

Flexible and accurate quantification of astrocyte activity reveals novel single-cell and population-level physiology

Yizhi Wang^{1*}, Nicole V. DelRosso^{2*}, Trisha V. Vaidyanathan³, Michelle K. Cahill³, Michael E. Reitman³, Silvia Pittolo², Xuelong Mi¹, Guoqiang Yu^{1¶}, Kira E. Poskanzer^{2,3,4¶}

¹Bradley Department of Electrical and Computer Engineering, Virginia Polytechnic Institute and State University, Arlington, VA 22203

²Department of Biochemistry & Biophysics, University of California, San Francisco, San Francisco, CA 94143

³Neuroscience Graduate Program, University of California, San Francisco, San Francisco, CA, 94143

⁴Kavli Institute for Fundamental Neuroscience

*Equal contribution

¶Co-corresponding authors

Running title: Astrocyte quantitative analysis (AQuA)

Corresponding authors:

Kira E. Poskanzer: kira.poskanzer@ucsf.edu

Guoqiang Yu: yug@vt.edu

Abstract

Recent work examining astrocytic physiology centers on fluorescence imaging approaches, due to development of sensitive fluorescent indicators and observation of spatiotemporally complex calcium and glutamate activity. However, the field remains hindered in fully characterizing these dynamics, both within single cells and at the population-level, because of the insufficiency of current region-of-interest-based approaches to describe activity that is often spatially unfixed, size-varying, and propagative. Here, we present an analytical framework that releases astrocyte biologists from ROI-based tools. The Astrocyte Quantitative Analysis (AQuA) software takes an event-based perspective to model and accurately quantify the complex activity in astrocyte imaging datasets. We apply AQuA to a range of *ex vivo* and *in vivo* imaging data, and uncover novel physiological phenomena in each. Since AQuA is data-driven and based on machine learning principles, it can be applied across model organisms, fluorescent indicators, experimental modes, and imaging resolutions and speeds, enabling researchers to elucidate fundamental astrocyte physiology.

Introduction

With increased prevalence of multiphoton imaging and optical probes to study the physiology of astrocytes¹⁻³, many groups now have the tools to study fundamental functions that previously remained unclear. Recent work has focused on new ways to decipher how astrocytes respond to neurotransmitter and neuromodulator circuit signals⁴⁻⁷ and how the spatiotemporal patterns of their activity shape local neuronal activity⁸⁻¹⁰. Recording astrocytic dynamics with the goal of decoding their disparate roles in neural circuitry has largely centered on cell type-specific expression of genetically encoded probes to carry out intracellular calcium (Ca^{2+}) imaging using variants of GCaMP³. In addition, many groups have studied astrocytic function by performing extracellular glutamate imaging using GluSnFR², and several more recently developed genetically encoded fluorescent probes for neurotransmitters such as GABA¹¹, norepinephrine (NE)¹², ATP¹³, and dopamine¹⁴ are poised to expand our understanding of astrocytic circuit biology.

Compared to neuronal Ca^{2+} imaging, astrocytic Ca^{2+} imaging using GCaMP presents particular challenges for analysis due to their complex spatiotemporal dynamics. Thus, astrocyte-specific analysis software has been developed to capture these dynamics, including techniques that divide the cell into distinct subcellular regions corresponding to their anatomy⁴ or apply a watershed algorithm to identify regions-of-interest (ROIs)¹⁵. Likewise, GluSnFR imaging analysis techniques are based on manually or semi-manually selected ROIs, or by analyzing the entire imaging field together as one ROI^{2,6,8,16}. It is worth noting that these and most, although not all^{17,18}, other current techniques rely on the conceptual framework of ROIs for image analysis. However, astrocytic Ca^{2+} and GluSnFR fluorescence dynamics

are particularly ill-suited for ROI-based approaches, because the concept of the ROI has several inherent assumptions that cannot be satisfied for astrocytic activity data. Astrocytic Ca^{2+} signals, for example, can occupy regions that change size or location across time, can propagate within or across cells, and can spatially overlap with other Ca^{2+} signals that are temporally distinct. ROI-based approaches assume that for a given ROI, all signals have a fixed size and shape as specified by the ROI, and all locations within the ROI undergo the same dynamics, without propagation. Accordingly, ROI-based techniques may over- or under-sample these data, thus obscuring true dynamics and hindering physiological discovery in these cells. An ideal imaging analysis framework for astrocytes would take into account, and quantify, all of these dynamic features and be free of these ROI-based analytical restrictions. In addition, an ideal tool should be applicable to astrocyte imaging data across spatial scales, encompassing subcellular, cellular, and population-wide fluorescence dynamics.

In this work, we set out to design an image analysis toolbox that would capture the complex, wide-ranging fluorescent signals observed in most dynamic astrocyte imaging datasets. We reasoned that a non-ROI-based approach would better describe the observed fluorescent dynamics, and applied probability theory, machine learning, and computational optimization techniques to generate an algorithm to do so. We name this resulting software package Astrocyte Quantitative Analysis (AQuA) and validate its utility by applying it to simulated datasets that reflect the specific features that make analyzing astrocyte data challenging. We next apply AQuA to three experimental two-photon (2P) imaging datasets—*ex vivo* Ca^{2+} imaging of GCaMP6 from acute cortical slices, *in vivo* Ca^{2+} imaging of GCaMP6 in primary visual cortex (V1) of awake, head-fixed mice, and *ex vivo* extracellular glutamate, GABA, and NE imaging. In these test cases, we find that AQuA accurately detects fluorescence dynamics by capturing fluorescence events as they change in space and time, rather than the activity from a single location in space, as in ROI-based approaches. AQuA outputs a comprehensive set of biologically relevant parameters from these datasets, including propagation speed, propagation direction, area, shape, and spatial frequency. Using these detected events and associated output features, we uncover neurobiological phenomena that have not been previously described in astrocytes.

A wide variety of cellular and circuit functions have been ascribed to astrocytes, and a key question currently under examination in the field is whether certain types of Ca^{2+} activities observed in these cells correspond to particular neurobiological functions. However, current techniques with which to classify these observed dynamics remain inadequate since they do not capture many of the dynamics recorded in fluorescent imaging of astrocytic activity. The framework we describe here allows for a rigorous, in-depth dissection of astrocyte physiology across spatial and temporal imaging scales, and sets the stage for a comprehensive categorization of heterogeneous astrocyte activities both at baseline and after experimental manipulations.

Results

Design principles of the AQuA algorithm

To move away from ROI-based analysis approaches and accurately capture heterogeneous astrocyte fluorescence dynamics, we set out to design an algorithm to decompose raw dynamic astrocyte imaging data into a set of quantifiable events (Fig 1a, Supp. Movie 1, Supp. Fig. 1–3). Here, we define an event as a cycle of a signal increase and decrease that coherently occurs in a spatially connected region, but this region is defined by the fluorescence dynamics, not *a priori* by the user or the cell morphology.

Algorithmically, this definition is converted to the following two rules: 1) the temporal trajectory for an event contains only one peak (single-cycle rule, Fig. 1b) and 2) adjacent locations in the same event have similar trajectories (smoothness rule, Fig. 1b). The task of the AQuA algorithm is to detect all events, and, for each event, to identify the temporal trajectory, the spatial footprint, and how the signal propagates within the footprint. Briefly, our strategy of event-detection is to a) explore the single-cycle rule to find peaks, which are used to specify the time window and temporal trajectory, b) explore the smoothness rule to group spatially adjacent peaks, whose locations specify the footprint, c) apply machine learning and optimization techniques to iteratively refine the spatial and temporal properties of the event to best fit the data, and d) apply statistical theory to determine whether a detected event is true or due to noise (Fig. 1).

Full statistical and computational details are provided in the Methods, but we want to highlight one technical innovation (Graphical Time Warping [GTW])¹⁹ and one new concept (the single-source rule) that jointly enable a nuanced analysis of astrocyte fluorescence dynamics as shown below in application to experimental datasets. With GTW, we are able to consider fluorescent signal propagation as integrated into each modeled event. To the best of our knowledge, signal propagation has never been rigorously accounted for and has been considered an obstacle to analysis. With GTW, we can estimate and quantify propagation patterns in the data. With the introduction of the single-source rule (Fig. 1b), each event only contains a single initiation source and we can separate events that are initiated at different locations but meet in the middle. The single-source rule also allows us to divide large-scale activity that can occur across an entire field-of-view into individual events, each with a single initiation location.

The output of the event-based AQuA algorithm is a list of detected events, each associated with three categories of parameters: 1) the spatial map indicating where the event occurs, 2) the dynamic curve corresponding to fluorescence change over time (dF/F), and 3) the propagation map indicating signal propagation. For each event, we use the spatial map to compute the event area, diameter and shape of the domain it occupies (Fig 1c). Using the dynamic curve, we can calculate maximum dF/F, duration, onset-time, rise-time and decay-time. Using the propagation map, we extract event initiation location, as well as propagation path, direction, and speed. In addition, AQuA computes features involving more than one

event, such as the frequency of events at a position, and the overall number of events in a specified region or cell. A complete list of features is in the Methods section.

Validation of AQuA using simulated data

To validate AQuA, we designed three simulation datasets so that we know the ground truth for the dynamics of each event. These three datasets independently vary the three key phenomena observed in astrocyte imaging datasets that cause ROI-based approaches to misanalyze the data: size-variability, location-variability, and propagation. While these three phenomena usually co-occur in real datasets, we simulated each phenomenon independently to examine their individual impact and test AQuA's performance relative to other fluorescence image analysis tools, including CaImAn²⁰, Suite2P²¹, CaSCaDe¹⁵, and GECI-quant⁴. CaImAn and Suite2P are widely used for neuronal Ca^{2+} imaging analysis while CaSCaDe and GECI-quant were designed specifically for Ca^{2+} activity in astrocytes. We should note that although CaSCaDe uses the term "events" to describe Ca^{2+} transients, all four methods are ROI-based. In our analysis of these simulated datasets, as elaborated in the Methods, we optimally tuned the ROI-detection for all methods for an objective comparison of the best performance of each method. We also systematically changed the signal-to-noise-ratio (SNR) to examine the effect of noise on each analysis method.

To evaluate the performance on all simulated datasets, we used two measures: IoU and a map of the event counts. IoU (intersection over union) measures the consistency between detected and the ground-truth events, and takes into account both the spatial and temporal accuracy of detected events. IoU ranges from 0 to 1, where 1 indicates perfect detection and 0 indicates a complete failure in detection. The map of the event counts is obtained by counting the number of events at each pixel in the field, and is used to visually assess the accuracy of event-detection results by a comparison to the ground-truth map.

We first studied the impact of size-varying events (Fig. 2a), in which multiple events occurred at the same location and the event centers remain fixed, but sizes changed across different events. The degree of size change is quantified using size-change odds (see Methods) where a size-change odds of 1 indicates events with the same size, while an odds of 5 is the largest size change we simulated. For instance, when we set the odds at 5, we simulate events with sizes randomly distributed between 0.2 and 5 times the baseline size, with an SNR of 10dB, chosen to closely match the noise level in real experimental data. When there was no size change (odds=1), all methods, as expected, performed well with IoUs near 0.95 (Fig. 2a). When the degree of size change was increased, AQuA still performed well (IoU=0.95), while all other methods quickly drop to 0.4–0.5. We then changed our analysis to study the impact of different SNRs on performance by varying SNR, but fixing the size-change odds. AQuA performed better with increasing SNR and achieved nearly perfect detection accuracy (IoU=1) at 20dB. In comparison, all other

methods had an IoU less than 0.6, even at high SNR (Fig. 2a). We also examined the results by visualizing event counts at each pixel (Supp. Fig. 4–5). With size change odds of 3 or 5, the map of ground truth event-counts did not show clear ROI boundaries, because events from the same ROI had various sizes, and because events from different ROIs can overlap at some spatial locations (Supp. Fig. 4). It is clear from these maps that AQuA reported faithfully the events under various SNRs but all other methods had erroneous event counts and produced artificial patterns.

We next focused on the impact of shifting the event locations. In these simulated datasets, event size was fixed but event location changed, and degree of change was represented by a location change score (Fig. 2b). A value of zero indicates no location change and greater values represent larger degrees of change. Here, results are similar to changing size, as above. AQuA models the location change well and its performance is not affected by degree of location change. Likewise, AQuA reached near perfect results when SNR was high. In contrast, all other analysis methods performed poorly with changing locations. In particular, the other astrocyte-specific methods (CaSCaDe and GECI-quant) missed many signals. Even though the overall conclusion is similar for both the size- and location-changing events, the peer methods had more variation of IoU performance among themselves and the event count map showed distinct patterns (Supp. Fig. 4).

In our third simulated dataset, we asked how the phenomenon of fluorescence signal propagation impacts the performance of AQuA compared to the other methods. Two propagation types—growing and moving—were simulated in this dataset (Fig. 2c), although they were also separately evaluated (Supp. Fig 6). Propagation frame number denotes the difference between the earliest and latest onset times within a single event. When propagation frame number is zero, all signals within one ROI, but not necessarily across ROIs, are synchronized and there is no propagation. Similar results to the two scenarios discussed above were obtained here, with AQuA out-performing all the other methods by a large margin. These results indicate that AQuA can handle various types of propagation well, while the performance of other methods degrades rapidly when propagation is introduced.

In summary, when any of the three ROI-violating factors—size-variability, location-variability, and propagation—is introduced, other methods do not accurately capture the signal dynamics of the simulated data, and AQuA outperforms them by a large margin. We expect that the performance margin on real experimental data is larger than those quantified in the simulation studies here, since real data exhibits multiple ROI-violating factors and the performance of the ROI-based methods is over-estimated in our simulations (see Methods). However, these IoU analyses and the event count visualizations informed us about different types of errors observed in ROI-based methods. To focus on astrocyte-specific methods, CaSCaDe tends to over-segment, as it is based on watershed segmentation. GECI-quant, especially its soma-segmentation step, is particularly challenged by noise, causing many signals to be lost (Supp. Fig.

5). This accounts for the result that GECI-quant is not able to detect anything when the SNR is low (Fig. 2b, right). We note that although all events are constrained in the same ROI in the propagation simulations, propagation caused ROI-based approaches to quickly decline in performance. Here, GECI-quant was influenced by noise level, while CaSCaDe's assumption of synchronized signals did not allow accurate capture of the event dynamics.

AQuA enables identification of single-cell physiological heterogeneities

To test AQuA's performance on real astrocyte fluorescence imaging data and ask whether AQuA could be used for classifying Ca^{2+} activities observed in single cells, we first ran AQuA's event-detection on Ca^{2+} activity recorded from astrocytes in acute cortical slices from mouse V1 using 2P microscopy. We used a viral approach to express the genetically encoded Ca^{2+} indicator GCaMP6f³ in layer 2/3 (L2/3) astrocytes. Unlike ROI-based approaches, AQuA detects both propagative and non-propagative activity, revealing Ca^{2+} events with a variety of shapes and sizes (Fig. 3a, left). Further, since AQuA not only detects Ca^{2+} events' spatial footprint but also their time-course, we can apply AQuA to measure the propagation direction each event travels over its lifetime. Imaging single cells, we used the soma as a landmark, and classified events as traveling toward the soma (pink), away from the soma (purple), or static (blue) for the majority of its lifetime (Fig. 3a, right). We used AQuA's automatic feature-extraction and combined multiple measurements (size, propagation direction, duration, and minimum proximity to soma) into one spatiotemporal summary plot (Fig. 3b). Since astrocytes exhibit a wide diversity of Ca^{2+} activities across subcellular compartments^{6,22,23}, plotting the signals this way rather than standard DF/F transients highlights these heterogeneities, allows us to map the spatial location of the Ca^{2+} signals, and enables a quick, visual impression of a large amount of complex data (Supp. Fig. 7). We note that while the expression of GCaMP6 in these experiments enabled us to analyze events within single cells, some probes do not allow clear delineation of single cells. However, a secondary fluorophore (such as TdTomato) often serves the purpose of defining the morphology of single cells, and the AQuA software has been designed to overlay morphological masks on the dynamic fluorescence channel.

We next asked whether some subcellular regions of astrocytes have more dynamic activity than others across all analyzed cells (n=11 cells). Although we detected more static events than dynamic overall (Supp. Fig. 8a), we observed a higher proportion of dynamic events than static events in the soma (59%, Fig. 3c, Supp. Fig. 8b). We then characterized events by propagation direction and event initiation location (Fig. 3d). Events that begin close to the soma (≤ 50 th percentile) and propagate away (purple) were on average larger than the events propagating toward the soma (pink, two-tailed t-test). Similarly, those events that began close to the soma (≤ 50 th percentile) and propagated away had on average a longer duration than events propagating toward the soma (two-tailed t-test, Fig. 3e, Supp. Fig. 8).

One of AQuA's strengths is its ability to automatically extract a large number of features. These features can be used to form a comprehensive Ca^{2+} measurement matrix, where each row represents an event and each column an extracted feature, and which includes all events for each cell (Supp. Fig. 9). Dimensionality reduction applied to this matrix can, in turn, be used to visualize each cell's Ca^{2+} signature (Supp. Fig. 9, white rows separate each individual cell). To do this, we applied t-distributed Stochastic Neighbor Embedding (t-SNE)²⁴, followed by k-means clustering to assign the cells to groups (Supp. Fig. 9), revealing clusters marked by cells with large differences in median frequency (Fig. 3f). Astrocytic Ca^{2+} frequency is commonly measured as the number of transients that occur over time within an ROI. Here, we instead define frequency from an event-based perspective in two ways: 1) for each event, the number of other events that overlap in time, and 2) for each event, the number of other events that overlap in space. We used these two measures (temporal and spatial overlap) and several other extracted measures (Supp. Fig. 9) to construct the matrix used for t-SNE visualization and clustering. We next tested how well our AQuA-specific features perform at clustering the heterogeneity among cells compared to two ROI-based methods (Fig. 3g), and found that the AQuA-based method outperformed the others. In fact, even when we only use AQuA-specific features for this analysis—area, temporal overlap, spatial overlap, and propagation speed—and remove all features that can be extracted from ROI-based methods, AQuA still significantly outperforms in clustering cells (Supp. Fig. 9g–i). AQuA-extracted features that correspond only to those that can be obtained by ROI-based methods—standard frequency, amplitude, duration—do not allow clustering significantly better than the ROI-based approaches themselves (Supp. Fig. 9g–i), suggesting that the AQuA-specific features are those that best capture dynamic fluorescence features that vary among single cells. This indicates that AQuA may be used to extract data from existing *ex vivo* Ca^{2+} imaging datasets to reveal previously uncovered dynamics and sort cells into functionally relevant clusters.

In vivo astrocytic Ca^{2+} bursts display anatomical directionality

Recent interest in astrocytic activity at the mesoscale has been driven by population-level, multicellular astrocytic Ca^{2+} imaging^{1,5,7,8,25–27}. To test the power of AQuA-based event detection, we next applied it to populations of *in vivo* astrocyte Ca^{2+} activity. Previous studies have described temporal details of astrocyte activation^{4,5,7,8,25}, yet have left largely unaddressed the combined spatiotemporal properties of Ca^{2+} activity at the circuit-level, across multiple cells. Here, we explored whether AQuA can uncover spatial patterns within populations of cortical astrocytes in an awake animal, and carried out head-fixed, 2P imaging of GCaMP6f activity in V1, L2/3 astrocytes. Populations of *in vivo* cortical astrocytes exhibit both small, focal, desynchronized Ca^{2+} activity²⁵, and large, coordinated activities^{4,5} that we refer to as bursts in this context. Importantly, AQuA detected both of these classes of Ca^{2+} activity in

the same *in vivo* imaging datasets, suggesting that it could be a powerful tool to investigate both kinds of activity—and the interactions between them—even within a single dataset (Movie 2, Fig. 4a). Similar to previous studies, we observed many (but not all) of the bursts co-occurring with locomotion periods (Fig. 4b, pink), and many events within these burst periods displayed propagation (Fig. 4c, top). These propagative events were larger in area and had greater propagation distances compared to the events that occurred during the inter-burst periods (Fig. 4c, bottom). Here, to test whether AQuA could help us understand these large bursts and discover discrete features of this phenomenon, we next focused our investigation on all the events occurring during the burst periods (Supp. Fig. 10).

To analyze the structure of these burst-period Ca^{2+} events, we investigated fluorescence propagation across multiple spatial scales: at the level of individual events, of subregions of the imaging field encompassing multiple events, and of the entire imaging field. At the level of individual events within a single burst, plotting the individual event direction within the entire field of view did not reveal a consistent propagation direction (Fig. 4d). However, when we divided our field-of-view into equivalently sized, subregional tiles (Fig. 4e), we observed more consistent propagation direction within single subregions (Fig. 4f). When we plot the cumulative count of the percentage of bursts with regions that propagate in the same direction, we indeed observe that this curve is right-shifted compared to a simulated random assignment of majority regional propagation direction (Fig. 4g), suggesting that there does exist regularity in the propagation pattern within bursts, but that this only becomes apparent at a spatial scale larger than individual events.

Since the consistency of propagation directionality increased with increasing spatial scales, we next explored whole imaging field dynamics during Ca^{2+} burst events. We noted that the percentage of the active field of view varied across burst periods (Fig. 4b), with a wide variability from few to hundreds of events (Fig. 4h). To control for number and size of events, we used the difference between each event's onset time to calculate a single burst-wide propagation direction (Fig. 4h, black arrow). Doing so revealed a consistent posterior-medial directionality of population Ca^{2+} activity in L2/3 V1 astrocytes (Fig. 4i). Although Ca^{2+} bursts have been previously observed using GCaMP6 imaging in awake mice^{4,5}, consistent spatial directionality with respect to the underlying anatomy has never been described. This observed posterior-medial directionality may be revealing anatomical and physiological underpinnings of these bursts, and since they have been shown to be at least partly mediated by norepinephrine^{5,7}, they could be reflective of the response of groups of cortical astrocytes to incoming adrenergic axons originating in locus coeruleus. Regardless of burst mechanism(s), these results suggest that *in vivo*, astrocytic Ca^{2+} propagation dynamics differ depending on the spatial scale examined, which may explain previously described discrepancies in dynamics.

AQuA-based detection of extracellular molecular dynamics

We next asked whether AQuA could be used to detect astrocytic fluorescent activities with distinct spatiotemporal dynamics than we observe when measuring intracellular Ca^{2+} . We decided to perform imaging of extracellular-facing probes, including GluSnFR², to measure extracellular glutamate dynamics, since it has been widely used for glutamate imaging^{2,6,8} and astrocytes regulate extracellular glutamate concentration. In addition, GluSnFR dynamics are much faster than GCaMP dynamics, which causes detection to be very susceptible to low SNR. This can be an additional challenge and thus much previous GluSnFR analyses has relied on averaging across multiple trials. While GluSnFR has been expressed both in astrocytes and in neurons previously^{2,8,16,28}, how cell type-specific expression and morphology—particularly relative to synaptic and extra-synaptic glutamate release—determines its fluorescent dynamics has not been fully explored^{16,28}. No previously applied analytical tools have been reported to automatically detect GluSnFR-based glutamate events to accommodate differential event sizes and shapes. Here, we explored whether application of AQuA could be used to detect cell type-specific differences in glutamate dynamics that may be based on heterogeneous underlying morphologies and cell biological mechanisms.

We expressed GluSnFR in either astrocytes or neurons using cell type-specific viruses² and carried out 2P imaging of spontaneous GluSnFR activity in acute cortical V1 slices from L2/3. Distinct morphological differences between astrocytic and neuronal expression of GluSnFR were evident, as has been observed previously^{8,29,30} (Fig. 5a). We applied AQuA to these datasets to detect significant increases in GluSnFR fluorescence, and were able to detect events that were too small and dim to detect by eye (Movie 3); AQuA-detected events were confirmed by *post hoc* ROI-based analysis. Indeed, 62% of astrocytic events had an area less than the size of a single astrocyte ($100 \mu\text{m}^2$), and 8% of astrocytic and 35% of neuronal glutamate events had a small maximum dF/F (less than 0.5). Because GluSnFR events have previously been detected by spatially averaging within a single cell or across broader areas of tissue, *or* by manual detection, the events that AQuA detects are most likely missed by ROI-based methods^{6,8,16} (Supp. Fig. 11). Because AQuA is designed to detect events independently from shape or size, events of heterogeneous size and shape were revealed during this analysis (Fig. 5a–b). A large proportion of these spontaneous GluSnFR events changed size over the course of the event, with 42% of total astrocytic and 32% of total neuronal glutamate events exhibiting changes in area. On average, astrocytic GluSnFR events were significantly larger ($274 \pm 39.56 \mu\text{m}^2$) than neuronal events ($172 \pm 57.06 \mu\text{m}^2$), sometimes encompassing an entire astrocyte (Supp. Fig. 11). Neuronal GluSnFR events were significantly more circular (Figure 5b–d), perhaps reflecting morphological differences between cell type somata. We also found that between cell types, GluSnFR events exhibited different size dynamics (Fig. 5b–c). While there was no difference in the rate of increase in event size between astrocytes and neurons,

we did observe that the rate of size decrease of astrocytic events between frames was larger than that of neuronal events (Fig. 5c), which may reflect differential synaptic and extrasynaptic glutamate dynamics in proximity to subcellular compartments of each cell type.

Once we found that AQuA-based detection was effective for quantification of spontaneous GluSnFR events, we wanted to test its performance on more spatially and temporally precise glutamate events, since GluSnFR can be used to measure synaptic release of glutamate when imaged at fast frame rates^{2,31}. To do this, we performed fast (~100Hz) GluSnFR imaging while photoactivating a caged glutamate compound (RuBi-glutamate^{10,32}) with a second laser beam. In these experiments, we tested uncaging pulses at various durations (25–150ms), and applied AQuA to detect these small-scale, fast events (Fig. 5d, right). Although we did observe an increase in detection accuracy (identification of event at the time and location of the laser uncaging pulse) of uncaging-defined GluSnFR events at longer uncaging durations, AQuA detection still showed high accuracy levels at shorter durations, with a minimum of 96% average accuracy across durations (Fig. 5d, right; n = 5 cells, 3 replicates/cell). These results indicate that AQuA works well for fluorescent event detection at fast frame rates.

We lastly wanted to demonstrate that AQuA can be used for other extracellular-facing probes that are relevant for astrocyte-neuron physiology. To do this, we imaged and analyzed two recently developed genetically encoded probes that sense extracellular neurotransmitters: GABASnFR¹¹ and GRAB-NE¹², which report GABA and NE dynamics, respectively. We expressed GABASnFR (Fig. 5e, left) in cortical astrocytes and GRAB-NE in cortical neurons (Fig 5f, left), and performed *ex vivo* 2P imaging before and after bath application of either NE or GABA. In both cases, we used AQuA to detect events as the neurotransmitter contacted the fluorescent sensors. For GABASnFR expression in individual astrocytes, we observed that each event increased in both fluorescence amplitude and area with cell-specific dynamics (Fig. 5e, right). While the widespread neuronal expression of GRAB-NE did not allow for cell-specific analysis, it did allow us to observe waves of NE as it was bathed over the slice. Here, AQuA enabled detection of the dynamic spatial location, amplitude, and area of these waves as they progressed across the slice (Fig 5f, right), indicating that AQuA may be useful to quantify propagating wavefronts in other contexts. Together, results in this section suggest that AQuA-based detection can be used to quantify the dynamics of extracellular molecules at a range of speeds and spatial spreads, across multiple cell types and expression patterns.

Discussion

With the development and application of a flexible event-based analysis tool for astrocyte imaging datasets, we hope to enable many research groups to accurately quantify observed fluorescence dynamics, including those that are un-fixed, propagative, and size varying. Here, we demonstrate that AQuA

performs better than many other image analysis methods—including those designed for astrocytic and neuronal applications—on simulated datasets, and describe fluorescent event detection in several types of datasets, using the genetically encoded GCaMP, GluSnFR, GABASnFR, and GRAB-NE indicators. Because AQuA is data-driven, it can be applied to datasets that have not been directly tested here, including those captured under different imaging magnifications and spatial resolutions, as well as confocal or wide-field imaging systems. In addition, since the AQuA algorithm functions independently from frame rate, datasets captured with faster or slower frame rates^{17,25} are also just as amenable to an event-based analysis with AQuA as those shown here. Further, AQuA is applicable to fluorescent indicators other than the ones tested here, particularly those that exhibit complex dynamics.

We envision the AQuA software and its underlying algorithm as enabling problem-solving for a wide range of astrocyte physiological questions, both because AQuA accurately captures dynamics exhibited by commonly used fluorescent indicators and because there are more features extracted by AQuA that can be analyzed than those extracted by existing methods. Since AQuA-specific features were able to capture heterogeneities we observed among single cells when imaging Ca^{2+} , we posit that these features may be more physiologically relevant than the standard measurements (amplitude, frequency, duration) used to describe astrocytic physiological differences, although these standard measures are also features extracted by AQuA. In the current work, we use these multiple features to describe the spontaneous astrocyte activity in cortex, but with varying spatial scale, fluorescent probe, and experimental preparation. In future work, we and others can apply AQuA-based analyses to other brain regions and layers to describe potential functional heterogeneities among astrocytes³³. Beyond baseline differences, we expect that AQuA will be a powerful tool to quantify physiological effects of pharmacological, genetic, and optogenetic manipulations, among others. These manipulations and subsequent analyses would allow researchers to examine both astrocyte-intrinsic and -extrinsic physiology, depending on whether astrocytes, neurons, or another brain cell type is being changed.

Significant disagreement in the field remains about basic physiological functions of astrocytes. Perhaps the most outstanding issue is whether astrocytes undergo vesicular release of transmitters such as glutamate. While we don't address this controversial topic in the current work, we expect that the heterogeneous activities that we uncover using an AQuA-based analysis of GluSnFR may be key in determining different sources of glutamate in neural circuits under different conditions, and could help untangle some of the conflicting data in this arena. Our tool enabled us to identify extracellular glutamate changes not only by cell type, but also by event size and shape dynamics, demonstrating an in-depth analysis of GluSnFR data. The event-based analytical tools presented here may be particularly useful as next-generation GluSnFR variants become available and make multiplexed imaging experiments increasingly accessible³¹.

When surveying astrocyte Ca^{2+} imaging data, experimental regimes can largely be grouped into two categories: single-cell, usually *ex vivo* imaging and population-wide, *in vivo* imaging focusing on large-scale activity of many cells. Experimental data and neurobiological conclusions from these two groups can differ quite widely, or even conflict. This may be due, in part, to the large, population-wide bursts observed with the onset of locomotion *in vivo*. Many techniques used to analyze these bursting events—all ROI-based—can under-sample events that occur between bursts by swamping out smaller or shorter signals. Here, we present a technique that can be used to sample small- and large-scale activity in the same dataset or across datasets, in order to bridge spatiotemporal scales these datasets. As such, this event-based analysis tool has the potential to aid researchers in resolving outstanding physiological problems, while also tackling new ones.

As demonstrated by its utility with Ca^{2+} , glutamate, GABA, and NE datasets, AQuA also has the potential to be applied to many other fluorescence imaging datasets that exhibit non-static or propagative activity. Although we designed AQuA specifically to study dynamic astrocyte fluorescence, it is open-source and user-tunable, and we anticipate that experimentalists will find it advantageous in other contexts in which neuronal or non-neuronal cells exhibit non-static or propagative fluorescence activity. For example, recently described Ca^{2+} activity in oligodendrocytes displays some similar properties to that in astrocytes^{34,35} and AQuA-based analysis may be useful. Likewise, subcellular compartments in neurons, such as dendrites or dendritic spines, have also been shown to exhibit propagative, wave-like Ca^{2+} signals³⁶ and large-scale, whole-brain neuronal imaging can capture burst-like, population-wide events³⁷.

While we predict that the potential applications are wide, it is also important to note the limitations of AQuA, and be clear about when it will not be the most effective approach. Since AQuA detects local fluorescence changes as events, it is not well suited to strictly morphological dynamics, such as those observed in microglia, and it does not improve on the many excellent tools built for analyzing somatic neuronal Ca^{2+} activity^{20,21}, where ROI assumptions are well satisfied. In addition, AQuA was optimized for and tested on 2D datasets, as these comprise the majority of current imaging experiments in the field. As techniques for volumetric imaging rapidly advance, an extension to accommodate 3D imaging experiments will be necessary. We expect that AQuA is expandable to 3D datasets, based on the fact that the algorithmic design is not restricted to 2D assumptions. Although developing a full version of a 3D extension is beyond the scope of this paper, we have built a 3D prototype to test the feasibility of extending events from 2D to 3D (Supp. Fig. 12). Using a simulated 3D dataset based on published¹⁷ astrocytes, the prototype performed well on 3D data, including detection of various event sizes and signal propagation rates (Supp. Fig. 12). These results suggest that a full 3D AQuA extension—including the optimization of computational efficiency and visualization—will work on real 3D datasets in the future.

In addition, the results demonstrate that AQuA is a flexible and robust platform that can accommodate new types of data without large changes to the underlying algorithm.

Acknowledgements

The authors acknowledge members of the Poskanzer and Yu labs for helpful discussions and comments on the manuscript. We thank Gregory Chin and Sae Yokoyama for excellent technical assistance. K.E.P. is supported by NIH R01NS099254, NIH R21DA048497, NSF 1604544, Brain Research Foundation Frank/Fay Seed Grant, E. Matilda Ziegler Foundation for the Blind, and the Bold & Basic grant from the QBI at UCSF. G.Y. is supported by NIH R01MH110504 and NSF 1750931.

Author Contributions

K.E.P and G.Y. conceived and designed the study. Y.W. and G.Y. designed and implemented the AQuA algorithm, software, and simulations. N.V.D. analyzed all imaging data and provided critical conceptual input to software design. T.V. carried out GluSnFR experiments; M.K.C. performed *ex vivo* Ca^{2+} , GABASnFR, and GluSnFR experiments; M.R. performed *in vivo* Ca^{2+} imaging experiments; and S.P. performed GRAB-NE experiments. X.M. implemented the Java version of the software, and built and tested the 3D prototype. K.E.P supervised the experimental team. G.Y. supervised the computational team. Y.W., N.V.D., K.E.P, and G.Y. wrote the manuscript with input from all authors.

Methods

Viral injections and surgical procedures

For slice experiments, neonatal mice (Swiss Webster, P0–P4) were anesthetized by crushed ice anesthesia for 3 minutes and injected with 90nL total virus of *AAV5-GFaABC1D.Lck-GCaMP6f*, *AAV5-GFaABC1D.cyto-GCaMP6f*, *AAV1-GFAP-iGluSnFR*, *AAV1-hsyn-iGluSnFR*, *AAV2-GFAP-iGABASnFR.F102G*, and *AAV9-hsyn-NE2.1* at a rate of 2–3nL/sec. Six injections 0.5 μm apart in a 2x3 grid pattern with 15nL/injection into assumed V1 were performed 0.2 μm below pial surface using a UMP-3 microsyringe pump (World Precision Instruments). Mice were used for slice imaging experiments at P10–P23.

For *in vivo* experiments, adult mice (C57Bl/6, P50–P100) were given dexamethasone (5mg/kg) subcutaneously prior to surgery and then anesthetized under isoflurane. A titanium headplate was attached to the skull using C&B Metabond (Parkell) and a 3mm diameter craniotomy was cut over the right hemisphere ensuring access to visual cortex. Two 300nL injections (600nL total virus) of *AAV5-GFaABC1D.cyto-GCaMP6f* were made into visual cortex (0.5–1.0mm anterior and 1.75–2.5mm lateral of bregma) at a depth of 0.2–0.3mm and 0.5mm from the pial surface, respectively. Virus was injected at a

rate of 2nL/s, with a 10min wait following each injection to allow for diffusion. Following viral injection, a glass cranial window was implanted to allow for chronic imaging and secured using C&B metabond³⁸. Mice were given at least ten days to recover, followed by habituation for three days to head fixation on a circular treadmill, prior to imaging.

Two-photon imaging

All 2P imaging experiments were carried out on a microscope (Bruker Ultima IV) equipped with a Ti:Sa laser (MaiTai, SpectraPhysics). The laser beam was intensity-modulated using a Pockels cell (Conoptics) and scanned with linear or resonant galvonometers. Images were acquired with a 16x, 0.8 N.A. (Nikon, *in vivo* GCaMP and *ex vivo* GRAB-NE) or 40x, 0.8. N.A. objective (Nikon, *ex vivo* GCaMP, GluSnFR, and GABASnFR) via a photomultiplier tube (Hamamatsu) using PrairieView (Bruker) software. For imaging, 950nm (GCaMP), 910nm (GluSnFR and GABASnFR), or 920nm (GRAB-NE) excitation and a 515/30 emission filter was used.

Ex vivo imaging

Coronal, acute neocortical slices (400 μ m thick) from P10–P23 mice were cut with a vibratome (VT 1200, Leica) in ice-cold cutting solution (in mM): 27 NaHCO₃, 1.5 NaH₂PO₄, 222 sucrose, 2.6 KCl, 2 MgSO₄, 2 CaCl₂. Slices were incubated in standard continuously aerated (95% O₂/5% CO₂) artificial cerebrospinal fluid (ACSF) containing (in mM): 123 NaCl, 26 NaHCO₃, 1 NaH₂PO₄, 10 dextrose, 3 KCl, 2 CaCl₂, 2 MgSO₄, heated to 37°C and removed from water bath immediately before introducing slices. Slices were held in ACSF at room temperature until imaging. Experiments were performed in continuously aerated, standard ACSF. 2P scanning for all probes was carried out at 512x512 pixel resolution. Acquisition frame rates were 1.1Hz (GCaMP), 4–100Hz (GluSnFR), 6Hz (GABASnFR), and 1.4Hz (GRAB-NE). For GluSnFR imaging and RuBi-glutamate uncaging experiments, GluSnFR imaging was performed at 950nm excitation to ensure that no RuBi-glutamate was released during scanning. Acquisition rates were between 95–100Hz, using resonant galvonometers. 300 μ M RuBi-glutamate was added to the circulating ACSF and using a second MaiTai laser tuned to 800nm, five uncaging points were successively uncaged at each cell at durations indicated in the figure and at power <3mW that were shown in control experiments to cause no direct cell activation.

In vivo GCaMP imaging

At least two weeks following surgery mice were head-fixed to a circular treadmill and astrocyte calcium activity was visualized at ~2Hz effective frame rate from layers 2/3 of visual cortex with a

512x512 pixel resolution at 0.8 microns/pixel. Locomotion speed was monitored using an optoswitch (Newark Element 14) connected to an Arduino.

AQuA algorithm and event detection

Overview of the AQuA algorithm

Astrocytic events are heterogeneous and varying with respect to many aspects of their properties. In AQuA, we extensively applied machine learning techniques to flexibly model these events, so that our approach is data-driven and physiologically relevant parameters are extracted from the data instead of imposing *a priori* assumptions. Probability theory and numerical optimization techniques were applied to optimally extract fluorescent signals from background fluctuations. Here, we delineate the eight major steps in AQuA (Supp. Fig. 1), discuss motivations behind the algorithm design, and describe key technical considerations in further detail.

Step 1: data normalization and preprocessing. This step removes experimental artifacts such as motion effects, and processes the data so that noise can be well approximated by a standard Gaussian distribution. Particular attention is paid to the variance stabilization, estimate of baseline fluorescence, and variance. Step 2: detect active voxels. Step 3: identify seeds for peak detection. Step 4: detect peaks and their spatiotemporal extension. These three steps work together to achieve peak detection. To detect peaks we start from a seed, which is modeled as a spatiotemporal local maximum. However, since random fluctuations due to background noise can also result in local maxima, we need to detect active voxels such that only the local maxima on the active voxels are considered as seeds. Here, active voxels are those likely to have signals. Step 5: cluster peaks to identify candidates for super-events. Temporarily ignoring the single-source requirement, the set of spatially-adjacent and temporally-close peaks is defined as a super-event. However, clustering results of spatially adjacent peaks are not super-events themselves, because a peak group may consist of noise voxels and temporally distant events. Step 6: estimate the signal propagation patterns. Step 7: Detect super-events. To get super-events from peak clusters, we compute the temporal closeness between spatially adjacent peaks by estimating signal propagation patterns. The propagation pattern for each event is also important for its own sake, by providing a new way to quantify activity patterns. Step 8: split super-event into individual events with different sources. A super-event is split into individual events by further exploiting propagation patterns. Based on propagation patterns within a super-event, the locations of event initiation are identified as local minima of the onset time map. Each initiation location serves as the event seed. Individual events are obtained by assigning each pixel to an event based on spatial connectivity and temporal similarity.

Step 1 (Data normalization and preprocessing): We correct for motion artifacts in the *in vivo* dataset using standard image registration techniques³⁹ before applying AQuA. However, AQuA does not necessarily require motion correction because it performs event-based analysis, which is localized temporally and thus less prone to motion artifacts.

We perform data normalization and preprocessing to approximate noise by a Gaussian distribution with mean=0 and standard deviation=1. For those who want to directly modify our code, pseudocode is available (Supp. Table 1). To achieve this normalization, we first apply a square-root transformation to the data to ensure that the noise variance after transformation is not related to the intensity itself, an operation also known as variance stabilization. Variance stabilization is important so that events with bright signals are not over-emphasized and events with dim signals not neglected. Next, the noise variance of the transformed data for each pixel is estimated as half of the median of the square of differences between two adjacent values in the time series at the pixel. Mathematically, denote X_i the time series at the i th pixel, where $X_i[t]$ is the value of the t th time point. Then, the noise variance σ_i^2 at the i th pixel is estimated as

$$\sigma_i^2 = \frac{1}{2} \text{median}_{t=2, \dots, T} \{(X_i[t] - X_i[t-1])^2\}.$$

We do not use the conventional sample variance $\frac{1}{T} \sum_{t=1}^T (X_i[t] - \text{mean}(X_i))^2$ as the estimate, which tends to inflate the variance when signals exist in the time series. Third, to estimate the baseline fluorescence F_0 for each pixel, we compute the minimum of the moving average of 25 time-points in a user-specified local time window (default=200 in our experiments). We do not use the full time series to identify the minimum, in order to be robust to image degradation or other long-term trends. Considering the minimum is a biased estimate of the baseline fluorescence, we add a pre-determined quantity to the minimum to serve as the estimate of F_0 . Here, the pre-determined quantity depends on the extent of the moving average and the size of the time window, and is found through simulation. Denote $V_i[t]$ as the value of i th pixel at the t th frame in the raw video data. In the following, all modeling is performed on the normalized data,

$$Z_i[t] = \frac{\sqrt{V_i[t]} - \sqrt{F_{i0}}}{\sigma_i},$$

where the subscript i denotes that the baseline fluorescence and noise variance are location- and pixel-specific. It is worth noting that the data normalization and preprocessing is used to build an accurate statistical model so that signals can be reliably distinguished from noise. When we extract activity-related features such as rise and decay time, raw data is used.

Step 2 (Detecting active voxels): A voxel is defined as a pixel of a certain frame. For example, voxel (x, y, t) denotes the pixel at location (x, y) in the t th frame in the movie. An active voxel is the voxel that

contains an activity signal. If a voxel is not associated with any event, it is not considered active. Since an event often occupies multiple pixels and extends several frames, we first apply 3D Gaussian filtering to smooth the data to reduce the impact of noise. Then, we calculate the z-scores for each voxel in the smoothed data. Here, z-score is computed as the value of the voxel divided by its standard deviation, which can be estimated as in the normalization procedure above, but now on the smoothed data. All voxels that have z-scores larger than a given threshold are considered tentative active voxels. A liberal threshold is used here to retain most signals, often at a z-score of 3. We next calculate the size of groups of connected tentative active voxels, with spatially connected tentative active belonging to the same group and a minimum size threshold (often 4). If a group of tentative active voxels is less than the threshold, all voxels in this group are removed, resulting in a final list of active voxels. Pseudocode is presented in Supp. Table 1.

Our event detection can be roughly described as finding bumps in space-time, and it could be asked why we don't simply set a threshold on the ΔF data and identify all components larger than the threshold as events. This method is essentially equivalent to our Step 2 (*Detecting active voxels*), but there are motivations to go beyond this step. Even if simple thresholding detected individual events perfectly, many features of each event—such as propagation patterns—cannot be directly derived. More importantly, no threshold we have found can faithfully detect real events (Supp. Fig. 3). Three types of commonly observed complexities lead to this insufficiency in thresholding (Supp. Fig 3b). First, not all pixels drop below the threshold because of residual signals. Second, neighboring events occurring at different times can be spatiotemporally connected. Third, two events can be temporally separated when they appear, but meet with each other after propagation.

Step 3 (Identifying seeds for peak detection): Similar to the detection of active voxels, we apply 3D Gaussian smoothing to the normalized data and then find all local maxima, defined as connected components of pixels with a constant intensity value and with all neighboring pixels having a lower value. Considering our time-lapse images as three-dimensional arrays (2D space plus 1D time), each single pixel has 26 neighbors. Although each local maximum generally occupies one pixel due to random fluctuation inherent in the data, this definition allows a local maximum to occupy multiple pixels of the same intensity value. This is helpful for the case in which some pixels have saturated values. Because pure random fluctuation can also lead to local maxima, we restrict the search of local maxima to active voxels only (see pseudocode in Supp. Table 1). The resultant local maxima are considered seeds for the purpose of peak detection, the subsequent step in the algorithm. We start peak detection by identifying the local maxima as seeds because the maxima are likely to contain the strongest signal and thus have a better SNR

than other points. The 3D Gaussian smoothing is used to further improve SNR, motivated by the fact that an event occupies multiple pixels and spans multiple time points.

Step 4 (Detecting peaks and their spatiotemporal extent): We partially and temporally extend each seed detected above to all voxels that are potentially associated with each event. We call the collection of the seed and its extended voxels the *super voxel*. Seeds are processed one-by-one, with higher intensity seeds processed first. Each seed is first extended temporally, then spatially.

The spatiotemporal index (x_0, y_0, t_0) denotes the seed. When we temporally extend the seed backwards and forwards (Supp. Fig 2b, Supp. Table 2), we encounter two main scenarios. In the first, a voxel before (x_0, y_0, t_0) has a value close to the baseline F_0 , and a voxel after (x_0, y_0, t_0) also is close to F_0 . If a voxel has an intensity $<20\%$ of the seed value, it is defined as close to baseline. In this scenario, the seed is extended temporally until it reaches these two voxels. In the second scenario, extension in at least one of two directions never meets a voxel with value that is considered close to the baseline before meeting another seed. This scenario could happen when one event begins before the previous event drops completely to the baseline fluorescence level. To determine whether these two seeds should be merged and thus belong to the same event, we denote V_{\min} the minimum value between the two seeds and calculate the difference between V_{\min} and value at the seed (x_0, y_0, t_0) . If the difference is larger than the threshold Δ_{tw} , which is $2\sigma_0$ by default for most data, the minimum is considered the end of the extension. Otherwise, these two seeds are merged and the extension continues. For very high peaks, this threshold is too low for perceptually meaningful separation. To split two adjacent high ΔF peaks, a strong decrease between them is needed and the threshold is changed to $\Delta_{tw} = \max(0.3\Delta F(x, y, t_0), 2\sigma_0)$.

Once each seed i is temporally extended from (x_i, y_i, t_i) to a peak $(x_i, y_i, (t_i - a_i):(t_i + b_i))$, $(t_i - a_i):(t_i + b_i)$ denotes a time window spanning from $t_i - a_i$ to $t_i + b_i$. We define reference curve c_i as the average of nine pixels around (x_i, y_i) in that time window. To spatially extend each peak to cover most signals-of-interest, each seed becomes a set of voxels $\bigcup_{k=1}^K (x_{ik}, y_{ik}, (t_i - a_i):(t_i + b_i))$ after extension, where the corresponding spatial footprint, K pixels $\{(x_{ik}, y_{ik}), k = 1 \dots K\}$, is spatially connected. During this process, each seed is associated with two sets. The first is Ω_i , which are pixels already associated with seed i . Each pixel in this set, (x_{i1}, y_{i1}) , e. g., corresponds to a set of voxels: $(x_{i1}, y_{i1}, t_i - a_i: t_i + b_i)$. The second set is Θ_i , the set of pixels to avoid. Initially, $\Omega_i = \{(x_i, y_i)\}$ and Θ_i is empty.

The spatial extension operation for each seed is repeated a maximum of 40 rounds. For each seed, the set Ω_i is spatially dilated with a 3x3 square, thus only testing pixels adjacent to the Ω_i boundary. Next, Θ_i is removed from the dilated region. We then test whether each new pixel should be added to Ω_i or not. Because for each given new pixel and each time window we have a time series, we can calculate the Pearson correlation coefficient between this time series and c_i . The correlation coefficient is converted to

a z-score using the Fisher transform. If the z-score is higher than the user-defined given threshold, the pixel is added to Ω_i . Otherwise, it is added to Θ_i . Because all seeds are local maxima, no time alignment is needed here.

During the extension process, different super voxels can meet. We want to stop the extension process of one super voxel only when it meets the bright part of other super voxels (50% rising to 50% decaying). For example, we have two peaks from two seeds: $(x_1, y_1, (t_1 - a_1): (t_1 + b_1))$ and $(x_2, y_2, (t_2 - a_2): (t_2 + b_2))$. Assume the first seed has already occupied pixel (x_3, y_3) . When the second seed tries to determine whether it should extend to (x_3, y_3) or not, we calculate whether $(t_1 - a_1: t_1 + b_1)$ and $(t_2 - a_2: t_2 + b_2)$ sufficiently overlap. Two peaks sufficient overlap if the 50% rising to 50% decay ranges of the two peaks overlap. Thus, if $(t_1 - a_1: t_1 + b_1)$ and $(t_2 - a_2: t_2 + b_2)$ sufficiently overlap, seed two will not include pixel (x_3, y_3) and it is added to Θ_2 . Otherwise, it is added to Ω_2 . After spatial extension is complete, we remove super voxels with $\Omega_i < 4$ pixels or total voxels < 8 pixels.

Step 5 (Clustering peaks to identify candidates for super-events): A super-event is defined as a group of events connected spatially but originating from different initiation locations. This step was motivated by the frequent observation in real data that multiple events can be spatially connected at some time point. One example is a large burst in the *in vivo* dataset, where multiple events start at different places but merge as a burst at later stage. Another example is a set of two events originating from different places, propagating and meeting each other in the middle (Supp. Fig. 3). Thus, in a spatial direction, we may encounter multiple events within the super-event. However, we never encounter two or more events in the temporal direction, which guides the following algorithm. To identify candidates for super-events, we next cluster peaks, but these results are not identical to super-events, because voxels extended to be associated with the peak may have some errors. As discussed below, the candidate super-event must be purified to resolve the final super-event.

Since each super-voxel extends from its seed (representative peak), we also call the process of clustering super-voxels as clustering peaks for conceptual convenience. If two super-voxels are connected and their rise-time difference is less than a given threshold (as discussed below), they are considered neighbors. For two super-voxels, if 10% of pixels of either super-voxel is also occupied by the other, they are a conflicting pair. For each super-voxel, we list all its neighbors and conflicting counterparts. To cluster peaks/super-voxels (Supp. Fig. 2b, Supp. Table 3), we begin with the earliest occurring super-voxel and check each of its neighbors. If a neighbor is not conflicting with that super-voxel, it is combined with the super-voxel. This process is repeated until no new super-voxel can be added. Then we move to the next earliest super-voxel that is not added to any others, and repeat this process. An iterative approach prioritizes events that are close to each other. Supposing the largest rising time difference for

super-voxels that is allowed to be neighbors is 10, we start the procedure with the allowed difference as 0 and merge the super voxels. Then we increment the allowed difference by 1 and repeat the step above, until the rising time difference allowed reaches 10.

Step 6 (Estimating signal propagation patterns): For each spatial location/pixel, an associated time series indicates the signal dynamics. Estimation of propagation patterns is formulated as a mathematical problem of time alignment between the time series at each location and a representative/reference time series. Time alignment results directly relay the delay of a given pixel at a time frame with respect to the representative dynamics. Conventionally, time alignment is accomplished by dynamic time warping (DTW)⁴⁰. However, DTW is notoriously prone to noise, which leads to unreliable propagation estimation. Since two adjacent pixels have more similar propagation patterns than two distant pixels, we impose a smoothness constraint on neighboring pixels using our recently developed mathematical model—Graphical Time Warping (GTW)—to explicitly incorporate this constraint¹⁹. However, since we do not have a representative time series at the very beginning, our strategy is to guess a reference time series from the data and align time series at each pixel to this reference. Then, we use alignment results to obtain an updated reference, and iterate the process of alignment and update of reference until it converges (Supp. Fig. 2b, Supp. Table 3).

To initialize the reference time series, we search for the voxel with the largest $\Delta F/F$ value and record that voxel's location. The initial reference is then estimated as the average time series of the pixels in the 5x5 square around that location, with the square size a user-tunable parameter. The voxel with the largest $\Delta F/F$ value is used because it has the best signal-to-noise ratio. We do not use the time series at a single pixel to initiate the reference because it is noisy, nor do we use the average time series over all the pixels, because the average would be a large distortion to the representative dynamics due to signal propagation. Next, we supply the neighborhood graph and the reference time series to GTW to calculate the time alignment between all pixels and the reference. For each pixel, we consider the 8 pixels around the 3x3 grid as neighbors. A GTW parameter controls the balance between fitness and smoothness of the alignment. We empirically found 1 to be a good value. To control computational complexity, GTW has another parameter corresponding to the maximum time delay allowed. In all our experiments, we found no time delay induced by propagation is larger than 11 frames. So, we set that parameter to 11.

Step 7 (Detecting super-events): Once the time alignment between the representative dynamics and the time series at each pixel is obtained, we refine candidate super-events to obtain final super-events. A super-event is defined as detected when the representative dynamics and all voxels are associated with the super-event. Since each voxel is jointly specified by spatial location and time frame, we next determine

which pixels and time frames jointly belong to the super-event. Since representative dynamics are already obtained in the previous step of propagation estimation, here we focus on how to determine which pixels and which time frames are covered by the super-event.

Because each pixel corresponds to a time series, if a pixel belongs to a super-event, its time series should be highly correlated to the representative dynamics of the super-event. Note that the correlation is calculated based on the aligned time series to account for the time distortion due to signal propagation. Thus, we first obtain a new time series for each pixel based on the time alignment obtained previously. Then, we calculate the Pearson correlation between each new time series and the representative dynamics, leading to a correlation map. We further convert the Pearson correlation to z-score using Fisher's transform. Here, we do not use a threshold for each z-score to determine whether that pixel is statistically significantly associated with the super-event because that ignores the neighborhood information in the correlation map and is less statistically powerful. Instead, incorporating the information from the neighboring pixels, we apply our recently developed order-statistics-based region-growing method to determine which pixels should be associated with the super-event (Supp. Fig. 2b, Supp. Table 3)⁴¹.

To determine which time frames are associated with the super-event, we now examine the representative time series, calculating the maximum intensity along the curve and considering all time frames with intensity $>10\%$ of the maximum to be associated with the super-event. Different pixels may have different time frames associated with the super-event. We use the time alignment results above to identify the time frames associated with the super-event for each pixel. A time frame at a given pixel is associated with the super-event as long as its corresponding time frame in the representative curve is associated with the super-event.

Step 8 (Splitting super-events into individual events with different sources): For each super-event, we have a 2D map of rise-time for each pixel by re-aligning the super-event using GTW. The local minima in this map are potential originating locations for events in this super-event. However, noise may produce random local minima, which do not correspond to true originating locations and are removed by merging with spatially adjacent local minima. We use rise-time to determine whether two local minima should be merged. This idea can be illustrated with the following 1D example: [1 2 4 2 2]. The two local minima are the first and the last pixel (pixel i and j , respectively), occurring at time 1 (t_0) and 2 (t_1), respectively. To determine whether they should be merged, we find all paths connecting them. In this example, there is only one path and the pixel with the latest rise-time in this path is the third pixel (rise-time=4 (t_m)). The distance between pixel i and pixel j induced by this path is therefore defined as $\max(t_m - t_0, t_m - t_1)$. If the distance induced by any path is less than the given threshold, these two local minima are merged.

We next separate super-events into individual events by simultaneously extending all remaining local minima. Each remaining local minimum corresponds to one event. Pixels attached to a local minimum are defined as growing. With each iteration, we add the earliest-occurring pixel to a growing event. If the pixel under examination is adjacent to a growing event, it is done, and then we find the next earliest occurring pixels. Otherwise, we add it to the waitlist and continue with the next earliest occurring one. Each time a pixel is successfully added to a growing event, pixels in the waitlist are checked as to whether they can be added to growing events. When the growing process ends, all individual events are obtained. (see pseudocode in Supp. Table 3).

Run-time analysis of AQuA algorithm: In our implementation, $\sim 60\text{--}70\%$ of the total running time is spent on propagation estimation, and super-event detection and splitting. The super voxel detection step takes about 15% of the total time. Another 7% extracts features from each detected event. The event cleaning and post-processing takes $\sim 6\%$ of the total time, while it takes $<1\%$ of the time to load the data and $<2\%$ in the active voxel step, which finds baseline, estimates noise, and gets seeds. Tested on a desktop computer with Intel Xeon E5-2630 CPU and 128 GB memory using Windows and MATLAB version R2018b, the total running time ranges from several minutes to 1.5 hours, depending on file size and data complexity (Supp. Table 4).

Generation of simulation data sets

Spatial footprint templates: We built a set of templates for event footprints from real *ex vivo* data which serve as the basis for the ROI maps in the subsequent step. Footprints are processed by morphology closing, hole filling, and morphology opening to clean boundaries, with 1683 templates generated total.

ROI maps: 2D ROI maps generated from spatial footprint are used to generate events in subsequent steps. Different simulation types have a different preference for the size of the ROIs. Maximum number of ROIs is set at 100; ROIs are randomly chosen and placed onto a 2D map <5 pixels from existing ROIs.

Simulation dataset 1 (size-varying events): To simulate event size changes, we generate events for each ROI and then alter them to have different sizes so that each ROI in the 2D map will be related to multiple events whose centers are inside that ROI, but whose sizes are different. The degree of size change is characterized by the odds ratio (maximum = 5) between the maximum and the minimum allowable sizes of the events associated with that ROI. For example, with an odds ratio of 2, the size of the event will range from 50–200% of the ROI area. The chances for the event size to be larger or smaller than the area

of the ROI are the same. To achieve this, we generate a random number between 1 and 2, then randomly assign whether to enlarge size by multiplying or shrink by dividing by this factor. Event duration is four frames.

To determine the frames at which the event occurs, we first put the event 10–30 frames (randomly) after the ROI occurs. Spatial distance of this event from others must be ≥ 3 pixels and temporal distance ≥ 4 frames. Part of the event may be inside the spatial footprint of other ROIs, as long as its spatiotemporal distance to other events is larger than the threshold set above. Events are generated for each ROI; on average, we simulate 250 frames with 800 events on 90 ROIs.

Simulation dataset 2 (location-changing events): To simulate event location changes, we generate events with the same size for each ROI and shift them to nearby locations. Thus, each ROI (450–550 pixel size) is related to multiple events near to that ROI. Denote $dist$ the distance between the event center and the ROI center. Denote $diam$ the diameter of the ROI. The degree of location change is quantified by the ratio between $dist$ and $diam$. For example, if we set 0.5 as the maximum degree of location change, the distance of the center of a new event to the ROI will be 0–0.5 times the diameter of the ROI. If the ratio is 0, we simulate a pure ROI dataset. The new event may be located any direction from the ROI, randomly picked from 0 – 2π . Shapes of new events are randomly picked from the templates, so may be different from the ROI while size is constant. Event duration is four frames, and the remaining steps are the same as above. On average, we simulate 250 frames with 800 events on 90 ROIs.

Simulation dataset 3 (propagating events): We simulated two types of propagation: growing and moving, leading to three types of synthetic datasets: growing only, moving only, and mixed. These three types are generated similarly. The ROI map is generated as above, and ROI sizes are 4,000–10,000 pixels, with events generated inside each ROI. In comparison, events in the size-change and location-change simulations can be (fully or partially) outside their corresponding ROIs. We simulate only one seed (starting propagation point) in each ROI. For each event, we generate a rise-time map (for each pixel in the ROI) and construct event-propagation based on the map. We obtain this map by simulating a growing process starting from the seed pixel, with the seed pixel active at the first time-point. At the next time point, its neighboring pixels are active with a variable success probability. Growth continues until $\geq 90\%$ of pixels in the ROI are included in the event. Based on the rise-time map, we identify frames at which pixels become active in the event. To determine when the event ends, we treat growing and moving propagation differently. In growing propagation, all pixels are inactive simultaneously 2 frames after the last pixel becomes active. For moving propagation, the duration is 5 frames. Typically, we generate approximately 140 events in 14 ROIs for each synthetic dataset.

Simulate various SNRs: Gaussian noise is added to the synthetic data to achieve various SNRs. We define the signal intensity as the average of all active pixels in all frames. SNR is defined as

$$20 \times \log_{10} \frac{\text{average signal intensity}}{\text{noise standard deviation}}.$$

When we change the degree of location change, size change, and propagation duration, we add noise with 10 dB SNR. To study the impact of SNR on size changes, size-change degree is 3. For location changes, distance-change ratio is 0.5 while varying SNRs. For propagation, propagation duration is 5 frames. Seven SNRs are tested: 0, 2.5, 5, 7.5, 10, 15, 20 (all in dB).

Post-processing simulated data: We set the average signal intensity at 0.2, with a range from 0–1. Synthetic data is spatially filtered to mimic blurred boundaries in real data. The smoothing is performed with a Gaussian filter with a standard deviation of 1. Signals with intensity <0.05 after smoothing are removed. Remaining signals are temporally filtered with a kernel with a decay τ of 0.6 frames. The rising kernel is linear. For propagation simulation, data is down-sampled by five. Next, we perform a cleaning step. For each pixel in each event, we find the highest intensity (x_peak) across frames. For that pixel, we set signals that are <0.2 times of x_peak to 0. Finally, a uniform background intensity of 0.2 is added (except for GECI-quant, where no background is added; see below).

Application of AQuA and peer methods on the simulation data sets

Based on our knowledge about simulated datasets, we apply specific considerations for each analytical method in order to set optimal parameters for each. In this way, we aim to assess the methodological limit of each method, rather than suboptimal performance due to inadequate parameter-setting. We expect that the performance of the peer methods on simulation data is an overestimate of their performance on real experimental data, because here we take advantage of the ground-truth knowledge, which is not available for experimental astrocyte data.

Event detection using peer methods: AQuA and CaSCaDe report detected events, while other methods report detected ROIs. For a consistent comparison, we detect events from those methods that use ROIs. Once ROIs are detected, we calculate the average dF/F curve for each ROI, as follows: The curve is temporally smoothed with a time-window of 20. The minimum value in the smoothed curve is considered baseline. Assume the minimum value occurs at time t_{base} . The baseline is then subtracted to obtain the dF curve. The noise standard deviation σ is estimated using 40 frames around t_{base} . We then obtain a z-score curve as dF/σ . A large z-score indicates an event; we use a z-score threshold of z_0 . The value z_0 is set according to ground-truth knowledge, so that the smallest-size event in the simulation data is detected

by this threshold. Denote x_0 and s_0 the peak intensity and the size for the smallest event in the ground truth. We also denote the ground truth noise level as σ_0 . Then, the threshold is calculated as,

$$z_0 = \min\left(\frac{0.9x_0\sqrt{s_0}}{\sigma_0}, 10\right).$$

We clip the score to 10 to avoid setting large values for high SNR. For CaSCaDe, we supply this value as the peak intensity threshold parameter.

Using the z-score curves and threshold, we detect events from ROIs for CaImAn, Suite2P, and GECI-quant. For each z-score curve, we find all frames with values $>z_0$. Each frame is a seed for an event. Assume the z-score for that frame is z_1 and we search before and after that frame. If the intensity of the frame is $\geq 0.2z_1$, the frame is associated with the event. If we meet frames with intensities $< 0.2z_1$, we stop searching that direction. Once finished, we obtain all frames associated with the event. We continue with another seed frame to find another event. Note that if a frame is considered part of an event, we do not consider it as a seed for another event, even if it is $>z_0$. The spatial footprint is fixed for all frames in an event, based on the ROI detected. Combining spatial footprint and frames, we obtain events for each ROI and identify all voxels belonging to an event.

Parameter setting for AQuA: The parameters of AQuA are based on the ex-vivo-GCaMP-cyto preset with the following modifications: For different noise levels, we apply different smoothness levels. The smoothing is performed only spatially and values are empirically chosen. The smoothness parameter is the standard deviation of the Gaussian smoothing kernel used.

SNR (dB)	0	2.5	5	7.5	10	15	20
Smoothness	1	0.9	0.8	0.7	0.6	0.5	0.1

We do not simulate motion of the field-of-view, so we do not discard any boundary pixels, and we set regMaskGap = 0. We do not simulate Poisson Gaussian noise; we use additional Gaussian noise only, so PG = 0. Event sizes in the simulation are >200 pixels, so we set the minimum event size to be a value much smaller: minSize = 16. An event may not have more than one peak, so we set cOver = 0. We do not simulate temporally adjacent events, so we set thrTWScl = 4. We do not use proofreading, so we choose a more stringent z-score of events: zThr = 5.

Specific considerations for CaSCaDe: We use the following parameters for CaSCaDe: According to the duration and temporal distances of the simulated events, we can safely set peak distance p.min_peak_dist_ed = 3 and minimum peak length p.min_peak_length = 2. We set the spatial smoothing filter size in the 3D smoothing function (bpass3d_v1) according to the size of the event, so we set p.hb

equal to 2x median of the radius of the spatial footprint of all events. We use this setting because the default settings could not detect larger events on the simulation data sets. For temporal smoothing, we set `p.zhb=21`. We do not need to correct background, so we set `p.int_correct= 0`. The minimum peak intensity is `p.peak_int_ed = z0`, as discussed above. Minimum event intensity is `p.min_int_ed = min(2, p.peak_int_ed *0.2)`. We modified the low-frequency part of the watershed segmentation step to allow larger events to be detected, by changing the function `bpassW` inside the function `domain_segment`. We replaced the noise estimator in CaSCaDe (function `estibkg`) with the more robust one used by AQuA.

CaSCaDe uses a supervised approach to classify detected events. Instead of manually labeling a large number of events and training many SVM models, we directly use ground truth to perform training. For example, for each event detected by CaSCaDe, we check the ground-truth data to test whether it is (part of) a true event. If so, it is retained; otherwise, it is discarded.

Specific considerations for GECI-quant: GECI-quant requires user input at each step. Here, we describe how to automate these steps by taking advantage of ground-truth information. This allows us to test many conditions and repeat many times.

First, we do not add background signals to the synthetic data, so background subtraction is ignored. The domain- and the soma-detection steps require manual thresholding. We estimate the best threshold using the ground-truth data for each simulation. To do so, we scan 255 thresholds and use the one that leads to the best correlation between binarized data and ground truth. We next cleaned the binarized signals with sizes <4 pixels. The data here is also smoothed as it is in GECI-quant (3x3 spatial averaging). Events with spatial footprints < 1,000 pixels are treated as domains and others are treated as somas. The soma segmentation step also uses a threshold. We first process the data as in GECI-quant: for every three frames, a standard-deviation map is calculated so that each voxel in ground-truth data is associated with a standard deviation value. The average of all standard deviations from the ground-truth data is used as the threshold.

We next made the entire analysis pipeline automatic. Fiji is called from the command line in each step and parameters are passed as well. The final ROIs from Fiji are brought back into MATLAB. ROIs are > 15 pixels in area. All other parameters are unchanged, including those for the particle detector. Note that this modification cannot be used as an automated version of GECI-quant for real applications since it relies on ground-truth information.

Specific considerations for CaImAn: We experimented with different parameters for CaImAn and found the following set of parameters performed best on simulation data. As event size can be large, we enlarge the patch size, so `patch_size = [128,128]` and `overlap = [32,32]`. Components to be found is set to `K = 50`.

The standard deviation of the Gaussian kernel (half size of a neuron) is enlarged to tau = 16. Maximum size is 5,000 and the minimum size is 25. Decay time is 0.5. Other parameters are based on default settings. No spatial or temporal down-sampling is used. Adjusting these parameters does not impact results on our simulated data. We used the 5/5/2018 version downloaded from <https://github.com/flatironinstitute/CaImAn-MATLAB>.

Specific consideration for Suite2P: The most critical parameter for Suite2P is neuron size. We set db.diameter equal to the minimum between 50 and the median of the radius of the spatial footprint of all events. Setting the diameter too large leads to an out-of-memory issue. We bypass the registration step. We used the 6/4/18 version downloaded from <https://github.com/cortex-lab/Suite2P>.

Performance evaluation on the simulated data

To evaluate the accuracy of detected events, we quantify the intersection over the union (IoU). We consider all event voxels, not only pixels as in ROI-based methods. For each detected event i , we find all the ground-truth events that have common voxels with event i . For each such ground-truth event, e.g., event j , we calculate an IoU score (also known as Jaccard index) between this pair of events as the following,

$$IoU_{i,j} = \frac{\text{Number(Voxels in event } i \cap \text{Voxels in event } j)}{\text{Number(Voxels in event } i \cup \text{Voxels in event } j)}.$$

When a detected event can be perfectly matched with a ground-truth event, its IoU score is 1. A score of 0 indicates this pair of events has nothing in common. For each detected event i , we find the maximum IoU score among all pairs between this event and a ground-truth event. We denote this maximum score as IoU_i . Similarly, we can compute a score IoU_j for the ground truth event j . The final IoU score is obtained by averaging over all events, including detected and ground-truth events. Supposing we have I detected events and J ground truth events, where I and J are not necessarily equal, we compute the final score as the following,

$$IoU = \frac{\sum_{i=1}^I IoU_i + \sum_{j=1}^J IoU_j}{I+J}.$$

All simulation is performed on a workstation with 16 cores, 128 GB RAM and 6TB hard drive. We use MATLAB 2018a on Windows 10 Enterprise Edition. GECI-quant is run on Fiji with ImageJ version 1.52h. Each simulation is repeated 10 times. The mean and 95% confidence interval (CI) of IoU score is calculated and plotted. The CI is calculated as $[\mu_{sim} - 2\sigma_{sim}, \mu_{sim} + 2\sigma_{sim}]$, where μ_{sim} is the estimated mean and σ_{sim} is the estimated standard deviation (σ_{sim}) based on 10 repetitive runs.

Open-source software for analyzing and visualizing dynamic fluorescent signals in astrocytes.

Applying software engineering principles, we developed an open-source toolbox for astrocyte fluorescent imaging data with detailed user guidelines. The software not only implements the AQuA algorithm for detecting events, but also provides an integrated environment for users to see the results, interact with the analysis, and combine other types of data such as cell/region masks and landmarks. There are two versions of the software with the same functionality, based on MATLAB or Fiji. The software is freely available at <https://github.com/yu-lab-vt/aqua> where detailed documents and example applications can be found. A list of extracted features is shown in Supp. Table 5. Here, we highlight several important functions of the software.

First, the software implements AQuA and provides several options to export the event-detection results, including TIFF files with color-coded events, event features in Excel, and MATLAB or Java data structures to be used by other programs. Second, the software can display analysis results by adding color to the raw video, where color encodes the value of a user-defined extracted feature such as propagation speed. Users can specify which feature to be displayed, either an existing feature in AQuA or a user-designed feature based on features provided by AQuA. We provide several pre-defined colormaps, but allow users to manually define colormaps as well. AQuA also provides a side-by-side view, to simultaneously display two features or a raw video plus one feature. Third, the software provides a convenient way to interactively view detected events and their associated features. By clicking on an event, the dF/F curve for the event is shown in a separate panel below the video, and the time-frames during which the event occurs are highlighted in red. The values of several other features for that event are also shown in another panel. The software allows multiple events to be selected simultaneously, so that their curves and features can be plotted together and compared. Fourth, the software provides both automatic and manual ways to proofread the results. For automatic proofreading, events are filtered by setting desired ranges for features-of-interest. Alternatively, users can choose the ‘delete/restore’ button and manually click an event to remove it. Fifth, the software provides flexible ways to incorporate cell morphology or landmark information. Users can manually supply cell morphology or regional information such as the cell boundary, which can assign events to individual cells. Users can also provide landmark information such as the location of a pipette for pharmacological application. Users can also load cell, region, or landmark information from other data sources, such as another fluorescence channel that captures cell morphology. The software can extract landmark-related features for each event, including the direction of propagation relative to a landmark.

Figures and Figure Legends

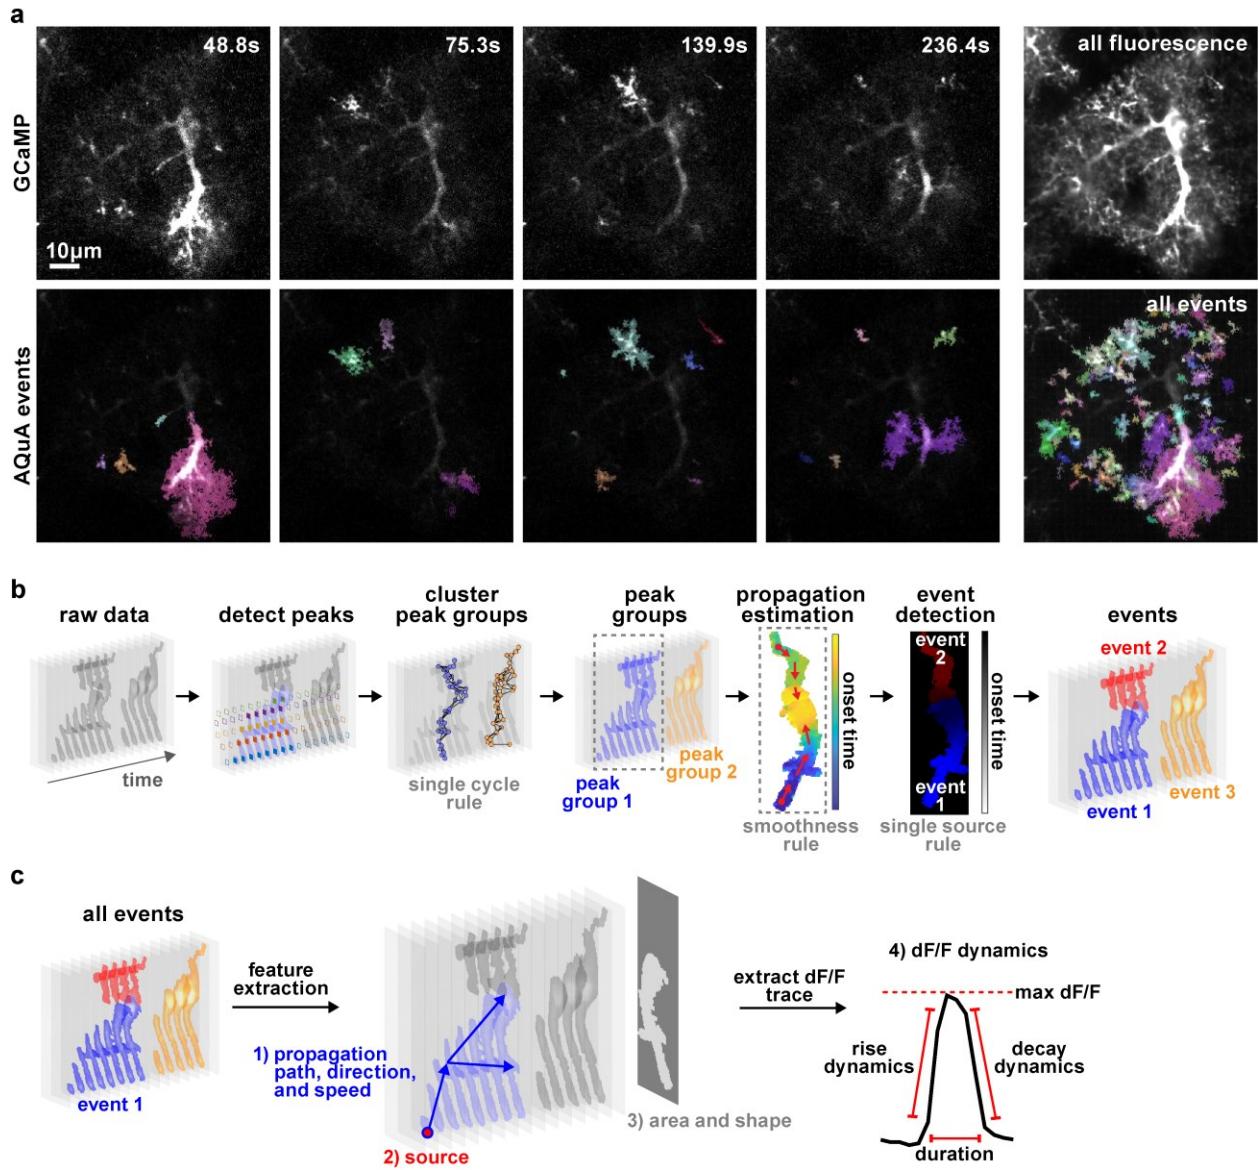


Figure 1. AQuA-based event detection. (a) Individual representative frames from 5-min *ex vivo* astrocytic GCaMP imaging experiment (top, Supp. Movie 1) with AQuA-detected events shown below. Each color represents individual event and is chosen at random. Right column shows the average GCaMP fluorescence (top) and all AQuA-detected events (bottom) from the entire movie. Note that contrast is different between rows to highlight events. (b) Flowchart of AQuA algorithm. Raw data is visualized as a stack of images across time with grey level indicating signal intensity. In the *detect peaks* panel, five peaks are detected and highlighted by solid diamonds, each color denoting one peak. Based on the single-cycle rule and spatial adjacency of the apexes (solid dots) of each peak, peaks are clustered into spatially disconnected groups. Based on smoothness, propagation patterns are estimated for each peak group. By applying the single-source rule, two events are detected for peak group 1. Three total events are detected. (c) Feature extraction. Based on the event-detection results, AQuA outputs four sets of features relevant to astrocytic activity: 1) propagation-related (path, direction, and speed); 2) source of events, indicating

where an event is initiated; 3) features related to the event footprint, including area and shape. Event 2 is plotted here; 4) features derived from the dF/F dynamics.

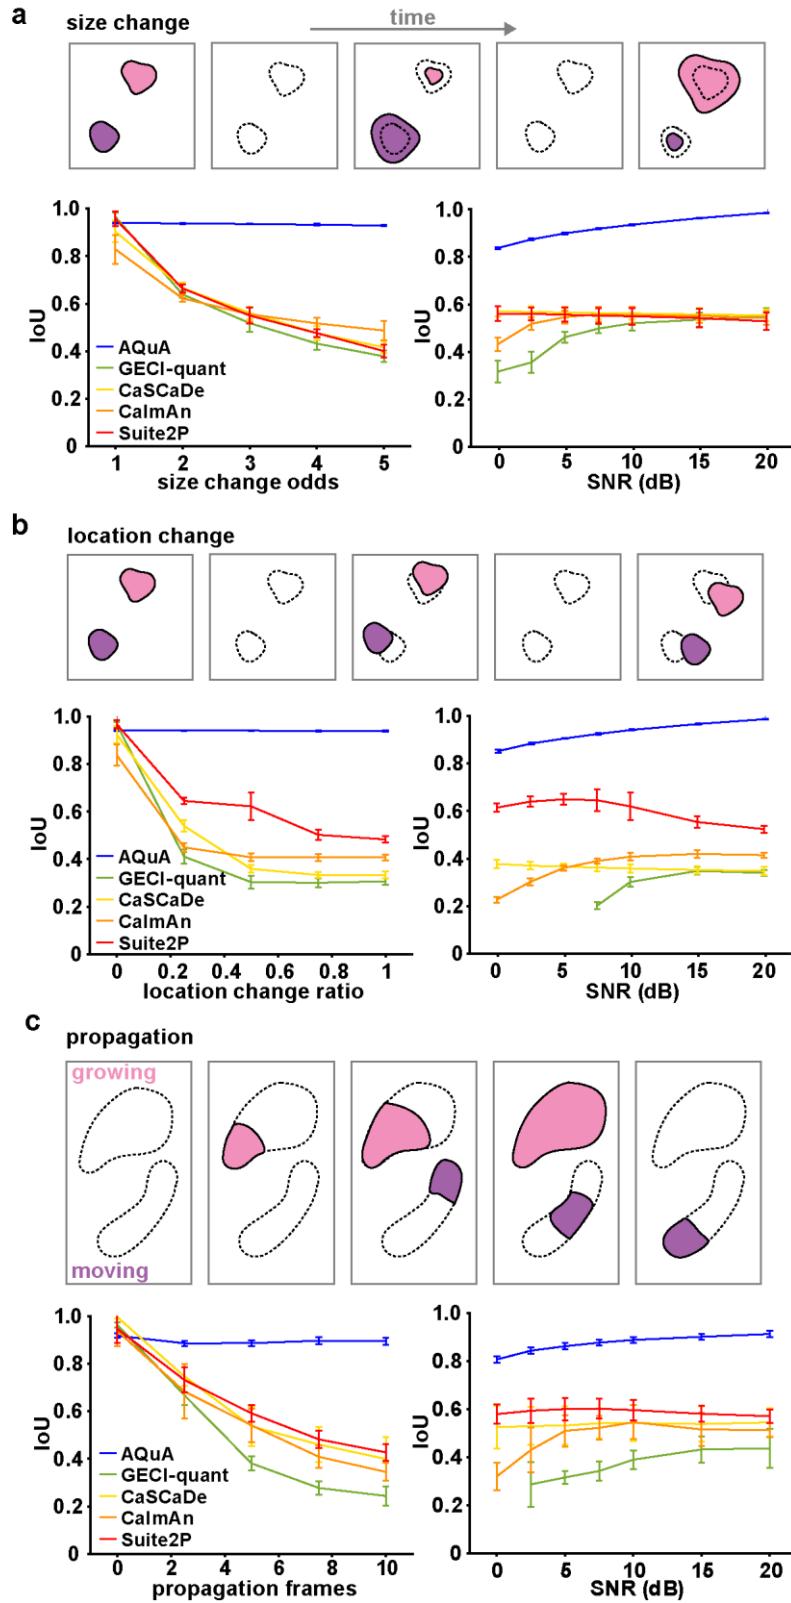


Figure 2. Performance comparison among image-analysis methods. (a–c) Schematic (top) and results (bottom) of performance of five image analysis methods (AQuA, GECI-quant, CaSCaDe, CalmAn, and

Suite2P) on simulated datasets, independently changing event size (a), location (b), and propagation duration (c). In results, change of independent parameter is shown in left panel, and varying SNR in right. For each result, the smallest value of the independent parameter corresponds to a simulation under pure ROI assumptions. The larger the values, the greater the violation of the ROI assumptions. IoU (intersection over union) measures the overlap between detected and ground-truth events. An IoU=1 is the best achievable performance, meaning that all detected events are ground-truth and all ground-truth events are detected. Error bars indicate the 95% confidence interval calculated from 10 independent replications of simulation, where each simulation contains hundreds of events.

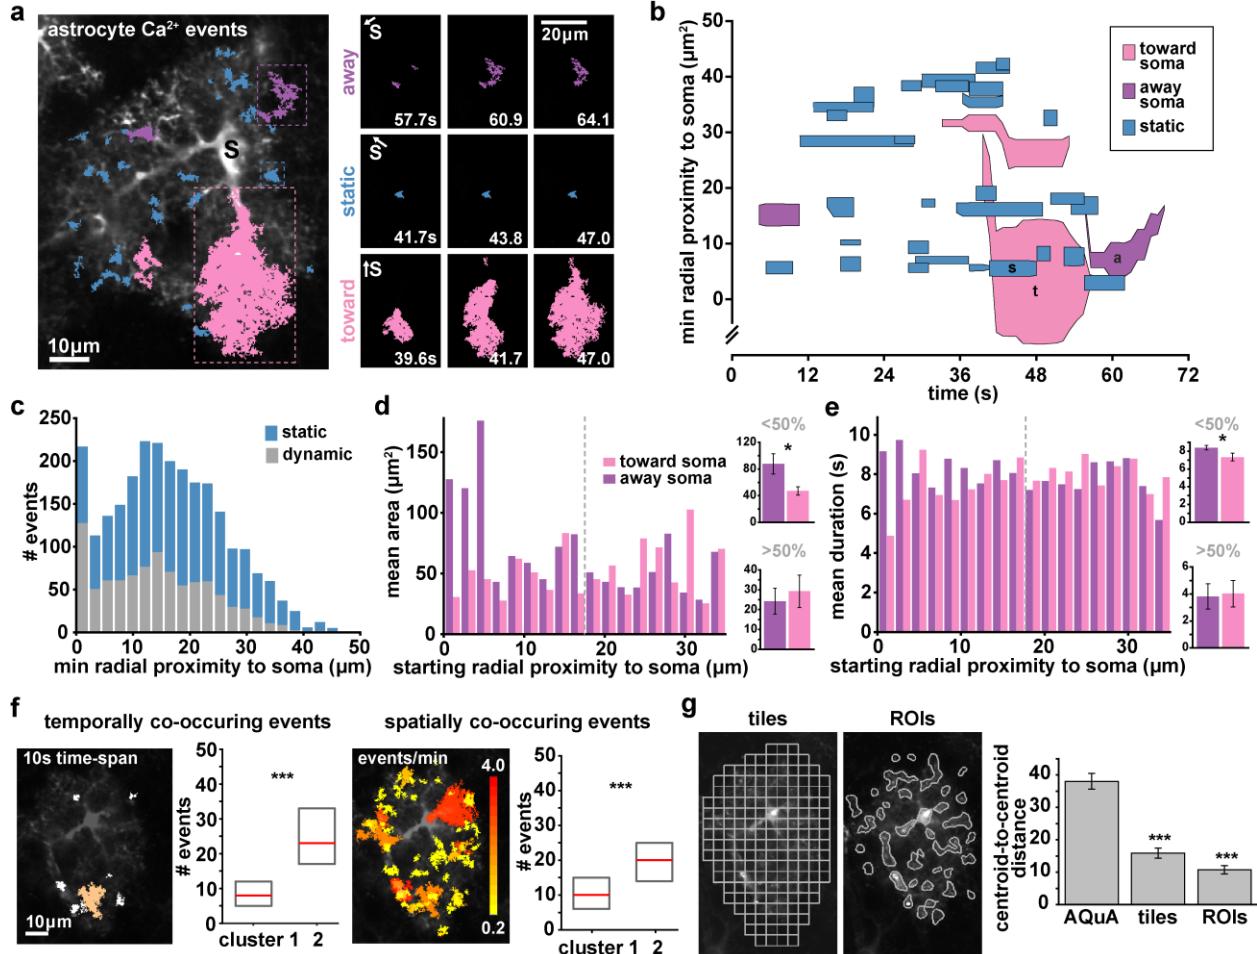


Figure 3. AQuA features capture heterogeneities among single astrocytes. (a) Representative GCaMP6f *ex vivo* image (left) with AQuA events overlaid from 1 min of a 5 min movie. Soma marked with black s. (Supplemental Movie 1). Right: Representative image sequence for each propagation direction class (blue = static, pink = toward soma, purple = away from soma. Soma direction marked with s and white arrow. (b) Spatiotemporal plot of Ca^{2+} activity from 1 min of movie. Each event is represented by a polygon that is proportional to its area as it changes over its lifetime. (c) Distribution of dynamic and static events as a function of minimum distance from soma (chi-square test, *** $p<0.001$, $n=5$ slices, 11 cells). All bin widths calculated by Freedman-Diaconis's rule. (d) Left: Propagative event size versus starting distance from soma, segregated by propagation direction. Dashed gray line denotes half the distance between the soma and the cell border. Right: Average event area for those that start <50% (top) and >50% (bottom) of the distance from the soma, (one-tailed paired t-test, * $p<0.05$). (e) Left: Event duration versus starting distance from soma. Right: Average event duration for those that start <50% (top) and >50% (bottom) of the distance from the soma (one-tailed paired t-test, * $p<0.05$). (f) Two event-based measurements of frequency: events with activity overlapping in time (left) and in space (right). Left: one example event (orange) co-occurs with six other events (white) within 10s. Right: event colors indicate event number/min (0.2–4) at each location. Median (red) and interquartile range (gray) from cells in each cluster in Supp. Fig. 9 (one-tailed Wilcoxon rank sum, *** $p<0.001$). (g) Quantification of centroid distances between cells from two clusters determined by t-SNE plots of Ca^{2+} activity using features calculated from ROIs and $5 \times 5 \mu\text{m}$ tiles (top), (bottom, one-tailed paired t-test, *** $p<0.001$).

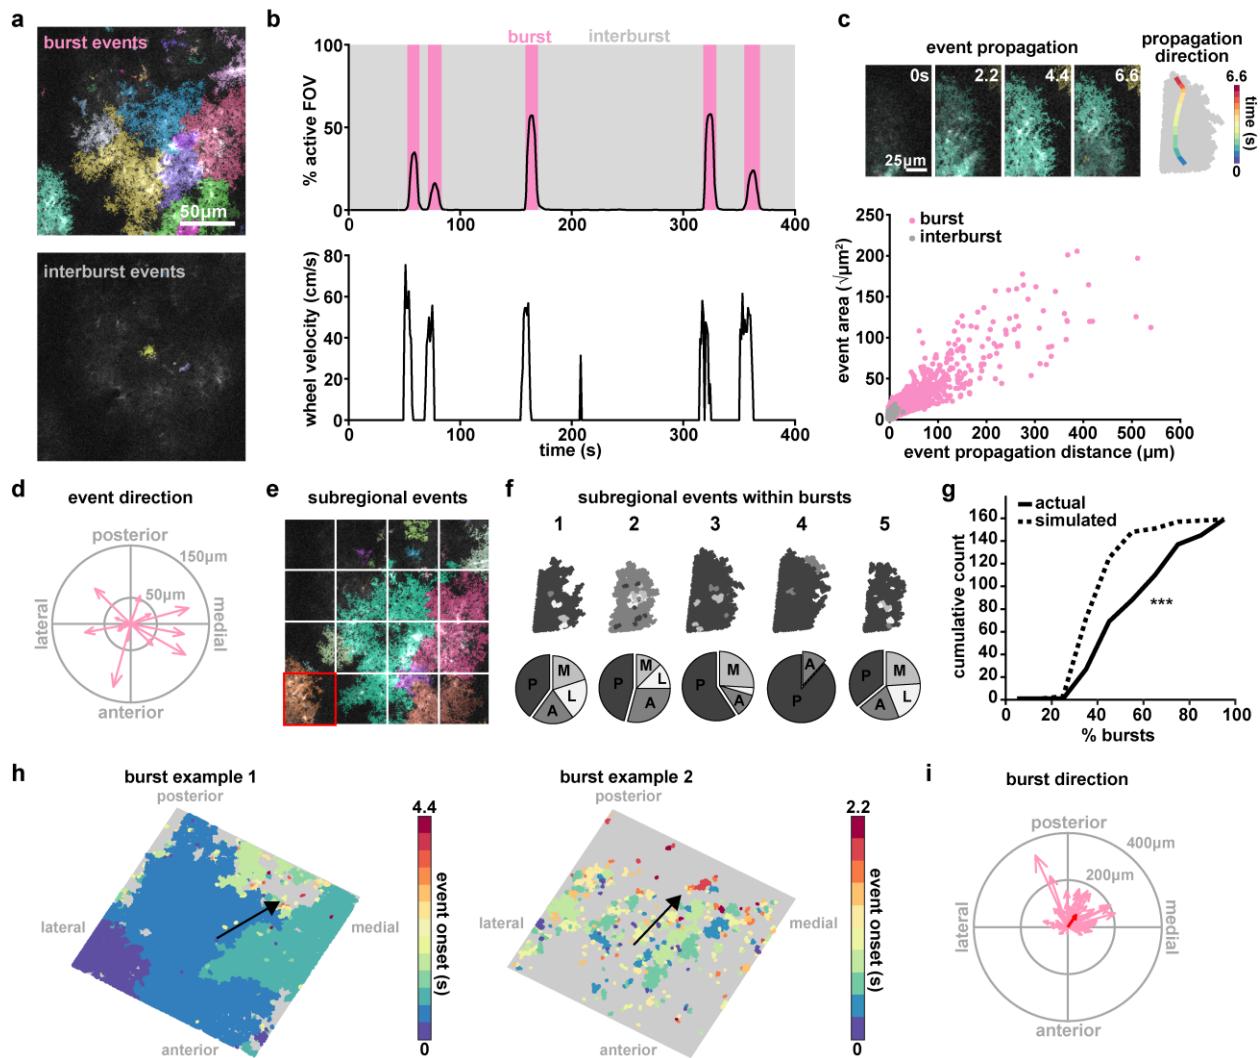


Figure 4. AQuA resolves astrocytic Ca^{2+} propagation directionality across scales. (a) Representative *in vivo* GCaMP6f images during a burst period (top) and inter-burst period (bottom) with overlaid AQuA-detected events. (b) Population Ca^{2+} events represented as percentage of the imaging field active as a function of time. Burst periods (pink) are identified when Ca^{2+} activity exceeds more than 1% of the active field of view and exceeds more than 10% of the maximum number of event onsets. (c) *In vivo* Ca^{2+} events propagate with specific directionality. Top: representative propagative event that occurred during the burst period in panel a. The propagation direction (change of centroid relative to its original location) for each frame is overlaid on the event (right). Bottom: Total propagation distance versus event size for all events within bursts (n=6 mice, 66 bursts, 14,967 events). (d) Event propagation direction from all events over the entire field in the burst shown in e. Length of arrow indicates propagation distance. (e) To test consistency of subregional directionality during bursts, sixteen 96x96 μm tiles are overlaid on images. (f) Top: All events within highlighted tile in d (red square) for five burst periods, color-coded by propagation direction (top). Bottom: Event propagation direction distributions (P=posterior; A=anterior; M=medial; L=lateral). (g) Cumulative distribution of percentage of bursts with events (within individual tiles/regions) propagating in the same direction in actual (solid) and simulated (dashed) data (one-tailed Wilcoxon rank sum, ***p<0.001) (h) Two representative maps of population burst propagation direction with each event color-coded by their onset time relative to the beginning of the burst period,

demonstrating variability of burst size. **(i)** Burst propagation direction calculated from onset maps in h ($n=66$ bursts). Event locations from the first 20% of the frames after burst onset are averaged together to determine burst origin. Event locations from 20% of the last frames after burst onset are averaged together and the difference between this and the origin determines burst propagation distance. Red arrow denotes average of all bursts.

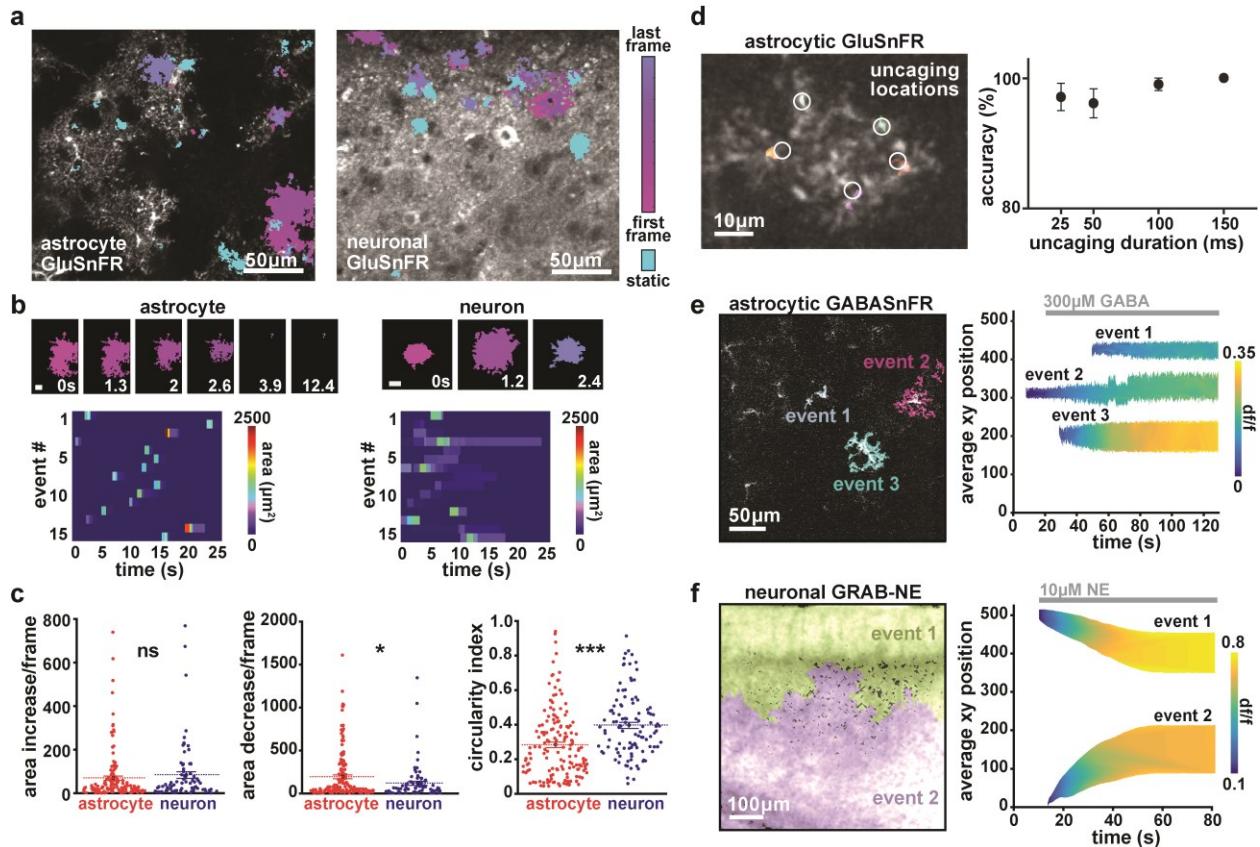
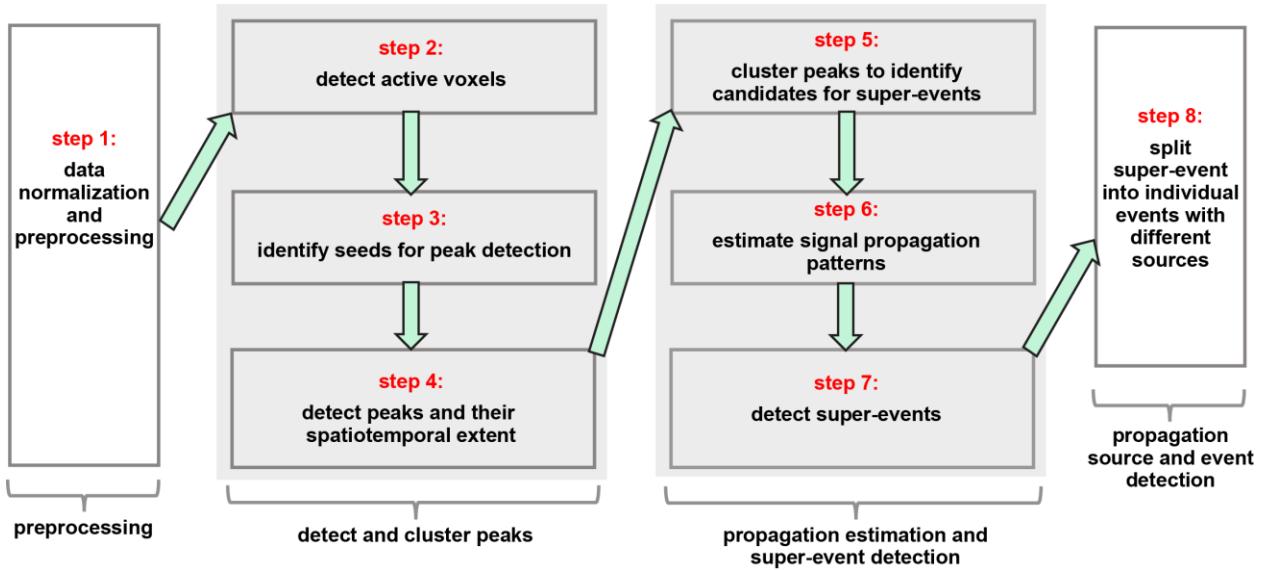
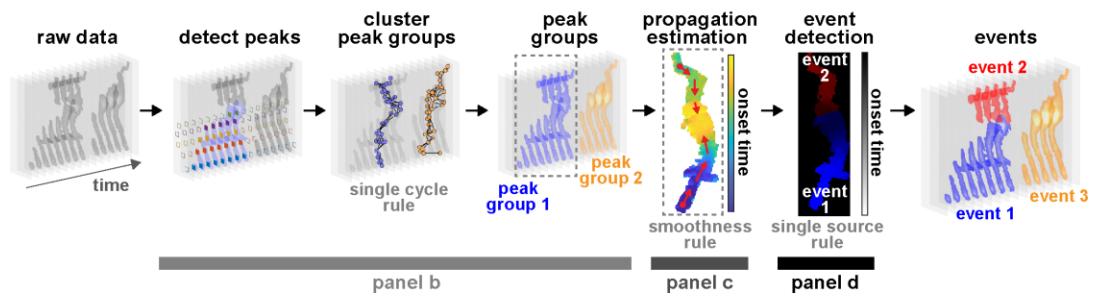


Figure 5. AQuA-based detection of extracellular dynamics via astrocytic and neuronal expression of genetically encoded neurotransmitter sensors. (a) Representative images of *ex vivo* slices with expression of astrocytic (left) or neuronal (right) GluSnFR. Color indicates detected events. Those with dynamic shape are shown in magenta, and static events in cyan. (b) Examples of timecourse of astrocytic (left, top) and neuronal (right, top) glutamate events. Scale bar = 10 μm. Raster plot of area of astrocytic (left, bottom) and neuronal (right, bottom) glutamate events. (c) Size dynamics (area increase [left] and decrease [middle] per frame) and shape (circularity index, right) of glutamate events when GluSnFR is expressed on astrocytes (red) or neurons (blue). (d) Left: Single astrocyte expressing GluSnFR, with AQuA-detected events (colors) with ~100Hz frame rate imaging and 25–150ms uncaging of RuBi-glutamate. Uncaging locations marked with white circles. Right: Percent correct events detected by AQuA, depending on duration of the laser uncaging pulse. (e) Example of three AQuA-detected events at single timepoint (97s) after addition of 300 μM GABA to slice with astrocytes expressing GABASnFR (left). Right: Detected events before and after addition of 300 μM GABA (gray bar) to circulating bath. Events are plotted to display spatial position in imaging field (y-axis), event area (height), and gradually increasing amplitude (color) over time. (f) Left: two detected events in cortical slice expressing GRAB-NE in neurons after addition of 10 μM NE. Right: Events plotted to display spatial position (y-axis), event area (height), and amplitude (color) dynamics over the course of the experiment. In (e) and (f), average xy position at each timepoint is calculated using the following equation: $((xLoc-1)*frameSize) + yLoc)/frameSize$.

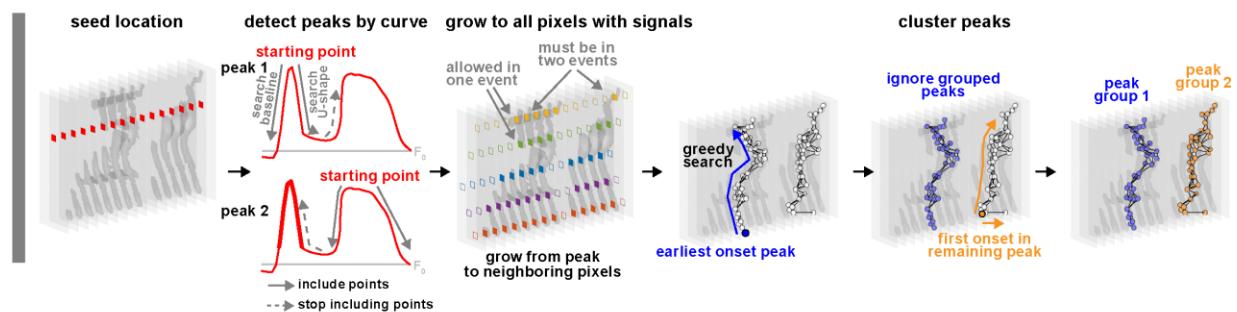


Supplemental Figure 1: Eight steps in the AQuA algorithm. The eight steps can be grouped into four modules indicated by brackets below panels. The last three modules are further illustrated in Supp. Fig. 2.

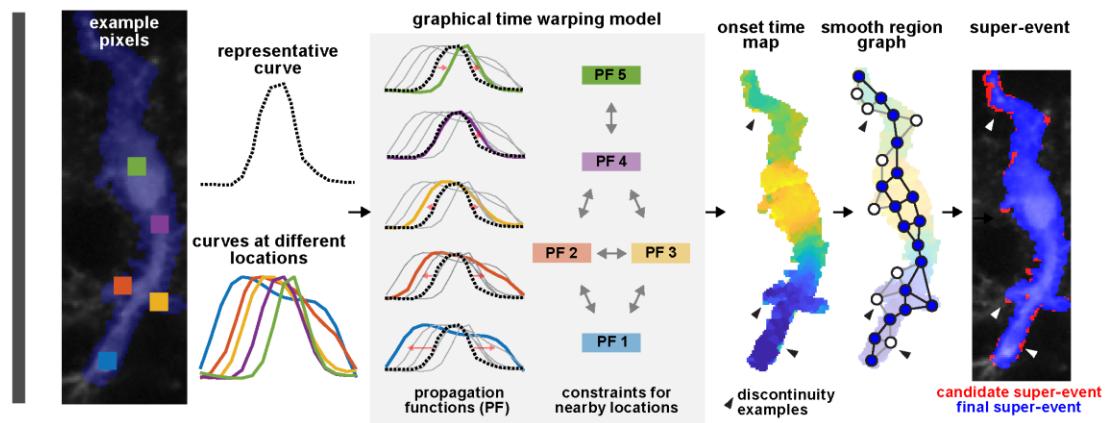
a AQuA flowchart



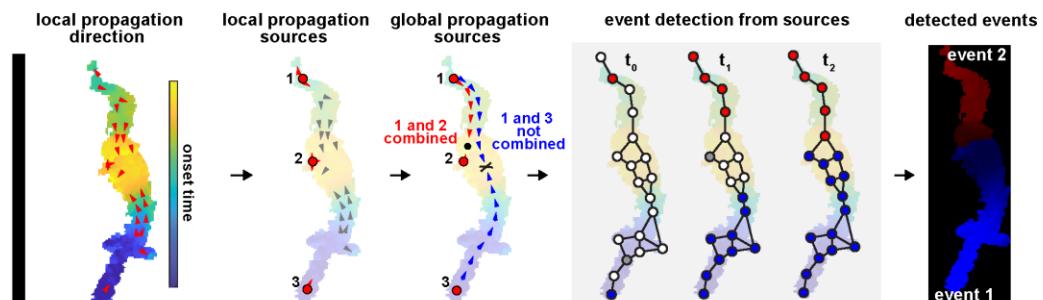
b detect and cluster peaks



c propagation estimation and super-event detection

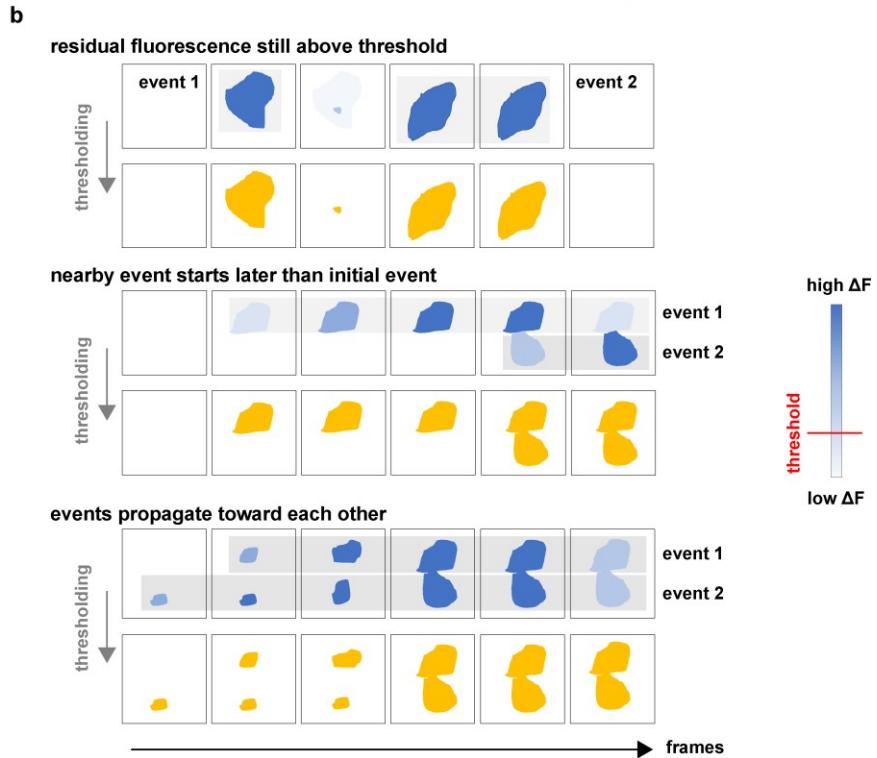
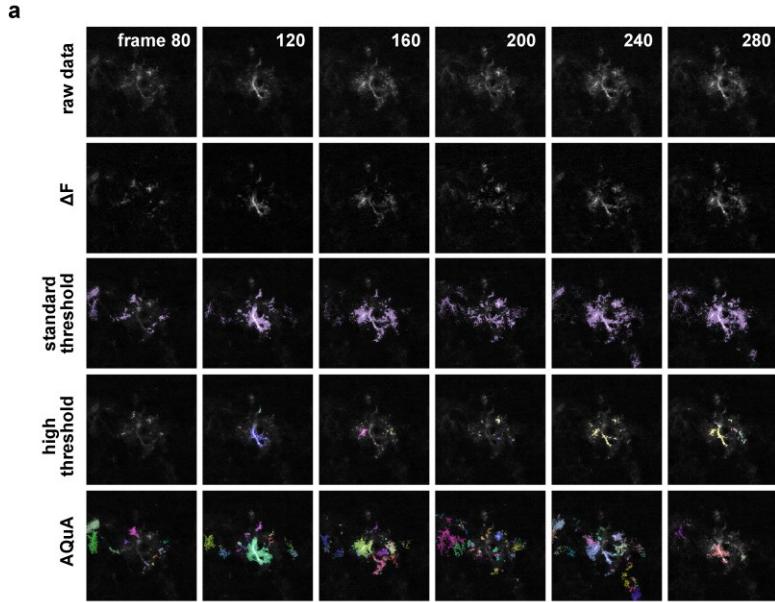


d propagation source and event detection



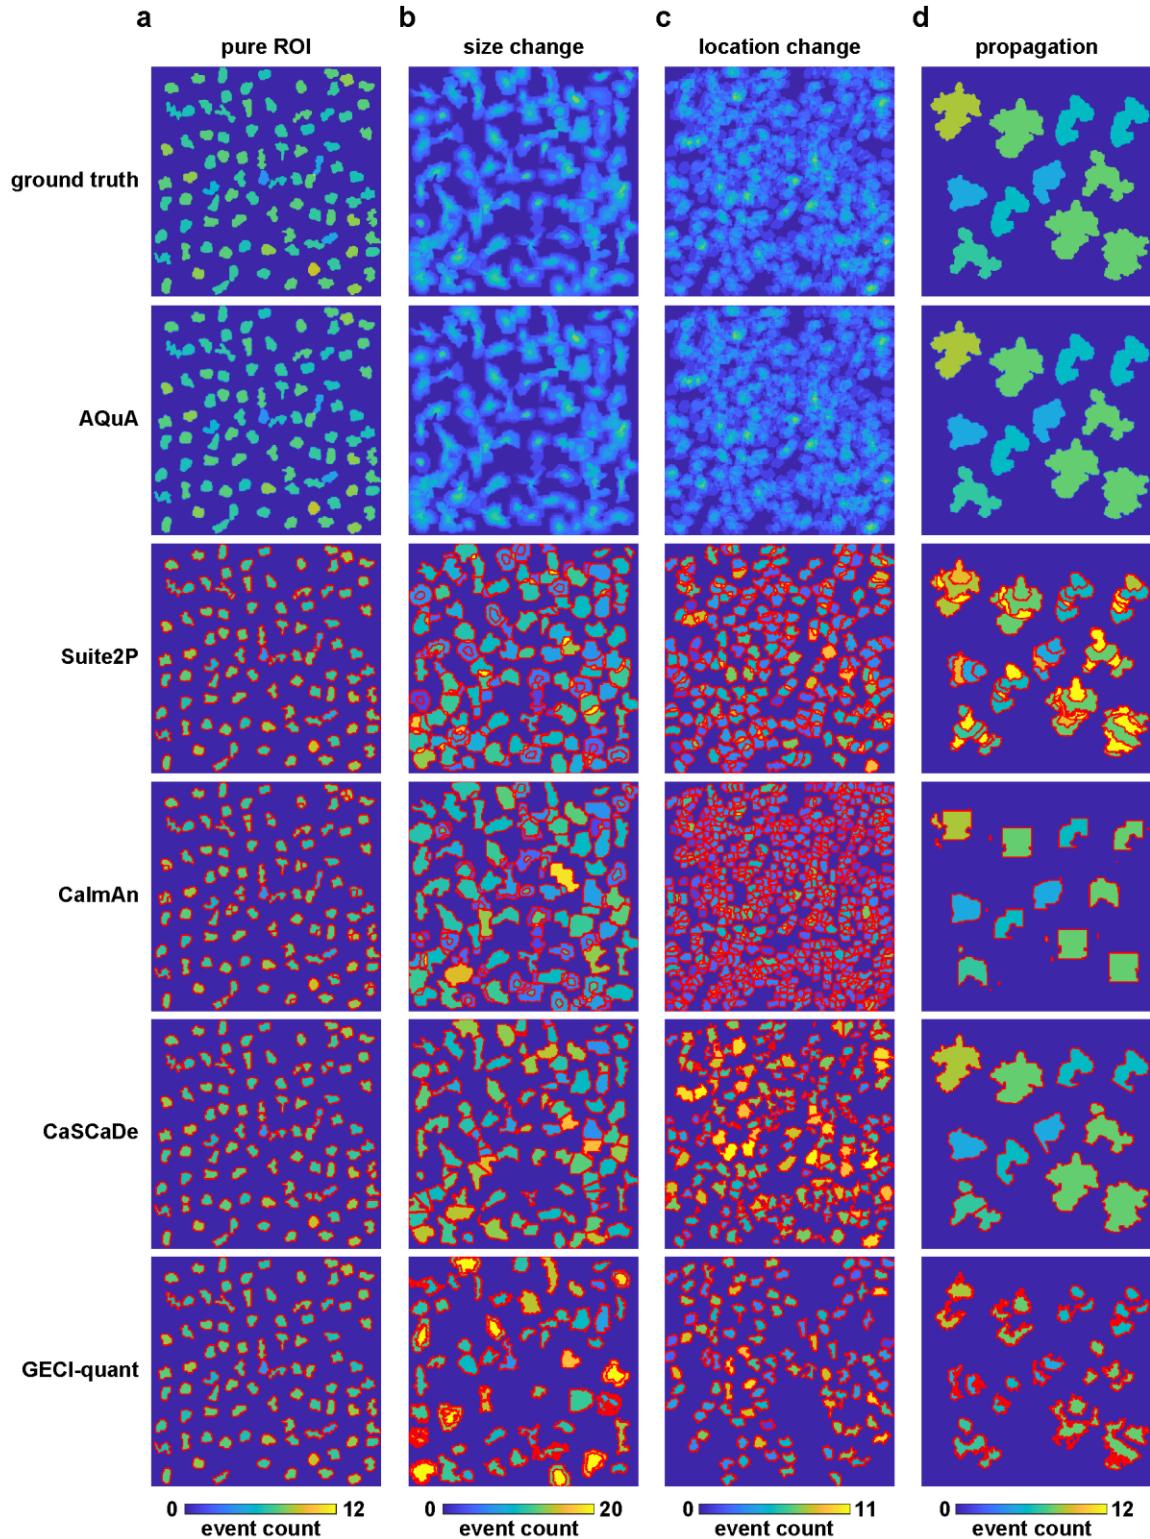
Supplemental Figure 2: Schematic illustration of three major modules in AQuA algorithm. Curves and regions taken from a real data set. **(a) AQuA flowchart**, with three gray bars below indicating where the three major modules are located with respect to the AQuA flowchart. **(b) detect and cluster peaks**: curves in the *detect peaks by curve* panel are associated with the location labeled by the red diamond in

the *seed location* panel. One curve may have multiple peaks, which are detected one-by-one. Once a peak is detected at a seed location, the peak is spatially extended to include its neighboring pixels as in the *grow to all pixels with signals* panel. Clustering of peaks starts from the peak with the earliest onset time and includes its spatially adjacent peaks based on the two inclusion rules shown in the *grow to all pixels with signals* panel. Two peaks at one location are never clustered into one group. Once the greedy search strategy can't find more peaks to include, it stops and one peak group is formed. Then, to find another peak group, the greedy search restarts from the first onset in the remaining peaks. The process is repeated until no peaks remain. **(c) Propagation estimation and super-event detection:** This module is applied to each peak group. The five colored curves are the dynamics of the five exemplar pixels with corresponding colors. The dashed curve is the representative or reference curve. In the *graphical time warping model* panel, red arrows indicate how the reference curve can be warped to represent the curve at each location. The graphical time warping model incorporates the information that nearby locations should have more similar curves than distant locations. A double-headed arrow between two functions informs the model that these curves should be warped similarly to the reference curve. As a comparison, if there is no double-headed arrow between curves, dissimilar warping functions are allowed. Once the warping function is calculated by the graphical time warping model, onset time is computed for each pixel, resulting in an onset time map. Note discontinuity of onset time examples at black triangles. These pixels are removed to obtain the final super-event, which may contain multiple events and are subject to the next operation. **(d) Propagation source and event detection:** Local propagation sources are obtained by finding local minima on the onset time map. According to the rules described in Methods, some local sources will be combined/merged, resulting in global propagation sources. Briefly, if the path between two local sources does not have to go through a location with a late onset time, these two local sources are combined. Then, each global propagation source leads to an event. Each event is obtained by growing each global source to include its neighboring pixels. In the *event detection from sources* panel, solid dots are pixels already assigned to an event, white dots are unexplored pixels, and grey dots are explored but await a later decision to be assigned to an event.

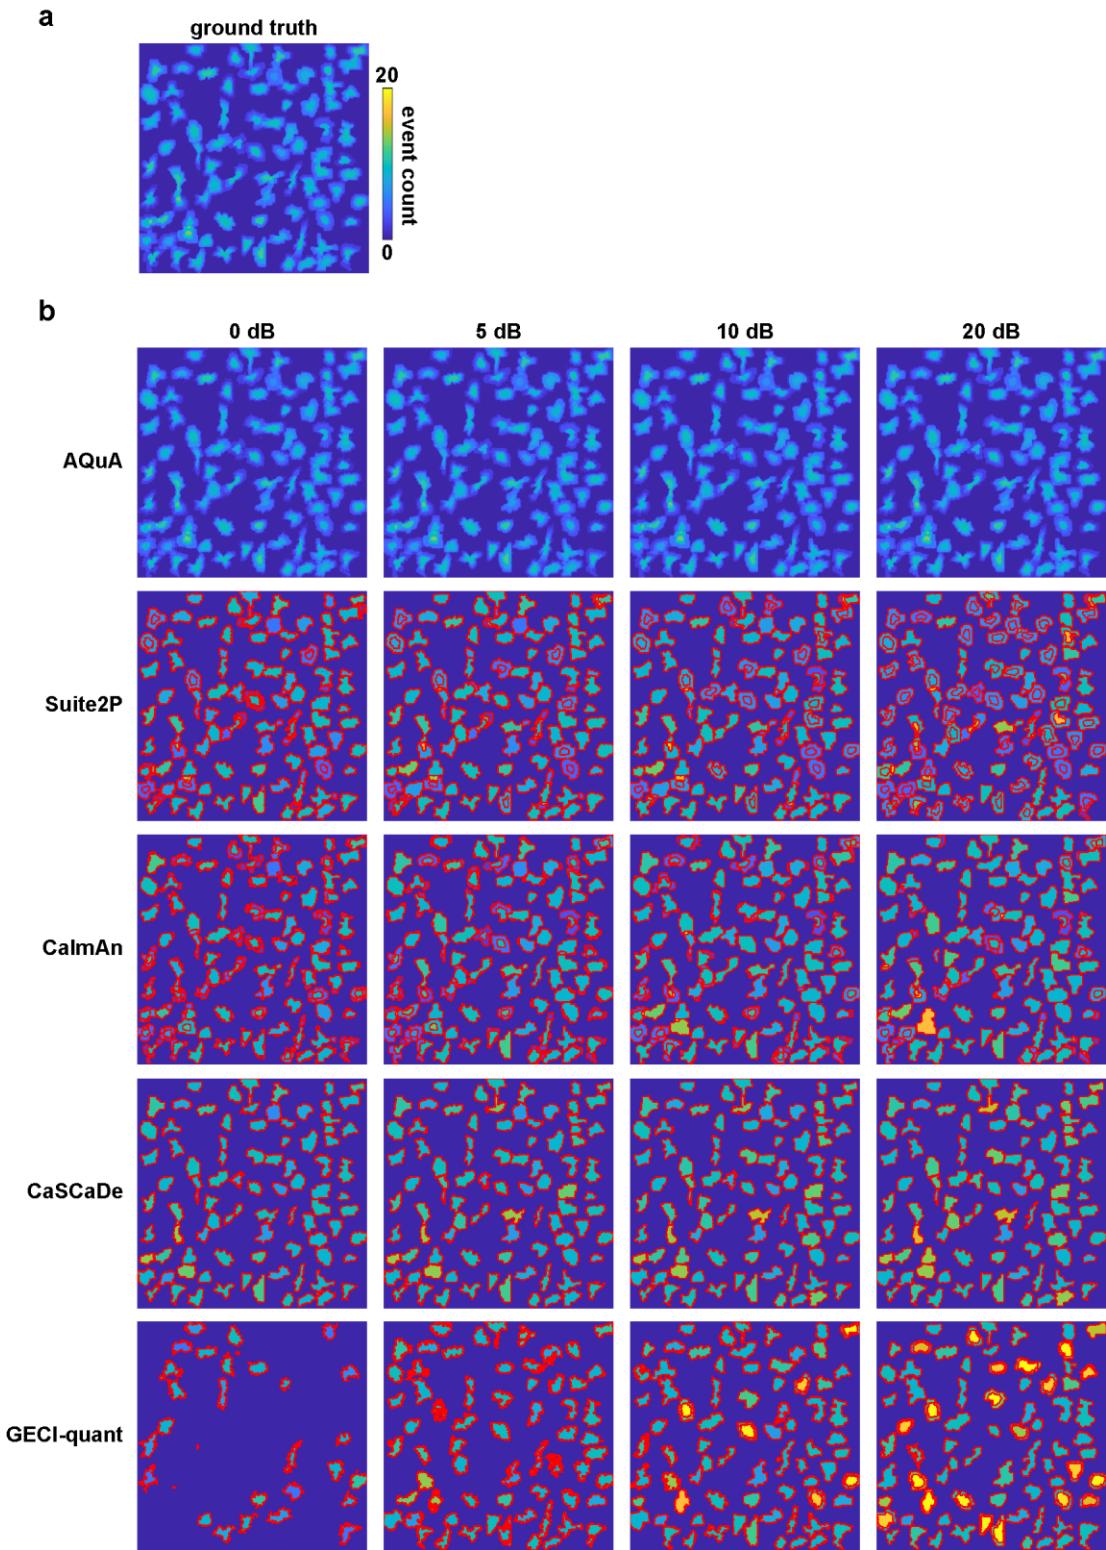


Supplemental Figure 3: Limitations of thresholding-based analysis. (a) *Top row:* Six example frames of *ex vivo* data smoothed with a Gaussian kernel of $\sigma=0.5$ after square root transformation. *Second row:* Baseline for each pixel is estimated with a 20-frame time window, noise level is estimated as σ_{noise} , and baseline is subtracted from raw data to obtain ΔF . *Third row:* Standard threshold is set at $3\sigma_{\text{noise}}$. Many individual events are erroneously detected as one very long and large spatiotemporal component, for reasons graphically explained in (b). *Fourth row:* A high threshold ($10\sigma_{\text{noise}}$) leads to loss of many true events, and many detected events are incomplete. Each color indicates an event. *Fifth row:* AQuA-detection avoids the pitfalls in threshold-setting and identifies each individual event. (b) Two events are incorrectly connected after thresholding (incorrect events in yellow in each sub-panel). Intensity color bar on right, with red indicating the threshold, refers to all panels. *Top:* Between multiple events in the same

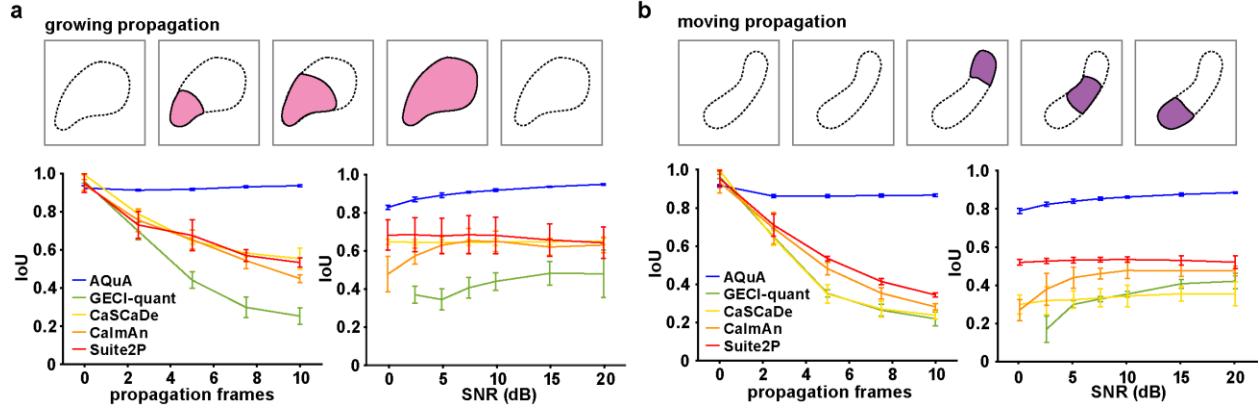
location, even though the intensity drops a lot, not all pixels will fall below the threshold. Each event is shown with a gray bounding box. The super-voxel step in AQuA solves this problem by finding a time window for each seed, and spatially extending the windows. *Middle*: Neighboring events are initiated at distinct times, but are spatiotemporally connected at a later time. If two regions have very different onset times, AQuA will treat them as different events in the super-event detection step. *Bottom*: Two events can be separated when they appear, but meet after propagating. In the event detection step, AQuA distinguishes these events based on the single-source rule.



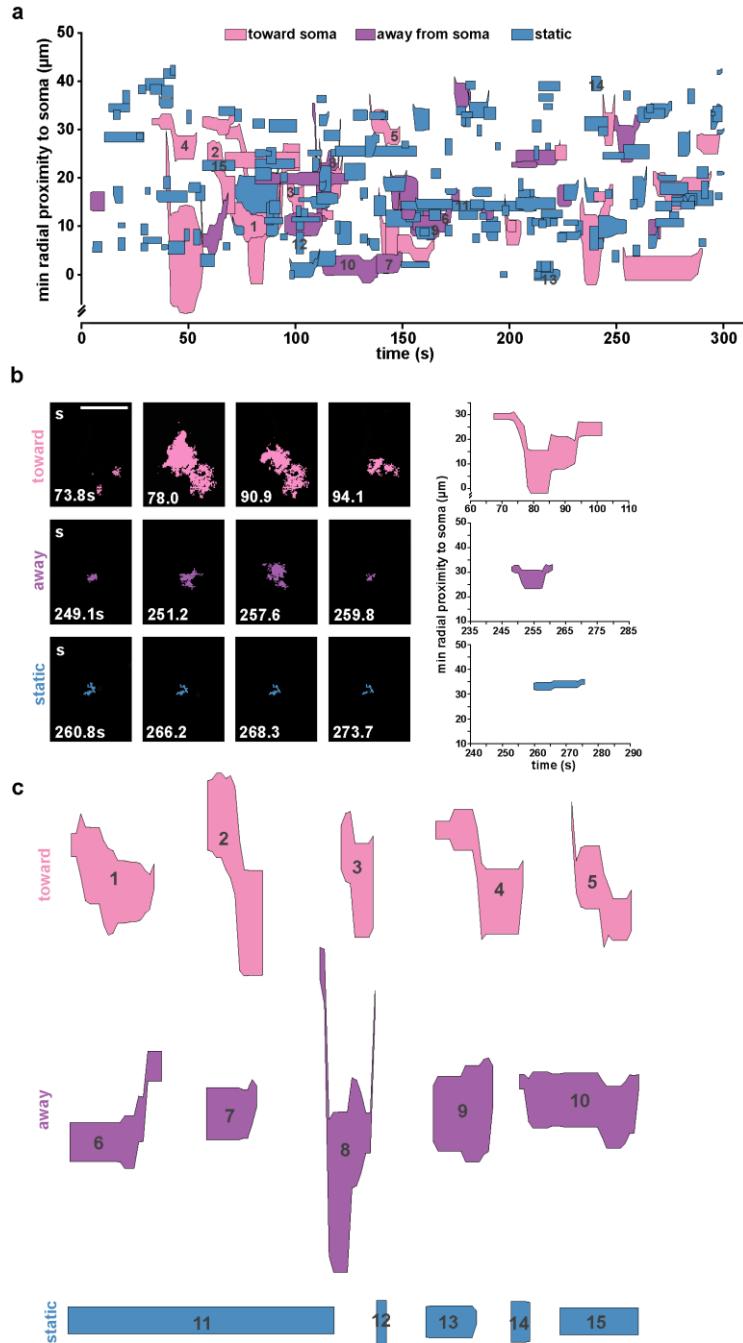
Supplemental Figure 4: AQuA detects ground truth events across three types of simulated data.
 Color represents event count for each pixel (note colors bars have different scales in each dataset). Red borders show ROIs detected by ROI-based methods. (AQuA does not detect ROIs.) (a) Pure ROI. (b) Size change odds of 5, indicating size changes 20–500% of ROI. (c) Location change ratio of 1. Average distance to the center of the ROI is 100% the ROI diameter. (d) Mixed propagation with 10 frames.



Supplemental Figure 5: Event counts under different SNRs. Study the impact of SNR change when size change ratio is 3. The color shows the count of events on that pixel. All plots share the same scale. The red lines are the boundaries of detected ROIs. **(a)** Ground truth event count and the color bar for all plots. **(b)** The event count for all methods under four different SNRs.

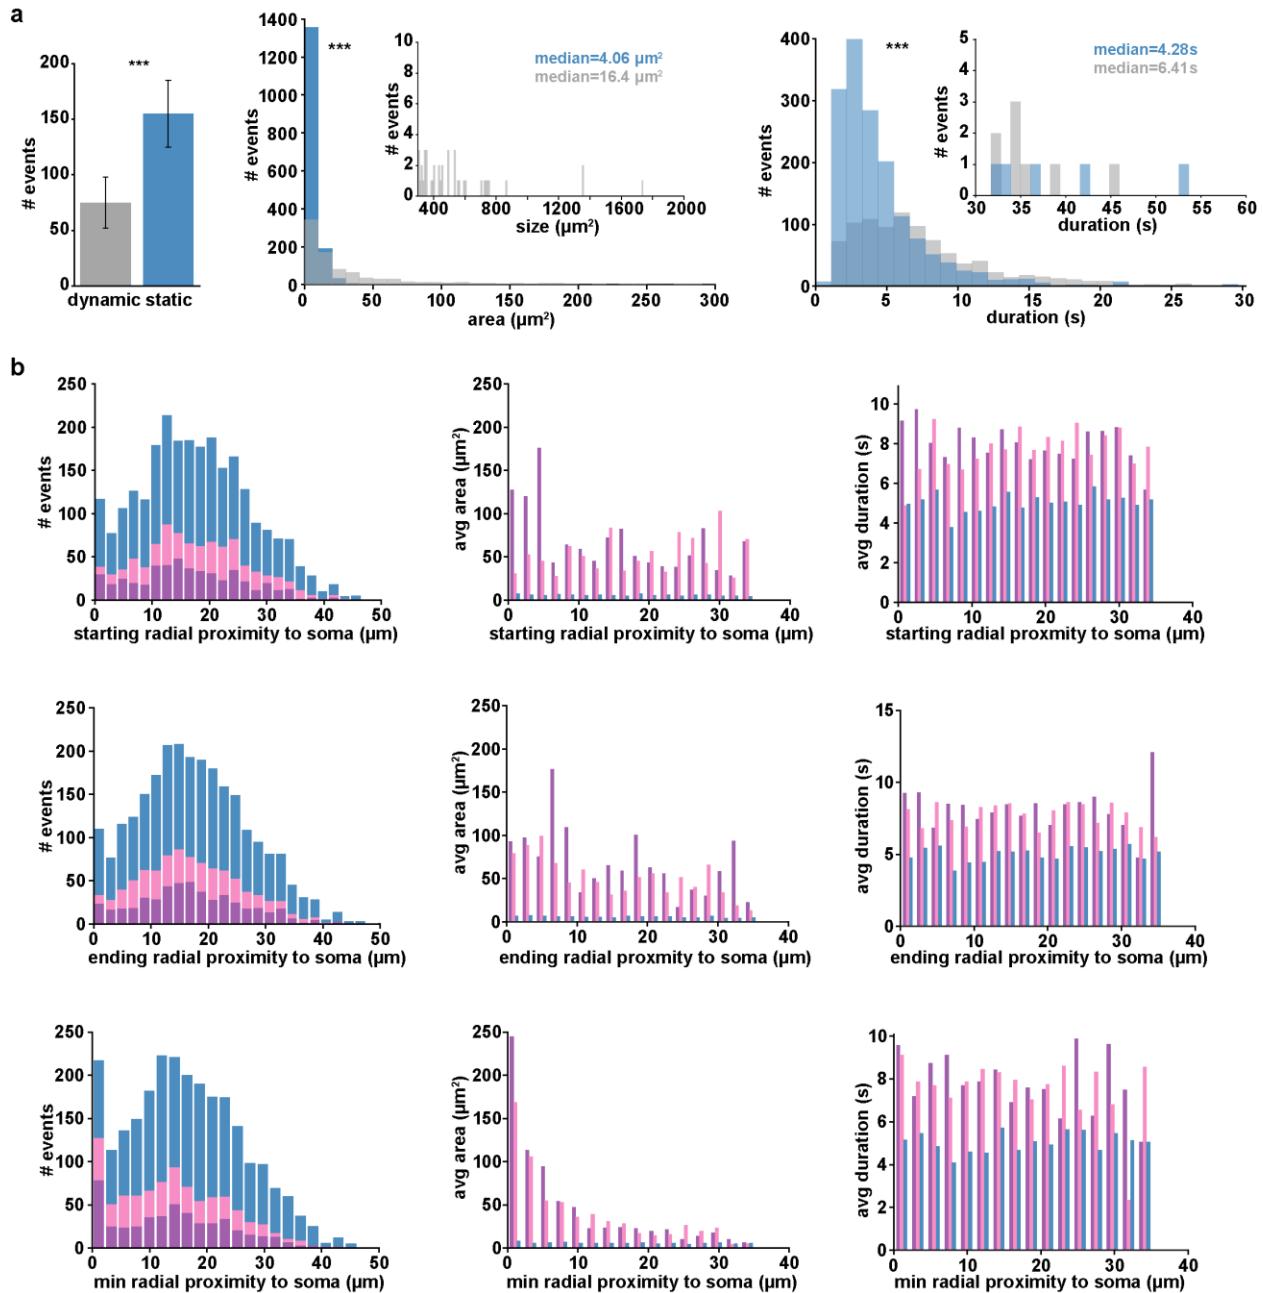


Supplemental Figure 6: Peer method performance on growing and moving propagation types.
 Schematic (top) and results (bottom) of performance of five image-analysis methods (AQuA, GECI-quant, CaSCaDe, CalmAn, and Suite2P) on simulated datasets with (a) growing propagation and (b) moving propagation. Change of the propagation frame number is shown in the bottom left panel, and varying SNR in the bottom right. When the number of propagation frames (not the event duration) is 0, the simulation is under pure ROI assumptions. IoU (intersection over union) measures the overlap between detected and ground-truth events. An IoU of 1 is the best performance achievable by any method, meaning that all detected events are ground-truth and all ground-truth events are detected. The bars on each curve indicate the 95% confidence interval calculated from 10 independent replications of simulation, where each simulation contains hundreds of events.

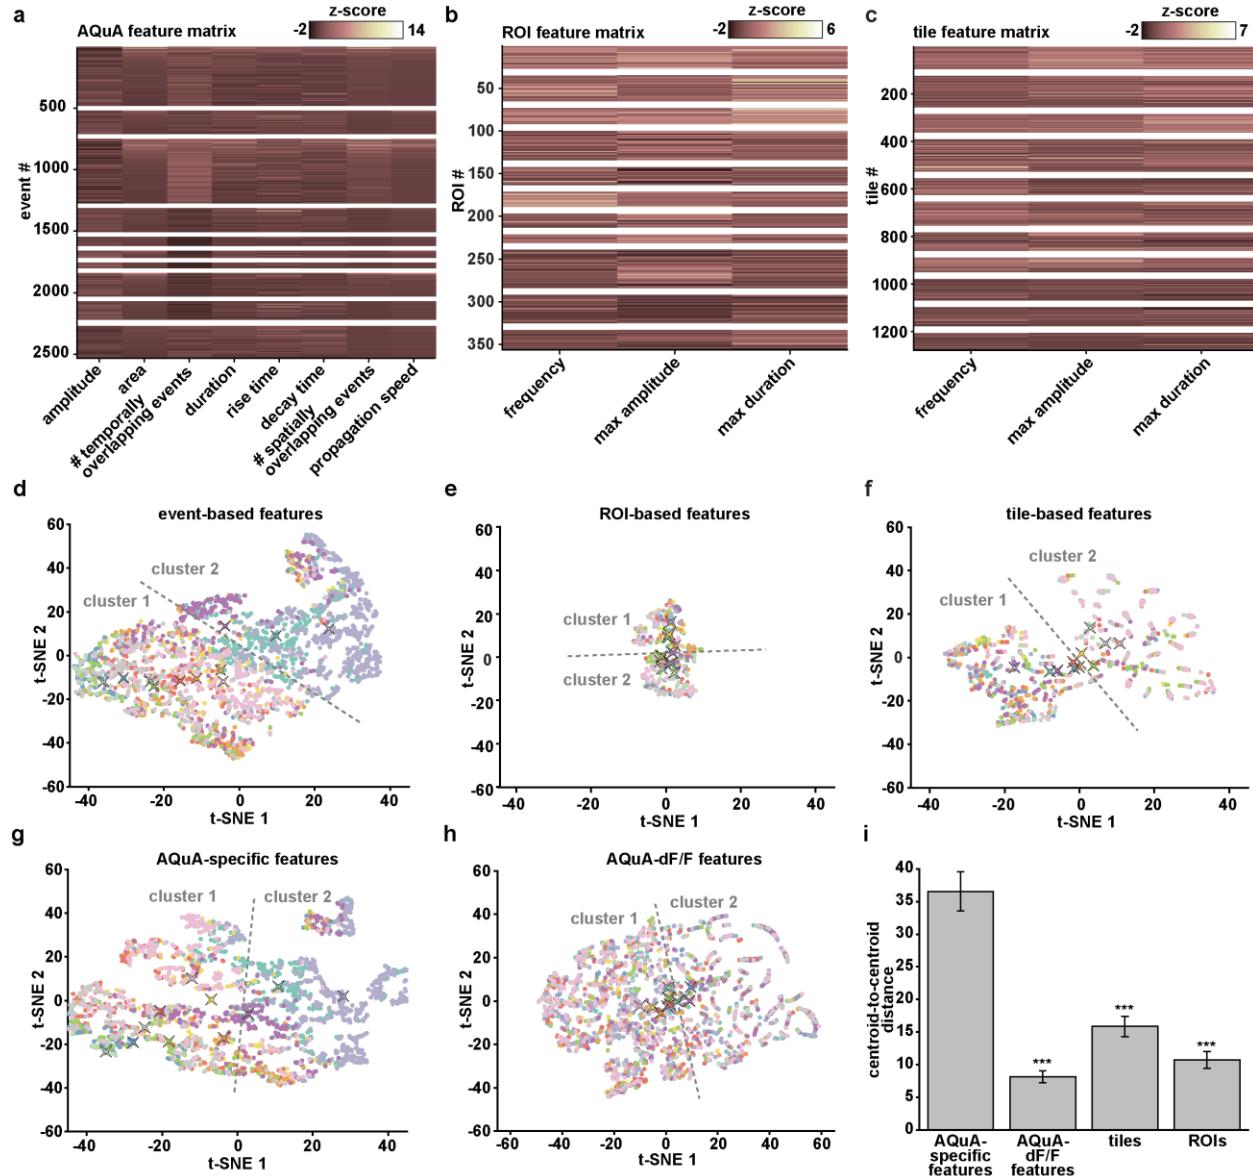


Supplemental Figure 7: AQuA features enable detailed Ca^{2+} activity plots. (a) Spatiotemporal plot of Ca^{2+} activity from a five minute movie (the first minute of which is shown in Fig. 3b). Each event is represented by a polygon that is proportional to the area of the event as it changes over its lifetime, and color-coded by propagation direction. (b) Example time series illustrating how propagation direction is determined (left). A propagation direction score is calculated for each event by multiplying the Euclidian distance between the event pixels' proximity to the soma at each frame by each pixel's intensity. The overall score is the summation of this weighted pixel intensity distance over the lifetime of the event. Therefore, if more pixels with higher intensity move toward the soma it will be classified as such (top). While some events appear in the plot as moving toward the soma, they are actually calculated as moving away from the soma (middle) since we are only displaying the minimum event proximity to the soma in the spatiotemporal plot, but calculate each pixel's proximity to the soma when generating propagation

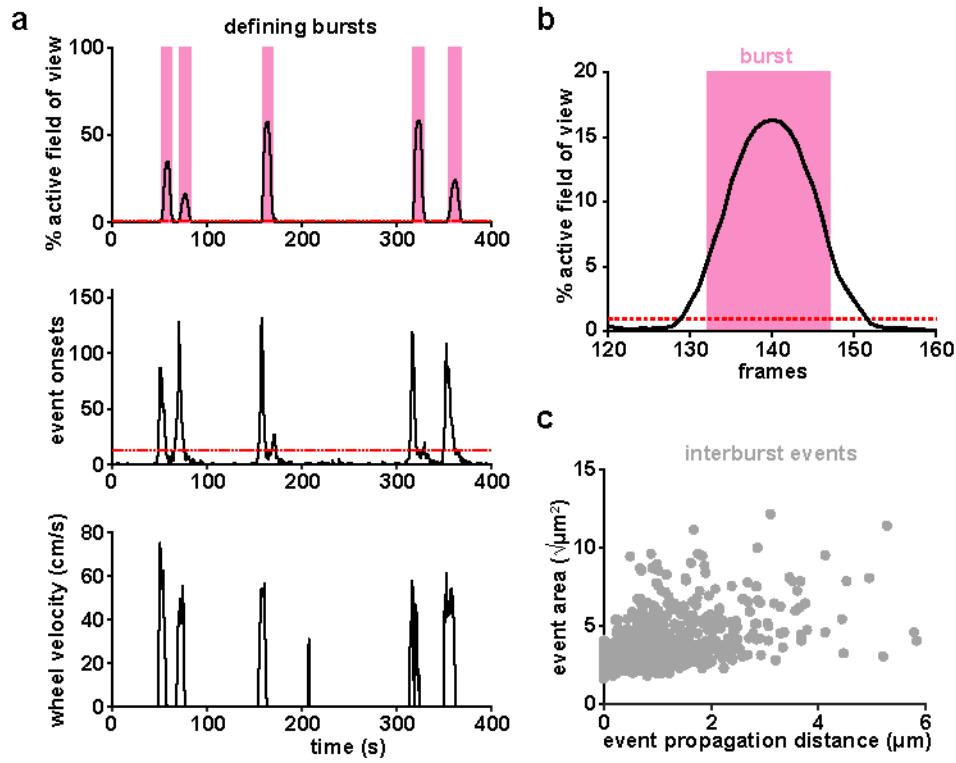
score (see Supplemental Movie 4). Further, pixel intensity is first thresholded at $0.3dF/F$. Therefore, events that move toward or away from the soma yet have pixel intensities below threshold (bottom) appear to have a propagation direction when plotted, yet have a zero propagation direction score when calculated. (c) Additional events plotted for each propagation direction category to demonstrate range of detected/plotted events. Scale is not equivalent to events shown in *b*, but is equivalent within entire group shown here.



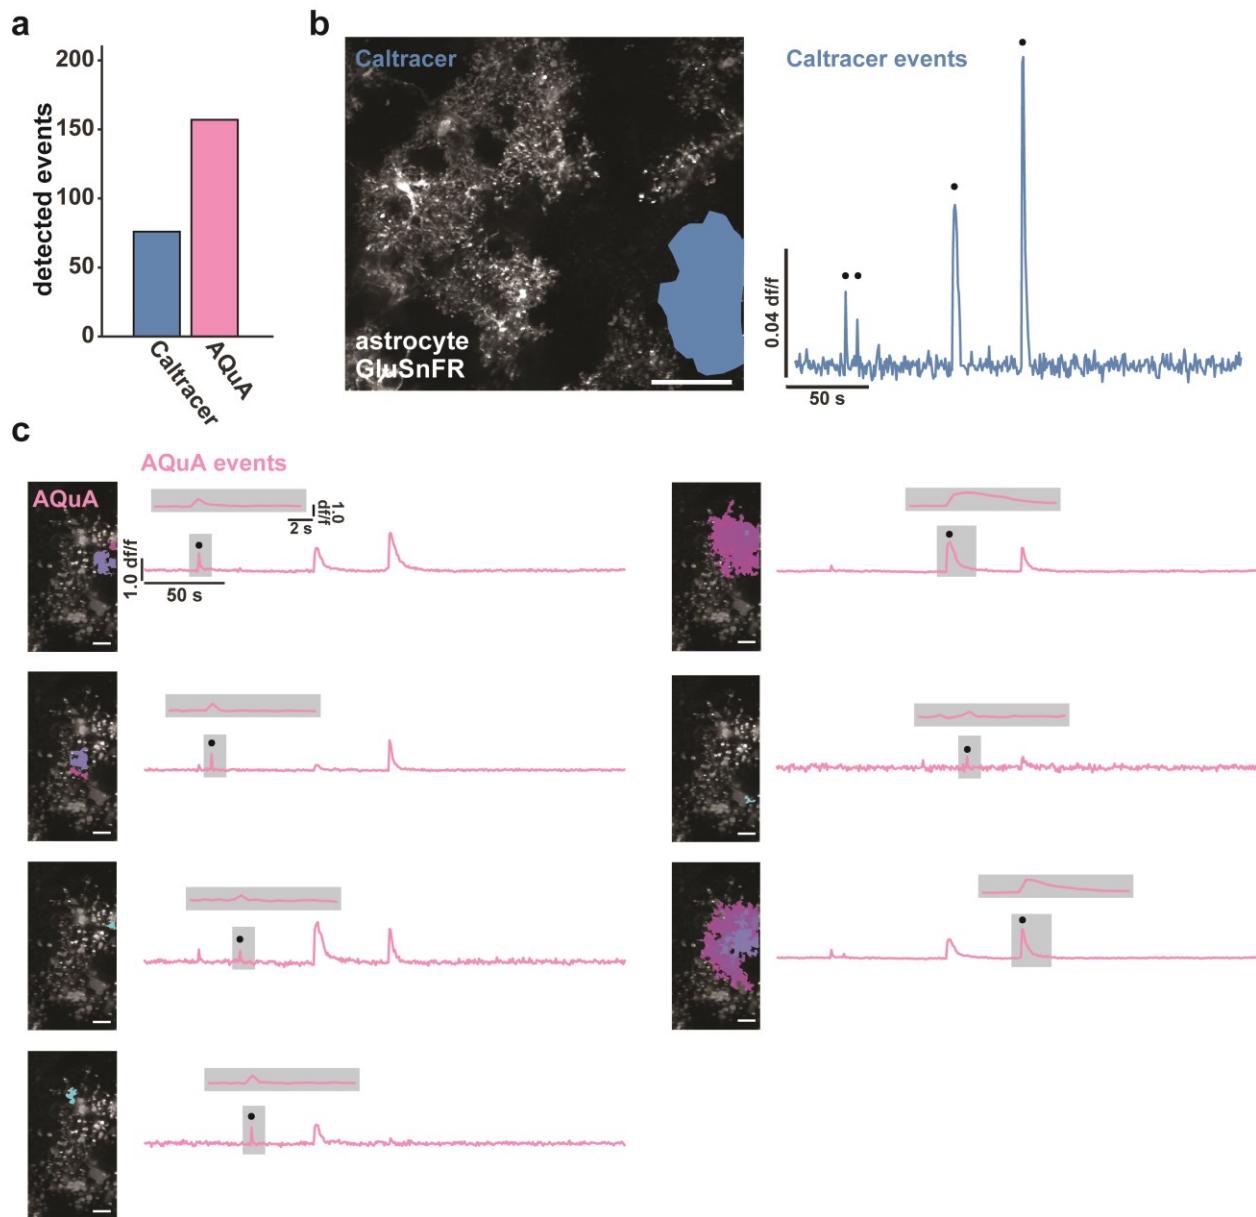
Supplemental Figure 8: Distribution of Ca^{2+} event features. (a) Left: total number of Ca^{2+} events that are dynamic (gray, propagation direction score > 0) and static (blue, propagation direction score = 0) within the 2D imaging plane, *** $p < 0.001$, $n=11$ cells, chi-square test for independence. Middle: distribution of Ca^{2+} event area for dynamic and static events, *** $p < 0.001$, one-tailed Wilcoxon rank sum test. Right: distribution of Ca^{2+} event duration for dynamic and static events, *** $p < 0.001$, one-tailed Wilcoxon rank sum test (right). (b) Distribution, average area, and average duration of events propagating toward soma (pink), away from soma (purple), and static events (blue) compared to starting distance from soma (top row), ending distance from soma (middle row), and minimum distance from soma (bottom row). Bin widths calculated by Freedman-Diaconis's rule.



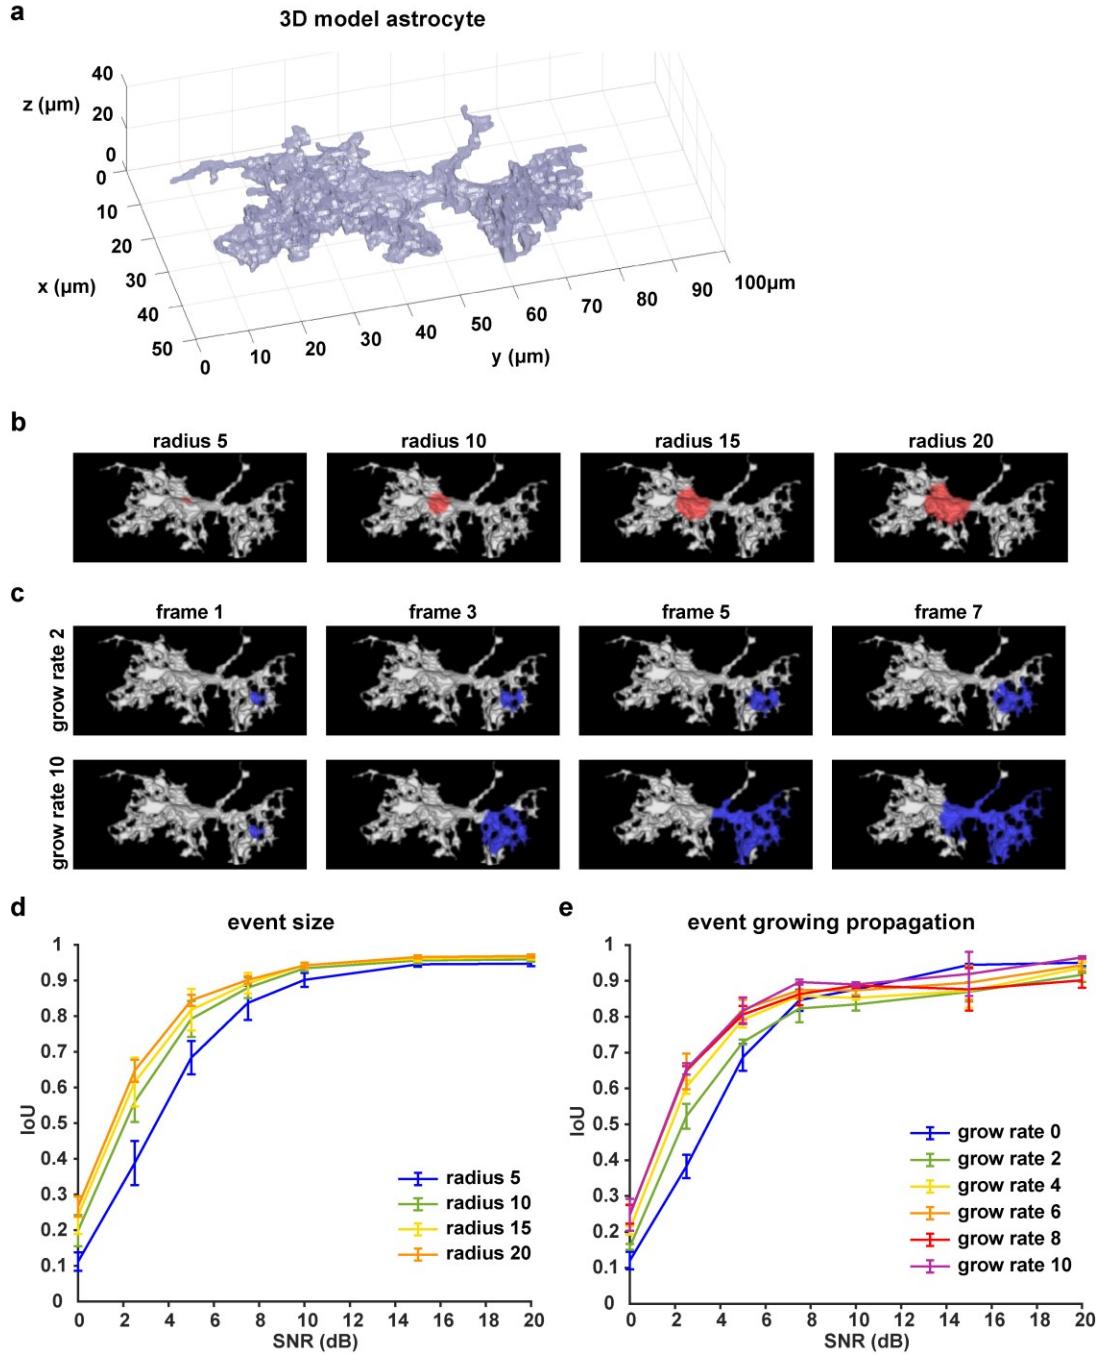
Supplemental Figure 9: Cluster analysis on features generated from three spatial footprint methods. (a) Heatmap of z -scores for eight AQuA features (x-axis) describing each event. White boxes demarcate events and features from individual cells. (b–c) Top: heatmaps of z -scores for three features describing the Ca^{2+} activity at each ROI (b) or tile (c) location. ROIs detected using average projection image with a $5\mu\text{m}$ square filter applied (for ROIs) or $5 \times 5\mu\text{m}$ tiles, based on fluorescence intensity and size. Ca^{2+} events with signals $> 0.03\text{dF/F}$ and two times the noise of each individual trace were selected. Pixels within each ROI or tile were averaged and dF/F was calculated by dividing each value by the mean values from the previous 25 seconds. (d–f) t-SNE visualizations of each cell's Ca^{2+} activity using features calculated using AQuA (d), ROIs (e), and tiles (f). High dimensional data (a–c) are reduced are displayed in two dimensions. Points that are clustered closer together can be interpreted as having more similar Ca^{2+} activity features. k-means clustering boundary denoted as dashed line. (g–h) t-SNE plots using only subsets of AQuA-calculated features from (a) and (d). (g) t-SNE plot of only the features specific to AQuA and not shared with ROI or tile analysis. (h) Plot using only AQuA-extracted features that correspond to those in ROI- or tile-based analyses. (i) Comparison of difference between two clusters generated from the t-SNE analysis followed by k-means clustering. Note increased separation using AQuA-specific features compared to others. (One-tailed paired t-test, *** $p < 0.001$)



Supplemental Figure 10: (a) Population Ca^{2+} events represented as two temporal traces: percentage of imaging field active (top) and number of AQuA event onsets (middle). Burst periods (pink) are defined from the top trace as periods when Ca^{2+} activity exceeds 1% of the active field of view (red dashed line, top), and exceeds 10% of the maximum number of event onsets (red dashed line, middle). Burst periods correlate with wheel velocity of the treadmill (bottom). (b) Burst onset is defined as the first frame in which 10% of the peak is exceeded and burst offset is defined as the last frame exceeding 10% of the peak. (c) The relationship between all interburst events' total propagation distance and size, similar to the burst events plotted in Fig 4c.



Supplemental Figure 11: Comparison of AQuA and Caltracer for event detection of astrocytic GluSnFR signals. (a) Applied to the same data sets, AQuA detects 157 events, while Caltracer^{2,9}, using a rising faces algorithm, detects 76 events with manually defined single-cell ROIs. (b) ROI example (left) and temporal trace with detected events marked by black circle using Caltracer software. Scale bar = 50 μ m. (c) AQuA-detected events from the same cell as in (b), and corresponding temporal traces (black dot, specific events shown above each trace). Scale bar = 10 μ m.



Supplemental Figure 12: AQuA performance on simulated 3D datasets. (a) The 3D model used for performance testing the 3D AQuA extension was extracted from real imaging data¹⁷. (b) Example of simulated data in which event size varies; each column represents a different event radius. (c) Examples of simulated events with slow (top row, growth rate 2) and fast (bottom row, growth rate 10) growing propagation rates. (d–e) Results of AQuA performance on simulated datasets with respect to varying SNR under independently varying event (d) size and (e) growing propagation rates. IoU (intersection over union) measures the overlap between detected and ground-truth events. An IoU of 1 is the best performance achievable, indicating that all detected events are ground-truth and all ground-truth events are detected. The bars on each curve indicate the 95% confidence interval calculated from 5 independent replications of simulation, where each simulation contains 10–20 events.

Supplemental Table 1: Pseudocode for preprocessing, detecting active voxels, and identifying seeds for super voxels. In the pseudocode, the operator colon is used. For example, $A(:)$ means all the elements in A , $A(j:k)$ means the elements $[A(j), A(j+1), \dots, A(k)]$, $A(j:end)$ means the elements $[A(j), A(j+1), \dots, A(n)]$ if $A \in R^n$, and $A(:, :)$ means all the elements in 2D matrix.

Algorithm: Preprocess, detect active voxels, and identify seeds for super voxels

```

1  Require: The original data  $datOrg \in R^{H \times W \times T}$ ; the spatial Gaussian filter parameter  $smoXY$ ;
2  the threshold scale  $thrARScl$ ; the region minimum size  $minSize$ .
3  procedure ACTTOP( $datOrg, smoXY, thrARScl, minSize$ ).
4      Calculate the squared difference of value of adjacent time points for each pixel,
5       $xx(:, :) = (datOrg(:, :, 2:end) - datOrg(:, :, 1:end - 1))^2, xx \in R^{H \times W}$ .
6      Get the square root of median in  $xx$ ,  $stdMap(:, :) = \sqrt{\frac{\text{median}(xx)}{0.9133}}, stdMap \in R^{H \times W}$ .
7      Do spatial Gaussian filter on  $stdMap$ .
8      Define and store the noise standard variance of raw data  $stdEst = \text{median}(stdMap)$ .
9      Define and store Gaussian blur matrix  $dat \in R^{H \times W \times T}$  by applying spatial Gaussian filter
10     on each frame of  $datOrg$  with parameter  $smoXY$ .
11     Get the background matrix  $F_0 \in R^{H \times W \times T}$  by calculating the minimum of moving
12     average of  $dat$ .
13     Calculate the dynamic matrix  $dF \in R^{H \times W \times T}$  by subtracting  $F_0$  from  $dat$ .
14     Construct a boolean matrix  $dActVox \in R^{H \times W \times T}$ , initialize all elements to false.
15     for  $x = 1:W$  do
16         for  $y = 1:H$  do
17             for  $t = 1:T$  do
18                 if  $dF(x, y, t) > stdEst \times thrARScl$  then
19                      $dActVox(x, y, z) = \text{true}$ .
20                 end if
21             end for
22         end for
23     end for
24     Get the 3D connected components of  $dActVox$ , remove the components whose 2D
25     projection size less than  $minSize$ , and set these components' voxels in  $dActVox$  as false.
26     Get all the local maxima located in the components according to value in  $dat$ . Store
27     these maxima as  $lmLoc$ .
28     Sort the  $lmLoc$  in descending order according to the values in  $dat$ .
29     return  $dat, dF, dActVox, lmLoc, stdEst$ 
30     end procedure

```

Supplemental Table 2: Pseudocode for peak detection and extension

Algorithm: Peak detection and extension

Require: The smoothed data $dat \in R^{H \times W \times T}$; the dynamic matrix $dF \in R^{H \times W \times T}$; the seeds of super voxels $lmLoc$; the temporal cut threshold scale $thrTWScl$; the seed growing threshold $thrExtZ$; the super voxel significance $thrSvSig$; the standard variance of noise $stdEst$.

1 **procedure** SPTOP($dat, dF, lmLoc, thrTWScl, thrExtZ, thrSvSig, stdEst$)

2 Construct label matrix $lblMap \in R^{H \times W \times T}$, initialize all values in it to zero.

3 Define nLm as the number of seeds contained in $lmLoc$.

4 **for** $i = 1: nLm$ **do**

5 Calculate the dat mean curve and dF mean curve of neighbor pixels (including seed) of the seed $lmLoc(i)$.

6 Get the time window of the curves calculated above around the peak with parameter $thrTWScl$, record the start time $t_0(i)$, the peak time $t_{peak}(i)$ and the end time $t_1(i)$ of this peak.

7 Initialize the 2D projection set belong to the super voxel of seed $lmLoc(i)$ as $S_i = \{lmLoc(i)\}$.

8 **end for**

9 **for** $k = 1: 40$ **do**

10 **for** $i = 1: nLm$ **do**

11 Find the pixels which are neighbor pixels of S_i . Store these pixels in set V_i .

12 **for** pixel x_0 in V_i **do**

13 Get the curve of x_0 in data as c_0 . Get the 2D location of x_0 as ih, iw .

14 **if** $\frac{\min(c_0(t_{peak}(i)) - c_0(t_0(i)), c_0(t_{peak}(i)) - c_0(t_1(i)))}{stdEst} \geq thrExtZ$ **then**

15 Update $S_i = S_i \cup \{x_0\}$.

16 For the 2D location (ih, iw), from $t_{peak}(i)$, search the time point t forward and backward until $lblMap(ih, iw, t) \neq 0$ or t reaches $t_0(i)$ or $t_1(i)$. Label all visited time points of this 2D location in $lblMap$ as i .

17 **end if**

18 **end for**

19 **end for**

20 **end for**

21 **end for**

22 **for** $i = 1: nLm$ **do**

23 Fill hole of the 3D component whose voxels labeled with i in $lblMap$.

24 Store the voxels labeled i in updated $lblMap$ as $svLst_i$.

25 Get the 2D projection of this 3D component, the number of pixels in it denoted as n_{2d} .

26 Calculate the dat mean curve of whole 3D component (do average spatially) as c_1 .

```

27   if  $\frac{\min(c_1(t_{\text{peak}}(i)) - c_0(t_0(i)), c_1(t_{\text{peak}}(i)) - c_0(t_1(i)))}{stdEst} n_{2d} < thrSvSig$  then
28     Set the  $svLst_i$  as  $\emptyset$ .
29   else
30     Calculate the rising time for  $i^{th}$  super voxel, store the rising time as  $riseX(i)$ .
31   end if
32 end for
33 return  $svLst, riseX$ 
34 end procedure

```

Supplemental Table 3: Pseudocode for estimating signal propagation, detecting super-events, and splitting super-events into individual events.

Algorithm: Estimate signal propagation, detect and split super-events.

Require: The smoothed data $dat \in R^{H \times W \times T}$; the dynamic matrix $dF \in R^{H \times W \times T}$; the super voxels set $svLst$; the rising time $riseX$; the slowest propagation $cDelay$; the rising phase uncertainty $cRise$; the GTW smoothness parameter $gtwSmo$.

1 **procedure** EVTTOP($dat, dF, svLst, riseX, cDelay, cRise, gtwSmo$)

2 Merge the super voxels: if two super voxels are adjacent and the difference of their rising time is small, merge them together as one super-event. Store the super-events as $seLst$.

3 Construct a 3D matrix $labelMap \in R^{H \times W \times T}$ and initialize all elements to zero.

4 **for** super-event $se0$ in $seLst$ **do**

5 Get 2D projection ihw of the super-event $se0$, where ihw is a list of pixels .

6 Calculate super pixels of this projection through super-pixel segmentation algorithm.

7 Align the curve of each super pixel using GTW algorithm of Y. Wang et al.(2016) with parameter $gtwSmo$.

8 Calculate the rising time of the normalized dF curve of each super pixel. Construct a 2D matrix $dlyMap \in R^{H \times W}$, assign the rising time to the corresponding locations in $dlyMap$.

9 Construct a 2D matrix $evtMap \in R^{H \times W}$ to represent the spatial distribution of events in $se0$, initialize all values in $evtMap$ to 0, $nSeed = 1$.

10 Split the $dlyMap$ into several parts according to parameter $cDelay$, where the rising time in each part is smooth enough. Denote these parts as a list $evtGroup$.

11 **for** each part $ihw0$ in $evtGroup$ **do**

12 Get the local minima of $ihw0$ according to $dlyMap$ as the potential event seeds.

13 Store them in a list $seedLm$.

14 Sort the $seedLm$ in descending order according to their rising time.

15 **for** seed $seed0$ in $seedLm$ **do**

16 Find all seeds in $seedLm$ whose rising time less than $rise_{seed} + cRise$, store them as a list $memCand$.

17 **for** seed mem in $memCand$ **do**

18 **if** mem and $seed0$ are connected in $dlyMap$ **then**

19 Merge mem and $seed0$. $seed0 = seed0 \cup \{mem\}$. Remove mem in $seedLm$.

20 **end if**

21 **end for**

22 Update $evtMap(seed0) = nSeed, nSeed = nSeed + 1$.

23 **end for**

24 Sort the remaining pixels in $ihw0$ (the pixels which are not local minima) in ascending order according to their rising time. Store as a list $pixels$.

25 Initialize $waitlist = \emptyset$.

26 **for** pixel $p0$ in $pixels$ **do**

27 **if** $p0$ is adjacent to a growing event in $ihw0$ **then**

28 Set $evtMap(p0)$ to the label of the adjacent event.

29 Check the $waitlist$, if the pixel in $waitlist$ is adjacent to an event in $ihw0$, remove the pixel in $waitlist$ and just to line 28.

30 **Else**

31 Update $waitlist = waitlist \cup \{p0\}$.

32 **end if**

```
33      end for
end for
34      Update the labelMap. Label the labelMap according to the 2D distribution evtMap
      of se0.
35      end for
36      Store the event set as evtList.
37      Calculate the basic features of each event as ftsLst.
38      Return evtList, ftsLst.
39      end procedure
```

Supplemental Table 4: AQuA running time. We ran AQuA on different datasets tested on a desktop computer with Intel Xeon E5-2630 CPU and 128 GB memory using MATLAB version R2018b and running on Microsoft Windows 10. The tests used the default parameter setting. All datasets are 512 x 512 pixels. Note that some datasets are stored as 32-bit files, so file size can be larger than those stored with 16 or 8 bits.

ID	type	frames	file size	event number	total running time (s)
1	<i>in vivo</i> Ca ²⁺	1090	545 MB	3,873	2,440
2	<i>in vivo</i> Ca ²⁺	1614	1.58 GB	31,829	5,678
3	<i>in vivo</i> Ca ²⁺	1614	1.54 GB	12,065	3,110
4	<i>ex vivo</i> Ca ²⁺	281	281 MB	3,824	2,020
5	<i>ex vivo</i> Ca ²⁺	281	281 MB	3,471	2,076
6	<i>ex vivo</i> Ca ²⁺	281	281 MB	4,158	1,800
7	GluSnFR	910	459 MB	6	140
8	GluSnFR	458	230 MB	55	355
9	GluSnFR	306	154 MB	33	182

Supplemental Table 5: AQuA-extracted features

Category	Name of feature	Unit	Description
Basic properties	Area	μm^2	Spatial area of each event
	Volume	μm^3	Spatial-temporal volume of each event
	Perimeter	μm	Perimeter of each event
	Circularity		Closeness to a circle for the shape of each event
Trace	P Value on max dF/F		Significance value of max dF/F for each event, which is related to the noise level in each trace
	Max dF/F		Peak of the dF/F curve for each event
	Duration 50% to 50%	s	Duration between 50% onset time to 50% offset time
	Duration 10% to 10%	s	Duration between 10% onset time to 10% offset time
	Onset duration 10% to 90%	s	Time taken for dF/F to rise from 10% to 90%
	Offset duration 90% to 10%	s	Time taken for dF/F to drop from 90% to 10%
	Decay time constant τ	s	Decay time constant by fitting exponential curve on dF/F
Propagation, onset	Overall	μm	Single event overall onset propagation score, which is the sum of such scores in four directions
	Anterior	μm	Single event onset propagation distance in anterior direction
	Posterior	μm	Single event onset propagation distance in posterior direction
	Left	μm	Single event onset propagation distance in left direction
	Right	μm	Single event onset propagation distance in right direction
	Anterior ratio	μm	Single event onset propagation distances in anterior direction, normalized by the sum of four directions
	Posterior ratio	μm	Single event onset propagation distances in posterior direction, normalized by the sum of four directions
	Left ratio	μm	Single event onset propagation distances in left direction, normalized by the sum of four directions
	Right ratio	μm	Single event onset propagation distances in right direction, normalized by the sum of four directions
Propagation, offset	Overall	μm	Single event offset propagation score, which is the sum of such score in four directions
	Anterior	μm	Single event offset propagation distances in anterior direction

	Posterior	μm	Single event offset propagation distances in posterior direction
	Left	μm	Single event offset propagation distances in left direction
	Right	μm	Single event offset propagation distances in right direction
	Anterior ratio	μm	Single event propagation offset distances in anterior direction, normalized by the sum of four directions
	Posterior ratio	μm	Single event propagation offset distances in posterior direction, normalized by the sum of four directions
	Left ratio	μm	Single event propagation offset distances in left direction, normalized by the sum of four directions
	Right ratio	μm	Single event propagation offset distances in right direction, normalized by the sum of four directions
Landmark	Event average distance	μm	Average distance of an event centroid to each landmark
	Event minimum distance	μm	Minimum distance of an event centroid to each landmark
	Event toward landmark		Score for propagating toward each landmark
	Event away from landmark	μm ³	Score for propagating away from each landmark
	Event toward landmark before reaching	μm ³	Score for propagating toward each landmark before it reaches the landmark
	Event away from landmark after reaching	μm ³	Score for propagating away from each landmark after it reaches the landmark
Region/Cell	Event centroid distance to border	μm	Distance from event centroid to the boundary of the region/cell
	Event centroid distance to border - normalized by region radius		Divide the distance to region boundary by the radius of the region
Network/ Interaction	Temporal density		Number of events that share spatial footprint with current event (including itself)
	Temporal density with similar size only		Number of events that share spatial footprint with current event and have similar size with it (50% to 200%)
	Spatial density		Number of events co-occurred with current event

References

1. Srinivasan, R. *et al.* New Transgenic Mouse Lines for Selectively Targeting Astrocytes and Studying Calcium Signals in Astrocyte Processes In Situ and In Vivo. *Neuron* **92**, 1181–1195 (2016).
2. Marvin, J. S. *et al.* An optimized fluorescent probe for visualizing glutamate neurotransmission. *Nat. Methods* **10**, 162–170 (2013).
3. Chen, T.-W. *et al.* Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295–300 (2013).
4. Srinivasan, R. *et al.* Ca(2+) signaling in astrocytes from *Ip3r2*(-/-) mice in brain slices and during startle responses in vivo. *Nat. Neurosci.* **18**, 708–717 (2015).
5. Paukert, M. *et al.* Norepinephrine controls astroglial responsiveness to local circuit activity. *Neuron* **82**, 1263–1270 (2014).
6. Haustein, M. D. *et al.* Conditions and constraints for astrocyte calcium signaling in the hippocampal mossy fiber pathway. *Neuron* **82**, 413–429 (2014).
7. Ding, F. *et al.* α 1-Adrenergic receptors mediate coordinated Ca2+ signaling of cortical astrocytes in awake, behaving mice. *Cell Calcium* **54**, 387–394 (2013).
8. Poskanzer, K. E. & Yuste, R. Astrocytes regulate cortical state switching in vivo. *Proc. Natl. Acad. Sci. U.S.A.* 201520759 (2016). doi:10.1073/pnas.1520759113
9. Ma, Z., Stork, T., Bergles, D. E. & Freeman, M. R. Neuromodulators signal through astrocytes to alter neural circuit activity and behaviour. *Nature* **539**, 428–432 (2016).
10. Poskanzer, K. E. & Yuste, R. Astrocytic regulation of cortical UP states. *Proc. Natl. Acad. Sci. U.S.A.* **108**, 18453–18458 (2011).
11. Looger, L. L. A genetically encoded fluorescent sensor for in vivo imaging of GABA. *bioRxiv* 322578 (2018). doi:10.1101/322578
12. Feng, J. *et al.* A Genetically Encoded Fluorescent Sensor for Rapid and Specific In Vivo Detection of Norepinephrine. *Neuron* (2019). doi:10.1016/j.neuron.2019.02.037
13. Lobas, M. A. *et al.* A genetically encoded single-wavelength sensor for imaging cytosolic and cell surface ATP. *Nature Communications* **10**, 711 (2019).
14. Patriarchi, T. *et al.* Ultrafast neuronal imaging of dopamine dynamics with designed genetically encoded sensors. *Science* **360**, eaat4422 (2018).
15. Agarwal, A. *et al.* Transient Opening of the Mitochondrial Permeability Transition Pore Induces Microdomain Calcium Transients in Astrocyte Processes. *Neuron* **93**, 587–605.e7 (2017).
16. Armbruster, M., Hanson, E. & Dulla, C. G. Glutamate Clearance Is Locally Modulated by Presynaptic Neuronal Activity in the Cerebral Cortex. *J. Neurosci.* **36**, 10404–10415 (2016).
17. Bindocci, E. *et al.* Three-dimensional Ca(2+) imaging advances understanding of astrocyte biology. *Science* **356**, eaai8185 (2017).
18. Wu, Y.-W. *et al.* Spatiotemporal calcium dynamics in single astrocytes and its modulation by neuronal activity. *Cell Calcium* **55**, 119–129 (2014).
19. Wang, Y. *et al.* Graphical Time Warping for Joint Alignment of Multiple Curves. *Neural Information Processing Systems* 3648–3656 (2016).
20. Giovannucci, A. *et al.* CaImAn: An open source tool for scalable Calcium Imaging data Analysis. *bioRxiv* 339564 (2018). doi:10.1101/339564
21. Pachitariu, M. *et al.* Suite2p: beyond 10,000 neurons with standard two-photon microscopy. (2016). doi:10.1101/061507

22. Poskanzer, K. E. & Molofsky, A. V. Dynamism of an Astrocyte In Vivo: Perspectives on Identity and Function. *Annu. Rev. Physiol.* **80**, 143–157 (2018).

23. Shigetomi, E. *et al.* Imaging calcium microdomains within entire astrocyte territories and endfeet with GCaMPs expressed using adeno-associated viruses. *J. Gen. Physiol.* **141**, 633–647 (2013).

24. van der Maaten, L. & Hinton, G. Visualizing Data using t-SNE. *Journal of Machine Learning Research* **9**, 2579–2605 (2008).

25. Stobart, J. L. *et al.* Cortical Circuit Activity Evokes Rapid Astrocyte Calcium Signals on a Similar Timescale to Neurons. *Neuron* **98**, 726–735.e4 (2018).

26. Kanemaru, K. *et al.* In vivo visualization of subtle, transient, and local activity of astrocytes using an ultrasensitive Ca(2+) indicator. *Cell Rep* **8**, 311–318 (2014).

27. Nimmerjahn, A., Mukamel, E. A. & Schnitzer, M. J. Motor behavior activates Bergmann glial networks. *Neuron* **62**, 400–412 (2009).

28. Parsons, M. P. *et al.* Real-time imaging of glutamate clearance reveals normal striatal uptake in Huntington disease mouse models. *Nature Communications* **7**, 11251 (2016).

29. Jiang, R., Diaz-Castro, B., Looger, L. L. & Khakh, B. S. Dysfunctional Calcium and Glutamate Signaling in Striatal Astrocytes from Huntington's Disease Model Mice. *J. Neurosci.* **36**, 3453–3470 (2016).

30. Hefendehl, J. K. *et al.* Mapping synaptic glutamate transporter dysfunction in vivo to regions surrounding A β plaques by iGluSnFR two-photon imaging. *Nature Communications* **7**, 13441 (2016).

31. Marvin, J. S. *et al.* Stability, affinity, and chromatic variants of the glutamate sensor iGluSnFR. *Nat. Methods* **15**, 936–939 (2018).

32. Fino, E. *et al.* RuBi-Glutamate: Two-Photon and Visible-Light Photoactivation of Neurons and Dendritic spines. *Front Neural Circuits* **3**, 2 (2009).

33. Chai, H. *et al.* Neural Circuit-Specialized Astrocytes: Transcriptomic, Proteomic, Morphological, and Functional Evidence. *Neuron* **95**, 531–549.e9 (2017).

34. Baraban, M., Koudelka, S. & Lyons, D. A. Ca 2+ activity signatures of myelin sheath formation and growth in vivo. *Nat. Neurosci.* **19**, 190 (2017).

35. Krasnow, A. M., Ford, M. C., Valdivia, L. E., Wilson, S. W. & Attwell, D. Regulation of developing myelin sheath elongation by oligodendrocyte calcium transients in vivo. *Nat. Neurosci.* **93**, 9887 (2017).

36. Higley, M. J. & Sabatini, B. L. Calcium signaling in dendritic spines. *Cold Spring Harb Perspect Biol* **4**, a005686–a005686 (2012).

37. Keller, P. J. & Ahrens, M. B. Visualizing whole-brain activity and development at the single-cell level using light-sheet microscopy. *Neuron* **85**, 462–483 (2015).

38. Goldey, G. J. *et al.* Removable cranial windows for long-term imaging in awake mice. *Nat Protoc* **9**, 2515–2538 (2014).

39. Pnevmatikakis, E. A. & Giovannucci, A. NoRMCorre: An online algorithm for piecewise rigid motion correction of calcium imaging data. *J. Neurosci. Methods* **291**, 83–94 (2017).

40. Müller, M. *Information Retrieval for Music and Motion*. (Springer Science & Business Media, 2007). doi:10.1007/978-3-540-74048-3

41. Wang, Y. *et al.* Automated Functional Analysis of Astrocytes from Chronic Time-Lapse Calcium Imaging Data. *Front Neuroinform* **11**, 48 (2017).