Part I

The Relationship Between Structure and Dynamics in a Supercooled Liquid

Introduction to Part I

As explained in Section 1.2.1, the transition to rigidity in glass-forming liquids involves the appearance of spatially heterogeneous dynamics. Using this as a handle, we examine the relationship between structure and dynamics in a supercooled glass-forming liquid. In particular, we address the following two questions: (i) Is there something in the structure that is responsible for dynamics that can vary by orders of magnitude from one region of the sample to another at T_g? and (ii) Can the mobile regions, as identified by anharmonic excitations or higher propensities, be associated with particular types of local structure, and if so what is the structural signature of these 'soft' spots?

There are at least two ways to analyse the spatial distribution of particle dynamics. The approach that we take is to consider the correlation between the initial particle configuration and the subsequent dynamics. The second, which has received considerable attention, is to consider the correlation between the mobile particles themselves. For completeness sake, and to provide further motivation for the work presented in this part, we provide a brief review of the latter approach.

Particle motion in dense liquids is, to a large extent, entrained so that particles follow along in each other's path. In 1998, Donati et al. [68] showed that displacements in a supercooled liquid exhibited a strong tendency to locally align. The dynamics associated with such correlations are generally quite complex. To appreciate this, consider, first, the *simple* scenario of a diffusing vacancy. The linear character of this pattern of displacements reflects a rough local conservation of free volume, that is, the volume left free by the motion of a particle is more likely to be filled by a single particle rather than the collective rearrangement of a number of particles. There is also a correspondence between spatial structure and temporal sequence so that one end of the resulting string of displacements represents the first step while the other end represents the last step. Ritort and Sollich [65] have recently reviewed the predictions of a number of diffusing defect models.

The diffusing defect picture, presented above, ignores the possibility that the propensity for motion lies distributed in a configuration and that relaxation is not a simple consequence of the transport of a rare fluctuation (even one more complex than a simple vacancy) but rather a sequence of unlocking events which add up, over the observation interval, to a linear path. Vogel et al. [69] have presented simulation

evidence of just this latter process. Delaye and Limoge [70], in an interesting study, considered the different fates of vacancies created in a model glass. The resulting behaviour was divided into three groups: those defects that remained stable and stationary, those that relaxed by propagation (the diffusing defects) and a third group that relaxed by being, essentially, 'absorbed' back into the surrounding disordered medium through a local collective rearrangement. The presence of this last process distinguishes the amorphous material from the crystalline.

In terms of a simple model that can capture this more complex collective behaviour, the class of facilitated spin flip models introduced by Fredrickson and Andersen [71] and extended by Jäckle [72] is particularly useful. The term 'facilitated' refers to the idea that the local state of a system can influence the kinetics of adjacent regions. To date, models of facilitated dynamics have all relied on the introduction of explicit kinetic constraints. Recently, Garraghan and Chandler [73] have proposed that this idea of 'facilitation' represents a general aspect of dense amorphous phases. The interesting unanswered question here is whether the implied general mapping from systems of interacting particles to systems governed by kinetic constraints exists. Central to this question is the need to understand the degree to which a particle configuration determines the propensity of particles to subsequently move. This is precisely the subject of Chapter 2.

The analysis of the correlation between particle displacements sketched here provides (i) a compact summary of the information presented by dynamic heterogeneities, (ii) an explanation of some observed features of relaxation functions and transport behaviour in terms of microscopic dynamics, and (iii) the prospect of identifying kinetic rules that govern relaxation in disordered systems. This approach, however, does not explain what it is about a configuration that permits motion in one region but not in another, nor how this distribution varies with temperature, composition, particle interactions, etc. One could imagine, for example, studying transport in crystalline solids via this description, amassing a considerable amount of phenomenological information about the dynamic heterogeneities and yet never arriving at a clear structural (and, hence, predictive) picture of vacancies and interstitials. For this reason, we would like to directly address the question of the relationship between structure and dynamics, which we investigate in detail in Chapter 3.

Chapter 2

Dynamic Propensity

In this chapter we tackle the question of the degree to which a given configuration determines the subsequent particle dynamics. We begin by demonstrating that not all of the dynamics in the supercooled liquid can be related to structure. Through the introduction of the isoconfigurational ensemble method and the dynamic propensity, we are able to provide explicit proof that something in the structure is *responsible for the development of spatially heterogeneous dynamics in this system. In particular, we find that as the liquid is cooled, the structure-determined propensity for particle motion becomes increasingly heterogeneous, both in magnitude and in spatial extent. We characterise the variation of the propensity distribution with temperature and configuration, study its convergence properties, and demonstrate that additional insight into glass-formation can be obtained via the analysis of other correlations within the isoconfigurational ensemble.*

2.1 Introduction

Over the last ten years, dynamic heterogeneity has become recognised as a general phenomenological feature of glass formation [11]. The existence of these long-lived kinetic fluctuations has been useful in rationalising some puzzling aspects of kinetics in supercooled liquids. These include non-Fickian diffusion [74, 75], deviations from classical nucleation theory [76], and the breakdown of the scaling between translational diffusion, on the one hand; and, on the other, rotational diffusion [77], shear viscosity [57], and structural relaxation [78]. Helpful as these developments are, they do not address the fundamental question of the relationship between structure and

kinetics in the supercooled liquid.

The *spatial* distribution of particle mobilities, however, does appear to offer a considerable amount of information on this very point. A number of papers have considered the local connection between dynamics and structure, the latter being characterised using topology [44], potential energy [69, 79], and free volume [80]. While most have reported some correlation, none have established a correlation of sufficient strength to indicate a causal link, i.e. that the local kinetics was determined by the selected aspect of the local structure. In a recent review [11], Ediger observed, "At present, it is an article of faith that something in the structure is responsible for dynamics that can vary by orders of magnitude from one region of the sample to another at T_g."

Instead of trying to directly address the question 'What aspect of the structure gives rise to the observed dynamic heterogeneity?', we will answer the related question, 'What aspect of the dynamic heterogeneity actually arises from the structure?' It is logically necessary to answer this question before attempting the first. As we show, it is also possible to answer the latter question without having to first identify the correct measure of the particle structure relevant to determining the subsequent dynamics.

The degree to which the liquid dynamics reflects a persistent influence of a configuration is related to the idea, introduced in Section 1.3, of the crossover between liquid-like and solid-like behaviour on cooling. One measure of this transition from the liquid to solid-like descriptions is the crossover temperature proposed by Goldstein in 1969 [81]. The crossover temperature marks the transition from the high-temperature liquid - where momentum transfer (as binary collisions and, collectively, in hydrodynamic modes) plays a dominant role - to the low-temperature liquid in which dynamics is said [82] to become 'landscape dominated', the landscape referring to the potential energy surface over the configuration space. The configuration space, in other words, has begun to break up into kinetically isolated domains.

We answer the question - 'what aspect of particle dynamics in a liquid is determined by the initial configuration?' - with the construction of a new measure of structure-related dynamics, which we term the *dynamic propensity*. The isoconfigurational method used to calculate the dynamic propensity, and the analysis of the resulting ensemble of runs, are the subject of this chapter.

The rest of this chapter is structured as follows. In Section 2.2 we describe the

model glass-former and algorithms used in this work and introduce the general isoconfigurational ensemble method. We begin, in Section 2.3, by testing the reproducibility of the dynamics, and demonstrate that not all of it can have a structural origin. Then we introduce the dynamic propensity as a measure of structure-dependent particle dynamics, and show that this new dynamic quantity is also spatially heterogeneous. The effect of varying configuration and temperature is examined, and the convergence properties of the propensity are also discussed. In Section 2.4 we show that the variance in individual particle motion from run to run represents a new piece of dynamic information, and discuss the consequences of this 'intermittent' particle motion. And in Section 2.5 we demonstrate how the isoconfigurational ensemble method can be used to obtain insight into the process of relaxation by looking at correlations between particle motion within the same run and within different runs of the isoconfigurational ensemble. Finally, we summarise the key results of this work in Section 2.6, discuss their implications, and suggest some areas for further study.

2.2 Model and Algorithms

We consider a two-dimensional (2D) glass-forming liquid consisting of an equimolar binary mixture of particles interacting via purely repulsive potentials of the form

$$u_{ab}(r) = \epsilon \left[\frac{\sigma_{ab}}{r}\right]^{12} \tag{2.1}$$

where $\sigma_{12} = 1.0 \times \sigma_{11}$ and $\sigma_{22} = 1.4 \times \sigma_{11}$. All units quoted will be reduced so that $\sigma_{11} = \epsilon = m = 1.0$ where m is the mass of both types of particle. Specifically, the reduced unit of time is given by $\tau = \sigma_{11}\sqrt{m/\epsilon}$.

A total of $N = 1024$ particles were enclosed in a square box with periodic boundary conditions. The molecular dynamics simulations were carried out at constant number of particles, pressure and temperature using the Nosé-Poincaré-Andersen (NPA) algorithm developed by Laird et al. [83, 84]. See Appendix A for further details of this method. The use of constant NPT constaints allows one to compare systems with different composition and particle interactions, as we do in Part II. The pressure ($P = 13.5$), was chosen so that our results would be directly comparable to those of Weeks et al. [85] for the single-component soft-disk system. The 'masses' of the Anderson piston and Nosé thermostat were $Q_v = 0.0002$ and $Q_s = 1000$, respectively, for all temperatures. The equations of motion were integrated using a generalised

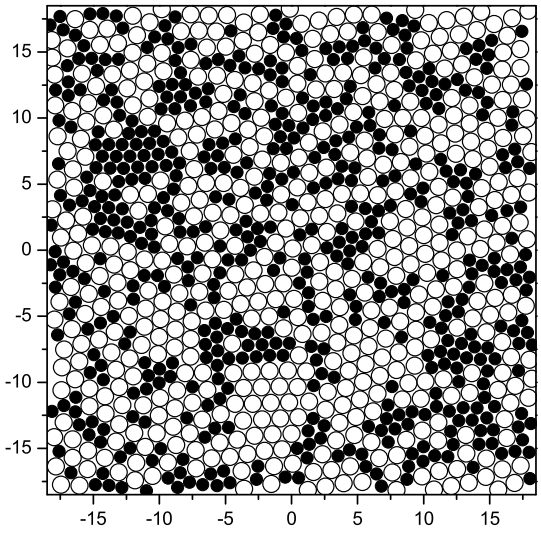

Figure 2.1: An equilibrated particle configuration at $T = 0.4$. The large and small particles are indicated by filled and open circles, respectively.

leapfrog algorithm [84], and are provided for 2D simulations in Appendix A. The time step employed was 0.05τ for $T > 1$, and 0.01τ for $T \leq 1$. For argon units of $\eta = 120k_B$, $m = 6.6 \times 10^{-23}g$ and $\sigma_{11} = 3.4\mathring{A}$, these time steps correspond to approximately 10 and 20 femtoseconds respectively.

The structural, dynamic and thermodynamic properties of this model glass-forming liquid have been characterised by Perera and Harrowell [44, 67]. For reference, the onset of the plateau region in the mean-squared displacement, and in the incoherent scattering functions, occurs near $T = 0.5$. The structural relaxation time and diffusion constants also have non-Arrhenius temperature dependence below $T = 0.55$. Together, these dynamic changes suggest the presence of at least two relaxation processes for $T \leq 0.5$.

All configurations investigated were equilibrated configurations taken from the study in ref. [67] and re-equilibrated with the NPA Hamiltonian. They represent amorphous stationary states in the sense that these states are stable over time scales at least an order of magnitude longer than the structural relaxation time and show no development of long-range correlations associated with established ordered phases. While the supercooled liquid state is strictly metastable, we will refer to such configurations as 'equilibrated'.

In this work we analyse correlations among the set of N-particle trajectories that

pass through a given configuration. To generate this *isoconfigurational ensemble* of runs at a given temperature T, we start with a configuration that has been equilibrated at T and for each run randomly assign the initial particle momenta from the appropriate Maxwell-Boltzmann distribution. A representative configuration at $T = 0.4$ is shown in Figure 2.1.

2.3 The Spatial Distribution of Dynamic Propensity

We examine the reproducibility of the dynamic heterogeneities in the supercooled liquid. Figure 2.2 shows the particle displacement vectors following three different runs starting from the same configuration - plotted in Figure 2.1 - of an equilibrated liquid. The three runs differ only in the random assignment of particle momenta from the Maxwell-Boltzmann distribution at the appropriate temperature. Each run was carried out at a pressure $P = 13.5$ and a temperature $T = 0.4$, which is below the onset-temperature of two-step relaxation.

Each plot in Figure 2.2 exhibits the now familiar features of dynamics in deeply supercooled liquids: large variations in the particle displacements, clear clustering of the 'slow' particles, and aggregation of the more mobile particles, sometimes in 'string-like' features. What is striking is that the spatial arrangement of particle

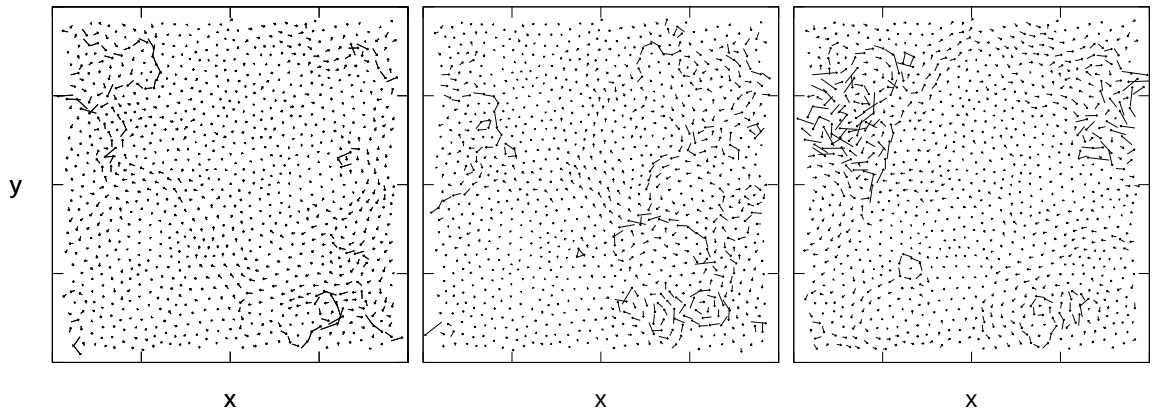

Figure 2.2: The particle displacements, indicated as vectors joining the initial to final particle positions, resulting from three MD runs of 1000τ at $T = 0.4$. All runs made use of the same initial configuration and differed only in the assignment of initial momenta to particles.

displacements differs markedly from plot to plot. While some particles exhibit a mobility that is reproducible from run to run, the dynamics of other particles varies substantially. We conclude that not all the dynamics can be determined by the initial configuration.

The variation in an individual particle's motion from run to run provides additional insight into the relationship between structure and dynamics. We explore this further in Section 2.4.

Now consider the possibility that there is no correlation at all between an initial configuration and the subsequent particle dynamics. In this case, each particle's squared displacement, averaged over many trajectories with the same initial configuration, would be the same as that of every other particle of the same species. This conclusion arises from the fact that the only point of connection between the different trajectories is the common initial configuration. It follows, therefore, that the magnitude of variation between the trajectory-averaged squared displacements of different particles of the same species is sufficient to establish the degree to which the initial configuration determines the dynamics.

To this end, we introduce the *isoconfigurational ensemble* consisting of N_{runs} separate simulation runs over a fixed time interval, all starting from the same particle configuration but with momenta randomly assigned from the Maxwell-Boltzmann distribution at the appropriate temperature. (Note that Doliwa and Heuer [86] have used multiple trajectories from a single configuration at *different* temperatures to establish the Arrhenius character of transitions between metabasins.) Let $f_i(\Delta r)$ be the ensemble distribution of the displacement of particle i over the fixed time interval. These distributions represent the ensemble characterisation of each particle's capacity for movement from a specific initial configuration. They are also invariant to time reversal. We shall refer to the second moment of $f_i(\Delta r)$, i.e. the ensemble mean of the squared displacement of particle i, $< \Delta r_i^2 >_{ic}$, as the *dynamic propensity* of particle i in the given initial configuration. The expression $< \cdots >_{ic}$ indicates an average over the isoconfigurational ensemble. We stress that 'propensity' as defined here should not be associated with the actual equilibrium distribution of trajectories that pass through a point in configuration space, since in constructing the propensity we have not taken into account the correlation between a particle's momentum and either the potential energy of that particle or the force acting on that particle.

To compare propensities from different temperatures T, we set the run time for a

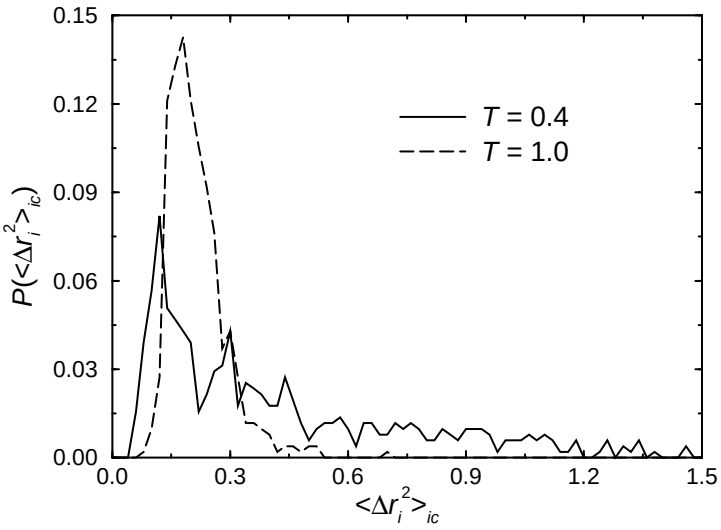

Figure 2.3: The distribution of small-particle propensities calculated using 1000 runs for single configurations at $T = 0.4$ and 1.0. Note the increase in width with increasing supercooling.

given trajectory to be 1.5 times the structural relaxation time τ_e (τ_e is defined in terms of the intermediate incoherent scattering function $F(k,t)$ such that $F(k_{max}, \tau_e) = 1/e$, where k_{max} is the wavevector of the Bragg peak and $e = 2.7183$, the base of the natural logarithm). This run interval was chosen to maximise the observed dynamic heterogeneities. If one was to choose run times much shorter or longer the dynamic heterogeneities would be unobservable - since they represent only a transient phenomenon - and all discussion of their relationship to structure would be obsolete.

Figure 2.3 shows the distribution of propensities for the small particles for configurations at $T = 1.0$ and 0.4, averaging over 1000 runs at each temperature. The width of the distribution increases substantially on cooling. As argued above, this increase in the range of the propensity distribution on cooling can only be the result of the increasing degree to which particle configurations determine the subsequent dynamics. With this result we can now replace the 'act of faith' of ref. [11] with the explicit demonstration of the heterogeneity of particle *propensities*, a feature that is completely determined by the initial configuration.

To visualise the spatial distribution of the propensity, it is useful to first remove the additional complexity of the configuration and use contour plots that contain no information about the location of individual particles. Any suitable graphing program can be used to generate a contour plot from regularly spaced data. In our case the

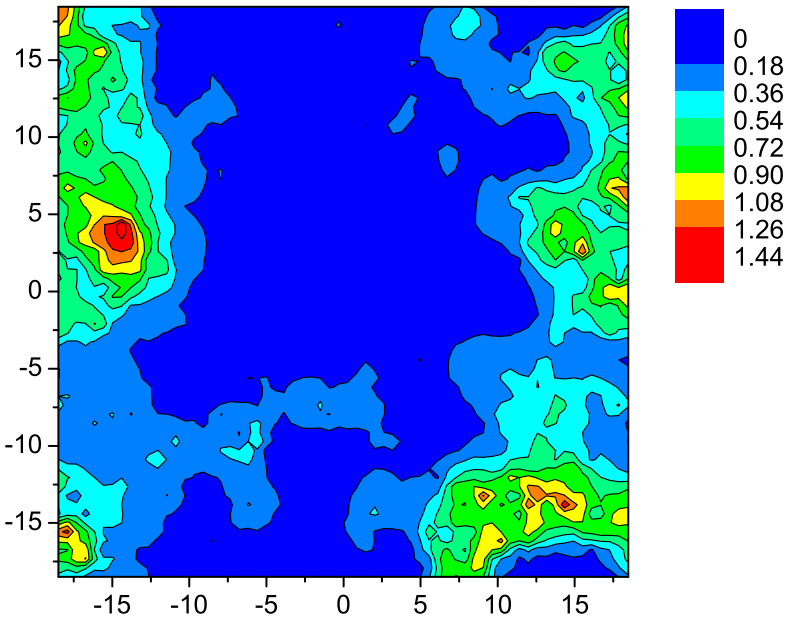

Figure 2.4: The spatial distribution of propensities at $T = 0.4$, calculated using 100 runs.

data points are generally located at irregularly spaced coordinates, namely the particle coordinates, so it is necessary to interpolate them. We found the modified version of Shepard's method [87] that is built into Origin [88] a useful algorithm for obtaining good fits to the data without introducing erroneous peaks and valleys. Interpolation parameters of 6 and 6 worked well. The occasional inconsistencies introduced by the interpolation near the periodic boundaries could be removed by fitting to a set of coordinates containing additional periodic images.

The spatial distribution of dynamic propensity for the same $T = 0.4$ configuration used in Figure 2.2 is mapped in Figure 2.4. Note the substantial spatial heterogeneity. There is a large sea of low propensity populated by islands of high propensity. On comparison of the propensity map with the individual trajectory maps, we find that the domains of high propensity are generally more compact than the often 'string-like' clusters of large displacements observed in individual trajectories. This suggests that the occasional string-like motion that is observed does not generally have a structural origin but rather must be the result of how mobility is transferred through the system. The comparison between particle trajectories and the propensity map also demonstrates that, as expected, motion occurs predominantly in regions of high propensity. We therefore conclude that the development of spatially heterogeneous dynamics in

this glass-former is due to the increasing influence of microscopic variations in structure on the motion of particles in the supercooled liquid.

An interesting question that arises is whether or not the local peaks in propensity represent 'defects'. We stress that the peaks indicate those particles with the highest average mobility over the runs examined, and may therefore not be the initiators of motion but rather some essential path through which relaxation is propagated. The propensity can provide only limited insight into the actual process of relaxation since it is the overlap of the individual relaxation processes. Other analyses of the isoconfigurational ensemble, however, can provide some insight into relaxation, e.g. the analysis of correlations between particle motion within the same run and within different runs. This is discussed further in Section 2.5.

2.3.1 The Increasing Heterogeneity upon Cooling

Because we are effectively probing single configurations, there will be variation in the spatial structure and degree of heterogeneity from run to run. We therefore study the effect of varying both the configuration and the temperature on the shape and spatial variation of the resulting propensity distribution. We generated ten configurations each at $T = 0.4$, 0.46, 0.5, 0.55, 0.6, 0.8 and 1 and calculated the propensities. The propensities were averaged over 100 runs, and the configurations were separated by $75\tau_e$ to ensure that they were significantly different from each other. The run time was kept constant at $1.5\tau_e$ for all temperatures. As a result, the run time ranged from 1000τ at $T = 0.4$ to 1.155τ at $T = 1$.

Table 2.1: Characteristic times for the 2D glass-former. τ_e is the structural relaxation time and t^* is the time at which the non-Gaussian parameter $A(t)$ (see text) reaches a maximum. The run times used to calculate the propensity have also been listed.

T	τ_e	run time	t^*
0.4	673	1000	200
0.46	51.7	76.8	65
0.5	13.6	20.2	32
0.55	4.31	6.4	16
0.6	2.91	4.32	8.5
0.8	1.19	1.77	3.5
1	0.775	1.155	1.4

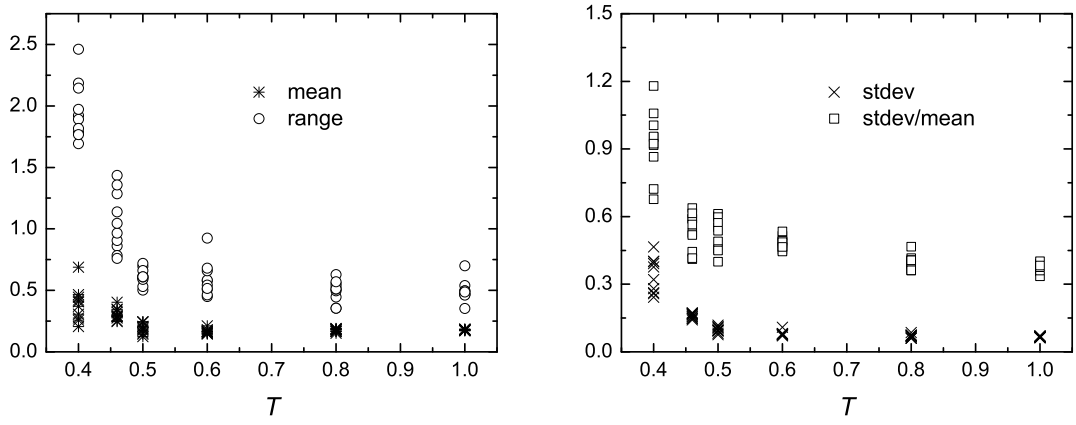

Figure 2.5: The mean, range, standard deviation (stdev) and the ratio stdev/mean for the propensity distributions calculated for ten configurations each at $T = 0.4$, 0.46, 0.5, 0.6, 0.8, and 1. At each temperature, the configurations were separated by $75\tau_e$ from each other, and the propensities were averaged over 100 runs of $1.5\tau_e$. Note the different y-axis scales.

We note that any choice of time scaling as a function of temperature is somewhat arbitrary given that different characteristic times can be chosen that scale differently with temperature. Another characteristic time that has previously been used [46] to study dynamic heterogeneities is the time t^* at which the non-Gaussian parameter $A(t)$ reaches a maximum (see Section 5.3.3 for the definition of $A(t)$). In Table 2.1 we list t^*, τ_e and the run time used to define the propensity as a function of temperature for our model glass-former, i.e. $1.5\tau_e$. The values of τ_e were taken from ref. [44] and t^* was calculated from data in ref. [89]. Note that t^* and τ_e scale quite differently. From a similar value at $T = 1$, t^* initially increases more rapidly with cooling. Later, τ_e increases more rapidly, until at $T = 0.4$ $\tau_e > 3t^*$. We note that at $T = 0.4$ the run time of 1000τ is several times longer than t^*; therefore the maximum heterogeneity in the propensity at $T = 0.4$ may well be larger than what we calculate here.

In Figure 2.5, we compare the mean, range, standard deviation (stdev) and the ratio stdev/mean for the propensity distribution calculated for each configuration. The most obvious change is a rapid increase in the range and standard deviation below $T = 0.5$, with the mean showing a smaller increase at low temperature. These changes are accompanied by an increase in the variation between different isothermal configurations, however a clear trend with change in temperature can still be distinguished. We conclude that below $T = 0.5$, there is a strong increase in the effect that

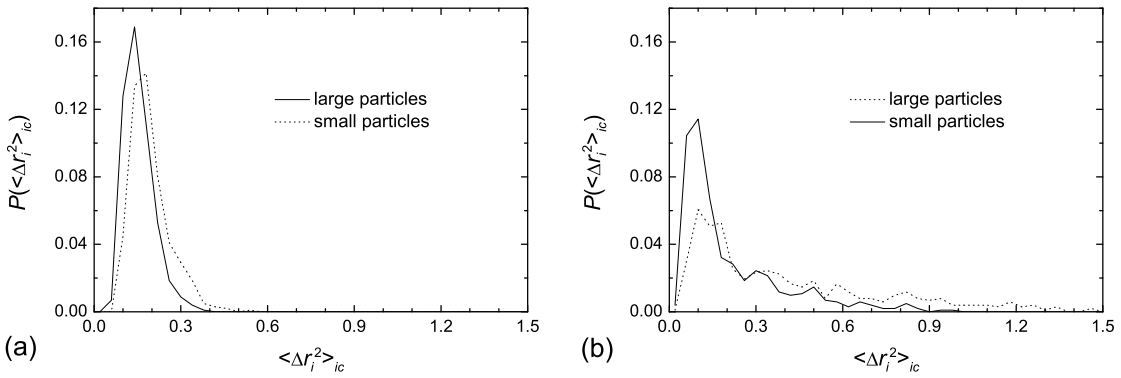

Figure 2.6: The propensity distributions over small and large particles for selected configurations at (a) $T = 1$ and (b) $T = 0.4$. The propensities were averaged over 100 runs of $1.5\tau_e$.

the structure has on the dynamics, and also in the effect that *fluctuations* in structure, both within and between different configurations, have on dynamics. In other words, the particle motion becomes increasingly sensitive to differences in structure.

In Figure 2.6 we plot the propensity distributions separately for small and large particles for individual configurations at $T = 0.4$ and 1. At $T = 1$ the distributions are very narrow and quite similar for both particle species, but as the temperature decreases the distributions become more spread out and the small and large particle distributions become less similar. The distributions still overlap, but on average the small particles have higher propensity than the large ones. This association between species type and propensity is discussed further in Section 3.2.1.

A change in temperature or configuration also affects the spatial variation of propensity. The propensity maps for four configurations at $T = 0.4$ are shown in Figure 2.7. As expected, the distribution of high and low propensity regions varies significantly from plot to plot. Particles that have low propensity in one configuration have high propensity in the next and vice versa, suggesting that sufficient change in structure has taken place to affect large changes in the spatial distribution of propensity. In Section 3.18 we present further analysis and discussion of the timescale over which structural changes affect the spatial distribution of mobility. Also note that although the mean propensity varies between configurations, i.e. some configurations are more mobile than others, their propensity distributions all have a high degree of spatial heterogeneity.

For comparison, Figure 2.8 shows the spatial distribution of propensity for selected

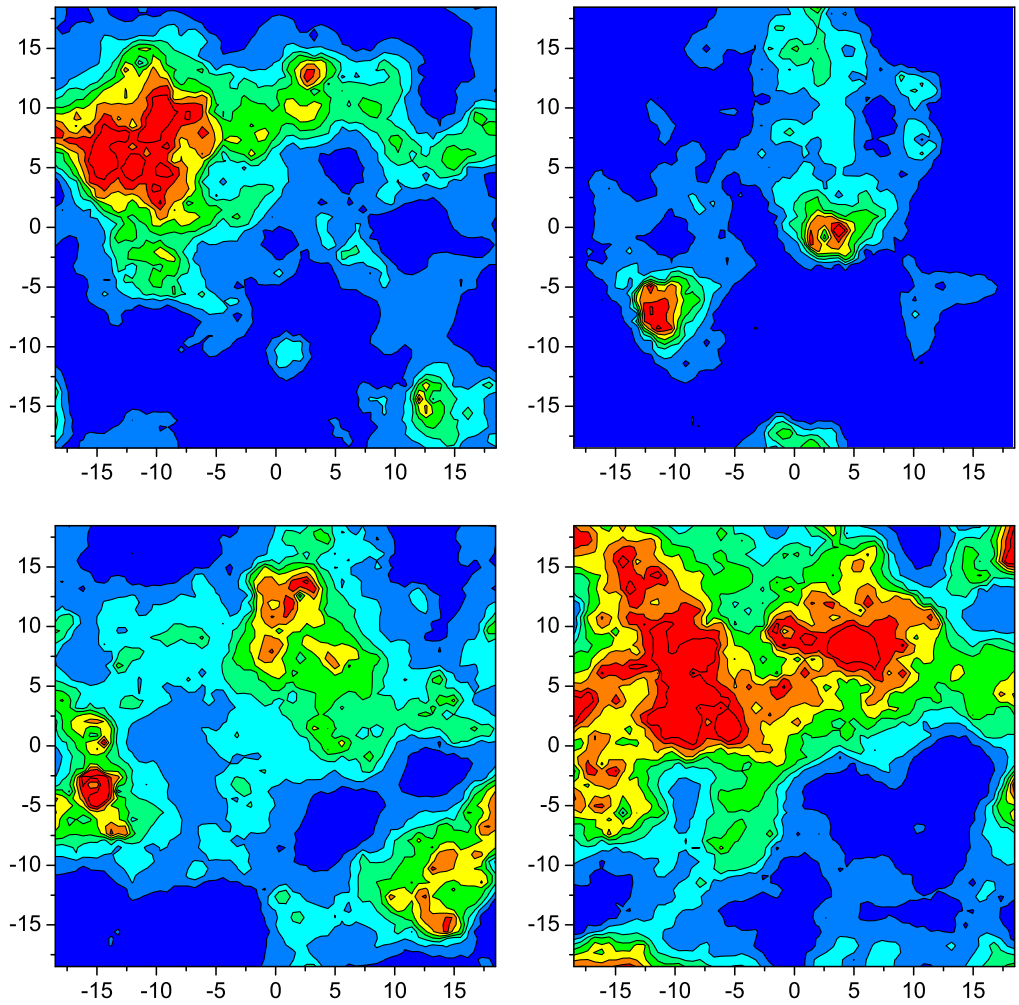

Figure 2.7: The spatial distribution of propensities at $T = 0.4$ for four configurations separated by $75\tau_e$. The propensities were averaged over 100 runs, and the scale is the same as in Figure 2.4.

configurations at $T = 0.46$, 0.5, 0.6 and 1. In order to observe the spatial variation at these higher temperatures, we have used different propensity scales than in Figures 2.4 and 2.7. Although the propensity range is similar for $T = 0.5$–1, there is a clear increase in spatial variation as the temperature is reduced (mainly due to an increase in the population of the extremes of the distribution), with further large increases at $T = 0.46$ and again at $T = 0.4$ (both of these mainly due to an increase in the width of the distribution).

There also appears to be an increase in the clustering of particles with similar mobility below $T = 1$. To better quantify this, we consider the aggregation of high

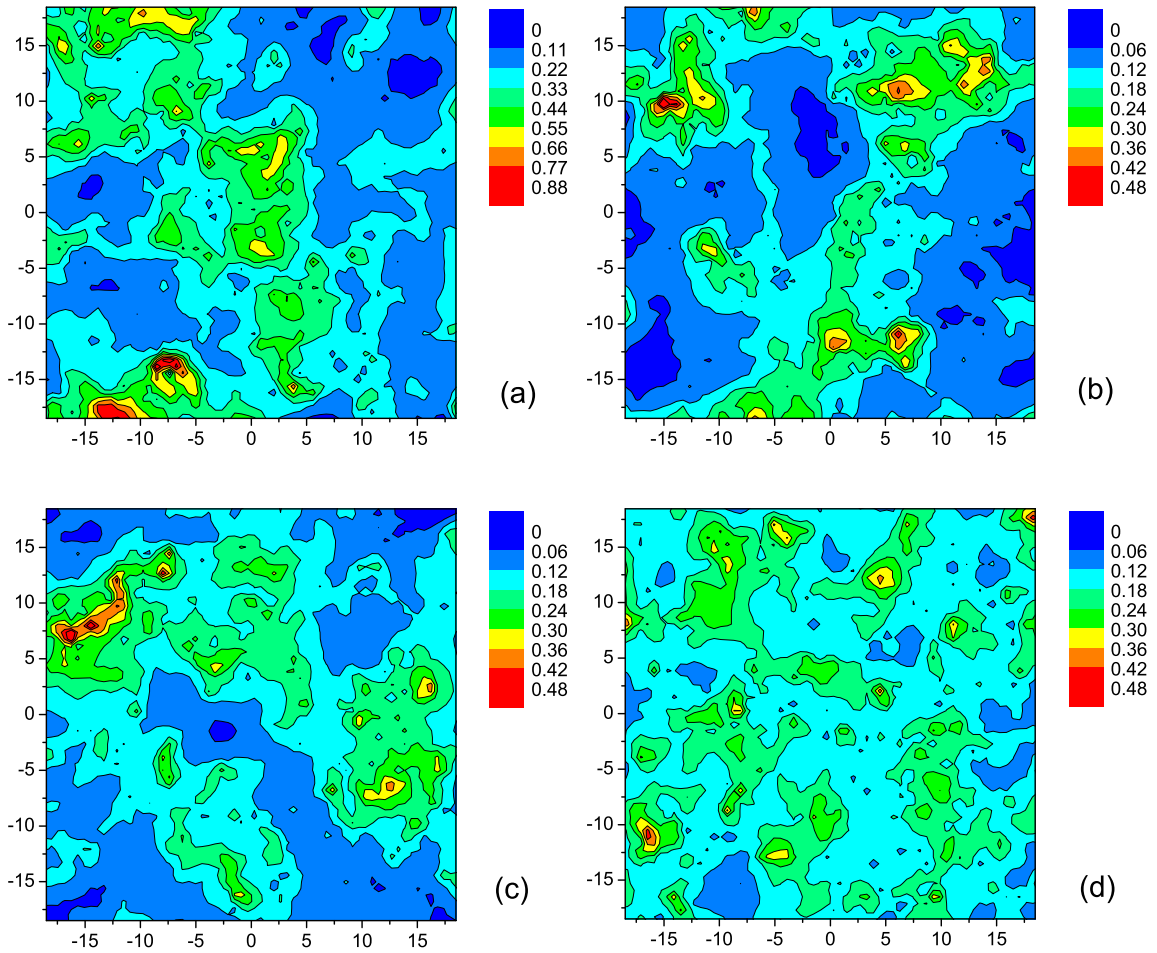

Figure 2.8: The spatial distribution of propensity for selected configurations at (a) $T = 0.46$, (b) $T = 0.5$, (c) $T = 0.6$, and (d) $T = 1$; where the propensities were averaged over 100 runs. Note the different propensity scales.

propensity particles using a cluster analysis that is described in general terms in Section 3.2.2. Basically, for each configuration we select the 10% of particles with the highest propensities and assign them to clusters depending on whether they are a nearest neighbour to another particle already in a cluster. When all particles have been assigned to clusters, we count the total number of clusters and the variance in cluster size, and use these two quantities to characterise the degree of spatial clustering. Figure 2.9 shows the results of the cluster analysis for ten configurations each at $T = 0.4$, 0.5, 0.6, 0.8 and 1. While there is considerable variation in the clustering of high propensity particles between different isothermal configurations and overlap between data points from different temperatures, it is clear that, on average, particles with high propensity cluster together more as the temperature is reduced.

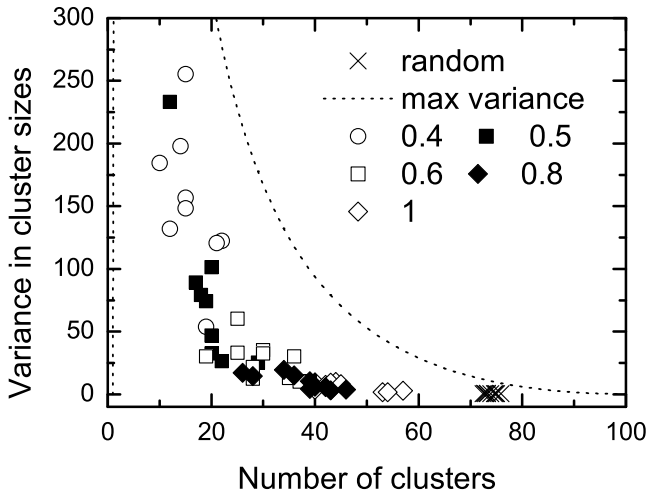

Figure 2.9: Cluster measures of spatial heterogeneity for particles with propensities in the top 10%. Data points are shown individually for ten configurations each at $T = 0.4$, 0.5, 0.6, 0.8 and 1. Statistics obtained using random values are shown for comparison. The dotted line represents the maximum variance possible for a given number of clusters (see text for more details).

In summary, we have established three ways in which the spatial heterogeneity of propensity increases upon cooling: (i) through a gradual increase in the clustering of particles with similar propensity, (ii) through an increase in the population of the extremes of the propensity distribution, and (iii) through a rapid increase in the range of the propensity distribution below $T = 0.5$. We also observed increasing variation between different isothermal configurations as the temperature was reduced. We therefore conclude that as the binary soft-disc liquid is cooled, variations in structure become increasingly important for the dynamic properties of the glass-former.

2.3.2 Statistical Convergence and Reliability

The need for the propensity as a measure of structure-related dynamics is directly due to the large difference in particle displacements from run to run, as discussed above. Given this high variability in particle motion, it is sensible to consider how the size of the isoconfigurational ensemble affects the propensity distribution and the comparisons that can meaningfully be made. In this section we investigate the uncertainty in the propensity distribution as a function of the number of runs, and discuss the practical consequences of our results.

The uncertainty in the propensity (mean squared displacement) of particle i is

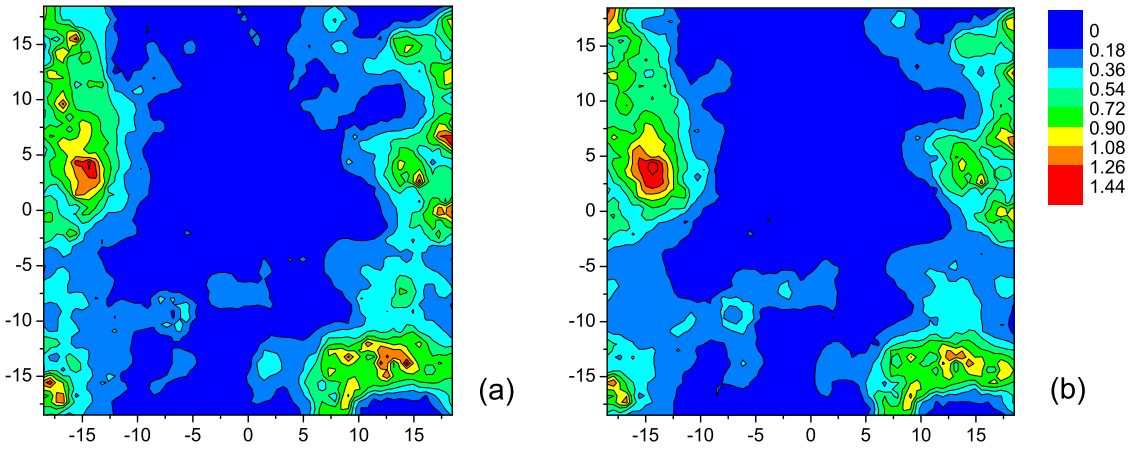

Figure 2.10: Convergence of the spatial distribution of propensity as a function of the number of runs for a configuration at $T = 0.4$. The propensities were calculated using (a) 50 runs and (b) 1000 runs. Note that there is little difference in the coarse grained spatial variation between the two plots.

measured by the standard error $\sigma_i/\sqrt{N_{runs}}$, where N_{runs} is the total number of runs and $\sigma_i = \sqrt{< {\Delta r_i}^4 >_{ic} - (< {\Delta r_i}^2 >_{ic})^2}$ is the standard deviation in the squared displacement distribution for particle i. Hence, the standard error decreases as a function of $1/\sqrt{N_{runs}}$, and the number of runs needed to further reduce the uncertainty increases rapidly. While this means that it takes many runs to reduce the uncertainty in a single propensity - for example after 200 runs some particles at $T = 0.4$ still have relative uncertainties of 60% - we find that both the shape of the total propensity distribution and its coarse-grained spatial variation converge far more rapidly.

In Figure 2.10 we compare the spatial distribution of propensity averaged over ensembles of 50 and 1000 runs for the same configuration at $T = 0.4$. Although there are minor differences between the two plots, it is clear that the coarse grained spatial variation has already converged by 50 runs, i.e. it is possible to distinguish all the regions of high, low and intermediate propensity after only 50 runs. We also find that it takes only 100 runs, at both $T = 0.4$ and $T = 1$, for the standard deviation of the total propensity distribution to converge to within 2% of the extrapolated limit at infinite number of runs. We therefore conclude that ensembles of 100 runs are large enough to determine the spatial distribution of propensity with a high degree of confidence.

To investigate the convergence of the individual propensities in more detail, we define the relative uncertainty in the propensity at the 95% confidence level and

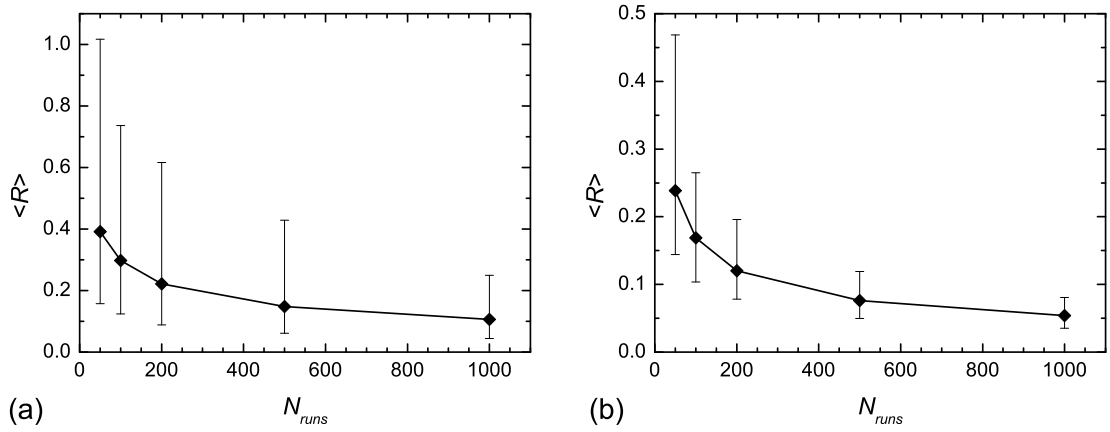

Figure 2.11: Convergence of the relative uncertainty in the propensity R (see Eq. 2.2) as a function of the total number of runs N_{runs} for configurations at (a) $T = 0.4$ and (b) $T = 1$. The error bars indicate the range of R values at a given number of runs, and the curve joins the mean values of R, where the average is taken over particles.

study its convergence as the number of runs increases. The $P\%$ confidence interval is defined as the interval in which there is a $P\%$ chance of finding the true population mean. To calculate the confidence interval for the propensity one should strictly use the two-sided Student's t-distribution [90] since the population mean and variance are unknown. However, in practice we find that the sample size, i.e. the number of runs, is sufficiently large that we can use the normal distribution instead (note that Student's t-distribution converges to the normal distribution as the sample size increases). The relative uncertainty R_i in the propensity of a particle i at the 95% confidence level as a function of the number of runs is therefore given by

$$R_i(N_{runs}) = 1.649 \frac{\sigma_i}{\sqrt{N_{runs}}} / < {\Delta r_i}^2 >_{ic} \tag{2.2}$$

where $< {\Delta r_i}^2 >_{ic}$ is the propensity of particle i and N_{runs} is the number of runs used to calculate the propensity. The normal and t-distributions can also be used to test whether the difference in propensity between two particles is significant or not.

In Figure 2.11, we plot the mean uncertainty $< R >$ (averaged over particles) as a function of the total number of runs for configurations at $T = 0.4$ and $T = 1$. The error bars indicate the range of R_i values at a given N_{runs}. At $T = 0.4$ [plot (a)], we find that while $< R >$ has decreased to about 0.2 (i.e. 20%) after 200 runs, the maximum value decreases much slower, e.g. there are still some particles

with $R_i = 0.6$. By 1000 runs $< R >$ has decreased to 0.1, but the largest relative uncertainties are still around 25%. In comparison, the uncertainty decreases much faster at $T = 1$ [plot (b)]. After 200 runs $< R >= 0.12$ and the maximum uncertainty is about 20%, and by 1000 runs $< R >= 0.06$ and the maximum is around 8%.

We therefore conclude that it is difficult to compare propensities that are similar in magnitude. How similar depends on how many runs one is willing to wait for, e.g. at $T = 0.4$ it will be rather impractical to compare propensities that are within 25% of each other. The reason why the *spatial* distribution of propensity converges far more rapidly is that the difference between high and low propensities, i.e. the range of the propensity distribution, is generally much larger than the mean. In fact, as shown in Figure 2.5(a), the range increases rapidly relative to the mean below $T = 0.5$. If this rate is faster than the rate at which the uncertainty in the propensity increases, which it appears to be, then the spatial distribution of propensity should converge even more rapidly at lower temperatures. It is an attractive idea that this may make propensity calculations practical at deeper supercoolings than we have studied here.

2.4 The Increasing Variance of the Individual Particle Motion

The variation in an individual particle's mobility between runs not only affects the convergence properties of the propensity, but represents another area where analysis of the isoconfigurational ensemble can provide new insight, this time into the *manner* in which the configuration influences the dynamics. To illustrate what we mean by this it is useful to consider that the same spatial distribution of propensity could be produced in many different ways. For example, a particle could move the same amount in every run, or its mobility could vary strongly from run to run, without necessarily changing its mean-squared displacement, i.e. its propensity. The shape of the displacement distribution therefore contains additional information about how the structure affects dynamics.

The increasing variation in particle mobility between runs, which we investigate here, indicates that there is considerable randomness or noise in the manner in which the configuration influences the dynamics at low temperature. We argue that this could be interpreted in terms of stick-release events. While there is a higher probability of a release event occurring in a high propensity region, both high and low

propensity regions are capable of sticking the particles, i.e. not allowing them to move in a given run. In other words, the configuration expresses its character intermittently.

The results in this section provide: (i) the reason why the propensity is needed to characterise the effect of structure on dynamics; (ii) physical information on the process by which the configuration influences the dynamics; and (iii) a view of dynamics that is consistent with recent experiments describing the intermittency of relaxation events in colloidal clays.

2.4.1 Variance versus Propensity

We quantify the variation in the ith particle's mobility between different runs using the standard deviation, σ_i, of the propensity, where $\sigma_i = \sqrt{< \Delta r_i{}^4 >_{ic} - (< \Delta r_i{}^2 >_{ic})^2}$. As shown in Figure 2.12, σ_i at $T = 0.4$ is significantly larger relative to the propensity $< \Delta r_i{}^2 >_{ic}$ than one would have expected from a continuum random walk in 2D. The inset shows the same results at $T = 1$.

To understand the significance of this finding we provide some background information. Freely diffusing particles in a liquid can be modelled as a continuum random

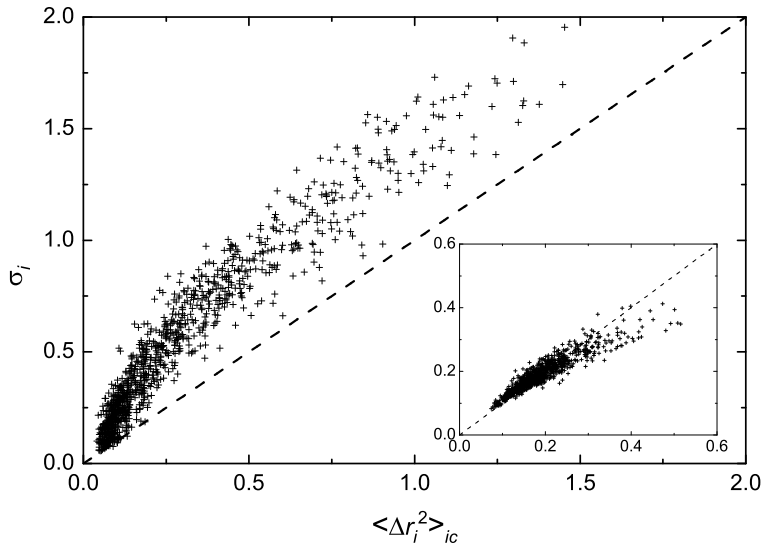

Figure 2.12: A scatter plot of the standard deviation σ_i (calculated over 1000 runs at $T = 0.4$) of the squared displacement of each particle plotted against its propensity. The inset shows the same data obtained for a configuration at $T = 1$. Note the difference in scale. The dashed line is the expected relation for a 2D random walk where each point along the line can be interpreted as arising from a different value of the diffusion constant.

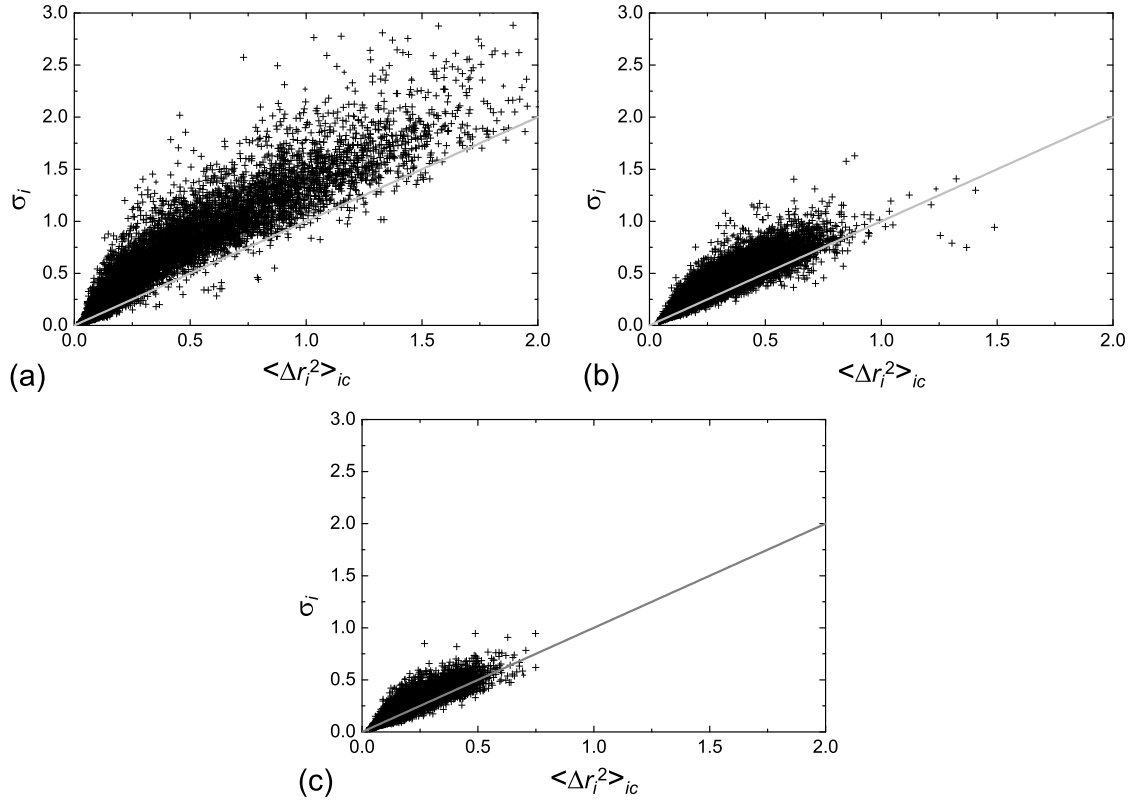

Figure 2.13: A scatter plot of the standard deviation σ_i of the squared displacement of each particle plotted against its propensity for ten configurations each at (a) $T = 0.4$, (b) $T = 0.46$, and $T = 0.5$. All quantities have been calculated using ensembles of 100 runs, and the solid grey line is the expected relation for a 2D random walk where each point along the line can be interpreted as arising from a different value of the diffusion constant.

walk. After a large number of independent steps in the random walk, i.e. after a time that is long compared to the mean collision time in the liquid, the particle's position will be given by a probability distribution that is equal to a normal distribution [91]. The isoconfigurational distribution of particle displacements for each particle in the 2D binary mixture should therefore be given by a normal distribution if the particle motion can be described by simple diffusion. In 2D the standard deviation and mean of a normal distribution are equal. Therefore each point along the line of slope equal to one in Figure 2.12 can be interpreted as arising from a different value of the diffusion constant. Hence, our results indicate that at low temperature the heterogeneity in the propensity cannot be described simply by a scenario in which particles are undergoing simple diffusion, but with different diffusion constants.

We investigate the relationship between the variability in particle motion, as measured by σ_i, and the propensity as a function of temperature and configuration. Figure 2.13 shows the relationship between σ_i and the particle propensity for ten configurations each at $T = 0.4$, $T = 0.46$ and $T = 0.5$. All quantities were calculated using ensembles of 100 runs. The data points at $T = 0.4$ show greater scatter compared to the results in Figure 2.12, partly due to the greater convergence of the variance and propensity in Figure 2.12 where these quantities were calculated over 1000 runs, and partly due to some variation from configuration to configuration. There is, however, a clear increase in variability below $T = 0.5$. At $T = 0.5$ the majority of data points are still clustered around the expected relation for a 2D random walk. The change in the variability of particle motion is therefore a strong characteristic of the change in dynamics as the liquid is supercooled.

We conclude that the increasingly large variation, upon supercooling, in an individual particle's movement from run to run represents an important piece of kinetic information, distinct from the propensities and their spatial distribution.

2.4.2 The Single Particle Non-Gaussian Parameter

The large variances of the individual particles are typically associated with highly asymmetric $f_i(\Delta r)$ distributions, with a peak at a low value of Δr and a long tail extending to large displacements, as shown in Figure 2.14 for a representative particle at $T = 0.4$.

This asymmetry can be quantified as a deviation from a Gaussian form through the use of a non-Gaussian parameter α_i for particle i given by

$$\alpha_i = \frac{< {\Delta r_i}^4 >_{ic}}{2(< {\Delta r_i}^2 >_{ic})^2} - 1 \tag{2.3}$$

The quantity α_i equals zero for a Gaussian distribution. The α_i distributions for configurations at $T = 1.0$ and 0.4 are plotted in Figure 2.15. While all the individual $f_i(\Delta r)$ distributions are close to Gaussian at high temperature, the supercooled sample exhibits a broad distribution of α_i values with most particles exhibiting a significantly non-Gaussian distribution of displacements.

Note that this non-Gaussian parameter is quite distinct from that discussed previously in the context of supercooled liquids [74, 92]. The α_i introduced here refers to the variety of displacements achieved by a *single* particle over the ensemble of

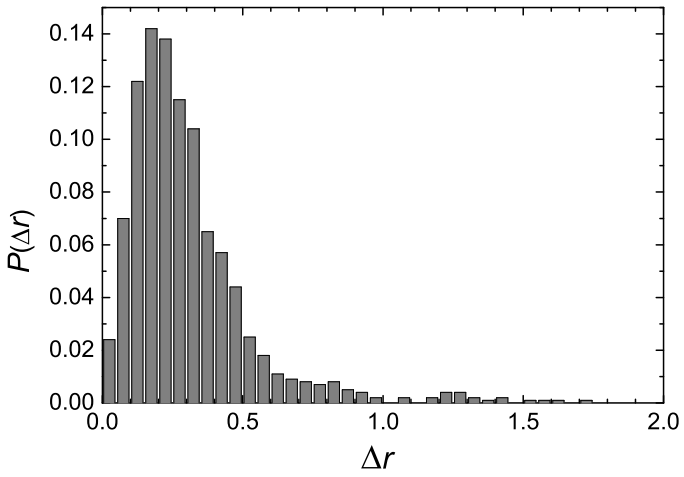

Figure 2.14: The distribution of particle displacements, $f_i(\Delta r)$, over an isoconfigurational ensemble of 100 runs for a single particle at $T = 0.4$. Note the highly asymmetric and non-Gaussian shape of the distribution.

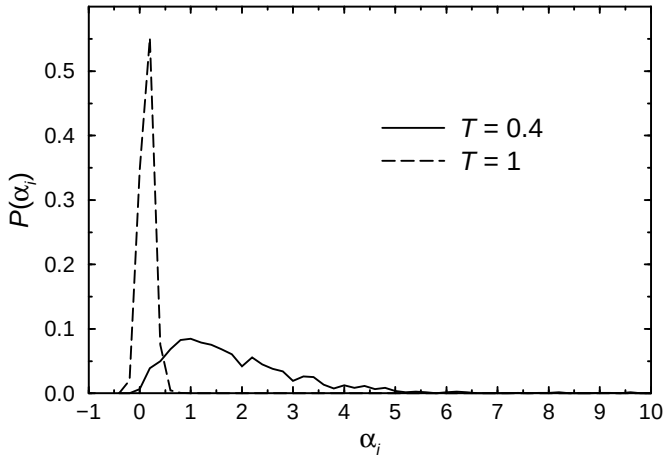

Figure 2.15: The distribution of single particle non-Gaussian parameters α_i (see Eq. 2.3) for configurations at $T = 0.4$ and 1.0, calculated using ensembles of 1000 runs.

trajectories, as opposed to the variety of displacements achieved by *different particles* in a single trajectory. In the language of the jump model of particle motion [93], the propensity characterises the average waiting time and jump length, while the non-Gaussian character of the $f_i(\Delta r)$ distributions is a result of either displacement correlations between successive jumps and/or non-Poisson statistics for the number of jumps within the observation time. We explore this question further in the next section.

2.4.3 The Jump Model of Particle Motion

We investigate the origin of the non-Gaussian character of the $f_i(\Delta r)$ distributions in terms of the jump model of particle motion [93]. In particular, we define a jump, compare the propensity to the average waiting time, and investigate whether the number of jumps per run can be described by a Poisson distribution. This analysis is incomplete as we do not study temperatures below $T = 0.5$, where the variability in particle dynamics changes most dramatically. Instead, it serves as a useful introduction to some conceptual views of dynamics in supercooled liquids and the methods by which they can be investigated.

The dynamics in supercooled alloys is often discussed in terms of 'caging' and the 'escape' from the cage. In this picture the plateau that develops in the intermediate scattering function upon supercooling (see Section 1.2) is ascribed to particles being trapped in the cage formed by their neighbours. The scattering function only starts to decay again once particles begin to escape from this cage. It is this picture of diffusive motion via large discrete jumps, between which the atoms oscillate as in a solid, that has inspired the jump model and large number of similar models and simulations (see, for example, refs. [94,95]). The common assumption is that jumping is the main process that explains the dynamics. Starting, generally, from a random walk, the anomalous diffusion is then incorporated via a distribution of waiting times or jump lengths.

Now imagine the simple scenario in which the waiting time, for a given particle, is the same for every jump. Then the resulting distribution of jumps per run should be given by a Poisson distribution (see Eq. 2.4 on page 41). Conversely, if the distribution of jumps per run is not Poisson distributed, then the waiting time cannot be the same for every jump. In this conceptual framework, the non-Gaussian character of the $f_i(\Delta r)$ distributions must be due to either displacement correlations between successive jumps or a tendency for mobile particles to continue being mobile, i.e. waiting times that are smaller for successive jumps. For this reason, we investigate whether the number of jumps per run can be described by a Poisson distribution.

To study the particle dynamics in this context, we first need to define a jump length. Ideally this jump length will maximise the amount of non-vibrational motion that is captured while minimising the amount of vibrational motion that is measured. One possibility is to define each jump individually using a relative criterion, where for each particle a jump is defined as motion that is of large amplitude relative

to short-time (vibrational) fluctuations in position. This method has been used to study motion in a Lennard-Jones (LJ) glass-former below T_g [95]. However, it appears unlikely that this method will be suitable for the present system. We are interested in motion that occurs above the glass transition, and it is not obvious that there will be a clear distinction between vibrational and non-vibrational motion. Even for the LJ glass studied in ref. [95], near T_g the amplitude of short-time fluctuations (0.14–0.16u) is not much smaller than the smallest (reversible) jumps (0.2u). A different jump criterion that has been used involves defining a minimum hopping distance [96, 97]. This criterion should provide us with sufficient dynamic information for the present question. Provided that we choose a distance that means that the local environment of the particle must have changed, this should be large enough that successive jumps or moves are statistically independent in time and direction for a normal liquid.

In this work we define a 'move' as occurring when a particle has moved a distance of R_{move} from its previous position. We choose $R_{move} = 1$ for all particles. Perera and Harrowell found that this distance maximised the degree of spatial segregation of the dynamics in the present glass-former when it was used to define a local relaxation time [44]. Values of 0.5 and 0.7 were also investigated, but appear to capture much vibrational motion. Single configurations at $T = 0.5$ and $T = 1$ were investigated using 1000 runs of $1.5\tau_e$ and 100 runs of 100τ at each temperature. We note that additional analysis below $T = 0.5$ would be interesting since the variability of particle motion increases rapidly below this temperature.

We consider the number of moves per run k for each particle, and define p as the probability of 'moving' into a new position in a given time interval. By increasing the number of time intervals N per run we can make p arbitrarily small. If p is the same for every move, then the distribution of moves per run $P(k)$ should be given by a Poisson distribution, i.e. we expect

$$P(k) = \frac{a^k}{k!}e^{-a} \tag{2.4}$$

where k is the number of moves per run, and $a = Np$ is the mean value of k

We define $f_i(k)$ as the distribution of the number of moves per run for particle i. To quantify how far this distribution deviates from Poisson we define a non-Poisson parameter as $nonP = u_2/u_1 - 1$, where $u_n = \sum_i (k_i - < k >)^n$ is the n-th central moment of the the $f_i(k)$ distribution. For a Poisson distribution $nonP$ will equal

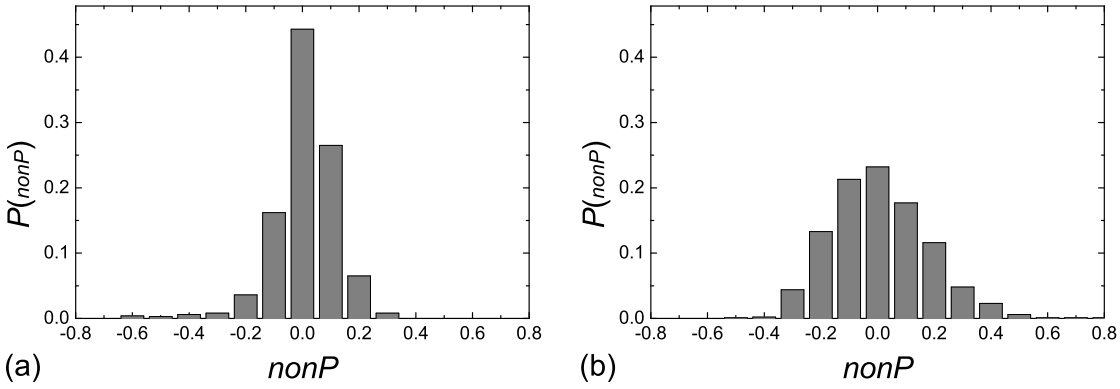

Figure 2.16: The distribution of non-Poisson parameters $nonP$ (see text) for particles in a configuration at $T = 0.5$. The distributions were calculated using (a) 1000 runs of duration $1.5\tau_e = 20.2\tau$, and (b) 100 runs of duration 100τ each.

zero, since $u_1 = u_2 = a$. To calculate an unbiased estimator for $nonP$ we use k-statistics [98], i.e. we calculate

$$nonP = \frac{k_2}{k_1} = \frac{\frac{n_{runs}}{n_{runs}-1} m_2}{\mu} \tag{2.5}$$

where k_n is the n-th k-statistic, n_{runs} is the total number of runs (i.e. the sample size), $m_2 = \frac{1}{n_{runs}} \sum_{i=1}^{n_{runs}} (k_i - \mu)^2$ is the sample variance, and $\mu = u_1$ is the sample mean.

The results are plotted in Figures 2.16 and 2.17. At $T = 0.5$, most individual particles have $P(k)$ distributions that are approximately Poisson distributed ($-0.4 \leq nonP \leq 0.4$ for most particles) for both choices of run length. At $T = 1$ we obtain similar results, although the longer 100τ runs have $nonP$ values that range from -0.6 to -0.2 rather than being centered about 0. The fact that we get Poisson statistics for the number of moves per run tells us that P, the probability of 'moving' into a new position, is roughly the same for every move. This suggests that there is little tendency for moving particles to continue moving above $T = 0.5$. The higher $nonP$ values obtained for the longer $100\tau = 86.6\tau_e$ runs at $T = 1$ are probably due to the fact that particles are able to sample both fast and slow regions over this very long timescale and thus have P values that change significantly over the course of the run.

We conclude that the waiting time per particle is roughly the same for every move above $T = 0.5$, at least over times that are of the same order of magnitude as the structural relaxation time. Similar analysis at $T = 0.4$ would help to address whether

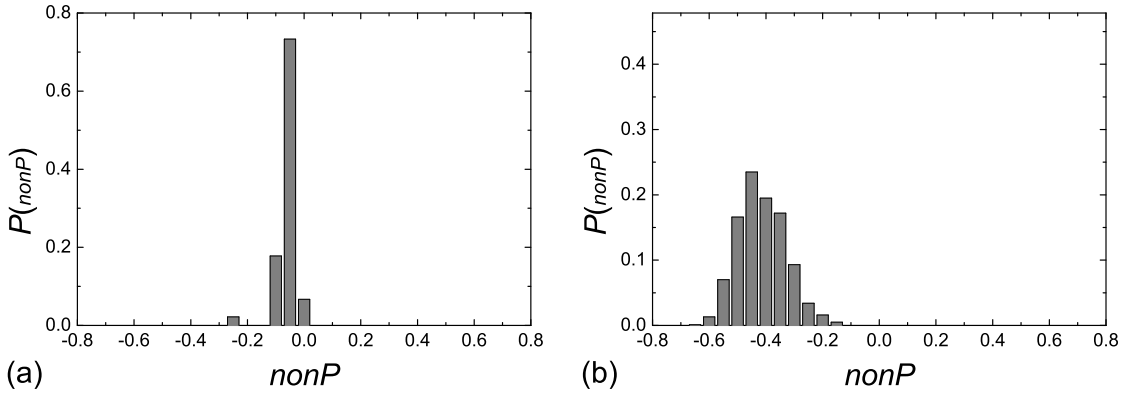

Figure 2.17: The distribution of non-Poisson parameters $nonP$ for particles in a configuration at $T = 1$. The distributions were calculated using (a) 1000 runs of duration $1.5\tau_e = 1.155\tau$, and (b) 100 runs of duration 100τ each.

it is correlated jump directions that produce the non-Gaussianity of the single particle displacement distributions observed at this temperature. Direct measurement of such directional correlations are recommended to verify this in the case that non-Poisson statistics are found for the number of moves per run.

Finally, we consider the distribution of move times, where the move time is defined as the time taken for a particle to move a distance of $R_{move} = 1$. Previously, we hypothesized that the propensity characterises the average waiting time and jump length. Here we test this hypothesis by considering the correlation between propensity and average move time with the jump length fixed at $R_{move} = 1$. We note that the propensities will only contain information about the average move time of those moves that are able to occur in the time period used to calculate the propensities. Hence it is only valid to compare the average move time obtained from runs of duration $1.5\tau_e$ with the propensities, since the same run times were used in calculating both quantities.

In Figure 2.18 we plot the propensity against the average move time $< t_{move} >$ for a configuration at $T = 0.5$. Only data for the 796 particles that moved in at least 50 of the total 1000 runs are plotted. This was done to ensure that there were adequate statistics to calculate the average move time for each particle. We find a moderate negative correlation between the propensity and the average move time. This provides some support for our previous hypothesis that the propensity characterises the average waiting time and jump length, but the broad scatter suggests that this is too simple a description to completely describe the data. It would be

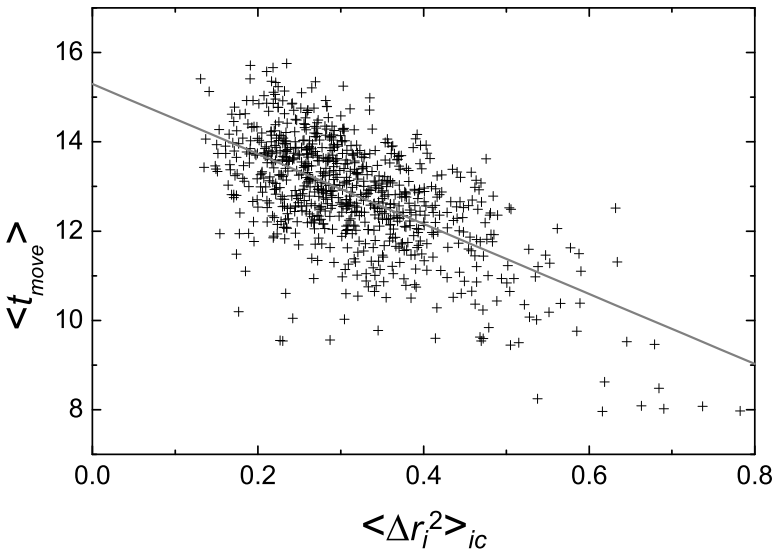

Figure 2.18: The relationship between propensity and average move time $< t_{move} >$ for 796 particles at $T = 0.5$. Properties were calculated using an ensemble of 1000 runs, and data points have only been plotted for particles that moved in at least 50 runs. The solid grey line is a linear fit to the data.

interesting to see if the correlation improves at lower temperature. At $T = 1$ we did not observe any correlation but this may have been due to a lack of data points (only 45 particles moved in at least 50 runs).

2.4.4 The Spatial Distribution of the Single-Particle Non-Gaussian Parameter

Given that the non-Gaussian $f_i(\Delta r)$ distributions represent a new piece of kinetic information, the spatial distribution of single-particle non-Gaussian parameters (α_i, defined in Eq. 2.3) may offer additional insight into the manner in which the configuration influences relaxation. Here we compare this distribution with the spatial distribution of dynamic propensity and discuss what insight may be gained from such analysis. Those particles with motion that varies the most from run to run can be thought of as having the least structural constraint on their mobility.

Figure 2.19 shows the spatial distribution of α_i and propensity for a configuration at $T = 0.4$. Regions with high propensity and particles with L06 topology (filled circles) tend to have low α_i, and particles with high α_i tend to have low propensity. Lines of high α_i appear to represent paths for rare motion in regions of low propensity,

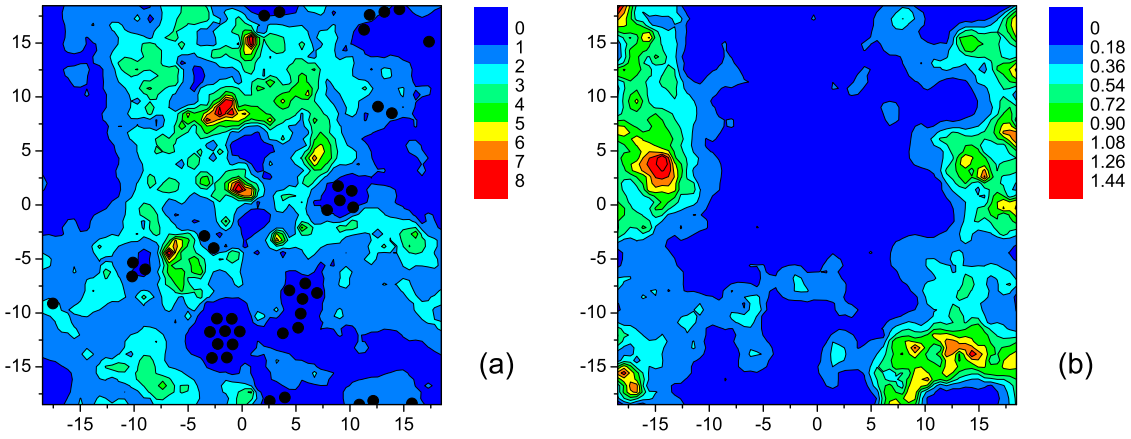

Figure 2.19: (a) The spatial distribution of the single particle non-Gaussian parameter α_i for a configuration at $T = 0.4$. The black circles indicate the positions of large particles with six large neighbours. (b) The propensity map for the same configuration used in (a). Quantities were calculated using ensembles of 100 runs.

and therefore a study of these may provide insight into mechanisms for relaxation of the slow regions.

2.4.5 Intermittency

One consequence of the high variability in particle motion from run to run, and the accompanying asymmetry of the displacement distributions, is that rare events (large displacements) have a significant impact on the mean squared displacement, i.e. the propensity. Other physical phenomena in which rare events have a significant influence on some mean property have previously been described as 'intermittent'. In particular, the term *intermittency* has been used to to describe distributions in which maxima in space or time are widely spaced and rare, but make a dominating contribution to the physical quantity of interest. See, for example, 'The Almighty Chance' [99] for an introduction to the concept of intermittency. The term was introduced by Batchelor and Townsend [100] for the patched temperature distribution in turbulence, and has been applied to the distribution of matter in space among other physical phenomena. The principal and characteristic property of intermittency is the abnormal relationship (when compared to the Gaussian one) between the consequent statistical moments.

In the present case, the rare events are not scattered in space or time, but over

the isoconfigurational ensemble, i.e. over the space of possible futures. However, the abnormal relation between the statistical moments is similarly non-Gaussian. Another way of interpreting this is to say that the configuration expresses its character, or effect on the dynamics, intermittently.

The concept of intermittency has also recently been used to describe some experimental results on relaxation in glasses. Ciliberto et al. [101] have reported intermittent voltage noise signals characterised by rare large noise spikes above the regular fluctuations during dielectric studies of a colloidal glass of clay particles called Laponite. The result is a non-Gaussian distribution for the voltage noise which, based on numerical work, has been interpreted in terms of activated and spontaneous relaxation events [102]. However, discrepancies have been found between different experiments, and between simulations and theory, and this area appears to need further work to rationalise the contrasting findings. A review of recent experimental, numerical and theoretical work on the intermittency of relaxation in glassy soft matter can be found in ref. [15].

In terms of the present work, we suggest that another way of interpreting the intermittent manner in which the configuration affects the dynamics, and possibly also the intermittent voltage noise, is in terms of 'stick' and 'slip' events. Even high propensity regions are able to 'stick' particles, i.e. to not allow significant motion to occur, but occasionally 'slips' occur, i.e. large displacements take place during a run. The difference between high and low propensity regions in this picture is that the frequency (or probability) of slips is higher in the high propensity regions. Of course, this is a rather crude picture in the sense that the distribution of displacements for a particle tends to be continuous rather than forming two discrete peaks, but the basic picture of an increasing variability of particle motion as the system is cooled is consistent with our results.

2.5 Correlations in Particle Motion

The isoconfigurational ensemble technique can also be used to explore the character of the dynamics in the supercooled liquid. This can be done by exploring correlations in particle motion, for example between different particles within the same run or between the same particle (self-correlation) in different runs. Information about the lengthscale over which particle motion is correlated can also be obtained.

2.5.1 Self-Correlation within the Isoconfigurational Ensemble

Analysis of correlations between the motion of a particle and itself in different runs can give information on the degree to which the original configuration confers directionality on particle motion. By directionality we mean a tendency to move in a preferred direction. In particular, we investigate the following question: is the variation in propensity due to the tendency of some particles to have a strong directional preference?

We define the directionality d_i of a particle i as the mean dot product over all pairs of displacement vectors normalised by the propensity, i.e.

$$d_i = \frac{\frac{1}{N_{\alpha\beta}} \sum_\alpha \sum_{\beta>\alpha} (\Delta\vec{r}_{i,\alpha} \cdot \Delta\vec{r}_{i,\beta})}{< \Delta {r_i}^2 >} \tag{2.6}$$

where α and β are run indices, $N_{\alpha\beta} = {}^{N_{runs}}C_2$ is the number of distinct pairs of runs in the isoconfigurational ensemble, $\vec{r}_{i,\beta}$ is the displacement vector of particle i in run β, and $< \Delta {r_i}^2 >$ is the propensity of particle i. For a random distribution of displacements, the vector pairs will be evenly distributed in magnitude and in intervector angle. Therefore, we should have $d_i = \frac{1}{\pi} \int_0^\infty \cos\theta \, \mathrm{d}\theta = 0$. And if the particle moves in the same direction in every run, we should have $d_i = 1$.

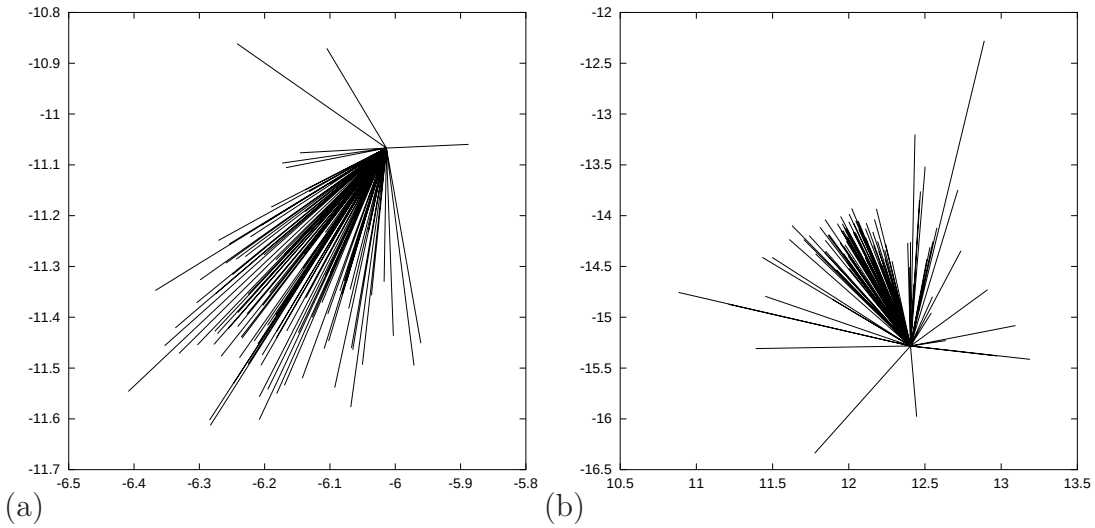

Figure 2.20: Displacement vectors for selected particles at $T = 0.4$ with high directionality and either (a) low or (b) high propensity. The vectors are from isoconfigurational ensembles of 100 runs.

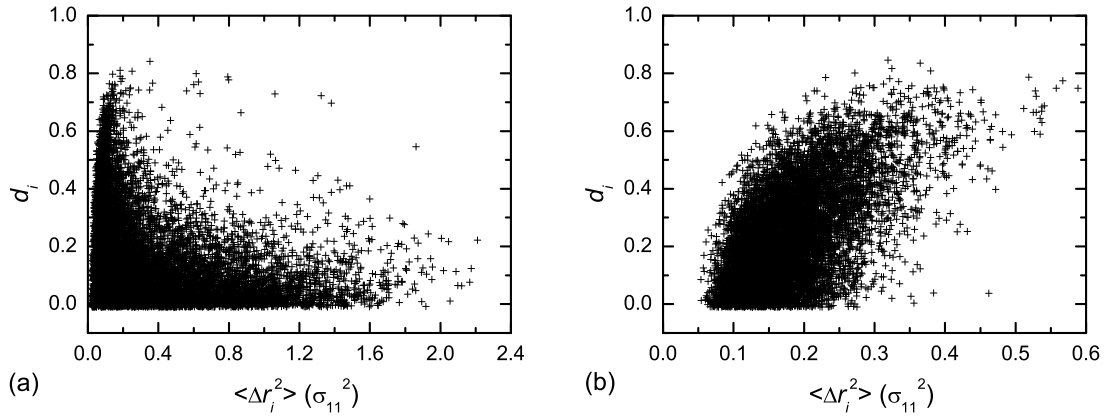

Figure 2.21: The particle directionality d_i as a function of propensity for ten configurations each at (a) $T = 0.4$ and (b) $T = 1$. Quantities were calculated using 100 runs. Note the different x-axis scales.

Figure 2.20 shows the displacement vectors for selected particles at $T = 0.4$. In particular, we considered an ensemble of 100 runs and selected particles with high directionality and either high or low propensity. We find that there are particles with $d_i > 0.8$, independent of propensity, whose displacement vectors fall almost exclusively within a 60° angle, i.e. that have strong directionality conferred upon their motion by the initial configuration.

An obvious question to ask is what role do particles with high directionality have in determining the propensity distribution. In Figure 2.21 we plot d_i against propensity, using data pooled from ten configurations each at $T = 0.4$ and $T = 1$. The configurations were separated from one another by $75\tau_e$, and the propensities and directionalities were calculated over ensembles of 100 runs. At $T = 1$, particles with high directionality have high propensity. This suggests that at high temperature the most mobile particles are those that have a high degree of directionality conferred upon their motion by the initial configuration. In contrast, at $T = 0.4$ the particles with high propensity generally have low directionality, suggesting that any directionality conferred by the initial configuration is rapidly 'forgotten' as a particle moves away from its initial position. Of course, this difference may be a consequence of how we have scaled the run time for the isoconfigurational ensemble. The mean collision time for this glass-former is 0.10τ for all $T \leq 1$ [89], therefore the run time at $T = 1$ is only one order of magnitude longer than this, compared to four orders of magnitude longer at $T = 0.4$.

The large number of particles with high directionality and low propensity at $T = 0.4$ could be explained by the following picture. At low temperature many particles remain trapped in the 'cage' formed by their neighbours. Since we are studying instantaneous configurations it is likely that some of these particles will be far from the centre of their local potential energy minimum (due to the cage) when the runs are begun, i.e. in the initial configuration. There is therefore a high probability that they will undergo mainly vibrational motion during the run and will consequently be found closer to the local potential energy minimum at the end of the run, thus resulting in a high directional bias in their motion.

We conclude that directional bias on particle motion due to the initial configuration is insufficient to explain the heterogeneity in the spatial distribution of propensities at low temperature. We do, however, note that some particles with high propensity also have high directionality. Given the increased clustering of high propensity particles at low temperature (see Figure 2.9), it is possible that these rare particles, with both high propensity and high directionality, have a role to play as initiators of motion in the high propensity regions. This could be an area for future research.

2.5.2 Correlation Between Motion of Unlike Particles

An analysis of correlations between the motion of unlike particles can yield information about the cooperativity of particle dynamics. For example, one can ask whether neighbouring particles move in the same direction, or investigate the lengthscale over which particle motion is correlated. This is an area that, as discussed in the introduction to Part I, has already been studied in much detail. The unique feature of the present analysis is that we average such measures over an isoconfigurational ensemble. This makes it possible, for example, to calculate the distance over which particle motion is correlated for individual particles in a given configuration.

Before discussing the distance over which particle motion is correlated, we present an analysis that can be used to identify particles whose motion is highly correlated with that of their neighbours. We define the flow f_i of a particle i as the mean dot product between the displacement vector of i and that of all its nearest neighbours, where the vectors are normalised and the isoconfigurational average is taken, i.e.

$$f_i = \left\langle \frac{1}{n_i} \sum_{j=1}^{n_i} \left(\frac{\Delta \vec{r}_{i,\alpha}}{|\Delta \vec{r}_{i,\alpha}|} \cdot \frac{\Delta \vec{r}_{j,\alpha}}{|\Delta \vec{r}_{j,\alpha}|} \right) \right\rangle \tag{2.7}$$

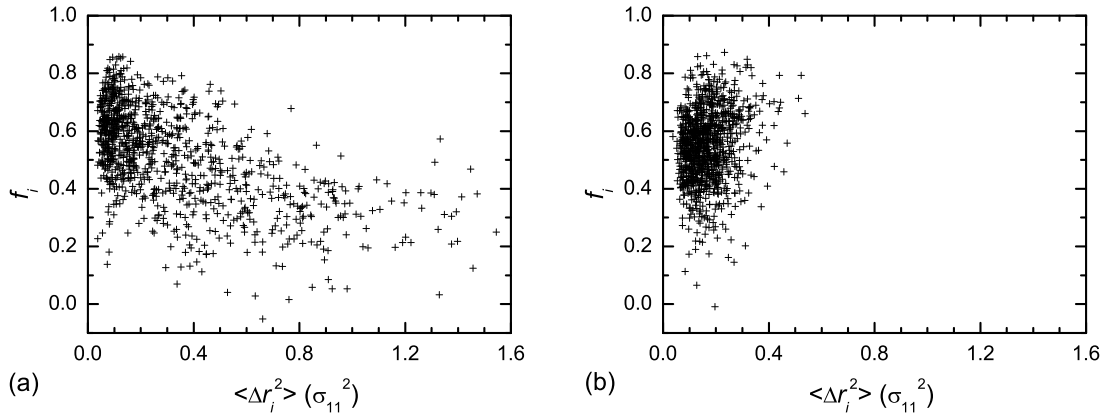

Figure 2.22: The particle flow f_i as a function of propensity for single configurations at (a) $T = 0.4$ and (b) $T = 0.6$. Quantities were calculated using 50 runs.

where α is the run index, n_i is the number of nearest neighbours of particle i, $\vec{r}_{i,\alpha}$ is the displacement vector of particle i in run α, and the angular brackets $<>$ represent an average over runs. If the motion of a particle i is uncorrelated with that of its neighbours then $f_i = 0$. On the other hand, if the motion of the particle is always in the same direction as that of its neighbours then $f_i = 1$.

In Figure 2.22 we plot f_i against propensity, for single configurations at $T = 0.4$ and $T = 0.6$. The propensities and flows were calculated over ensembles of 50 runs. While there is no strong correlation between propensity and flow, a few relations can be identified. Particles with high flow tend to be limited to those with $< \Delta r_i^2 >< 0.5$. The high flow of these low propensity particles may be due to normal modes or small hydrodynamic flows. There is also a strong change in the behaviour of high propensity particles upon cooling. At $T = 0.6$ there is a weak tendency for high propensity particles to have high flow, while at $T = 0.4$ there is a stronger tendency in the opposite direction, i.e. particles with high propensity tend to have lower flows on average than particles with low propensity. This relationship at $T = 0.4$ may be due to the large displacements and long run times involved compared to at $T = 0.6$. If several relaxation events take place in the same region over the course of many runs then the overall displacement vectors for these high propensity particles may no longer reflect the character of the motion during the individual relaxation events. This analysis may therefore be more useful at low temperature if applied to displacement vectors resulting from motion over shorter timescales.

Distance over which Motion is Correlated

Intuitively, it makes sense that if, in a dense supercooled liquid, one particle moves, then the particles nearby must also move to create the space needed for this particle to move into. One can ask over what distance such dynamical correlations persist. In this section we demonstrate how analysis of the isoconfigurational ensemble can be used to address this question.

Take a configuration and select a single particle i. Then, for each particle j, calculate the Pearson's correlation coefficient [90] K between the displacement magnitudes of particle i and particle j over an ensemble of 100 runs. Figure 2.23 shows the correlation coefficients for particles in a configuration at $T = 0.4$ as a function of their distance from particle i, where i is the high propensity particle at (-18,-15.5) in Figure 2.4. For distances greater than $5\sigma_{11}$ the correlation coefficient is scattered between -0.2 and 0.2 with an average value of about zero. The correlation coefficients are only greater than 0.2 for nearest neighbour particles, but the average correlation coefficient for particles at approximately the same distance (probably in the same coordination shell) remains greater than zero for distances up to at least $4\sigma_{11}$ and

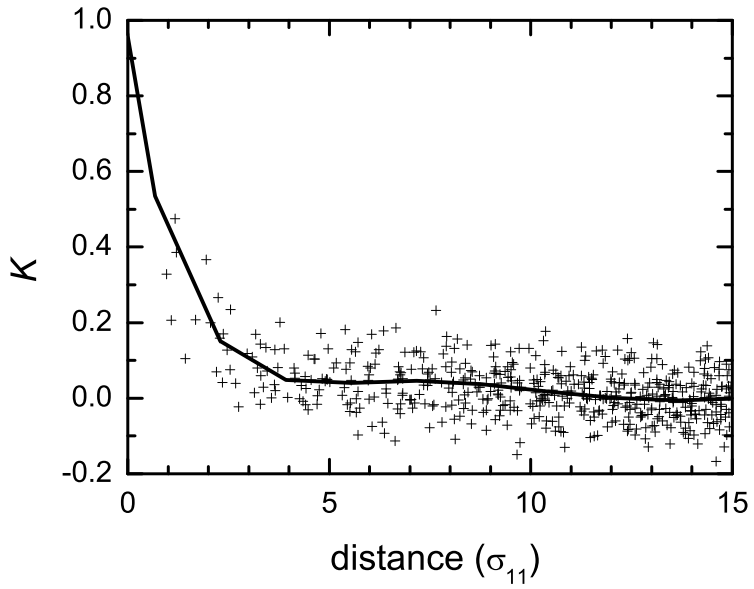

Figure 2.23: The correlation between the motion of a particle i and all other particles j as a function of the distance between i and j. The moving average has been indicated by a thick line, and the Pearson's correlation coefficient K between displacement magnitudes was calculated using data from an ensemble of 100 runs.

possible longer. Our results therefore suggests that the motion of i depends on the motion of at least some particles outside of its first coordination shell. This technique could be extended to calculate, for example, the mean distance over which particle motion is correlated for high and low propensity particles.

2.6 Discussion and Conclusions

Because the glass transition is defined by its dynamics, the task of establishing its structural origin requires us to begin with the dynamics and deduce what structures are responsible. This is an inversion of the usual problem in condensed matter physics and presents us with an important question, namely how trustworthy are the structural clues provided by the observed particle dynamics? In this chapter, we have demonstrated that there is considerable variation in the dynamical evolution of a specific particle configuration. We conclude that some aspects of the particle dynamics are not significantly correlated with the initial configuration and therefore cannot be 'explained' by reference to that configuration.

Through the introduction of the isoconfigurational ensemble, we have established that it is the spatial variation in the propensity for particle motion, rather than the motion itself, that is completely determined by the initial configuration. We find that, upon cooling, the spatial heterogeneity in the propensity increases both in amplitude and wavelength, i.e. the difference between high and low propensity increases and particles with similar propensity increasingly cluster together.

The increasingly large variation, upon supercooling, in an individual particle's movement from run to run also represents an important new piece of kinetic information, distinct from the propensities and their spatial distribution. At high temperature all particles have a variability in their motion from run to run, relative to their propensity, that is typical for a particle undergoing a continuum random walk in 2D. As the temperature decreases particles increasingly have a variability from run to run that is higher than would be expected if they were undergoing a random walk, characterised by highly non-Gaussian single-particle displacement distributions.

In summary, the increasing variability in particle motion that we have characterised provides: the reason why the propensity is needed to characterise the effect of structure on dynamics; physical information on the process by which the configuration influences the dynamics; and a view of dynamics that may explain recent

experiments describing the intermittency of relaxation events in colloidal clays. We have also outlined a number of ways in which further analysis of the variability of particle motion may provide additional insight into the process of relaxation and the role of structure in supercooled glass-formers.

The assignment of propensities to particles represents the major result of this chapter. This result provides a rigorous method for establishing the link between a given configuration and the subsequent dynamics. The remaining problem is to uncover the causal link between specific structural features of a configuration and the resulting propensities. A detailed account of the correlation between propensity and structure in the 2D mixture is presented in Chapter 3. Before moving on, however, we will discuss a number of other implications of the work presented in this chapter.

In the language of the jump model of particle motion [93], the propensity characterises the average waiting time and jump length, while the non-Gaussian character of the individual particle displacement distributions is a result of either displacement correlations between successive jumps and/or non-Poisson statistics for the number of jumps within the observation time. We have outlined how analysis of the isoconfigurational ensemble can directly test this conceptual view of dynamics in supercooled liquids.

A number of recent papers have characterised the transition in particle dynamics on supercooling as a transition from hydrodynamically-governed dynamics to landscape-dominated dynamics [86, 103]. The 'landscape' here refers to the potential energy surface over the configuration space. In this chapter we have arrived at an alternative description of this fundamental temperature dependent change, i.e. a transition, on cooling, from structure-independent (hence liquid-like) propensities for motion to structurally-determined propensities. One advantage of this new account over the landscape picture is that it refers directly to the behaviour of the liquid in real space rather than the abstract configuration space.

The propensity is related to the probability of a particle in a configuration undergoing a substantial displacement within a given time interval. Note that this is distinct from how far it is actually observed to move in a single trajectory. The propensity is therefore the starting point for several models of glass relaxation such as the facilitated spin models [71–73] and the cooperative lattice gas models [104–106]. In each of these, a set of rules determine the probability for movement based on the instantaneous configuration. In contrast, most models of molecular glass-formers are

defined by a Hamiltonian and structural constraints. Uncovering the relationship between particle configurations and the probability of particle motion in these models represents a major challenge. Having now defined and described the dynamic propensity, we have taken a first and necessary step. In the next chapter we address our ultimate goal, which is to predict the spatial pattern of dynamic propensity from a given configuration.

Chapter 3

Predicting Dynamic Propensity

Having established the dynamic propensity as the part of the dynamics that must be due to a property of the structure, we now search for its specific structural origins. In particular, we test the ability of reduced measures of structure to predict the spatial variation in dynamic propensity. We consider the local coordination environment, local free volume, local potential energy, and coarse-grained versions of these, among others. While we find some correlations between structure and dynamics, none of these measures are able to predict the spatial variation in propensity. We therefore turn to the short-time dynamics as a direct measure of structural 'looseness'. We define a single-particle Debye-Waller factor and find that it is able to predict the spatial distribution of propensity. We suggest that this provides an upper bound on the predictive ability of any structural measure. We then use the Debye-Waller factor to study the time evolution of dynamic heterogeneity and obtain results that are at odds with a simple picture of defect diffusion. Finally, we test whether there is a microscopic basis for a relationship reported between short-time dynamics and the geometric free volume.

3.1 Introduction

Some aspect of the structure in a glass-forming alloy determines the observed slow particle dynamics. For example, strong correlations are observed between the increase in shear viscosity and large angle scattering structure in metallic alloys following a temperature quench [107]. In the previous chapter we have shown that there is also something in the structure that is responsible for the *spatial* variation in dynamics.

More precisely, the spatial variation in dynamic propensity is determined by the initial configuration. But what aspect of the initial configuration is responsible for the spatial variation in propensity? In this chapter we test the proposition that reduced measures of structure are able to predict the spatial variation in propensity.

Variations in local environment have been cited as important in a number of studies of the local connection between structure and dynamics. For example, in a study of the same 2D glass-former considered in this part of the thesis, Perera and Harrowell found that large particles with high local six-fold orientational order had, on average, longer relaxation times than large particles with low six-fold orientational order [44]. And in studies of a Lennard-Jones (LJ) liquid, Kob et al. found differences in the radially averaged structure about particles with different mobilities [46, 79]. The latter work also found a correlation between the average potential energy and the average mobility of particles divided into subsets according to their mobility. However, no correlations were found of sufficient strength to indicate a causal link, i.e. that the selected aspect of the local structure determined the local kinetics.

As briefly discussed in the previous chapter, a number of recent papers have characterised the transition in particle dynamics on supercooling as a transition from hydrodynamically-governed dynamics to landscape-dominated dynamics, where the 'landscape' refers to the potential energy surface over the configuration space. Sastry et al. [103, 108] have shown that changes in various dynamic properties of a glass-forming LJ liquid can be related to changes in the energy of the accessible part of the energy landscape. This conceptual picture therefore provides a link between dynamics and structure, albeit in the abstract configuration space. Heuer et al. [86, 109] and La Nave and Sciortino [110] have also reported correlations between the dynamics of small systems of 60–120 particles and the potential energy of the inherent structure (IS). These calculations did not examine whether the correlations extend to the spatial distribution of the two quantities, an issue that we address here.

Free volume [34] is a another widely used concept to explain the relationship between structure and dynamics. The generation or disappearance of free volume has been invoked to explain shear banding [111] and positron annihilation [112] in metallic glasses. The concept of 'shear transition zones' [113] is also used in the context of non-equilibrium mechanical response. Intuitively, it is easy to conceive that less crowded regions will have more space for particle motion and will therefore be 'looser' or more 'mobile'. This intuitive picture has been formalised in free volume theories,

which are able to accurately describe some thermodynamic observations such as the temperature dependence of the viscosity (over more than 12 decades in magnitude for the viscosity in some cases [114]), and the dependence of T_g on quench rate [115]. A recent simulation study also found a strong correlation between the average free volume and the bulk averaged short-time mean squared displacement for monomers in a 'bead-spring' model of a glass-forming polymer over a range of temperatures [80]. Despite its popularity and the success of the free volume concept at a phenomenological level, there remains a persistent problem concerning the application of free volume to describe dynamics. In particular, what is the relationship between the geometric free volume - a quantity that can be well defined at the atomic scale - and the phenomenological free volume (a macroscopic quantity that is usually derived from the bulk density).

We begin by defining a measure of the local coordination environment and investigate the relationship between this and the propensity (Section 3.2.1). This is followed by a comparison of the spatial distribution of propensity and (i) the local potential energy (Section 3.2.2), and (ii) the local free volume (Section 3.2.3). In Section 3.2.4, we look at the effect of coarse graining the potential energy and the free volume on their spatial correlation with the propensity. We then consider a number of other structural measures (Section 3.2.5), including force networks and the proximity to crystalline clusters of large particles. In Section 3.3.1 we define a local Debye-Waller factor as a measure of structural 'looseness' and study its ability to predict the spatial distribution of propensity. We also use this quantity to study the time evolution of dynamic heterogeneities (in Section 3.3.2), and in Section 3.3.3 we investigate the ability of the geometric free volume to predict the spatial distribution of the short-time dynamics. Finally, we summarise the main results, draw some conclusions and suggest directions for future work (Section 3.4).

3.2 The Failure of Structural Measures to Predict the Spatial Variation in Propensity

Phenomenological correlations are the staple of the glass field. There is much evidence from experiments and simulations for correlations or anti-correlations between various bulk dynamic and thermodynamic properties (see, for example, ref. [116] and references therein). This is perhaps not surprising when one considers that, ultimately, it

must be the interaction potential that determines all dynamic, thermodynamic and vibrational properties of glass-formers. However, it is important to recognise that a correlation between two variables does not necessarily imply that there is a causal link between the two, i.e. that a change in one property is responsible for the change in the other. We propose that it is necessary that a strong microscopic correlation exist for there to be a causal link between two properties. In particular, we are interested in establishing whether or not structural measures are able to explain the spatial variation in propensity in a dense amorphous alloy.

3.2.1 Local Coordination Environment

We investigate the local coordination environment of each particle by defining a measure in terms of the number of neighbours it has of each type. The particles of species a that are nearest neighbour to a particle of species b are defined as those that lie within a distance cut_{ab}, the distance to the first minimum in the appropriate partial pair distribution function (PPDF) $g_{ab}(r)$. The PPDFs for this model can be found in ref. [67]. For reference, the cutoff distances, cut_{ab}, used to define nearest neighbours for $T = 0.4$–1 were (cut_{11}, cut_{12}, cut_{22})=(1.45, 1.65, 1.85).

We identify a particular neighbourhood, or local environment, with the following notation: A small particle with m small neighbours and n large neighbours is designated as Smn and the analogous large particle is indicated as Lmn. For example, S14 indicates a small particle surrounded by 1 small and 4 large neighbours, and L06 indicates a large particle surrounded by 0 small and 6 large particles. In this way, all the particles are divided into subsets.

For the following analysis, we pool the data from ten well-spaced configurations of the 2D binary mixture, each separated by a run time of $75\tau_e$ from the previous one. Propensities were calculated using a total of 100 runs for each configuration.

Figure 3.1(b) shows the populations of the different environments at $T = 0.4$. The small particles have either five or six neighbours, while the large particles find themselves in six- or seven-fold environments. Hence the four distinct clumps in the distribution. Also note the large number of different environments. The purely steric interactions of the 2D binary mixture do not lead to any significant chemical ordering in the liquid, unlike, for example, the Kob-Andersen model. The fractions in (a), (b) and (c) are relative to the total number of particles in each distribution.

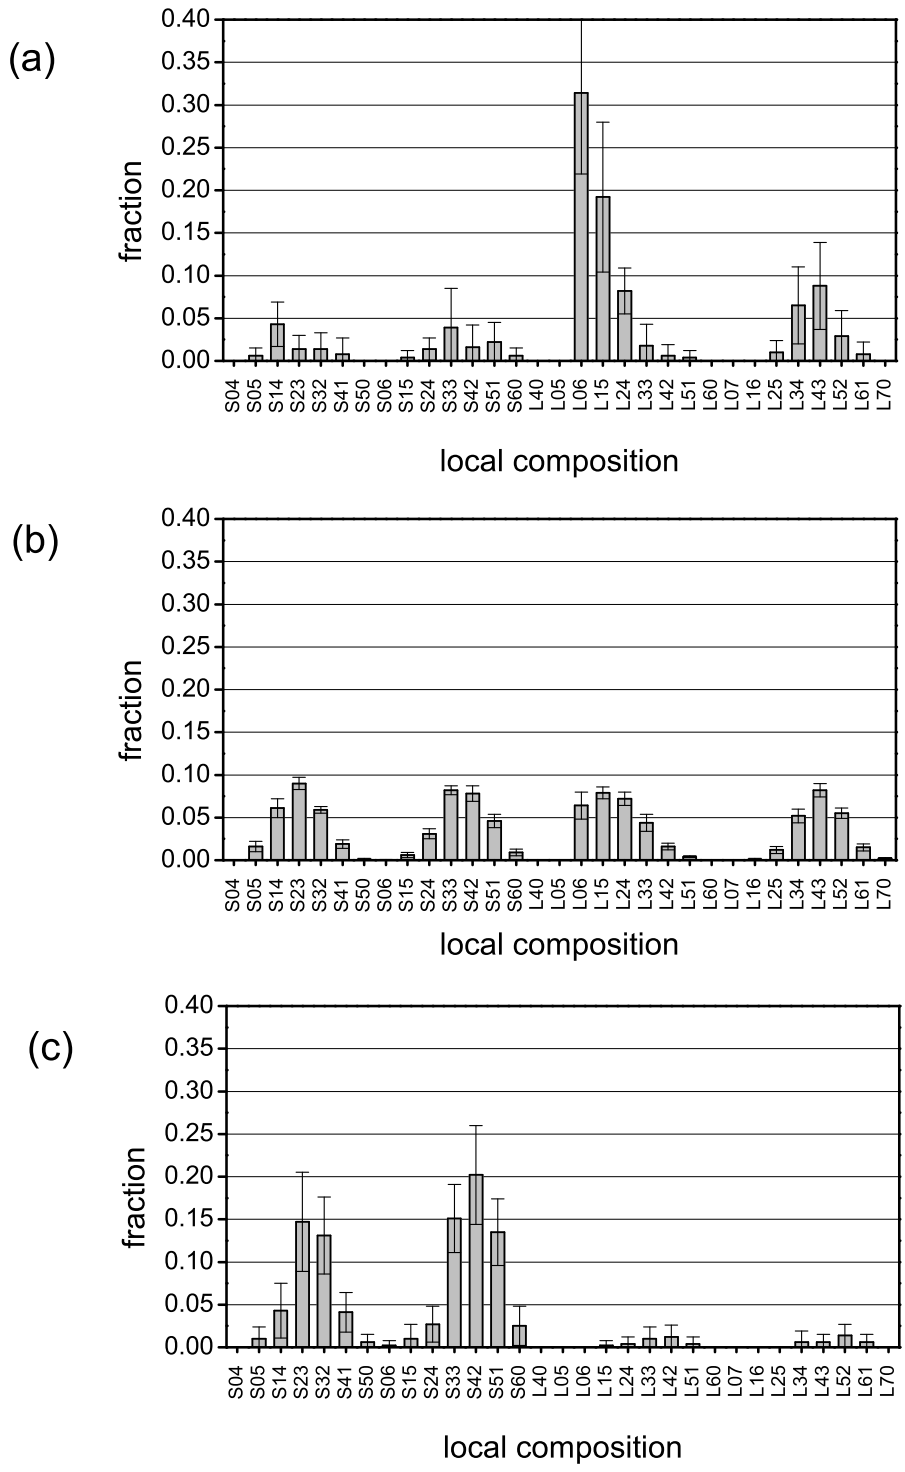

Figure 3.1: The distribution of local coordination environments in the binary mixture at $T = 0.4$ for (a) the slowest 5% of particles, (b) all the particles, and (c) the fastest 5% of particles. The environment code is explained in the text.

In Figure 3.1(a) we present the distribution of environments found among the slowest 5% of particles, i.e. the lowest 5% when ranked by their propensities. This subset is dominated by large particles, and in particular the L06 environment corresponding to local hexagonal packing of large particles. Particles with the L15 environment are also prominent. This finding is consistent with the earlier observation that large particles with high local six-fold orientational order had, on average, longer relaxation times than large particles with low six-fold orientational order [44].

More interesting perhaps, is the question of where motion will happen. Figure 3.1(c) shows the distribution of environments for the 5% fastest particles, i.e. the highest 5% when ranked by their propensities. This subset is clearly dominated by small particles, with a preference, relative to the total distribution in (b), for the more mixed environments: S23, S32, S33, S42, and S51.

Thus, we find some variation in the distribution of local environments between particles with high and low propensity, which is consistent with previous observations of differences in the radially averaged structure about particles with different mobilities [46,79]. However, consider a particle with a given coordination environment and it is clear that it may be present in either the high or the low propensity subsets, i.e. the local environment does not determine the variation in propensity. Another way to visualise this is to plot the distribution of propensities for the different environments.

The distribution of propensities for a selection of local environments are shown in Figure 3.2. While some environments, e.g. L06, provide a relatively strong constraint on particle motion, it is clear that most provide little constraint. Particles with low or high propensity can have almost any local environment. We conclude that some property other than the local coordination environment must be responsible for determining which of the particles exhibits high propensity.

While the local environment is unable to predict the spatial distribution of propensity, there is evidence for a general trend from lowest to highest average propensity as the number of large neighbours decreases and the number of small neighbours increases. Overall, particles with the L06 environment have the lowest propensity on average and particles with the S60 environment have the highest. That is, there is some correlation between local coordination environment and propensity when these variables are averaged over subsets of particles. However, we find no microscopic correlation of sufficient strength to indicate a causal link between the two.

For comparison, we plot the distribution of local coordination environments at

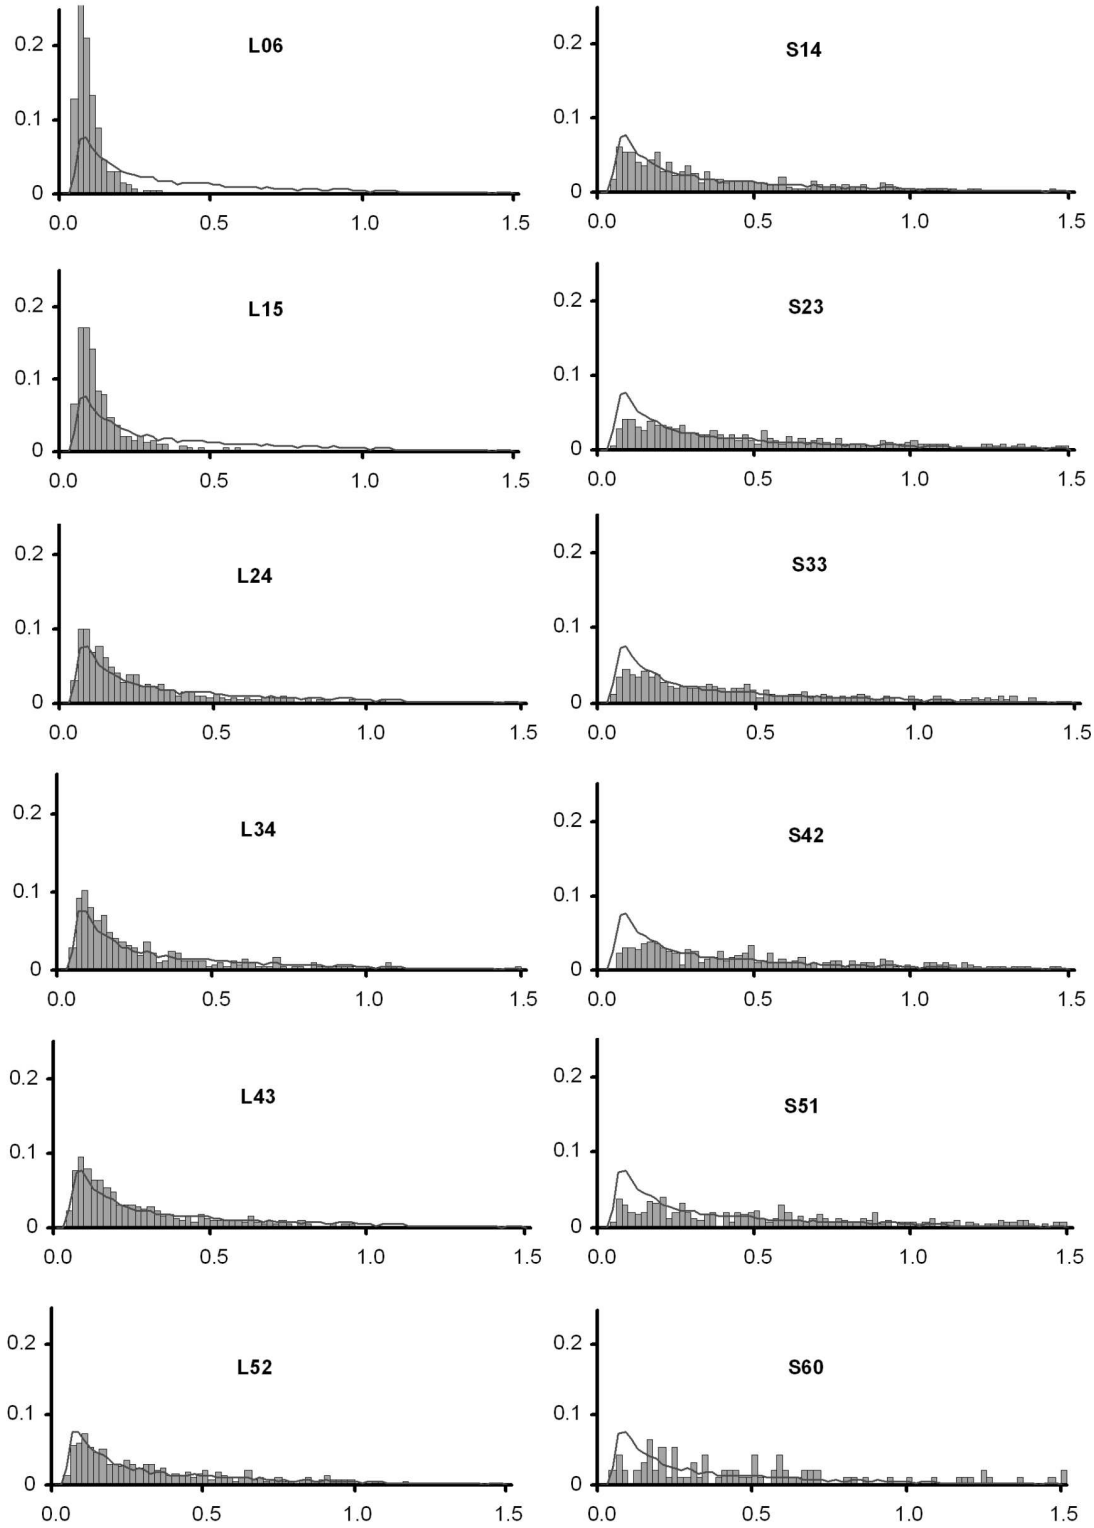

Figure 3.2: The distribution of particle propensities for selected local coordination environments at $T = 0.4$. The environment code is explained in the text, and the solid line indicates the propensity distribution over all particles.

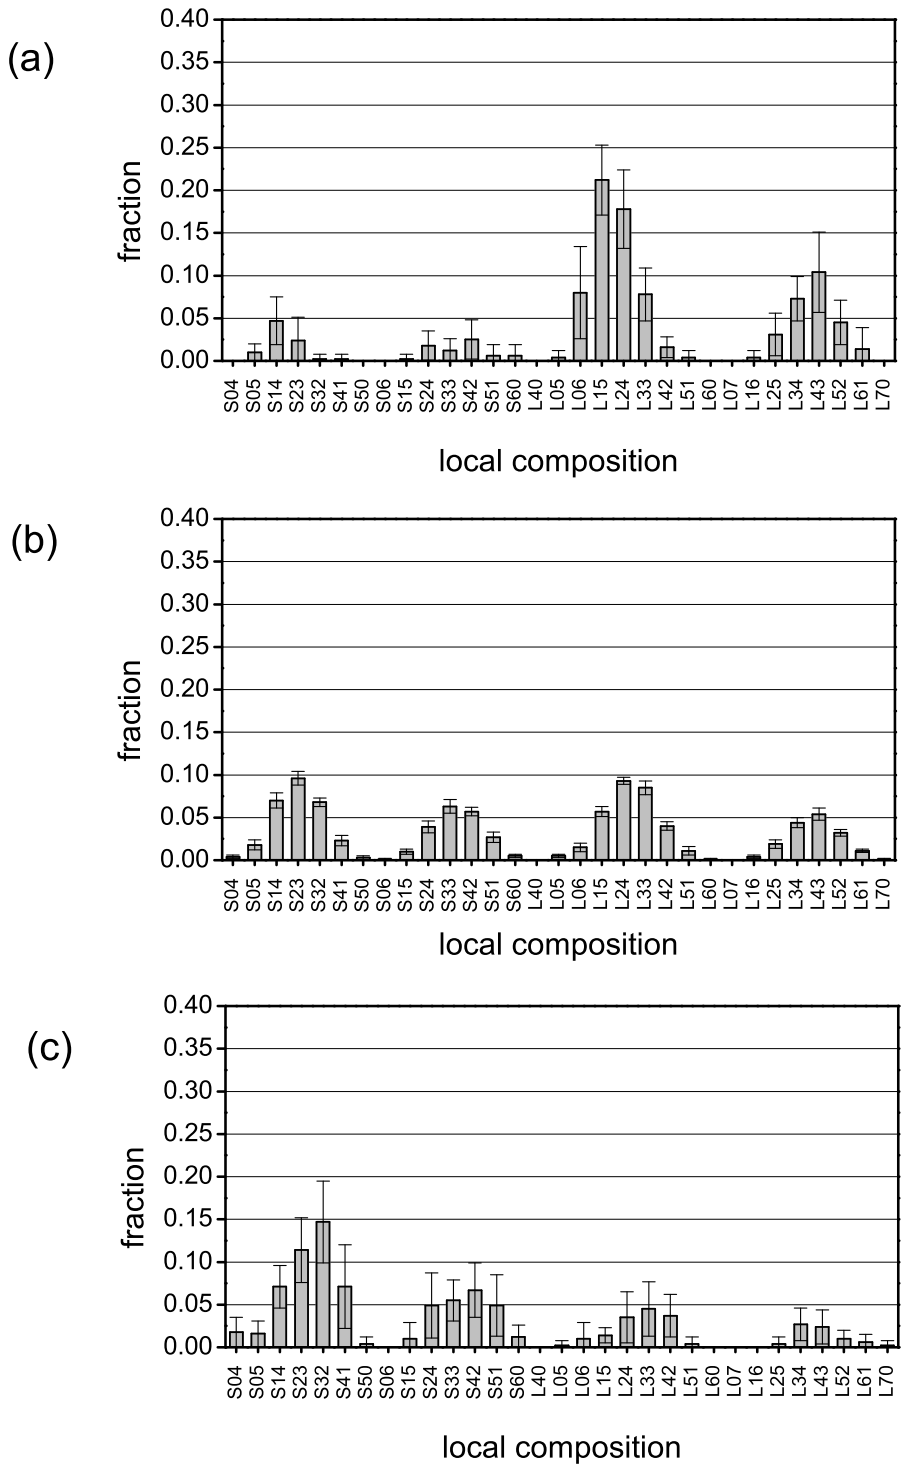

Figure 3.3: The distribution of local coordination environments in the binary mixture at $T = 1$ for (a) the slowest 5% of particles, (b) all the particles, and (c) the fastest 5% of particles. The environment code is explained in the text.

$T = 1$ in Figure 3.3. The proportion of L06 environments is now much smaller among the slowest 5% of particles, but the lowest and highest propensity subsets are still dominated by large and small particles, respectively, thus demonstrating that this is not a unique feature of low-temperature dynamics. It is also clear that the local environment does not provide a strong constraint on the propensity of a particle.

We therefore conclude that, while on average small particles tend to be more mobile than large particles, and the L06 environment is a good predictor of low propensity in the supercooled liquid, the local coordination environment is in general unable to predict the spatial distribution of dynamic propensity.

3.2.2 Potential Energy

We define the potential energy u_i of particle i as the sum over all potential interactions between it and its nearest neighbours (see Section 3.2.1 for the definition of nearest neighbours). We find that this local definition of the potential energy (PE) accounts for greater than 99.5% of the total potential energy in the soft-disc mixture at $T = 0.4$.

As described in the introduction, Doliwa and Heuer [86] and La Nave and Sciortino [110] have reported correlations between the dynamics of small systems and the inherent structure (IS) energy, but did not look at whether the correlation extends to the spatial distribution of these two quantities. We shall now examine the correlation between the spatial distribution of potential energy in the initial IS and the particle propensity. The inherent structures were obtained from the initial instantaneous configurations via a conjugate gradient minimisation of the energy and were used instead of the instantaneous configurations in order to remove the random effect of thermal fluctuations on the structure.

In order to aid in visualising the spatial variation of various quantities, we use contour plots. These are generated by interpolating the irregularly spaced particle properties onto a grid, as described for the propensity in Section 2.3. In Figure 3.4 we compare the spatial distribution of the IS potential energy and of the propensity for a configuration at $T = 0.4$. Note that the two maps are quite different, and that the spatial heterogeneities of the IS particle energies involve considerably shorter lengthscales than those of the propensity. We obtain similar results for nine other configurations at $T = 0.4$.

To quantify the spatial heterogeneity in a given property P, we use the following cluster analysis. The 10% of particles (102 particles in this case) with the largest

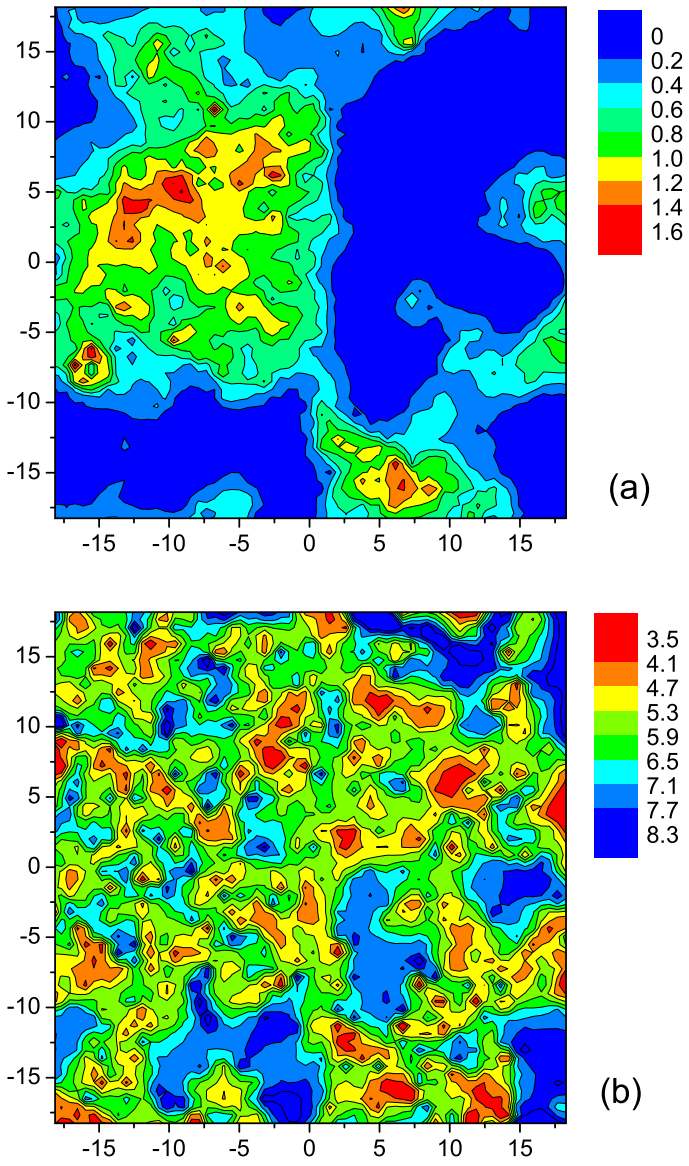

Figure 3.4: The spatial distribution of (a) dynamic propensity for a configuration taken from an equilibrated system at $T = 0.4$, and (b) potential energy per particle for the inherent structure of the same configuration used in (a). Propensities were averaged over 100 runs.

values of P are 'tagged'. Each tagged particle is then assigned to a cluster if it is a nearest neighbour to another tagged particle already in that cluster. When all tagged particles have been assigned to a cluster we count the number of clusters and calculate the variance in the number of particles per cluster. The maximum variance possible for a given number of clusters N occurs when $N - 1$ clusters consist of one particle

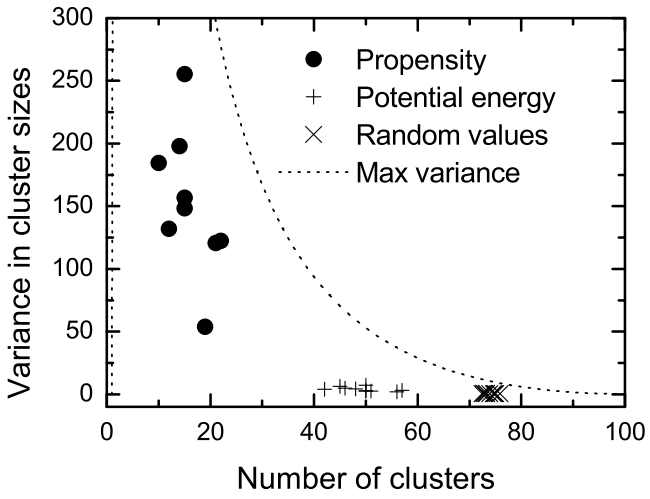

Figure 3.5: Cluster measures of spatial heterogeneity for particles with potential energies and propensities in the top 10%. Data points are shown individually for ten configurations at $T = 0.4$. Statistics obtained using random values are shown for comparison. The dotted line represents the maximum variance possible for a given number of clusters (see Eq. 3.1).

and one cluster consists of $102 - (N - 1)$ particles and is given by the relationship

$$\max(\sigma^2) = -205 - 10404/N^2 + 10608/N + N. \tag{3.1}$$

A random distribution without any spatial correlation will produce a large number of clusters with a correspondingly small variance, while a heterogeneous distribution will produce a smaller number of clusters.

In Figure 3.5 we plot the results of the cluster analysis for the propensity and the IS potential energy for ten configurations at $T = 0.4$. Particles with high potential energy cluster slightly more than an equal number of randomly selected particles, but significantly less than particles with high propensity.

These results highlight the absence of any significant spatial correlation between a particle's potential energy and its propensity. The apparent contradiction between our results and the previous reports [86,110] underscores the difficulty in interpreting correlations. A correlation between *average* values of two quantities does not necessarily mean that a microscopic, and hence causal, relationship exists. Conversely, if no microscopic correlation exists then we can explicitly rule out a direct causal link between the two properties. These are both important steps en route to understanding the macroscopic correlations. In the remainder of this section we demonstrate

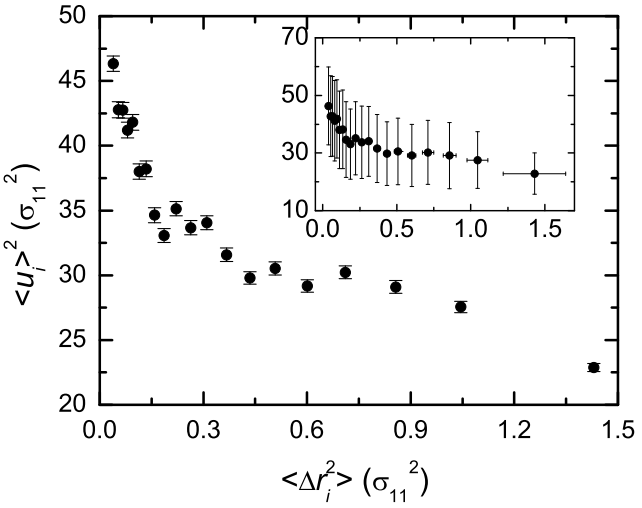

Figure 3.6: Potential energy per particle as a function of propensity. Data for ten configurations at $T = 0.4$ have been pooled together, and the particles divided into 20 subsets according to their propensities. Each subset is represented by a point in the graph. Note that we have plotted the square of the average potential energy $< u_i >^2$ in order to keep the units of the two axes the same. Error bars in the main graph represent one standard error. The inset shows the same data but with error bars corresponding to one standard deviation.

that correlations between suitably chosen averages also exist in the present system, despite their absence at a microscopic level.

If the particles are divided into 20 equal subsets according to their propensity, and the average propensity $< \Delta r_i^2 >$ and potential energy $< u_i >$ is calculated for each subset, then we do find a correlation. As shown in Figure 3.6, there is a clear increase in the average propensity as the average potential energy decreases. However, the size of the standard deviations (see inset) clearly demonstrates that it is impossible to predict the propensity of a particle from its potential energy. We have plotted the square of the average potential energy $< u_i >^2$ in order to keep the units of the two axes the same.

Kob et al. [79] have performed a similar analysis for an attractive Lennard-Jones liquid. They found that as the average potential energy increased, so too did the average mobility, and suggested that this was because particles with higher potential energy were able to find and move into more stable environments. For the repulsive 2D mixture that we have studied, the more stable environments - in the sense that they are less likely to be mobile - appear to be those with higher potential energy, i.e. the relationship between average potential energy and average mobility in the repulsive

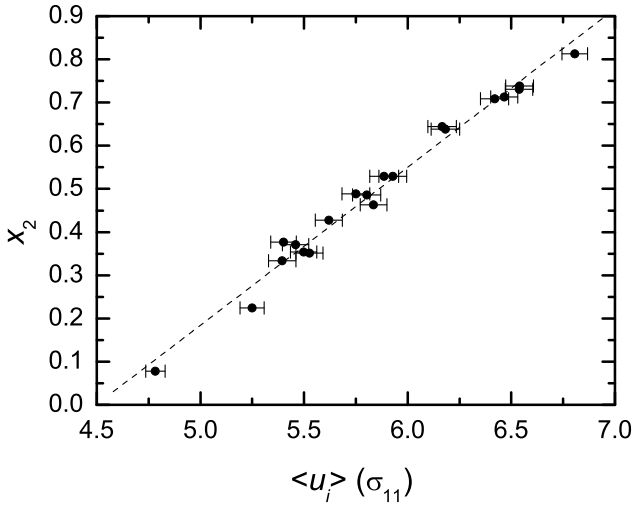

Figure 3.7: Average potential energy $< u_i >$ versus the fraction of large particles x_2 for the subsets considered in Figure 3.6. Each subset is represented by a point in the graph, and the error bars represent one standard error. There is a strong linear correlation between the two quantities as indicated by the dashed line.

soft-disc mixture is opposite to that in the attractive Lennard-Jones mixture.

We also find that, on average, small particles have lower potential energy than large particles, which is consistent with the observation in Section 3.2.1 that the subset of particles with the highest (lowest) 5% of propensities is dominated by the small (large) particles. To investigate this further, we again consider the propensity subsets but this time plot the average potential energy $< u_i >$ of each subset against the fraction of large particles x_2 in the subset. The results are plotted in Figure 3.7. There is a clear linear relationship between potential energy and composition, i.e. as the average potential energy increases so does the fraction of large particles in each subset.

In summary, correlations exist between average propensity and average potential energy, and between average potential energy and composition, when the particles are ordered by propensity and averages are taken over subsets. However, these correlations do not persist at a microscopic spatial level, and we conclude that local potential energy is unable to predict the spatial distribution of propensity. In Section 3.2.4 we consider the effect of coarse graining the potential energy, but first we investigate the spatial correlation between free volume and propensity.

3.2.3 Free Volume

Following Sastry et al. [117], we define the geometric free volume of a particle as the volume accessible to that particle with all its neighbours fixed. Our algorithm differs from that described in ref. [117] by considering overlaps between exclusion spheres rather than using the Voronoi construction. While this modified algorithm will not work well for systems with large voids, it provides an efficient method for calculating free volumes in the current system, and we expect it will work well for condensed phases consisting of mixtures of similarly sized particles. For the purpose of this calculation, we created a mapping to a hard-particle system by using a temperature dependent effective hard disc diameter corresponding to the distance of closest approach of two particles. We identified the closest approach by the distance at which the respective partial pair distribution function first exceeds 0.01. At $T = 0.4$ this corresponds to a distance of $0.9\sigma_{ab}$, where σ_{ab} is the lengthscale of the interparticle potential between particles of species a and b as listed in Section 2.2. We note that the relative ordering of particles by free volume, and in particular the spatial distribution of free volume, are both fairly insensitive to small changes in this effective hard disc diameter. Diameters of $0.85\sigma_{ab}$ and $0.8\sigma_{ab}$ resulted in changes in the relative ordering of free volumes by $+/-10\%$, but made little difference to their spatial distribution. And while values of $0.95\sigma_{ab}$ and σ_{ab} caused some particles to have zero free volume, those particles that still had non-zero free volumes were ordered similarly to what we found using other values.

We again consider the ten well-spaced configurations at $T = 0.4$ for which we calculated the propensity distributions. The distribution of free volume was calculated for each initial configuration and for its inherent structure, i.e. the local potential energy minimum obtained when the initial configuration is used as the start of a conjugate gradient minimisation of the energy. Slightly stronger correlations were observed using the inherent structures, therefore we present only the free volume analysis for these configurations.

To account for differences in particle size, we scale the free volume for each particle by $\pi.\sigma_{aa}^2/4$, where a is the particle species. The scaled free volume v_i has the added property that a very local correlation exists between this geometrical measure and the local potential energy (u_i) in both the inherent and instantaneous structures. In Figure 3.8 the raw and scaled free volumes are plotted against the local potential energies of particles in the inherent structure. Unlike the raw free volumes, the scaled

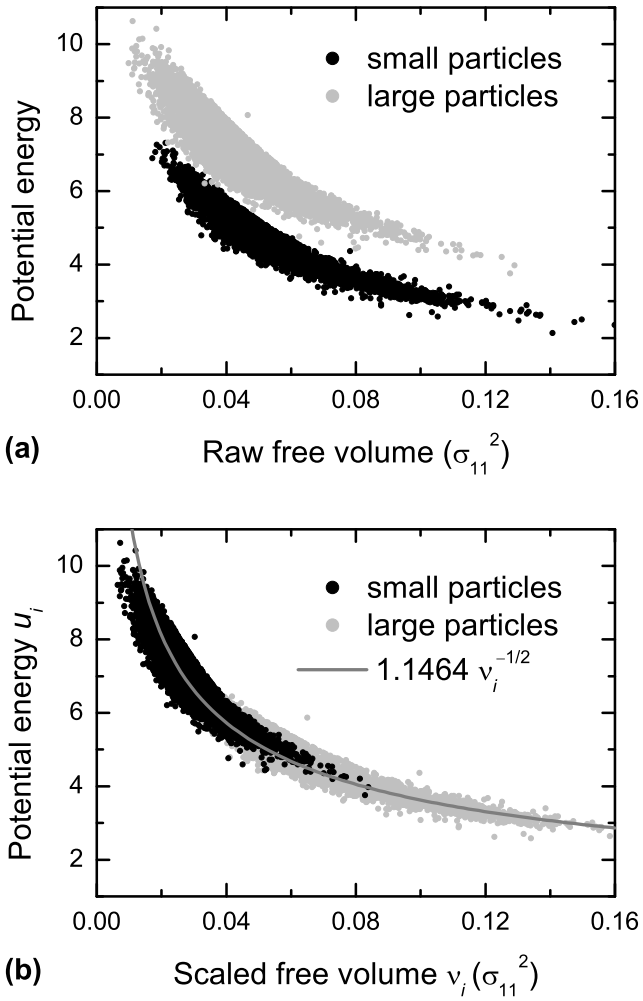

Figure 3.8: The relation between potential energy u_i and (a) raw free volume (b) scaled free volume v_i for particles in ten configurations at $T = 0.4$. Scaling the raw free volumes by $\pi.\sigma_{aa}^2/4$, where a is the species of each particle, collapses the data onto a single curve that is well described by the relation $u_i = 1.1464 v_i^{-1/2}$.

free volumes from the two particle species produce a single smooth curve when plotted against the energy. This curve is well described by the expression $u_i = 1.1464 v_i^{-1/2}$. As a consequence of this relation, most of the results described in this section are similar to those already presented for the potential energy. This will, of course, not necessarily be the case for all glass-forming liquids. From now on we consider only the scaled free volume, which we shall refer to simply as the free volume.

In Figure 3.9 we compare the spatial distribution of free volume and propensity for a configuration at $T = 0.4$. In spite of the coincidence of high propensity and

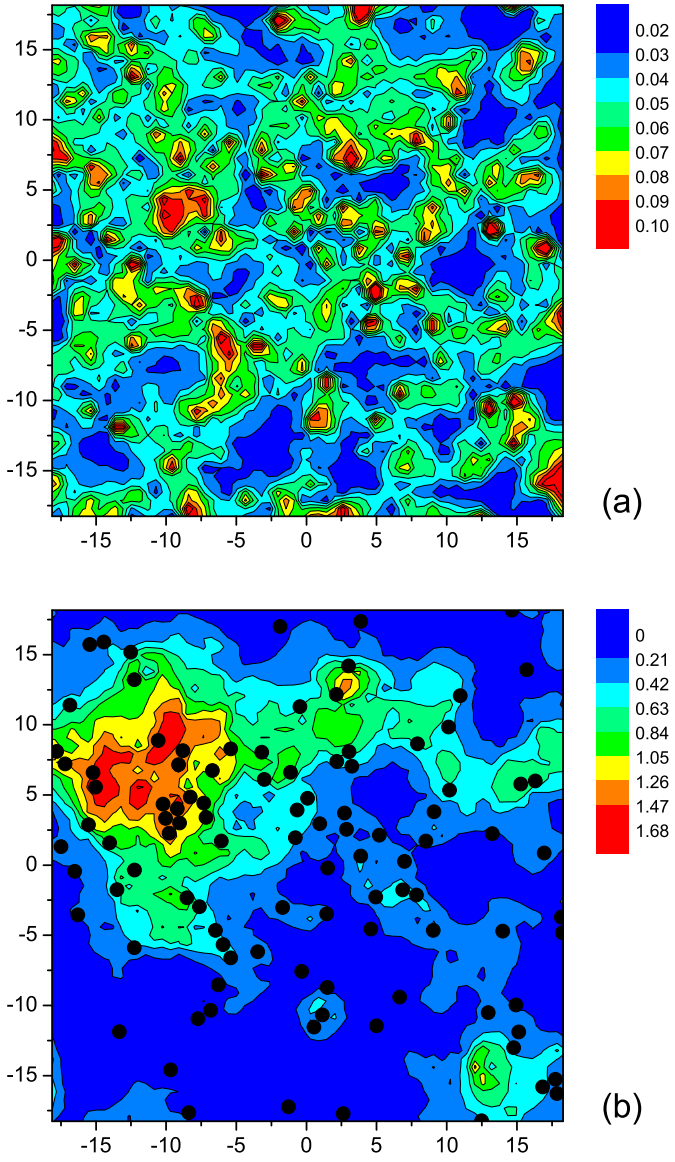

Figure 3.9: Contour plots of the spatial distribution of (a) free volume and (b) propensity for a configuration at $T = 0.4$. The black dots in (b) indicate the 10% of particles with the highest free volume. Propensities were averaged over 100 runs.

high free volume in some cases, there are clearly many 'false positives', i.e. particles with high free volume but low propensity. The spatial distribution of free volume and propensity are also clearly different. This is quantified in Figure 3.10(a) using the cluster analysis described in the previous section. Particles with high free volume cluster only marginally more than an equivalent number of randomly selected particles. For completeness, we repeat the subset analysis by dividing particles into 20 subsets

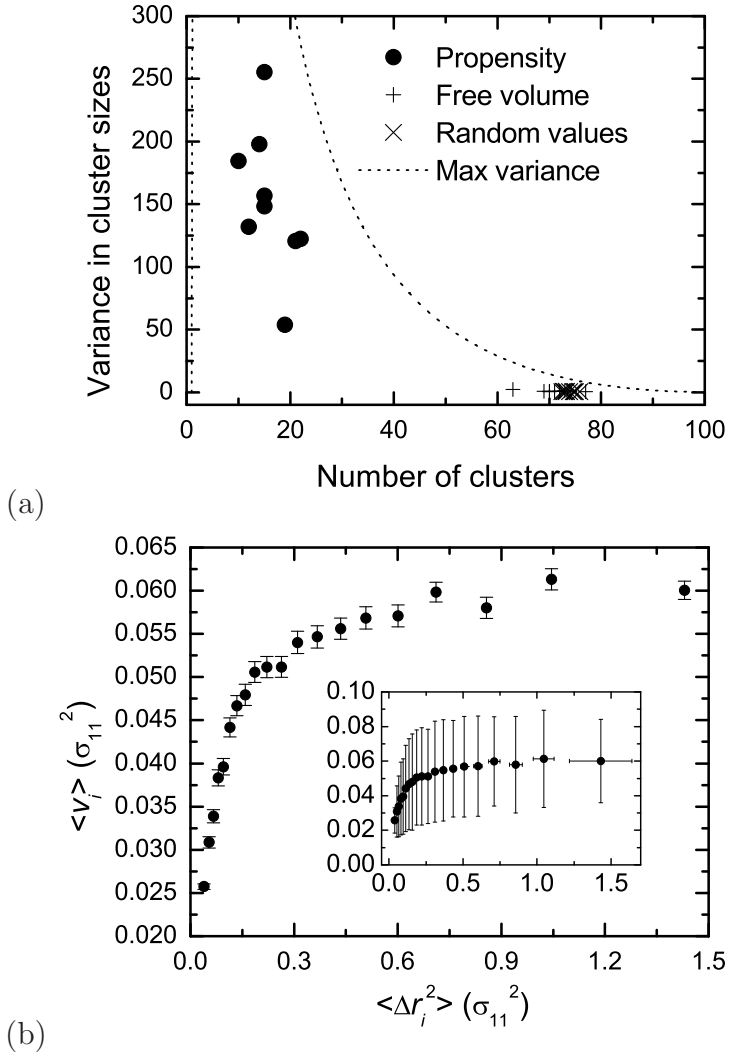

Figure 3.10: (a) Cluster measures of spatial heterogeneity for particles with propensities and free volumes in the top 10%. Data points are shown individually for ten configurations at $T = 0.4$. Statistics obtained using random values are shown for comparison. The solid line represents the maximum variance possible for a given number of clusters (see Eq. 3.1). (b) Free volume as a function of propensity. Data for ten configurations at $T = 0.4$ have been pooled together, and the particles divided into 20 subsets according to their propensities. Each subset is represented by a point in the graph. Error bars in the main graph represent one standard error. The inset shows the same data but with error bars corresponding to one standard deviation.

according to propensity and calculating the average free volume and propensity for each subset. The results are plotted in Figure 3.10(b). Similar to the potential energy analysis, there is an approximately linear correlation at low propensity followed by a rollover to a plateau with an asymptotic value of $v_i \approx 0.06$. However, the large

standard deviations (see inset) clearly demonstrate that there is considerable overlap between the free volume of particles in the different propensity subsets.

In spite of this general failure of the free volume to exhibit any strong spatial correlation with the propensity, there appears to be some cause for hope. Referring to the contour plot of free volume in Figure 3.9(a), we note the presence of a significant number of isolated particles with a high free volume. These 'rattlers' are the source of most of the false positives in Figure 3.9(b). It seems reasonable to expect that the degree to which a particle can 'utilise' a neighbour's free volume depends upon the size of its own free volume. Based on this argument, one may hope to distinguish 'useful' free volume from that which cannot contribute to relaxation based on an analysis of clustering of particles with significant free volume. We therefore study the effect of coarse-graining.

3.2.4 The Effect of Coarse-Graining

Given that the local potential energy and the geometric free volume fail as predictors of the spatial variation in propensity, it makes sense to ask if we can improve the correlation by using a suitable spatial averaging. For example, Qian et al. [50] found that there was an optimal local averaging length (a coarse-graining length) for which the Pearson's correlation coefficient of density and a residence time was maximised in simulations of the low molecular weight glass-formers propylene carbonate and salol.

We coarse-grained the free volume, the potential energy and the propensity by assigning to each particle the value of the relevant property averaged over the local values for that particle and of the particles lying within a distance r of that particle. Values of r in the range $0 \leq r \leq 10\sigma_{11}$ were used. The degree of clustering, as measured by the cluster analysis described in Section 3.2.2, increases steadily with increasing r. This is a trivial consequence of the coarse-graining. The clustering observed in the particle propensity is approximately reproduced in the coarse-grained free volume and potential energy for $r = 2$.

To measure correlation, we use Spearman's rank-order correlation coefficient K [118]. This calculates a linear correlation coefficient of ranks rather than values. We use this method because it is more robust than standard linear correlation methods such as Pearson's. Also, because the distribution of ranks is known (they are uniformly distributed), it is possible to calculate the significance of non-zero values. Except for rare cases where $|K| < 0.1$, we typically obtain non-zero K values with

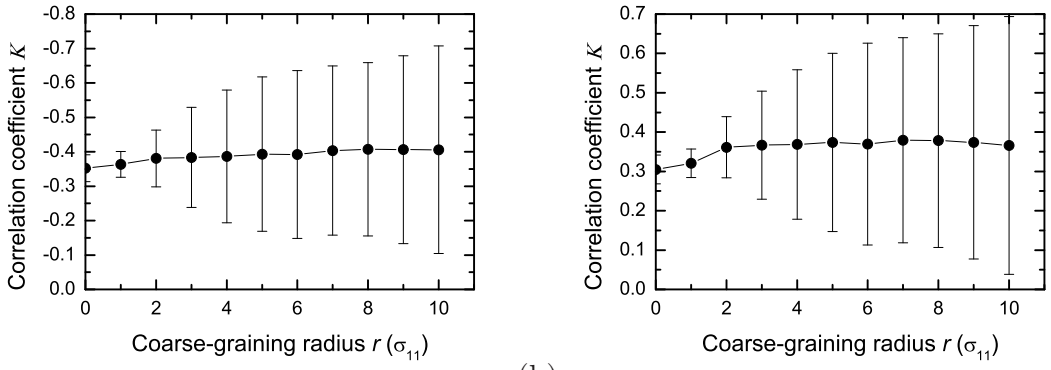

(a) (b)

Figure 3.11: Correlation between propensity and (a) potential energy, (b) free volume, as a function of coarse-graining radius r. Correlation coefficients (Spearman's rank-order correlation) have been averaged over ten configurations. Error bars represent one standard deviation.

greater than 99.9% significance. As for Pearson's correlation coefficient, K varies from -1 to 1, with 0 indicating that the two data sets are uncorrelated.

Figure 3.11(a) shows the correlation between potential energy and propensity as a function of coarse-graining radius. For the case without coarse-graining $K = -0.35$ averaged over the ten configurations studied (potential energy and propensity are anti-correlated). We find only a marginal increase in the average correlation on coarse-graining. The correlation between free volume and propensity as a function of coarse-graining radius is plotted in Figure 3.11(b). Without coarse-graining $K = 0.30$, averaged over the ten configurations studied, and there is only a small increase in the average correlation on coarse-graining. In particular, we note that while coarse-graining improves the correlation for some configurations it makes the correlation worse for others. This is evidenced by the growing standard deviation in the correlation coefficient with increasing coarse-graining radius, for both the potential energy and the free volume.

The average rank-order correlation coefficients that we find are equal to or larger than any of the (Pearson's) correlation coefficients obtained by Qian et al. [50]. Given that we find no spatial correlation between structure and dynamics in the present study, we suggest that it is unlikely that there will be any spatial correlation between the structural and dynamic measures considered in ref. [50].

In light of these results, and the analysis presented in the previous three sections, we conclude that neither the local coordination environment, local potential energy,

geometric free volume, nor simple spatial averaging of the latter two, are able to predict the spatial distribution of propensity. This is despite finding correlations between all of these quantities and dynamics, as measured by the propensity, when suitable averages are taken. We find that the L06 local environment provides a strong constraint on the propensity, and that *on average* particles with lower propensity have higher potential energy, lower free volume and are dominated by large particles. However, our spatial analysis clearly demonstrates that there is insufficient correlation between these reduced measures of structure and the propensity on a *microscopic* level to be able to argue that the particular aspect of structure is the cause of the spatial heterogeneity in the propensity. While it is possible that more elaborate coarse-graining/clustering schemes may improve upon the correlation between propensity and structure, we consider it unlikely that any of these structural measures alone will be sufficient to explain the spatial distribution of dynamics.

3.2.5 Other Reduced Measures of Structure

This section describes some less common descriptions of structure that we also examined. Some interesting results are obtained, but the results are similarly negative or inconclusive as regards predicting the spatial distribution of propensity. In Section 3.3 we describe a different approach to the problem of predicting propensity.

Proximity to Particles with the L06 Environment

A number of studies suggest that relaxation times in amorphous materials are affected by proximity to boundaries. In the case of free-standing polymer films T_g is found to decrease [58], indicating that the relaxation time must be reduced by proximity to a free boundary. And a study of the correlation length of cooperative motion in the facilitated kinetic Ising model [119] found that the relaxation time close to a rigid/free wall was increased/decreased relative to the bulk. In addition, there are plenty of experimental results that find confinement of glass-formers can influence T_g [120]. Ref. [121] reviews recent computational work on the effect of confinement and also discusses some experimental results.

Therefore, motivated by the observation that a large proportion of low propensity particles have the L06 environment, and that the number of L06 particles and their degree of clustering increases with cooling, we consider the following hypothesis: that the spatial variation in dynamic propensity is dominated by the proximity of particles

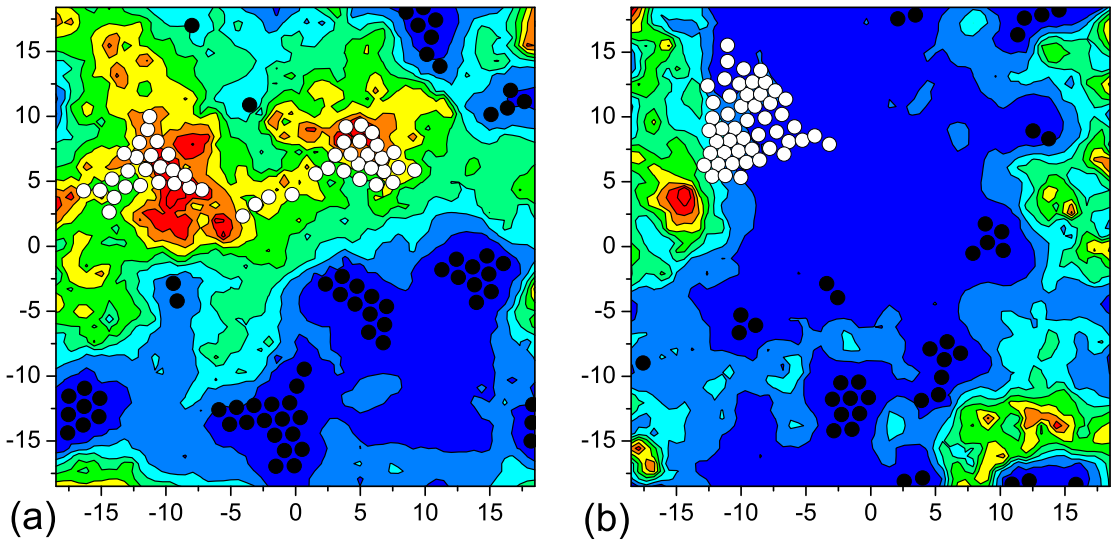

Figure 3.12: Distance from L06 particles as a predictor of the spatial distribution of dynamic propensity for two configurations at $T = 0.4$. The black circles indicate the positions of L06 particles, and the white circles indicate the 50 particles that lie furthest from these. The contour maps have been drawn to different scales to emphasise the spatial variation in propensity.

to crystalline L06 clusters, being higher in regions that are furthest away from these 'rigid walls'. We tested this hypothesis by plotting the location of L06 particles for ten configurations at $T = 0.4$ (the same configurations that were used previously) and compared their spatial distribution to the respective propensity maps. We also predicted the location of high propensity particles using several measures of distance from L06 particles. Finally, we considered the effect of including L15 particles in the analysis. The results shown here are for analysis of the inherent configurations, however very similar results are obtained from analysis of the instantaneous configurations since the local particle environments only change rarely during the conjugate gradient quench.

In Figure 3.12 we plot contour maps for the propensity for two configurations at $T = 0.4$ along with the location of particles with the L06 environment (black circles) and the location of the 50 particles (white circles) that lie furthest from *any* L06 particle. This measure does well at predicting the spatial distribution of propensity for the configuration plotted in (a), but poorly for the configuration plotted in (b). These results are typical of what we find for the other configurations. Sometimes distance from L06 particles is a good predictor of propensity and sometimes it is a

poor predictor. Including particles with the L15 environment in the analysis, by identifying those particles furthest from any L06 or L15 particle, improves the prediction of high mobility for some configurations but does not work in general, e.g. for the configuration in (b).

We also considered a more complex scoring function for predicting the location of high propensity particles in the following. The score for each particle i is calculated as

$$P_i = \sum_{j=1}^{N} \exp^{-r_{ij}/\eta} \delta(\vec{r}_i - \vec{r}_j) \tag{3.2}$$

where the sum if taken over all the particles with the L06 local environment (or the L06 and L15 environments), $r_{ij} = |\vec{r}_i - \vec{r}_j|$ is the distance between particles i and j, η is a weight parameter for the distance, and the delta function ensure the score is 0 if $i = j$. The higher the score the more mobile the particle is predicted to be. We considered η values in the range 1–5, but found no improvement with this measure. With $\eta = 1$ we obtain results that are very similar to the simple minimum distance criterion used above, and we expect the results to approach those above as η approaches 0. And with $\eta = 5$ we found that the high propensity particles were predicted to lie in almost the same position for all configurations.

In general, clusters of hexagonally-packed large particles influence the dynamics by strongly defining where regions of high propensity cannot be located, however they do not appear to strongly influence which of the remaining configuration will have high propensity. We therefore propose that at low temperature there are special configurations that impart high mobility, e.g. when high propensity particles are found very close to L06 clusters they often appear to be part of a group of particles rotating about a single central particle.

The Force Network

Alexander [122] has argued that the relaxation of internal stresses is the most important mechanism in 'selecting' the structure of the amorphous state and in determining its stability. With this in mind, Kustanovich, Rabin and Olami [123] have studied atomic bond tensions in 2D and 3D binary Lennard-Jones glasses at zero temperature and argued that they are not randomly distributed, but instead associated with directional correlations in the nearest neighbour interactions. While they do not explicitly consider dynamics, they make the interesting point that isotropic measures

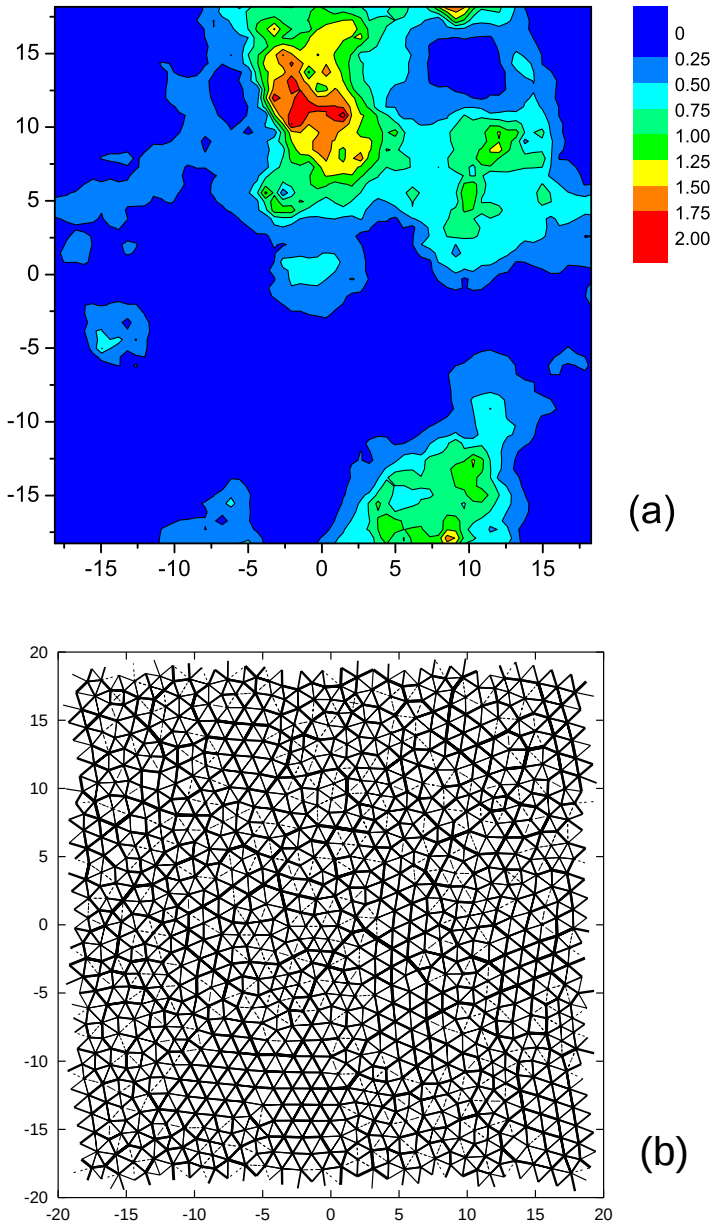

Figure 3.13: (a) The propensity map for a configuration at $T = 0.4$, and (b) the force network for the same configuration. The thickness of the lines connecting particles is proportional to the force between them. Propensities were averaged over 100 runs.

of disorder (which include free volume, local potential energy and local composition) will fail to take into account such correlations. We consider the relationship between the force network and the spatial distribution of propensity as a first step towards investigating whether anisotropic bond tensions are responsible for the spatial variation in dynamics. The results presented here do not form a complete study, but rather

serve to indicate some directions in which this work could be continued.

In the purely repulsive soft-disc model that we study all 'bond' tensions are negative. We therefore consider the force network instead, where the force between two particles i and j is given by

$$F_{ij} = -\frac{dU(r_{ij})}{r_{ij}} = -J_{ij} \tag{3.3}$$

$U(r_{ij})$ is the interparticle potential, r_{ij} is the distance between i and j, and J_{ij} is the 'bond' tension. We calculate the force between each pair of neighbouring particles for the inherent structure of a configuration at $T = 0.4$. The force network, shown in Figure 3.13(b), is represented by connecting neighbouring particles with a line whose thickness is proportional to the force between them. The spatial distribution of propensity for the same configuration is shown in plot (a). While there is no obvious correlation between the force network and the propensities, it is difficult to analyse the anisotropy visually and further numerical analysis may yield more insight and may be worthwhile. Interestingly, we found no correlation between the magnitude of the net force on a particle (in the instantaneous configuration) and its propensity.

3.3 The Single-Particle Debye-Waller Factor as a Measure of Structural 'Looseness'

As we have demonstrated, local structural measures fail to predict the spatial distribution of propensity. We therefore reconsider the question of what determines a particle's ability to move. Ultimately, this must be associated with the degree to which particles are constrained by their surroundings. The hypothesis that potential energy or free volume would correlate with local mobility rests on the expectation that these local scalar measures capture an essential aspect of this constraint. Having found this not to be the case, we now consider the nature of local constraints explicitly.

Thorpe [124,125], building on earlier work of Maxwell [126] and Phillips [127], has shown how the lack of full constraint in network glass-formers is manifest as floppy modes, an observation that appears to offer a connection between a configuration and its dynamic heterogeneity. However, this constraint counting has not been applied to glasses stabilised by dense packing (as opposed to directional bonds) because of the

unsolved problem of identifying local constraints in the former case. Brito and Wyart [128] have recently made some progress on this problem, reporting that the stability of a hard sphere glass can be related to that of networks of elastic bonds through the use of a time-averaged effective potential. On the other hand, Donev et al. [129] have developed an algorithm for testing whether a large continuous deformation of a hard-disc configuration is possible. These deformations, however, are highly collective and their algorithm is unable to single out dynamically significant motions associated with a small number of particles. We will sidestep this problem by looking directly for floppy modes, rather than trying to quantify the constraints responsible. The central question then becomes, is there a spatial correlation between the floppy modes (a measure of the short-time dynamics) and the propensity (a measure of the long-time dynamics)?

There is already experimental evidence that the short-time dynamics of particles can provide information about the long-time behaviour of a system. Buchenau and Zorn have reported that in selenium the mean-squared particle displacement, $< u^2 >$, scales with the viscosity, η, as

$$\eta = \eta_0 exp(C/[< u^2 >_{crystal} - < u^2 >_{liquid}]) \tag{3.4}$$

over many order of magnitude in η [130]. Subsequently, a range of polymeric and small molecule glass-formers have also exhibited strong correlations between the short-time fluctuations associated with the Debye-Waller (DW) factor and the viscosity [131]. And there are other reports in the literature of correlations between the DW factor and fragility [132], the non-ergodicity parameter and fragility [133], and between the shape of the interaction potential and both the fragility and the non-ergodicity, thus tying the latter two results together [116].

Here we provide evidence that such correlations between short- and long-time dynamics can exist at a *microscopic spatial level*. This is especially noteworthy considering that all the other macroscopic correlations that we have tested - between propensity and composition, free volume and potential energy - do not hold at the microscopic level.

3.3.1 Predicting Propensity on the Basis of Short-Time Heterogeneities

To characterise the short-time dynamics, we again consider an isoconfigurational ensemble of runs, all starting from the same 'equilibrated' configuration, with random initial momenta chosen from the appropriate Maxwell-Boltzmann distribution. However, in contrast to the propensity, we consider much shorter runs and study the variance in the particle positions. In this way, we define a particle Debye-Waller (DW) factor.

The DW factor is a standard measure of displacement fluctuations in solids, defined as the mean-squared deviation of an atom from its equilibrium position, averaged over all particles. Thus, one may write the DW factor as $<< (< \vec{r}_i > -\vec{r}_i(t))^2 >>$, where $\vec{r}_i$ is the instantaneous position of particle i, the inner angle brackets $<>$ and $< \vec{r}_i >$ refer to the time average and the outer brackets denote an average over particles. To calculate a DW factor for each individual particle in a configuration, we use a similar definition, except now the outer brackets indicate an average over an isoconfigurational ensemble of runs. We shall refer to this variance in the position of a single particle as the particle DW factor, i.e.

$$DW_i =<< (< \vec{r}_i >_{time} -\vec{r}_i(t))^2 >_{time}>_{ic}, \tag{3.5}$$

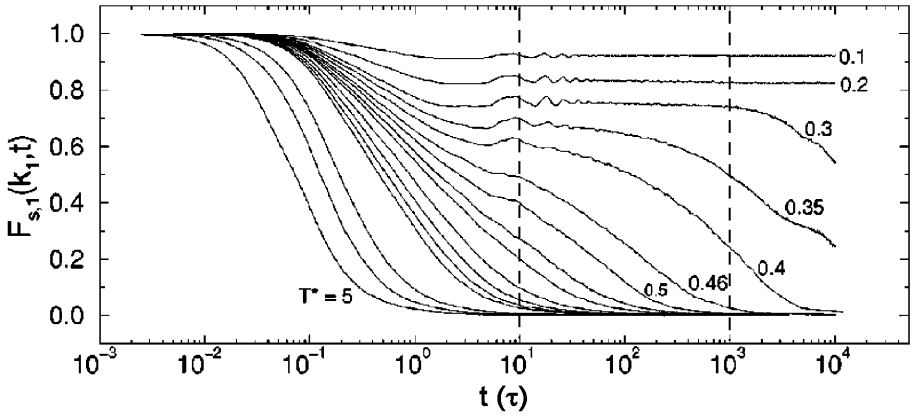

Figure 3.14: The incoherent scattering function as a function of temperature for the small particles in the 2D mixture. The times used to define the local DW factor (10τ) and the propensity (1000τ) at $T = 0.4$ are indicated by dashed vertical lines. Figure reproduced (with modifications) from ref. [44].

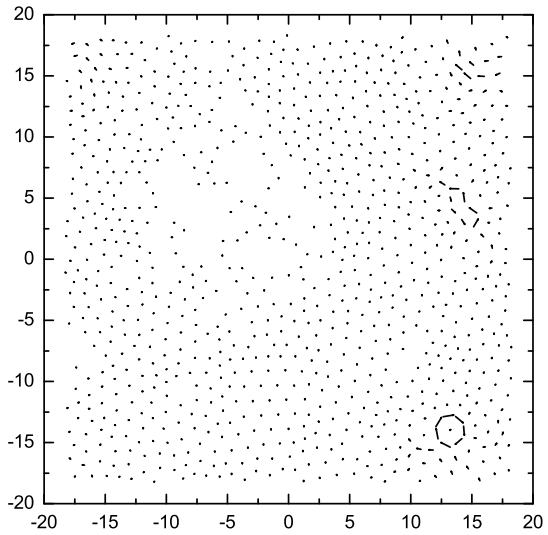

Figure 3.15: Particle displacement vectors connecting the inherent structures belonging to configurations at either end of a 10τ Debye-Waller run at $T = 0.4$. The vectors are drawn as arrows pointing from the initial to final particle positions.

where the time average is taken over the selected time interval and the isoconfigurational average is taken over the multiple trajectories. All results presented in this section are calculated using an ensemble of 100 runs.

We choose a runtime that lies in the middle of the characteristic plateau region in the log-log plot of the mean-squared displacement, and in the log-linear plot of the incoherent scattering function, and is therefore a characteristic time for fast β relaxation. At $T = 0.4$ we use a duration of 10τ to calculate the particle DW factors, which is two orders of magnitude shorter than the 1000τ runtime used to calculate the propensities. The incoherent scattering functions for the 2D mixture, with run times indicated, are shown in Figure 3.14. For reference, the structural relaxation time τ_e (the time at which the incoherent intermediate scattering function has decayed to $1/e$, and therefore a characteristic time for the primary α relaxation) is 673τ at $T = 0.4$.

From comparison of the inherent structure (IS) configurations at either end of the 10τ intervals, we find that this time roughly corresponds to the first 'escape' from the initial IS, involving a small localised reorganisation of particles. In Figure 3.15 we show the displacement vectors connecting IS configurations at either end of a DW run. Note the small localised rearrangements that consist mostly of displacements of less than one particle diameter. Analysing a small sample of runs we find they all end

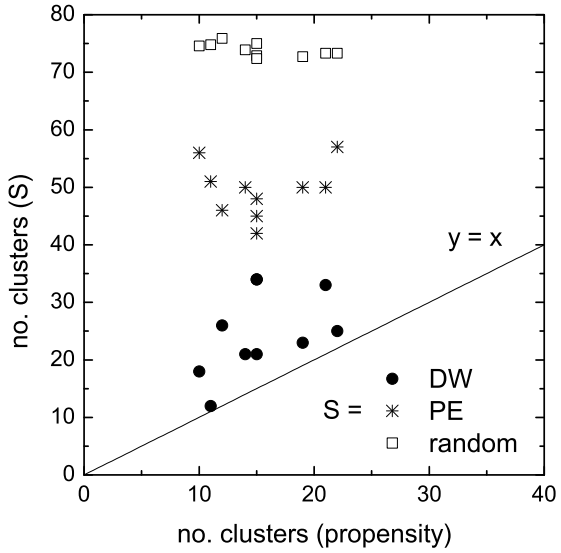

Figure 3.16: A cluster analysis of the spatial distribution of propensity, Debye-Waller factors (DW), and potential energy per particle (PE) for ten independent configurations at $T = 0.4$. 'Random' refers to the number of clusters generated by random samples of 102 particles.

up in different IS's to the initial configuration, and that most of these are unique.

We repeat the cluster analysis described in Section 3.2.2 for the particle DW factors, and compare the clustering of the DW in Figure 3.16. The DW factors result in a relatively small number of clusters, evidence of a substantial and non-trivial heterogeneity, and the number of these clusters is quite similar to that produced by the propensities and significantly smaller than that produced by the PE and random samples. Next we consider how well the local DW factors predict the spatial distribution of the propensity.

To use the DW factors as a predictor of propensity we shall require them to meet two criteria: they must exceed a Lindemann-like threshold of 0.035, which is similar to values obtained for a soft-sphere mixture near T_g [134], and the particles must be in a cluster of three or more particles. Our Lindemann-like threshold is inspired by the observation that crystals can sustain a maximum vibrational amplitude before they melt. This maximum vibrational amplitude, equivalent to the bulk DW factor relative to the interparticle spacing, is usually referred to as the Lindemann criterion [135]. Here we use a similar threshold with the aim of identifying local 'melting', in the sense of local particle mobility. We note that Stillinger [136] has suggested and studied an alternative definition of the Lindemann parameter, that can be extended into the

high temperature liquid phase, based on intrabasin vibrational displacements of the inherent structures accessible at a given temperature. Our minimum cluster criteria is to reflect the fact that any substantial displacement in a dense liquid will require more than one particle to be mobile. In 2D the smallest group of particles that can locally rearrange is three.

We find that the selected particles do very well at predicting the spatial variation of the propensity. In Figure 3.17 we compare the prediction of high propensity using the particle DW factors (black circles) with the propensity maps for six independent configurations. Most regions of high propensity are identified by the selected particles, and very few points lie in regions of low propensity. Compare this to the predictions of high propensity based on the potential energy or free volume maps and the improvement is obvious.

Our data supports the proposition that the high DW regions represent the precursors to the long-time motion, and that the subsequent propagation of the consequences of these 'seed' motions is not readily accessible from the initial configuration, hence the coarse-grained character of the DW factors' predictive success. Interestingly, in a study of a supercooled Lennard-Jones (LJ) liquid, Kob et al. [46] found that the *alpha*-relaxation time of particles with high mobility (measured on the timescale of the maximum in the non-Gaussian function) was on the order of the end of the *beta* relaxation time of the bulk. Thus, it appears likely that there will also be spatial correlation between the short-time and long-time dynamics in this attractive 3D glass-former. Additionally, Vollmayr-Lee et al. [137] have studied the same binary LJ system below the glass transition, characterising the particle mobility via a measure similar to our local Debye-Waller factor but averaged over multiple time intervals rather than multiple runs, and found that the dynamics was spatially heterogeneous. They found that those particles with high vibrational amplitude formed clusters that were compact relative to the lower-dimensional dynamic structures found for the same system above T_g. This is similar to our observation of relatively compact clusters of particles with similar mobility, as measured by the particle DW factor and propensity, relative to the structures found in single trajectories. Furthermore, Vollmayr-Lee and co-workers speculated that the clusters of highly mobile particles below T_g may be the nucleation point of the large scale motions found above T_g. Here, we have demonstrated that *above* T_g the short- and long-time propensities for motion are spatially correlated. We therefore conclude that it is likely that the compact

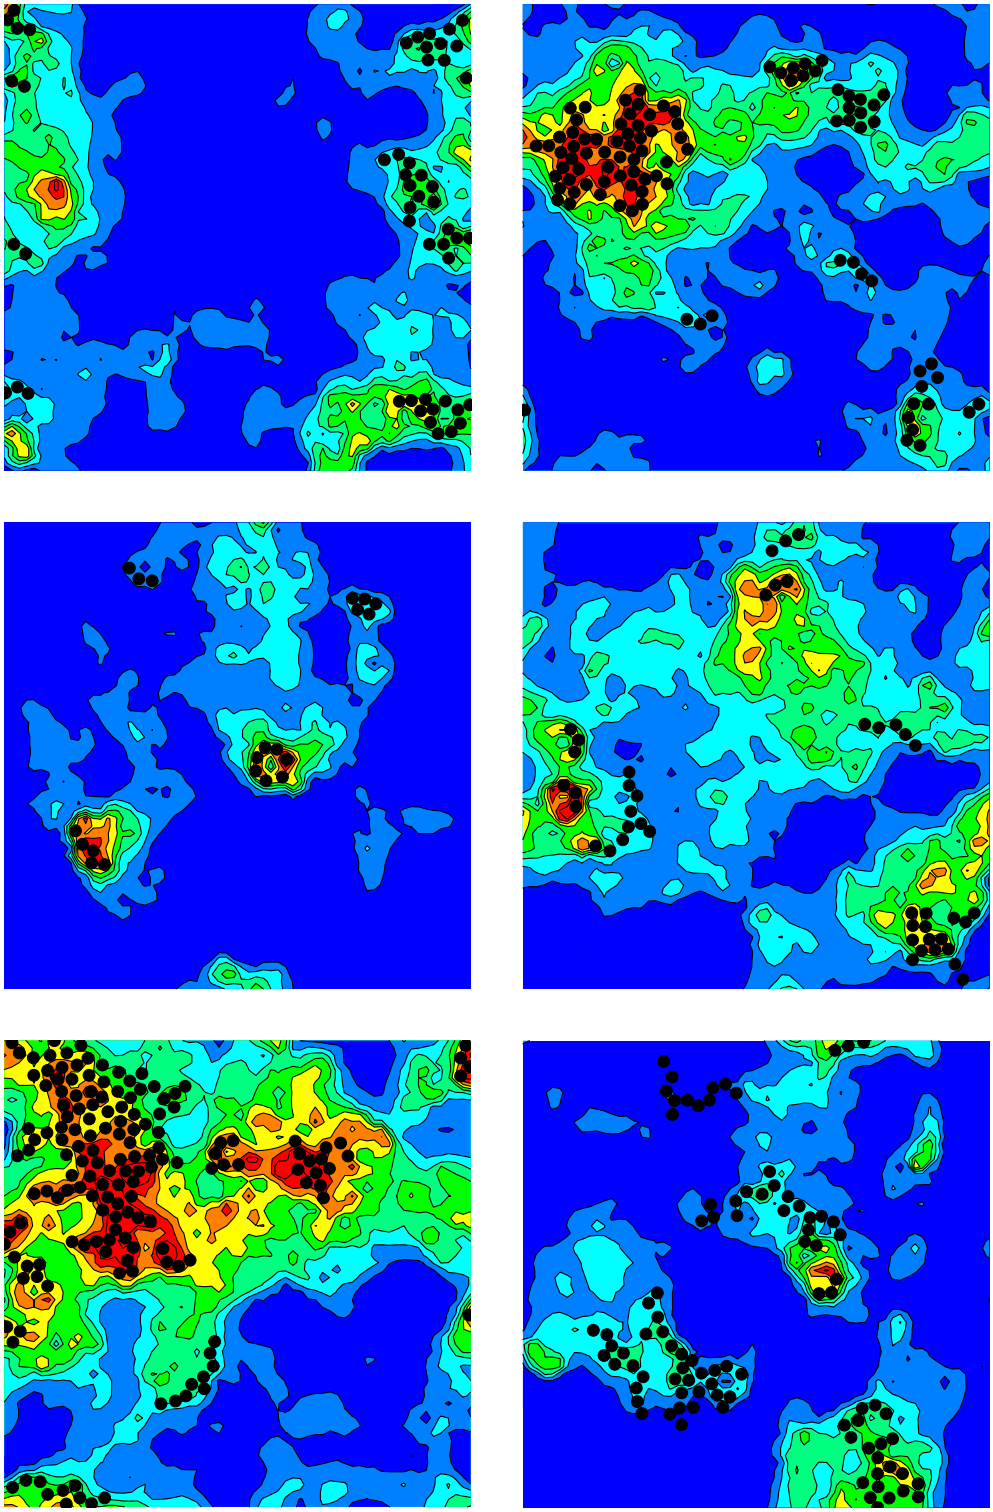

Figure 3.17: A comparison of the predictions of high propensity (filled circles) based on the local Debye-Waller data as described in the text with the actual propensity distributions for 6 independent configurations taken from an equilibrated system at $T = 0.4$. The colour scale is the same for all plots.

mobile clusters found in ref. [137] below T_g are indeed the nucleation point of the large scale motions found above T_g, and that the spatially heterogeneous distribution of these nucleation points is directly due to something in the structure. Laird and Schober [138] have reported spatial clustering of low-frequency modes in a quenched soft-sphere model glass, which suggests that the same relationship may also exist in one-component glass-formers. And Weeks et al. [139] have studied a colloidal glass-former and observed statistically similar clusters of mobile particles on the timescale of the β relaxation in *both* the supercooled liquid and glassy states. This could possibly be experimental evidence for the same relationship between structure and dynamics that we have described here.

3.3.2 Time Evolution of the Debye-Waller Factor

While much effort has gone into studying dynamic heterogeneities in supercooled liquids, little work appears to have been devoted to the time evolution of the dynamic heterogeneities in space. Doliwa and Heuer [47] have studied the time evolution of dynamic heterogeneities in a polydisperse hard-disc system. They found that clusters of slow particles can persist for very long times (up to $1300\tau_e$), that the dynamics of slow clusters are sub-diffusive and highly restricted in space, and that particles leave and join the cluster during its lifetime. In contrast, few particles leave and join clusters of fast particles during their comparatively short lifetime, and they undergo relatively little translational motion.

We have used our definition of a local DW factor to study the time evolution of dynamic heterogeneities at $T = 0.4$. In Figure 3.18, we plot the spatial distribution of DW factors for six configurations separated from each other by the timescale of a DW run (10τ). What is striking is that there are substantial changes in the spatial distribution of mobility over this short timescale. 'Loose' red regions disappear from one configuration to the next, sometimes to reappear again later. And new red regions appear that are not next to any existing red regions. Some configurations, such as the one in (f), have very few 'loose' regions. The timescale over which the distribution of DW factors changes is also significant. Although the system is able to move from one inherent structure (IS) to the next over 10τ, little diffusion of particles is able to take place during this time. Therefore, whatever it is that allows a susceptible subset of particles to become mobile must be able to be transmitted through space without the intervening particles moving substantially. This strongly suggests that the transfer

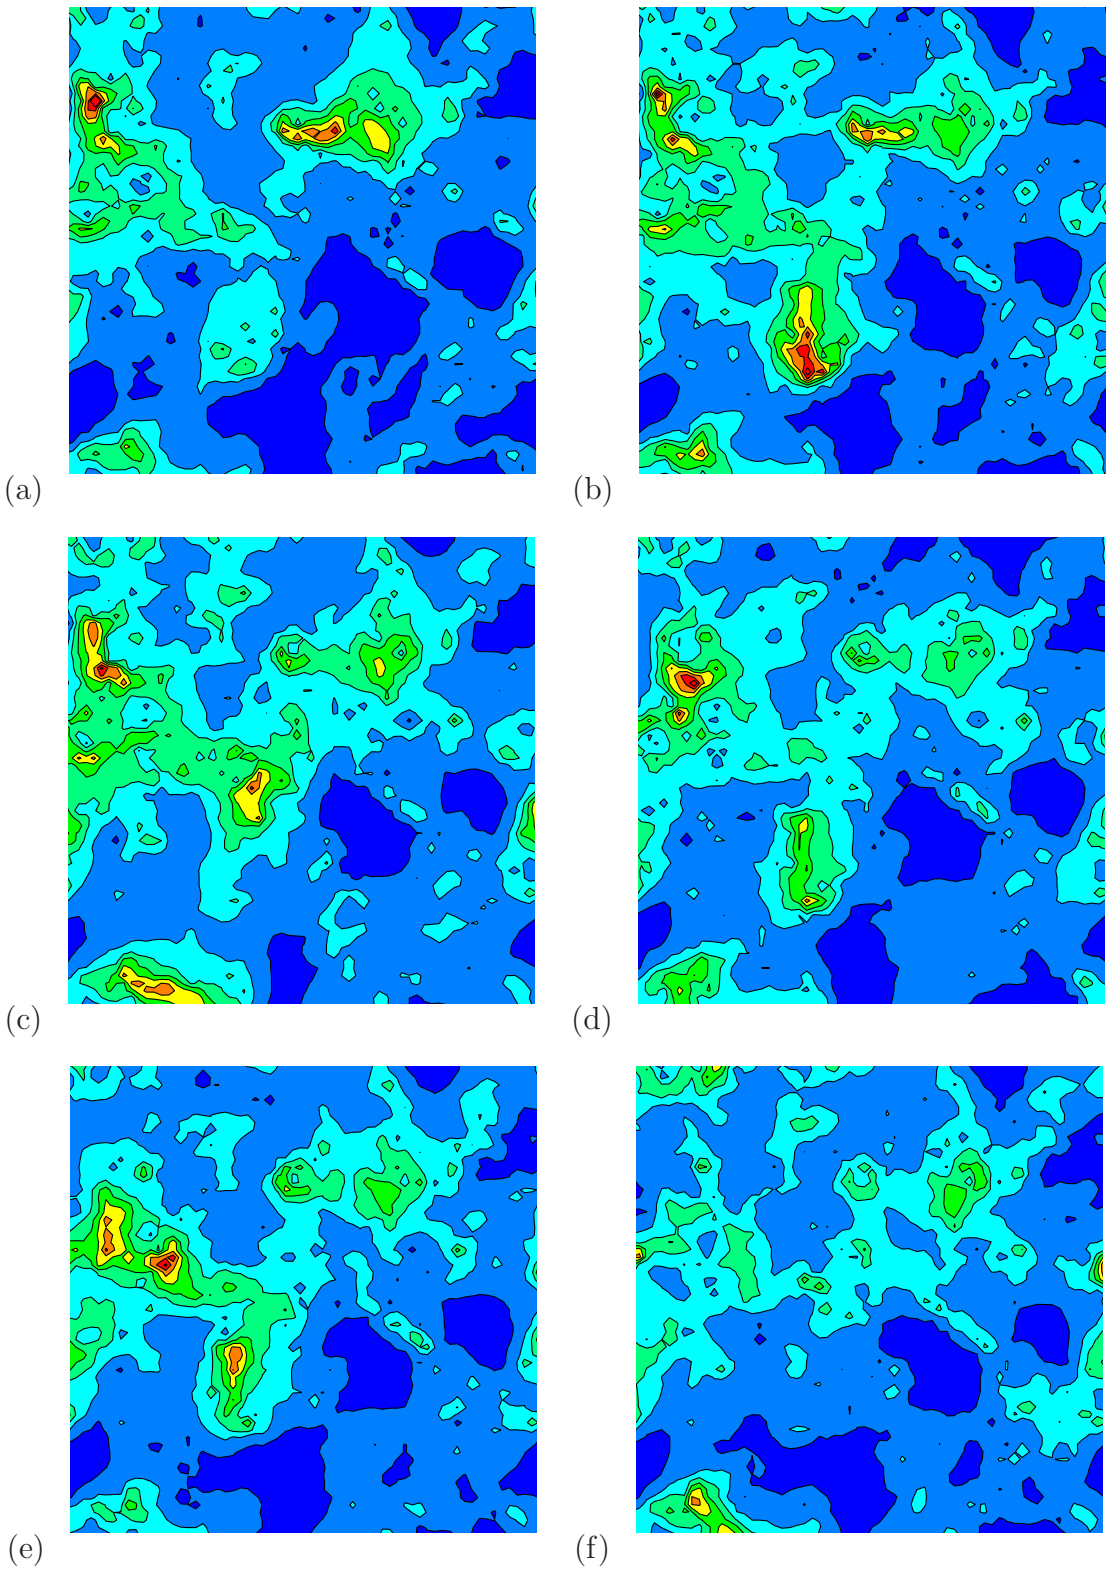

Figure 3.18: Time evolution of the spatial distribution of DW factors at $T = 0.4$. The configurations used to generate plots (a)–(f) are separated from each other by 10τ, i.e. the timescale of a DW run. Ensembles of 100 runs were used to calculate the DW factors.

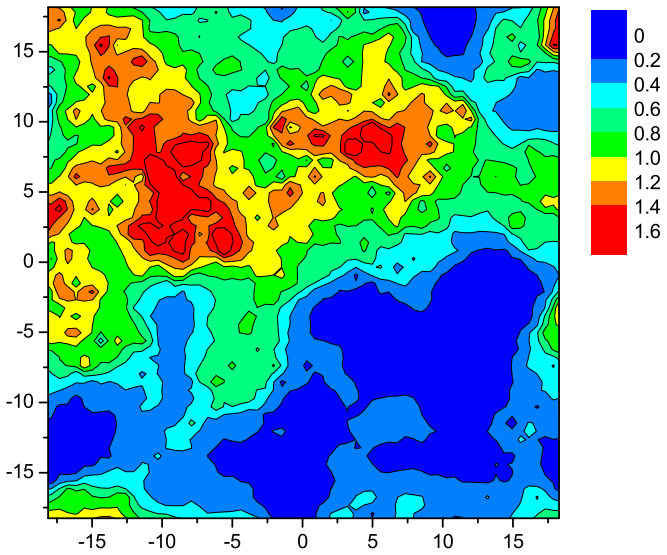

Figure 3.19: The spatial distribution of propensity at $T = 0.4$ for the configuration used to start the time series of DW maps shown in Figure 3.18. Propensities were calculated using an ensemble of 100 runs.

of mobility through the system cannot be explained simply in terms of the concept of defect diffusion [140], i.e. the transfer of some structural defect that allows for particle mobility. Note that it is this latter mechanism that operates in crystalline materials, in the form of the diffusion of vacancies and interstitials [141], and which, in a more abstract sense, is the picture built into facilitation models [71–73].

From the rapid change in the location of mobile regions in the DW maps it seems likely that the effect of the initial configuration on mobility changes faster than the timescale used to calculate the propensity. An obvious question to ask, then, is why does the DW factor predict the spatial distribution of propensity as well as it does? We suggest that this is because: (i) large regions of the structure, with low probability for motion, remain largely unchanged during most runs even after the longer α-relaxation time, and therefore still retain 'memory' of the initial configuration; (ii) as a result, the fast motion is confined to the remaining parts of the structure, resulting in the larger length scale of spatial heterogeneity in the propensity maps compared to the DW maps; and (iii) the propensity will retain 'memory' of the effect of the initial configuration on dynamics long after the instantaneous dynamics has forgotten about the initial structure, since the propensity will not become uniform until all regions have moved the same amount on average, i.e. the initial motion is included in the

propensity and needs to be averaged out by motion that no longer remembers the initial configuration and this won't happen until some time *after* the slowest regions become fast. For comparison, we have plotted the propensity map for the initial configuration used to start the time series of DW maps in Figure 3.19.

Considering the time evolution of the DW maps and the propensity maps for different configurations at $T = 0.4$, a natural dynamic hierarchy of domains can be identified. On short times, these are those that never become mobile (e.g. regions of hexagonally packed large particles), those that are mobile at any given moment, and those that have the potential to become mobile. On the longer timescale of the propensity maps this hierarchy can be extended. The slowest regions to relax are the crystalline L06 clusters, which remain largely intact over the $675\tau_e$ that separate the configurations studied at $T = 0.4$. Then there are non-L06 regions that remain slow on the timescale of the propensity ($1.5\tau_e$), but which relax on a timescale shorter than the $75\tau_e$ between configurations. Next come the regions in the DW maps that have the potential to be 'loose' over short times, and finally those that have a high probability of significant relaxation over longer times (the high propensity regions). The similarity between the DW and propensity maps described in Section 3.3.1 tells us that although the spatial distribution of short-time mobility changes rapidly, the 'loose' spots remain confined to certain regions of the sample over much longer times.

In light of these results, the following questions still await a satisfactory explanation: (i) what determines the susceptibility of a particle region to mobility? and (ii) what is it that can allow a susceptible subset of particles to become mobile and that can be transmitted long distances through space without significant particle motion? Answers to these questions may also help to explain why some non-L06 regions remain slow on the timescale of the propensity.

Continuing the idea of 'stick' and 'slip' introduced in Section 2.4, we suggest that at $T = 0.4$ the most mobile regions can be viewed as marginally stable states in the sense that they are able to 'stick' the configuration, i.e. to not allow diffusional motion in a given run, but more often that not become mobile, i.e. 'slip'. Since their stability is not optimised their structure is also not likely to be unique. This wealth of structural possibilities may help to explain why none of the simple structural measures that we have tested have been able to predict the spatial distribution of propensity. The hierarchy of dynamic domains would then imply a hierarchy of states of different stabilities. This picture of marginally stable states would also explain why mobility is

able to occur, and to be transferred, via small-scale rearrangements in the supercooled liquid. In contrast, in crystalline materials - which have unique optimised structures - motion can only occur via large-scale rearrangements. We note that Nagel and co-workers have also discussed the idea of marginal stable states in the context of granular materials and glasses. See ref. [142] and references therein.

Given that the DW factor is perhaps a better reflection of the effect that an instantaneous configuration has on the dynamics, one may ask whether the structural measures that we have considered are able to predict its spatial variation. The simple answer is, not in general. In the next section we present a direct comparison of the DW factor and free volume, since this directly addresses a correlation reported in the literature between the bulk averages of these two quantities. Additional spatial plots of propensity, DW factors, free volume, potential energy and inherent structures at $T = 0.4$ can be found in Appendix B.

3.3.3 Free Volume cannot Explain the Spatial Heterogeneity of Debye-Waller Factors

As described in the introduction, the free volume concept has been both popular and successful at a phenomenological level, yet there remains a persistent problem concerning the application of free volume to describe dynamics. In particular, what is the relationship between the geometric free volume - a quantity that can be well defined at the atomic scale - and the phenomenological free volume (a macroscopic quantity that is derived from the bulk density)? We address this question here.

For hard disc and sphere systems, and for systems that can be mapped to these, it is possible to define the local free volume unambiguously, as the volume accessible to a particle with all its nearest neighbours fixed in place [117]. Using this method, Starr et al. [80] recently found a power law relation between the average free volume and the bulk averaged short-time mean-squared displacement for monomers in a 'bead-spring' model of a glass-forming polymer over a range of temperatures. This success of the free volume idea was qualified by their failure to find any significant correlation between the mobility of a specific monomer and its local free volume. In this section we use our definition of the local DW factor to test for a causal relationship between the geometrical free volume and the short-time dynamics. We then discuss the relationship between the geometrical and the phenomenological free volume in the context of our results.

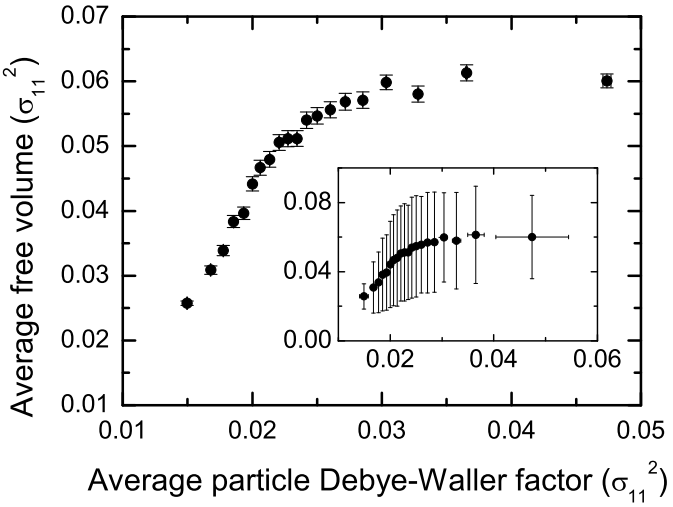

Figure 3.20: Free volume as a function of particle Debye-Waller factor. Data for ten configurations at $T = 0.4$ have been pooled together, and the particles divided into 20 subsets according to their DW factors. Each subset is represented by a point in the graph. Error bars in the main graph represent one standard error. The inset shows the same data but with error bars corresponding to one standard deviation.

We calculate the particle DW factors and local free volume, as described previously, for ten well-spaced configurations at $T = 0.4$. As shown in Figure 3.20, we find a smooth monotonic relation between the particle DW factor and the average free volume when particles are divided into 20 subsets according to their DW factors. For values of the DW factor up to about 0.022 this relationship is linear, a result analogous to that found in Starr et al. [80]. In addition, our work identifies an upper bound on the average particle free volume of about 0.06. This is visible as a plateau for DW factors above 0.03. We note that the maximum value of the DW factor, and hence the length of this plateau, increases with increasing time interval used to calculate the DW factor. For example, compare this plot to the one in Figure 3.10(b). Our results certainly support the phenomenological results of a strong connection between free volume and dynamics. The data also appear to support the idea that there is a well-defined threshold value of free volume or particle DW factor, above which large amplitude displacements occur.

This interpretation fails, however, when applied on a particle-by-particle basis. To see this, consider the standard deviation of the free volumes shown in the insert in Figure 3.20. Clearly, the free volume in a given dynamically defined subpopulation exhibits substantial fluctuations. Particles with wildly varying free volumes can

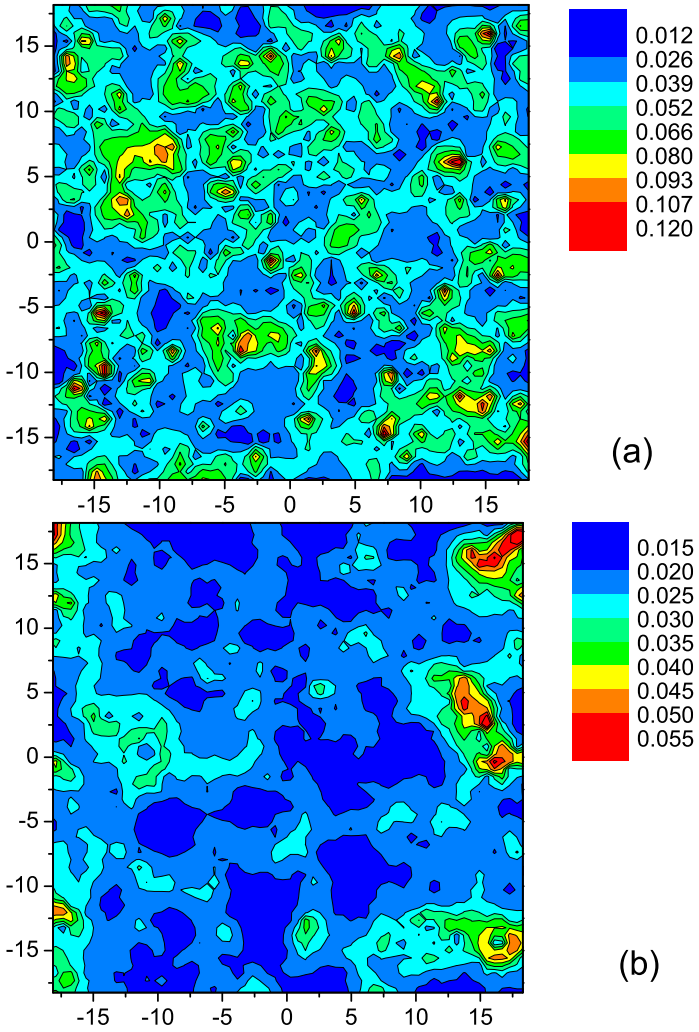

Figure 3.21: Contour plots of the spatial distribution of (a) free volume and (b) particle Debye-Waller factors for a configuration at $T = 0.4$. There is some spatial correlation between regions with low free volume and low DW factors but not in general between regions of high free volume and high DW factors.

exhibit similar values of the DW factor. We conclude that a particle's mobility, as characterised here by the DW factor, is not the result of its geometric free volume. As the amplitude of a particle's DW factor is a measure of the degree to which it is constrained by its surroundings, we conclude that the geometric free volume of that particle can only provide a haphazard glimpse of the degree of that constraint.

Our conclusion, that the variations between particles in terms of their geometric free volume cannot explain the variations observed in their DW factors, is supported by consideration of the spatial distribution of the two quantities. In Figure 3.21 we show contour maps for the free volume and the DW factor for a configuration

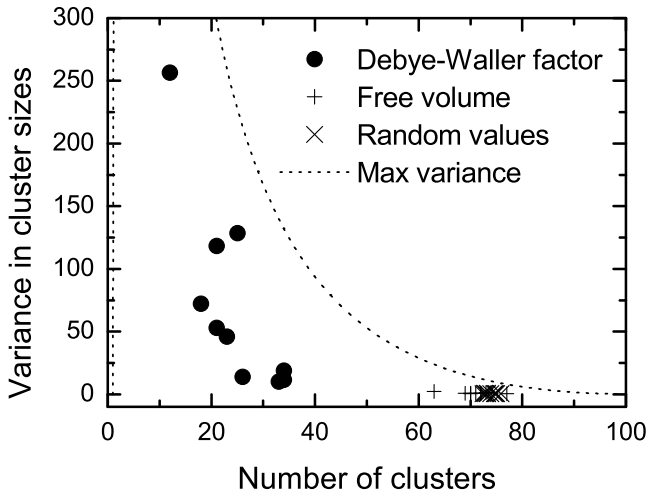

Figure 3.22: Cluster measures of spatial heterogeneity for particles with Debye-Waller factors and free volumes in the top 10%. Data points are shown individually for ten configurations at $T = 0.4$. Statistics obtained using random values are shown for comparison. The dotted line represents the maximum variance possible for a given number of clusters (see Eq. 3.1).

at $T = 0.4$. There is a clear difference in the characteristic length scales of the distributions with the DW factors exhibiting significantly stronger clustering than the free volume.

We have quantified this observation using the cluster analysis described previously (see Section 3.2.2). Figure 3.22 shows the results of the cluster analysis for the free volume and the particle DW factor for the ten configurations. Particles with high free volume show no significantly greater clustering than an equal number of randomly selected particles. In contrast, particles with high DW factor show significantly more clustering. These results, in addition to highlighting the absence of any significant correlation between a particle's free volume and its DW factor, point to the source of the problem. The clear spatial clustering of those particles with large DW factors is evidence of the cooperative character of even this short-time dynamics. The geometrical free volume fails to capture the subtle configurational features that result in enhanced local motion.

If the geometric free volume fails as a predictor of the local dynamic heterogeneity because the latter relies strongly on non-local correlations, can we improve the relevance of the free volume by using a suitable spatial average? For example, as mentioned in Section 3.2.4, Qian et al. [50] found that there was an optimal local

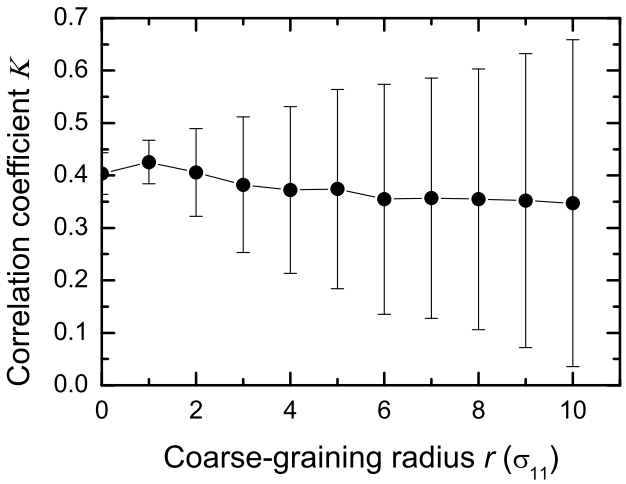

Figure 3.23: Correlation between free volume and particle Debye-Waller factor as a function of coarse-graining radius r. Correlation coefficients (Spearman's rank-order correlation) have been averaged over ten configurations. The error bars represent one standard deviation.

averaging length (a coarse-graining length) for which the Pearson's correlation coefficient of density and a residence time was maximised.

We have coarse-grained the free volume and the local DW factor, in the same way that we coarse-grained the propensity and local potential energy, by assigning to each particle the value of the relevant property averaged over the local values for that particle and of the particles lying within a distance r of that particle. We consider values of r in the range $0 \leq r \leq 10\sigma_{11}$. The degree of clustering as measured in Figure 3.22 increases steadily with increasing r. This is a trivial consequence of the coarse-graining. The clustering observed in the particle DW factor is approximately reproduced in the coarse-grained free volume for $r = 2$.

To measure correlation, we again used Spearman's rank-order correlation coefficient K [118]. This calculates a linear correlation coefficient of ranks rather than values. For the case without coarse-graining we find a value of $K = 0.40$, averaged over the ten configurations studied. Readers are reminded that we have already demonstrated that there is no strong correlation between the scaled free volume and the particle DW factor through the comparison of spatial maps and the cluster analysis. As shown in Figure 3.23, we find no increase in the average correlation between the free volume and the particle DW factor on coarse-graining. The short-comings of geometric free volume as a predictor of dynamics, we conclude, are not to be corrected

by simple spatial averaging.

We conclude that, in this 2D glass-former, having a larger free volume does not cause a particle to exhibit larger amplitude fluctuations in position. Rather, as the correlation between averages over subsets indicates, it only increases the likelihood that the reduced local constraints necessary for large amplitude motion might apply. Even over the short timescales studied here, collective (i.e. non-local) processes are important and these are not well correlated with a purely local measure such as free volume. For this reason, we believe that the results reported here are likely to be common to many glass formers.

In most of its popular usages, however, the phenomenological free volume refers, not to an explicit geometrical volume, but to a reduction of mechanical constraints on particle motion. In this sense, the single particle Debye-Waller factor that we have defined in this chapter probably provides the better match, since it is an explicit measure of particle constraint, even if it lacks a purely geometric definition. If one accepts this proposition - that it is the particle DW factor rather than the geometrical free volume that provides the better microscopic expression of the phenomenological free volume - then the outstanding question for developing a microscopic treatment of dynamics in glassy materials is to see if there exists a method for predicting the particle Debye-Waller factors from a given configuration that is algorithmically simpler than the dynamic averages presented in Section 3.3.1.

3.4 Discussion and Conclusions

In this chapter we have sought to identify a measure of the configuration that can predict the spatial pattern of propensity for the 2D binary mixture. We have examined the local coordination environment, local potential energy and local free volume, and coarse-grained versions of these, as well as the proximity to 'rigid' L06 clusters. In addition, we have used our isoconfigurational approach to examine whether the correlations observed in simulations and experiments between bulk averaged quantities (structural versus dynamic, and dynamic versus dynamic) persist at a microscopic level, in the sense that they are spatially correlated. Such microscopic correlations, we argue, must be present for any causal relation to exist between the two quantities.

While we found that high and low propensity was clearly associated with small and

large particles, respectively, and that the L06 local environment provides a strong constraint on the propensity, some property other than local coordination environment must also be involved. Most environments provide little constraint on the particle propensity, and particles with a given propensity can be found among most of the 23 different types of local environment. We also found that *on average* particles with lower propensity have higher potential energy, lower free volume and are dominated by large particles. However, our spatial analysis clearly demonstrates that there is insufficient correlation between these reduced measures of structure and the propensity on a *microscopic* level to be able to argue that any of these aspects of the structure are the cause of the spatial heterogeneity in the propensity.

This inability of local measures to predict the spatial distribution of propensity is perhaps not surprising. With increasing supercooling the particle motion becomes increasingly collective. The degree to which a particle is constrained must therefore depend not just on the arrangement of neighbours, but also on the degree to which those neighbours themselves are constrained, which in turn requires consideration of the neighbours' neighbours, and so on. Naively, one would expect coarse-graining to provide a non-local extension of, for example, the free volume that reflects this cooperativity and distinguishes 'rattlers' from those particles whose free volume is available for collective reorganisation. Yet we found that simple coarse-graining does not in general improve the correlation between structure and propensity. While the prediction of propensity improved for some configurations it become worse for others.

We therefore conclude that neither the local composition, local potential energy, geometric free volume, nor simple spatial averaging of the latter two, are able to predict the spatial distribution of propensity. The best we can say is that there are definite packings that in general inhibit motion, which can be topologically defined as L06 clusters. There also appear to be definite things in the structure that in general make motion more likely, however these cannot be simply understood in terms of accumulation of free volume, low potential energy, special environments, or distance from 'rigid' clusters.

With this failure of reduced measures of structure to predict the spatial distribution of propensity, we defined a direct measure of structural looseness in the form of the local Debye-Waller (DW) factor. Invoking criteria inspired by a heterogeneous extension of the Lindemann melting criterion for amorphous materials, we have shown that the DW factors provide an excellent predictor of the spatial distribution of the

high propensity domains in each configuration studied. This success is the more striking when compared with the absence of any strong correlation between propensity and quantities like the local energy or free volume. We conclude that the initial configuration determines the local DW distribution (corresponding to the β process), which in turn is the precursor to the subsequent dynamic propensity (characteristic of the α process). These results extend the growing evidence from experiments [130–133] and simulations [116] for correlations between high and low frequency response to the spatial heterogeneities of the two processes. We also conclude that it is likely that the soft spots observed in several different glass-formers below T_g [137–139] are indeed the nucleation point of the large scale motions found above T_g, and that the spatially heterogeneous distribution of these nucleation points is directly due to spatial variations in the structure.

Given the subtlety of the collective mechanical constraints probed by the short-time dynamics, it is very unlikely that any measure of the initial configuration will provide a better prediction of the dynamic propensity than that provided by the spatial distribution of the DW factors. Subsequent answers may improve the algorithmic efficiency in mapping between configuration and the selected DW map, but it is also unlikely that they will improve upon the quality of the answer. If this proposal is accepted then one has, in this work, a sense of the limits one should expect in the answer to the core problem of the glass transition, i.e the causal connection between structure and dynamics. Wolfram [143, 144] has pointed out that there are phenomena in complex systems that are irreducible, in the sense that the future behaviour cannot be obtained by an algorithm more efficient than the solution of the equations of motion. Recently, Israeli and Goldenfeld [145] have qualified this observation by arguing that prediction is possible for suitably coarse-grained versions of the outcome. Our results certainly support the idea that judicious coarse-graining of the structure-dynamics problem is an important part of obtaining a satisfactory solution.

We also studied the time evolution of the DW factor and found that its spatial distribution changed significantly over times almost two orders of magnitude faster than the structural relaxation time. 'Loose' regions appeared and disappeared from one configuration to next, and were not present at all in some configurations. This suggests a hierarchy of dynamic regions on short times: those that are never mobile, those that have the potential to become mobile, and those that are mobile at any given time. These results indicate that whatever it is that allows a susceptible region to

become mobile can be transferred through space without significant particle motion. This picture of dynamics is significantly different from that built into models of defect diffusion and facilitation.

Continuing the idea of 'stick' and 'slip' introduced in Section 2.4, we suggest that a distribution of states with differing stability may provide a better conceptual picture of the low-temperature dynamics. The most mobile regions can be viewed as marginally stable states, in the sense that they are able to 'stick' the configuration, i.e. to not allow diffusional motion in a given run, but more often that not become mobile, i.e. 'slip'. Since their stability is not optimised their structure is also not likely to be unique. This wealth of structural possibilities may help to explain why none of the simple structural measures that we have tested were able to predict the spatial distribution of propensity. The hierarchy of dynamic domains would then imply a hierarchy of states of different stabilities. And this picture of marginally stable states would also explain why mobility is able to occur, and to be transferred, via small-scale rearrangements in the supercooled liquid. In contrast, in crystalline materials, which have unique optimised structures, motion can only occur via large-scale rearrangements.

We also conclude that the success of the DW factor in predicting the spatial distribution of propensity is because large regions of the structure, with low probability for motion, remain largely unchanged during most runs even after the longer α-relaxation time, and therefore still retain 'memory' of the initial configuration. As a result, the fast motion is confined to the remaining parts of the structure, resulting in the larger length scale of spatial heterogeneity in the propensity maps compared to the DW maps. Furthermore, the propensity will retain 'memory' of the effect of the initial configuration on dynamics long after the instantaneous dynamics has forgotten about the initial structure, since the propensity will not become uniform until all regions have moved the same amount on average, i.e. the initial motion is included in the propensity and needs to be averaged out by motion that no longer remembers the initial configuration and this will not happen until some time *after* the slowest regions become fast.

Finally, we examined the microscopic relationship between the short-time dynamics and the geometrical free volume and found no correlation of sufficient strength to indicate a causal link. As discussed in Section 3.1, free volume has been invoked to

explain a variety of dynamic changes in supercooled liquids, and at least one numerical study [80] has found a relationship between the average geometrical free volume and the bulk averaged short-time mean-squared displacement. While having a larger free volume might increase the likelihood that the reduced local constraints necessary for large amplitude motion will exist, we found that in this 2D alloy it does not cause a particle to exhibit larger amplitude fluctuations in position. Even over the short timescale of the DW factor, collective (i.e. non-local) processes are important and these are not well correlated with a purely local measure such as free volume. For this reason, we believe that the results reported here are likely to be common to many glass-formers. In most of its popular usages, however, the phenomenological free volume refers not to an explicit geometric volume, but to a reduction of mechanical constraints on particle motion. In this sense, the single particle Debye-Waller factor, that we have introduced in this chapter, probably provides the better match since it is an explicit measure of particle constraint, even if it lacks a purely geometric definition. If one accepts this proposition, i.e. it is the particle DW factor rather than the geometric free volume that provides the better microscopic expression of the phenomenological free volume, then the outstanding question for developing a microscopic treatment of dynamics in glassy materials is to see if a method exists for predicting the particle DW factors from a given configuration that is algorithmically simpler than the dynamic averages presented in this chapter.

There may be some value in the idea that the stability of the different regions is related to the symmetry of the bond tensions or forces around each atom. Certainly, the most stable L06 domains have the most symmetric local force networks. It is also possible that the relationship between structure and dynamics may be simpler in other glass-formers, for example those with strong chemical ordering, such as the Kob-Andersen model or the 2D glass-formers that we characterise in Part II. It will certainly be interesting to apply the tools developed in this part to the study of structurally different glass-formers.