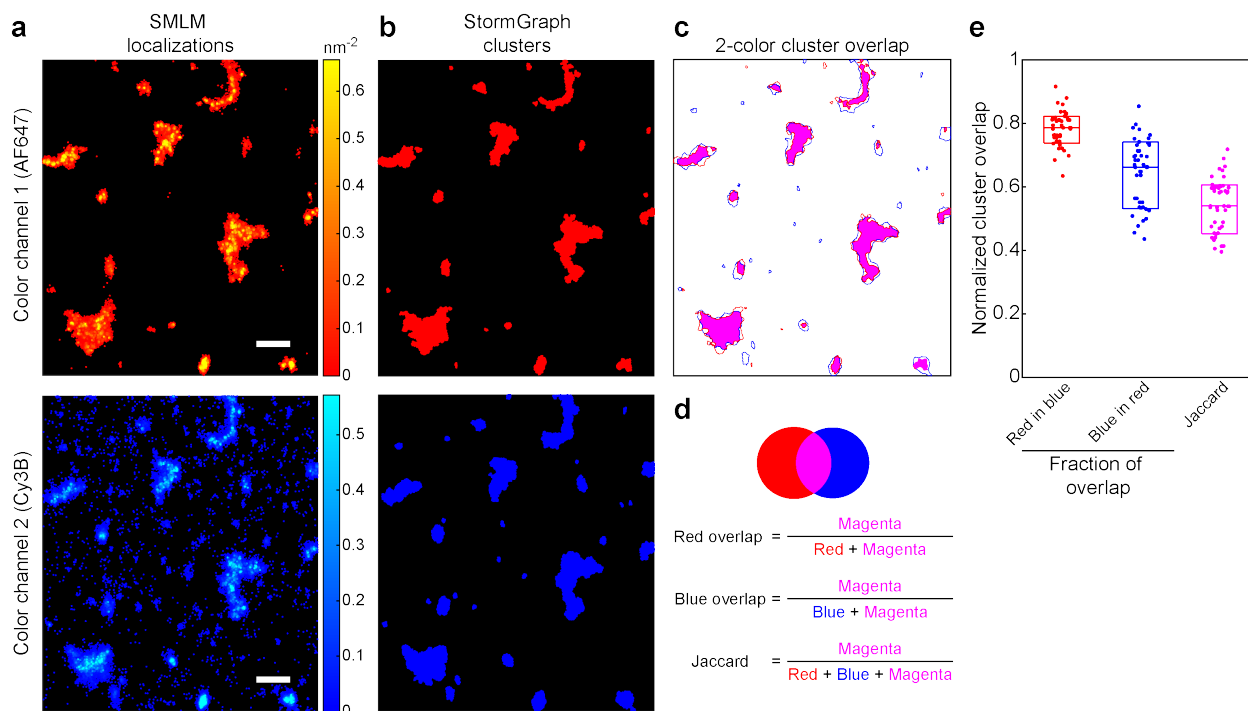
Figure 4.



Scurll *et al.*, Figure 4

Figure 4: **StormGraph results are not sensitive to the global average density of localizations.** (a) StormGraph cluster results (color bar = cluster area (nm²)) for a dSTORM dataset with 0%, 25%, 50% or 75% of the localizations randomly removed (left panels, color bar = density (nm⁻²)). Localization uncertainties were either used (+) or not used (-) by StormGraph during clustering. The value of r_0 was set using either the kNN method with $k = 15$ or the heuristic method (auto.). Scale bar = 500 nm. (b) Cluster areas quantified by each implementation of StormGraph for each of the four datasets, showing that the distribution of quantified cluster areas was not significantly affected by the random removal of localizations ($p > 0.05$ as determined by two-sample Kolmogorov-Smirnov tests). Boxes show medians and interquartile ranges.

Figure 5.



Scurll *et al*, Figure 5

Figure 5: Two-color cluster overlap analysis using StormGraph. (a) IgG-isotype B-cell antigen receptors (IgG-BCRs) on A20 B cells were labeled simultaneously with bivalent anti-IgG antibodies that were conjugated to either AF647 (top panel; red) or Cy3B (bottom panel; pseudo-colored blue) and imaged using dSTORM. Bivalent antibodies were used to induce clustering, since each antibody can bind to up to two IgG-BCRs. Figure shows the IgG-BCR dSTORM localizations in a region of interest (ROI) from one representative cell analyzed using StormGraph. Scale bar = 500 nm, color bars = density (nm⁻²). (b) Binary images of the AF647 (top) and Cy3B (bottom) clusters identified by StormGraph in the ROI shown in (a). (c) Merged image of the outlines of the AF647 clusters (red) and Cy3B clusters (blue) identified by StormGraph, with the overlapping area colored in magenta. (d) Pictorial description of the three area-based cluster overlap scores calculated by StormGraph, in the same order as the columns in panel (e). (e) Cluster overlap scores calculated using the formulae in panel (d) for 31 StormGraph-analyzed ROIs from multiple A20 cells imaged in the same experiment. Each ROI contributes one dot to each column. Boxes show medians and interquartile ranges. These scores determine the maximum observable overlap that could be expected for clusters of IgG-BCRs and a different molecule labeled using these same two fluorophores on A20 cells, imaged using the same imaging setup and analyzed by StormGraph.

Figure 6.

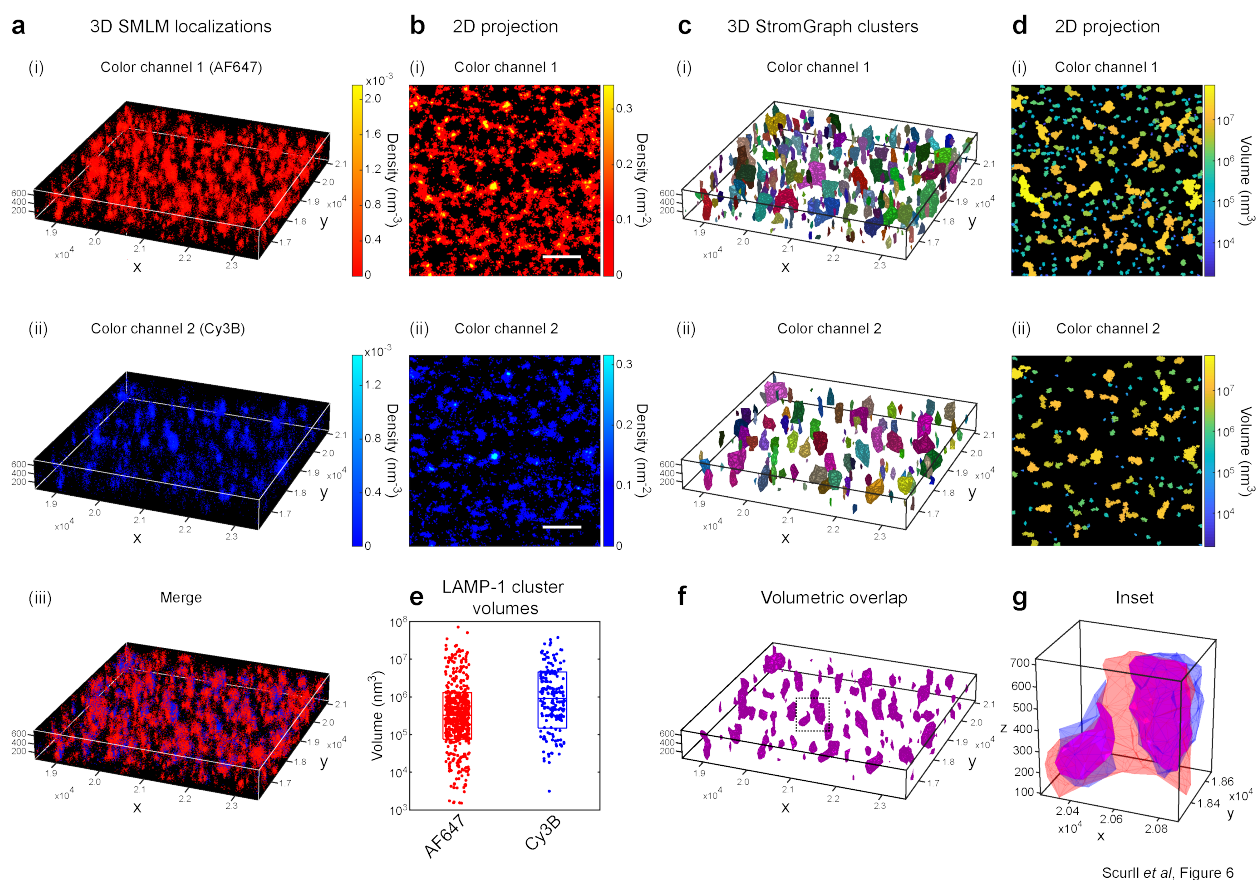


Figure 6: **Analysis of 3D SMLM data using StormGraph.** (a) Localizations of intracellular LAMP-1 detected by simultaneous staining with anti-LAMP-1 antibodies that were visualized using the same secondary antibody labeled with either (i) AF647 (red) or (ii) Cy3B (pseudo-colored blue). A murine B16 melanoma cell was imaged by two-color, 3D dSTORM and a $5 \mu\text{m} \times 5 \mu\text{m} \times 700 \text{nm}$ ROI was selected for analysis. Color bars = density (nm^{-3}). (iii) Merge. (b) 2D projections of the (i) AF647 and (ii) Cy3B localization data onto the xy -plane. Color bars = density (nm^{-2}). Scale bar = $1 \mu\text{m}$. (c) (i) AF647 and (ii) Cy3B clusters found by StormGraph using localization uncertainties, $k = 15$, $\alpha = 0.1$, and a minimum cluster size of 5 localizations. Clusters of localizations that could not be confidently distinguished from a single, multiply counted fluorescent probe were removed. Colors distinguish different clusters. (d) 2D projections of the clusters in panel (c) colored by their volumes (nm^3). (e) All volumes of AF647 and Cy3B clusters detected by StormGraph, as in (d). Boxes show medians and interquartile ranges. (f) Overlapping volumes (magenta) of the AF647 and Cy3B clusters in panel (c). (g) Enlarged region showing overlap (magenta) between one AF647 cluster (red) and two Cy3B clusters (pseudo-colored blue).