

Tracking the diffusion of GABA_A receptors

Supervisors: Josef Kittler (Physiology), Lewis Griffin (Computer Science)

Word count: 4,016

Abstract

This essay discusses the experimental process of tracking the diffusion of GABA_A receptors in the cell membrane of neuronal cells by fluorescent confocal microscopy, and the mathematical issues arising from the interpretation of the results – i.e. the motion tracking task of converting a video of the movement of blurry ‘blobs’ into estimates of underlying diffusion coefficient distribution.

The first section gives the biological context for the problem including definitions for terms such as GABA_A receptor. Section two describes typical experimental methods used to carry out fluorescent microscopy on receptors. Section three discusses the mathematical process of motion tracking as applied to the experimental results, and discusses some of the inherent limitations imposed by the experimental procedure.

It should be noted that although this essay is based upon the study of GABA_A receptors, the experimental procedures and mathematical issues of motion tracking are applicable to a wider range of biological entities including other neuronal receptor types and other sub-cellular structures.

1. Biological context

The vertebrate brain is composed of an extremely complex network of nerve cells called neurons. The exact morphology and topology of a neuron varies across the different regions of the brain, however in general a neuron can be said to be composed of a set of branching dendrites acting as inputs to the neuron, a soma (cell body), and an axon (also with a branching structure) forming outputs from the neuron. Each neuron is therefore connected to thousands of other neurons both as ‘inputs’ and ‘outputs’, and the relationship between inputs and outputs of a given neuron is analogous to some highly complex non-linear form of signal processing. Inputs can be excitatory or inhibitory; that is the inputs can make it either more or less likely for the neuron to produce an output (by means of an ‘Action Potential’).

Signals between nerve cells are passed by the means of chemical stimulation of one neuronal cell by the other at junctions called synapses. The synapse is a sub-micron gap between the cells where neuro-transmitter is secreted from the presynaptic membrane of one neuron and diffuses rapidly to the postsynaptic membrane of the other neuron where specific receptor proteins bind to the chemical ‘agonist’. Activated receptors will generally trigger the flow of some specific ionic species (often Ca²⁺, Na⁺ or Cl⁻) through an ion-channel embedded in the cell membrane of the neuron. ‘Ionotropic’ receptors are those receptors incorporating an ion-channel domain such that the binding of the agonist causes a conformational change allowing ions to flow through the ion-channel part of the

protein. The electrical signal produced by the flow of ions is transmitted down through the dendritic network where it is integrated with the other input signals being passed to the cell.

The receptor proteins are highly specific in which chemical molecules may be bound as agonists. Similarly the receptor may bind specific antagonist molecules which deactivate the receptor and modulator molecules which enhance or decrease the sensitivity of the receptor to the agonist/antagonist molecules. Clearly an understanding of how chemical molecules interact with receptor proteins is important for designing drugs targeted at modifying neuronal activity.

This essay is primarily concerned with a sub-class of Gamma-aminobutyric acid (GABA) receptors, called GABA_A receptors. GABA_A receptors are inhibitory ionotropic receptors – binding of the agonist GABA produces a conformational change in the receptor protein allowing passage of Cl⁻ ions whose net effect on the neuron is to reduce the likelihood of producing an overall output.

GABA_A receptors are pharmacologically and physiologically interesting since they are among the receptors whose activity is modulated by drugs such as benzodiazepines and barbiturates. These drugs have the effect of altering the frequency and duration of receptor ion-channel opening. Making the ion-channels open more often and for longer will produce an overall decrease in the activity of the neuronal network (since these receptors are inhibitory receptors) which can be registered in the patient as a decrease in alertness or tension, whereas modulator chemicals which decrease the activity and effectiveness of GABA_A receptors will produce the opposite effect. Too much influence in either direction could lead to the possibility of global over-excitement of the brain (resulting in seizures) or global under-excitement of the brain (resulting in unconsciousness / coma).

In addition to the effect of drugs on the action of GABA_A receptors it is also important to understand the spatial distribution of receptor proteins in the synaptic region. Previous studies have established that receptor proteins are not randomly spread around the entire cell membrane of the neuron but are quite precisely located around the synapses with only a relatively small number of receptors in the extrasynaptic region of the membrane [1].

Commonly used mathematical models for the receptor motion are free diffusion, anomalous diffusion, or confined motion [2]. Determining which regime is most appropriate, and in what region each regime applies, allows conclusions to be made about what sub-cellular structures contribute to receptor confinement. For example [3] shows experimentally that gephyrin is one membrane protein which is important in localising receptors to the synaptic region. Other cytoskeletal elements are also thought to help anchor the receptor proteins in place in the membrane [4].

In order to draw such conclusions about receptor dynamics at synapses it is necessary to be able to directly observe the receptor protein movement, and to measure the diffusivity

of individual receptors, or groups of receptors, both in the synaptic region and in the extrasynaptic region.

2. Experimental procedures

Thanks to recent advances in the area of fluorescence microscopy it is now possible to fluorescently tag individual receptor molecules and observe their movement using a confocal microscope with suitable laser illumination to excite the fluorophor. Fluorophors may be either artificially manufactured, as in the case of semiconductor quantum dots (QD) or may be naturally occurring fluorescent proteins, e.g. the green fluorescent protein (GFP) family. A large number of chemical fluorophors are readily available with varying characteristics (brightness, absorption/emission characteristics, toxicity); the selection of the 'best' fluorophor for a given experiment is part of the process of designing the experiment. Depending on the fluorophor chosen, there are a number of ways to chemically attach the fluorescent probe to the receptor protein to be studied.

Natural fluorophors such as GFP may be attached to the receptor protein by genetic manipulation of the DNA sequence for the receptor protein – the sequence for GFP is simply appended to the DNA sequence for the receptor, and then the modified gene sequence is inserted into a culture of neuronal cells, some of which will begin to express the new protein receptor which should have the GFP protein domain attached to it, allowing it to be tracked optically. The advantage of this technique is that the GFP protein is relatively stable and non-toxic to the cell, and the transfection process is now well established. Drawbacks include that the GFP protein itself is rather large and so may interfere with the working of the protein to which it is attached. Furthermore it is necessary to culture cells and transfect them with the modified protein sequence, not all of which will satisfactorily take up and express the new protein.

Quantum dots are semiconductor coated latex beads engineered to fluoresce at a specific wavelength. Advantages include good photostability and high efficiency (i.e. they fluoresce brightly) with a large range of colours and sizes being readily available commercially. In order to attach a quantum dot probe to a type of receptor it is necessary to employ immuno-cyto-chemistry, whereby antibodies for the receptor type are obtained and chemically attached to the QD. The fluorescently attached antibodies are then injected into the cell in order to allow the antibodies to bind with the receptors, resulting in fluorescently labelled receptors. A drawback with this method is that some antibodies may not bind with any receptors resulting in high background noise (compared to GFP which is always expressed alongside the receptor protein). Also depending on the antibody valency it may be possible to bind multiple antibodies to a single receptor giving misleading brightness results.

Using combinations of these two experimental procedures simultaneously allows multiple fluorophors to be attached to different biological entities of interest in the cell and for them to be imaged concurrently, thus maximising the information that may be obtained from each experiment.

3. Motion tracking

The results from these experiments will be video sequences of fluorescently tagged proteins, from which diffusion characteristics must be obtained. Although it is possible to draw qualitative conclusions from such video sequences ‘by eye’ (the human brain is quite good at carrying out motion tracking) it is technically challenging to use a computer to obtain quantitative information such as diffusion coefficients.

The problem can be split into two parts: segmentation and motion tracking. Segmentation is the determination of the co-ordinates of each receptor protein in a single frame. Since the fluorophors are smaller than the wavelength of light which they appear as ‘blobs’ in each frame rather than sharp points. Apart from introducing uncertainty as to exact location, this also has the effect of making all the points appear similar so any identifying features of individual receptors are not available to aid tracking. The ‘blobs’ in the frame are best modelled by a point-spread-function (PSF), or more simply by a Gaussian [5]. Although overlapping ‘blobs’ and variations in brightness complicate this process, for the purposes of this essay it will be assumed to be the less problematic part and henceforth it is assumed that a list of co-ordinates of receptors is available for each frame in the video.

Motion tracking is concerned with the ‘correspondence problem’: given two successive frames of a video, how to match up one particle in the first to another in the second and be able to say with high confidence that they are the same particle, given that each ‘dot’ has no uniquely identifying features. In addition there is no guarantee that each dot in a frame will have a matching dot in the next frame due to the issues listed below.

There are a number of physical limitations imposed by the experimental techniques including:

Limited field-of-view (FOV) of the microscope – receptors can migrate in/out of the field of view, thus there will not necessarily be a 1:1 correspondence between dots in each frame.

Detection errors / shot noise – the video images are blurry and the high gain amplifiers are susceptible to shot noise, resulting in some false identification of receptors in each frame which will not truly correspond to any other receptors in previous/successive frames, but which the motion tracking algorithm will nonetheless have to deal with.

Endocytosis / exocytosis - in addition to lateral motion through the cell membrane GABA_A receptor proteins also have a fixed lifetime before they are reabsorbed into the cell through endocytosis for recycling or destruction, while new receptors are carried from within the cell to the cell membrane by exocytosis. In motion tracking terms this means receptors will suddenly ‘appear from nowhere’ in the video, and similarly disappear.

Fluorophor blinking – the fluorophors may temporarily enter non-emitting highly excited ‘triplet’ states, or undergo photochemical reactions so they no longer function as fluorophors. In motion-tracking terms this will be similar to the situation with endo/exocytosis, except that in the case of photochemical breakdown the image will get progressively fainter as the fluorophors become ‘photobleached’. Practically this limits the length of time that a cell can be observed, and the maximum frame rate. (Use of a two-photon excitation system can mitigate against this at the expense of more complicated experimental setup.)

CCD integration time – the time between successive frames is limited by the integration time of the camera attached to the microscope (most likely a charge coupled device (CCD) camera). This may result in motion-blur of the images depending on how fast the receptors are moving - though this in itself could be useful when it comes to inferring the direction of motion, and thus most likely vector to corresponding dot in the next frame. More problematic will be that the higher the diffusivity, the further the receptors will be expected to have moved between frames leading to a number of problems: increased likelihood of a receptor leaving the field of view (leading to error), increased likelihood that the receptor will have moved sufficiently far that it may be confused with another nearby receptor in the next frame (getting worse the more densely packed the receptors in the field of view).

Hardware bandwidth – the bandwidth available in the hardware for collecting and recording data imposes a limit on the number of pixels per frame, and frames per second, which may be collected, though this is unlikely to be more of a limiting factor than the CCD integration time discussed above.

3.1 Least-squares algorithm

The most simple approach for motion tracking between frames is to make a number of assumptions: no blinking of fluorophors; no endo/exocytosis; no particles enter or leave field of view – allowing an assumption that there is 1:1 correspondence between particles in successive frames. In addition we assume that the particles carry out a ‘random walk’ following a Gaussian probability distribution, such that:

$$P(r) = \frac{1}{4\pi Dt} \exp\left(-\frac{r^2}{4Dt}\right)$$

Where $P(r)$ is the probability that the particle is displaced by r in time t , given diffusion coefficient D . The diffusion coefficient D may be calculated from the Mean Square Displacement (MSD) function (i.e. the mean of the squared displacements of all the particles in time t) where:

$$MSD = 4Dt$$

Under these assumptions, we can match each of N particles in the first frame to each particle in the second frame, and calculate the probability of each set of matching scenarios. We can then assume that pairing up of particles which gives the greatest overall probability is the correct configuration of matching of pairs of particles between successive frames – this is a Maximum Likelihood (ML) calculation. In fact for numerical convenience one could sum the natural logarithms of probabilities, which results in the maximum log-likelihood. In the case of a Gaussian probability distribution is the same thing as minimising the square of the displacements, which is computationally convenient.

As shown by the proof-of-concept Mathematica notebook given in Appendix, this method appears sufficient to carry out motion tracking with surprisingly good accuracy even when the diffusivity of the particles is relatively high. Furthermore it was demonstrated that underlying variations in diffusivity can be detected even though the least-squares approach as written assumes constant diffusivity. Using an iterative approach to solving the underlying diffusivity (by appropriately modifying the probability of how far a particle will move, based on its x,y co-ordinates) could allow accurate determination of the diffusivity distribution in the field-of-view.

In biological terms this numerical data for diffusivity leads to conclusions about the nature of the confinement of receptor proteins in the cell membrane. Concurrent imaging of different protein structures in the membrane could be used to link diffusivity measurements to the physical structures which help confine receptor proteins within the postsynaptic membrane.

However the algorithm suffers from severe deficiencies in scalability. The computation time scales with N factorial since every possible permutation of particles matching in successive frames is tested for optimality. The proof of concept Mathematica notebook is limited to approximately $N=8$ for reasonable computation time. (Obviously the algorithm scales linearly with number of frames, so at a certain point that too becomes a limiting factor).

3.2 Dense tracking problem

In the case of a densely packed field of view it also seems intuitive that the algorithm will begin to make mistakes, especially where the expected diffusion distance (per frame) of the receptors exceeds the mean inter-receptor spacing. It seems natural to assume that under these conditions the algorithm would incorrectly match points in consecutive frames because any point in one frame will always be a short distance from some other (unrelated) point in the next frame, at a distance less than the real diffusion length. Assuming that this is true for all points, one might expect that the least squares algorithm will match these nearby points rather than making the correct matches, which in terms of the original distribution would correspond to a 'narrower' Gaussian distribution (leading to a conclusion of lower diffusivity than is in fact the case).

The best way to test this hypothesis would be to simulate this situation using Mathematica and to observe any tracking mistakes. However due to the low value of N at which the problem becomes computationally intractable it is not feasible to test the hypothesis fully. Simulations with nine points and increasingly high diffusivity (i.e. high variance used in the Gaussian distribution used to generate the motion) have shown that the algorithm only rarely makes a mistake, even when particles come very close to each other. Although this is by no means conclusive proof that the least squares algorithm will not fail as suggested above, it indicates that the algorithm is more robust than ‘common sense’ would naturally assume.

A possible explanation is that the least square algorithm effectively computes the maximum likelihood function for a Gaussian distribution of point displacements, which is the true underlying pattern of displacements. Displacements produced by mistakenly tracking to other points are likely to be uniformly distributed since points are independent of each other, and so those displacements could be attributed to a superposition of a uniform distribution to the real Gaussian distribution. In essence the breakdown above suggests that the least squares algorithm will somehow start to turn into a ‘nearest neighbour’ algorithm at sufficiently high numbers of particles. Alternatively: the distribution of nearest neighbours to points in consecutive frames starts to obey a Gaussian distribution which the least squares algorithm maximises instead of the underlying diffusion pattern.

[5] suggests that the separability of statistical distributions of point displacements may actually be utilised to carry out motion tracking with $O(N^2)$. The basis of the method is to calculate the possible displacements for each point in one frame to each point in the next frame independently (there will be N^2 displacements in total, N of which will correspond to the real motion of the points, the rest being uniformly distributed tracking errors to independent and uncorrelated points). This distribution is posited to be a Gaussian distribution (the actual diffusion) superposed with a uniform distribution (the unrelated points), and it is suggested that the Gaussian portion should be easily identifiable (see Figure 1 (taken from Figure 5 of [5])).

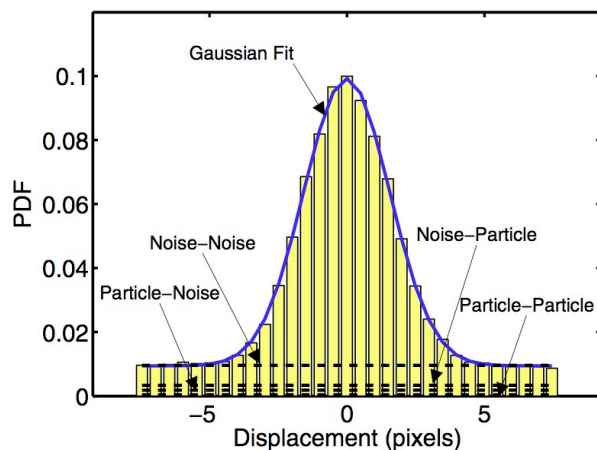


Figure 1 (from [5]) – Displacement distribution of simulated diffusion data including random tracking contributions.

Mathematica simulations of this idea (see Appendix) produced displacement data which was indeed very fast to compute even with high numbers of simulated receptors, however the mathematical alchemy required to produce a Gaussian describing the diffusion from this data remained elusive; useful results were not obtained (Figure 2).

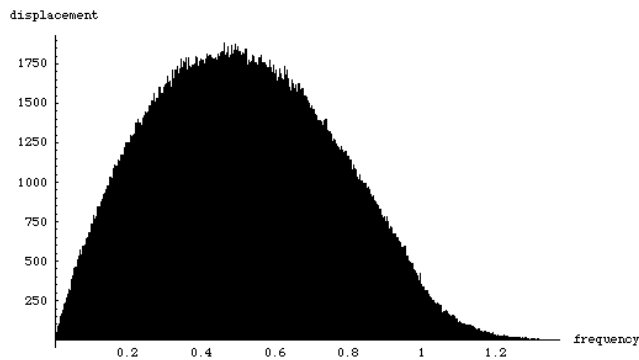


Figure 2 – Histogram of N^2 displacements for points from one frame to the next, acting under diffusion. Attempts to extract the Gaussian distribution of the actual diffusion from this data (cf. Figure 1, taken from [5]) failed miserably.

3.3 Improvements

The best approach to improving on the least squares algorithm is to reduce the number of permutations being calculated. The first priority should be a windowing approach based on some *a priori* estimate of the diffusivity such that only pairs of points within a physically plausible diffusion length are tracked – it is obvious that points on opposite sides of the field of view in successive frames are not correlated. In addition to being used as the basis of windowing, the *a priori* estimate could be used to validate the final distribution of displacements selected in that very small diffusivities could indicate that the tracking mismatching problem discussed above is occurring.

Experimentally one could try to validate the diffusivity value yielded by the least squares algorithm by using an alternative quantitative technique for finding the bulk diffusivity of the receptors, such as FRAP (Fluorescence Recovery After Photobleaching¹). The FRAP measurement could be used as the *a priori* estimate above.

Other improvements to the algorithm are also required to take into account the realities of particle movement into and out of the field of view, fluorophor blinking, and endo/exocytosis.

Alternative algorithms which allow for these complications are described in [6] and [7]. [6] seems to be a nearest neighbour algorithm where points are attributed to the nearest

¹ a measure of diffusivity by timing how long a region takes to show fluorescent activity after the fluorophors have been deliberately destroyed by photobleaching

track in the previous frame, taking into account the diffusivity already inferred for that track. The authors suggest that the algorithm scales well for large numbers of fluorophores and could be adapted for real-time tracking. How such a nearest-neighbour algorithm performs under 'dense tracking' conditions would be of interest; it is tempting to postulate that such an algorithm would be susceptible to the mismatching discussed above.

[7] uses a more mathematically sophisticated method, in that the fluorescent images are combined in a single stack to form a three dimensional cloud of points. Vertical tracks through this space correspond to motion of individual fluorophores and the authors attempt to reconstruct the tracks by considering the problems as one of 'minimal paths in a Riemannian metric'. Although the method also suffers from possible computational intractability ('combinatorial explosion') the authors claim that use of a locally optimal 'greedy algorithm' allows solutions to be computed in 'reasonable computational time'.

4 Conclusion

The experimental procedure for fluorescently tagging receptor proteins and observing their motion under a microscope is well established. Interpreting the results of these experiments, in order to draw conclusions about receptor movement in the cell membrane and their confinement by sub-cellular scaffold structures, requires solving a traditional motion tracking problem, with a few extra complications caused by the experimental limitations. This problem is analogous to problems in engineering areas such as robotic vision and flow velocimetry, and mathematical techniques for approaching this problem already exist, although there is no 'correct answer' and research is ongoing.

A least squares algorithm can be shown to work well in simulation, however the low number of points it may track before the computation time increases unacceptably limits its usefulness in practice. A modified least squares algorithm which makes some assumptions about underlying diffusivity will be more practical. The performance of least squares algorithms under dense tracking conditions remains unknown, however it is speculated that the algorithm will continue to work well in comparison to other possible approaches, like nearest neighbour.

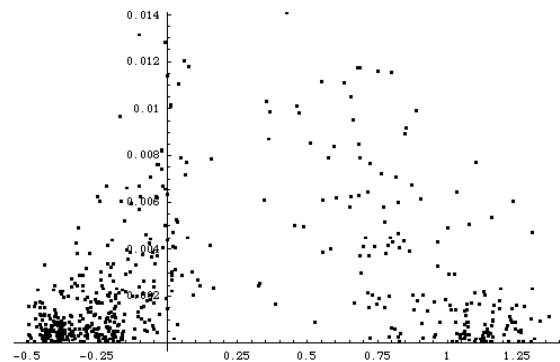
Attempts to simulate a statistical method for calculating diffusion characteristics (following [5]) failed. However as a general principle, it may be beneficial to consider approaches to tracking which do not rely on accurately tracking individual points, but yield some statistical estimate of diffusivity based on changes in density of points in successive frames, as this will give a more computationally tractable algorithm.

References

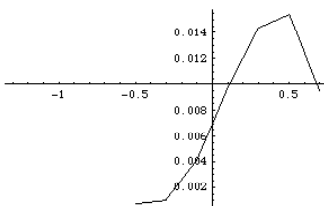
- [1] M. Dahan, S. Lévi, C. Luccardini, P. Rostaing, B. Riveau, A. Triller, “Diffusion Dynamics of Glycine Receptors Revealed by Single-Quantum Dot Tracking”, *Science*, **302** 442 (2003)
- [2] M. Saxton, K. Jacobson, “Single Particle Tracking: Applications to Membrane Dynamics”, *Annu. Rev. Biophys. Biomol. Struct.* **26** 373 (1997)
- [3] T. Jacob, Y. Bogdanov, C. Magnus, R. Saliba, J. Kittler, P. Haydon, S. Moss, “Gephyrin Regulates the Cell Surface Dynamics of Synaptic GABA_A Receptors”, *Journal of Neuroscence*, **25** 10469 (2005)
- [4] D. Choquet, A. Triller, “The Role of Receptor Diffusion in the Organization of the Postsynaptic Membrane”, *Nature Reviews (Neuroscience)*, **4** 251 (2003)
- [5] J. Guasto, P. Huang, K. Breuer, “Statistical particle tracking velocimetry using molecular and quantum dot tracer particles”, *Proceedings of IMECE 2005, 2005 ASME International Mechanical Engineering Congress and Exposition* (2005)
- [6] S. Bonneau, L. Cohen, M. Dahan, “A Multiple Target Approach for Single Quantum Dot Tracking”, *Biomedical Imaging: Macro to Nano, 2004. IEEE International Symposium* **1** 664 (2004)
- [7] S. Bonneau, M. Dahan, L. Cohen, “Single Quantum Dot Tracking Based on Perceptual Grouping Using Minimal Paths in a Spatiotemporal Volume”, *IEEE Transactions on Image Processing* **14** 1384 (2005)

Appendix

Mathematica notebook written by Lewis Griffin, for proof of concept of least squares algorithm.

[illegible]

```
In[120] = ListPlot[Table[{lim + 0.1, Mean[Select[tmp, (#[[1]] ≥ lim && #[[1]] ≤ lim + 0.2) &][[All, 2]]}], {lim, -1.4, 0.6, 0.2}],  
PlotJoined → True];
```



Modified notebook by David Fallaize.

```

In[1]:= << Statistics`ContinuousDistributions`
        << Graphics`Colors`
        << Graphics`Graphics`
        << Statistics`ClusterAnalysis`
        << LinearAlgebra`MatrixManipulation`

In[6]:= startingLocations[n_] := Array[{#, (Random[Real, {0, 1}], Random[Real, {0, 1}])} &, n]

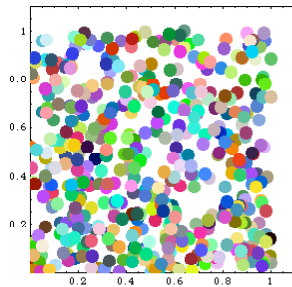
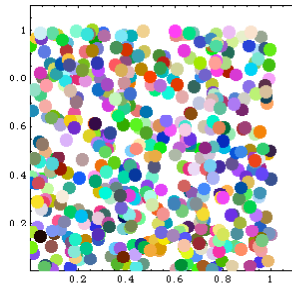
In[7]:= updateLocationsUniformMobility[pts_, e_] := Map[{#1, Mod[#2, 1] + {Random[NormalDistribution[0, e]], Random[NormalDistribution[0, e]]}. 1}] &, pts]

In[30]:= locationHistory = NestList[updateLocationsUniformMobility[#, 0.1] &, startingLocations[500], 1];

In[8]:= dspHistory[p_] := With[{fov = Max[Abs[p][All, All, 2]]}, colours = Array[RGBColor[Random[], Random[], Random[]] &, Length[p[1]]]];
        Map[Show[Graphics[{PointSize[0.05], MapThread[{#2, Point[#1[[2]]]} &, {#, colours}]}], PlotRange -> {1.1 fov {0, 1}, 1.1 fov {0, 1}},
            AspectRatio -> Automatic, Frame -> True] &, p];

In[31]:= dspHistory[locationHistory];

```

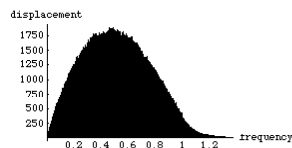


```

In[33]:= bareHistory = Map[{#1[[All, 2]]} &, locationHistory];
        firstFrames = Take[bareHistory, {1, -2, 1}];
        nextFrames = Take[bareHistory, {2, -1, 1}];
        pairsAcrossFrames = MapThread[Tuples[{#1, #2}] &, {firstFrames, nextFrames}];
        distances = Map[Map[EuclideanDistance[{#1[[1]], #1[[2]]}] &, #] &, pairsAcrossFrames];

In[32]:= Histogram[Flatten[distances], AxesLabel -> {"frequency", "displacement"}]

```



Out[32]= - Graphics -