


Using fluid structures to encode predictions of glassy dynamics

Tomilola M. Obadiya^{*} and Daniel M. Sussman[†]

Department of Physics, Emory University, Atlanta, Georgia 30322, USA

 (Received 31 October 2022; accepted 25 August 2023; published 3 November 2023)

Data-driven approaches that infer the local structures responsible for plasticity in amorphous materials have made substantial contributions to our understanding of the failure, flow, and rearrangement dynamics of supercooled fluids. Some of these methods, such as the “softness” approach, have identified combinations of local structural features in a supercooled particle’s environment that predict energy barriers associated with particle rearrangements. This approach also predicts the onset temperature, often characterized as the temperature below which the system’s dynamics becomes non-Arrhenius and above which local structures are no longer predictive of dynamical activity. We implement a transfer learning approach in which we first show that classifiers can be trained to predict dynamical activity even far above the onset temperature. We then show that applying these classifiers to data from the supercooled phase recovers the same essential physical information about the relationship between local structures and energy barriers that softness does.

DOI: [10.1103/PhysRevResearch.5.043112](https://doi.org/10.1103/PhysRevResearch.5.043112)

I. INTRODUCTION

At the microscopic scale glassy materials possess amorphous structures very similar to their dense fluid phase [1], but their dynamical properties are radically different [2]. Many different temperature scales have been identified in the transition from an equilibrium fluid to an out-of-equilibrium amorphous solid; the first one encountered upon cooling the system is the onset temperature T_0 , which marks a transition to “landscape-influenced” dynamics [3–6]. Above this temperature the particle-level dynamics are simple and diffusive and do not depend on local structure; below this temperature dynamics become both spatially and temporally heterogeneous, depend strongly on local structure, and (for fragile glass formers) begin slowing down superexponentially [7–11].

The role of local microscopic structure in local dynamics is a deep question, and a major theme of recent research has been the search for correlations between them [12–17]. Data-driven approaches, using numerical simulations to generate large data sets for the training of support vector machines (SVMs) [18–20], graph neural networks [21], and other techniques [22,23], have shown substantial promise for finding maximally correlative structures for the dynamics at different timescales in strong glass formers [24], fragile glass formers [19,21,25] and even anomalous glass formers modeled after biological systems [26]. Notably, approaches based on SVMs [19] have found *physically interpretable* classifiers.

This physical interpretation comes from training classifiers at low temperatures and studying how they behave when applied to data at other temperatures: by maximizing correlations in the training data in a compressed way, these classifiers are learning combinations of structural features that can be interpreted as a local energy barrier to particle rearrangements. However, relatively little is known about why these approaches work and whether or how one could use these results to help build a more robust theoretical description of glass-forming systems. Key unresolved questions include why these particular approaches lead to what is apparently a local order parameter for the supercooled liquids and how the learned energy barriers actually depend on the construction of the classifiers.

In this work we bridge the changing dynamical behavior above and below the onset temperature, on the one hand, and the physical interpretation of amorphous state classifiers, on the other. We first demonstrate that the same machine learning techniques that have successfully correlated structure and dynamics in the supercooled phase can be used to classify “extreme diffusive” events even far above the onset temperature. In the spirit of a transfer learning approach, we show that these liquid-state classifiers can statistically identify activated events in the supercooled phase, even though the character of the activated dynamics below T_0 changes dramatically. We further show that not only can accuracy on a classification task be maintained but also the physical interpretability is maintained: apparently, fluid-phase classifiers also learn energy barriers in the supercooled phase.

II. METHODS

A. Model and simulations

Our analysis is focused on a large set of molecular dynamics simulations of $N = 4096$ particles, using the standard

^{*}tomilola.medetin.obadiya@emory.edu

[†]daniel.m.sussman@emory.edu

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

80:20 Kob-Andersen model [27] (with a cutoff distance of 2.5) at a density of $\rho = 1.2$ in a cubic box with periodic boundary conditions. Throughout we report all quantities in dimensionless (reduced) units, using the standard Lennard-Jones (LJ) convention in which the base units are distance (measured in units of the large-particle diameter σ_{AA}), energy (in units of the interaction parameter ϵ_{AA}), and mass (in units of the particle mass m). For this model the dimensionless onset temperature is often reported as $T_0 \approx 0.87$ [7,19]; given the broad crossover in the dynamics, values between 0.8 and 1.2 are also reasonable estimates for this temperature scale [4,11]. Our simulations were done in the NVT ensemble for temperatures in the range $T \in [0.45, 2.0]$ in standard LJ units (i.e., ranging from the moderately supercooled to the supercritical fluid phase). We note that our use of a deterministic Nosé-Hoover thermostat [28] should not unduly influence the results reported below: in our classification task we do not include any information about the fictitious degrees of freedom used by the thermostat to maintain an average temperature, and we expect that, if anything, our quantitative results would improve upon conducting simulations in an NVE ensemble.

All simulations were conducted using HOOMD-BLUE [29], and in the following we focus on the behavior of the large particles. Our simulation configurations were generated as follows. We first equilibrated a system with random initial conditions 5000τ at a temperature of $T = 0.45$. We used the final configuration of this as an initial seed for our other simulations: a snapshot was loaded as the initial configuration for our other simulations, each of which was allowed to equilibrate for 1000τ at its target temperature. After this, data used in this study were saved at intervals of 1τ .

We begin by characterizing the local structural environment of particles in these simulations, using two-point radial structure functions $G_X(i; r, \delta)$ [30]. For a target particle i , these functions are defined as

$$G_X(i; r, \delta) = \sum_{j \in X} \exp\left(\frac{-(r - R_{ij})^2}{2\delta^2}\right), \quad (1)$$

where X denotes which of the components of the binary mixture is being considered, r is a parameter controlling the distance from which dominant contributions to the feature come, δ is a parameter controlling the width of the Gaussian shells, and R_{ij} is the distance between particles i and j . We characterize the local environment of particle i with a vector in a 100-dimensional feature space \vec{F}_i , with $\delta = 0.2$, $0 < r < 5$ in increments of 0.1, and $X = A, B$. Each feature is standardized [31] so that at the training temperature it has zero mean and unit variance.

We next choose a measure for the dynamics of particles; to be consistent with work on activated dynamics we use p_{hop} as introduced in Ref. [32]. We use an observational time window of 10 LJ time units, for which

$$p_{\text{hop}}(i, t) = \sqrt{\langle (\vec{r}_i(t) - \langle \vec{r}_i \rangle_{w_2})^2 \rangle_{w_1} \langle (\vec{r}_i(t) - \langle \vec{r}_i \rangle_{w_1})^2 \rangle_{w_2}},$$

where $w_1 = [t - 5, t]$, $w_2 = [t, t + 5]$, and thus $\langle \dots \rangle_{w_i}$ averages over one half of the observation window. We do not believe that using p_{hop} as a dynamical label is crucial—preliminary results indicate that choosing instead to measure particle dynamics using their cumulative displacement

over the same time window leads to qualitatively identical results—and we note that the length of the time window is optimized for detecting activated events in the supercooled regime.

B. Machine learning protocol

We train SVMs connecting structural features with dynamic observables largely following the “softness” methodology [19]. We build a training set by combing through molecular dynamics trajectories for examples of dynamically active (“rearranging”) and inactive (“nonrearranging”) particles and train a linear soft margin SVM (using the SCIKIT-LEARN package [33]) to classify these examples. We can then use the learned classifier (here, a hyperplane in feature space) to try to predict dynamics based on a particle’s instantaneous environment, and we define the softness of particle i at time t , $S_i(t)$, as the shortest distance between its vector of structural features and this classifying hyperplane.

The training set construction for our softness classifier closely followed the protocol outlined in Ref. [19]. We constructed a balanced 7600-sample training set using the coldest temperature considered ($T = 0.45$): 3800 rearranging samples and 3800 nonrearranging samples. We adopted the previously used convention of associating the structural data of a particle i at time $t - 2\tau$ with the dynamical state at time t . We defined a rearranging particle if, at time t , $p_{\text{hop}}(i, t) > p_c$, where $p_c = 0.2$. We defined a nonrearranging particle by requiring its $p_{\text{hop}}(i, t)$ value to remain less than a lower threshold of $p_l = 0.0085$ for at least 120τ duration of time. We then used the local structure of the nonrearranging particle in the middle of its time of low activity. Unlike in previous work, we take structure and p_{hop} values directly from the thermal configurations rather than quenching to the inherent states (in part because the fluid-phase simulations would be far from any minima). Unless otherwise stated, we used a soft-margin misclassification hyperparameter of $C = 10^{-2}$.

A major finding of Ref. [19] was that this signed distance—softness—encodes the probability the target particle would rearrange at a given temperature. The corresponding curves for the probability of rearranging at different values of S as a function of T all intersected at a common temperature, which in turn suggested the existence of an onset temperature above which structure was no longer predictive of dynamical events. The predicted value of T_0 was consistent with alternative definitions [4,7,19] and with the numerical values cited above. Before we return to this finding, we first ask, Can we learn to classify dynamical events based on structure not in the supercooled regime but at and even *above* the onset temperature?

For $T > T_0$ individual particle motion is diffusive rather than activated, and it is not clear that using p_{hop} as a dynamical label is the most natural choice. We continue to use it as an indicator function—it is still large for dynamical trajectories that move a particle far from its initial position and small for diffusive motions that stay near a particle’s initial position—and will show in a later work that this choice is not crucial to our results. To identify “extreme events” to classify, we select particles in high and low tails of the probability density function of p_{hop} at different temperatures. To have similarly sized training sets as in the case of softness, we choose lower

and upper cutoffs (p_l and p_u) that capture the most extreme 0.033% of low- and high-activity events, respectively. We identify particle i at a given time t as “extremely diffusive” if $p_{\text{hop}}(i, t) > p_u$ and associate it with the particle’s local structure at time $t - 2\tau$. Similarly, if $p_{\text{hop}}(i, t) < p_l$, the particle is identified as “extremely nondiffusive,” and its structure at time $t - 2\tau$ is included in the training set. The training set for each temperature we considered above T_0 contained 10 400 balanced samples.

Aside from this difference in choosing “rearranging” and “nonrearranging” labels, we follow the methodology above: we find a linear soft-margin SVM that best classifies a labeled training set and then apply this classifier to new data. To distinguish it from softness we call the distance of a point in feature space to such a classifying hyperplane the “fluidity,” and we use \mathcal{F}_i^T to denote the fluidity of particle i with respect to a classifier trained from data at temperature T . Given any of our classifiers, one can compute a particle’s softness or fluidity by computing its feature vector (which depends only on the instantaneous structure around the particle) and evaluating $\alpha_i(t) = \vec{w}_\alpha \cdot \vec{F}_i(t) - b_\alpha$, where \vec{w}_α is the normal vector, b_α is the bias defining a classifying hyperplane, and α refers to either softness S or fluidity \mathcal{F}^T . We note that the term “fluidity” was previously used to describe an average rate of plastic events in models of soft glassy rheology [34]; while our definition is different, we will see that highly “fluid” particles have more active dynamics at high temperatures and, indeed, are more likely to undergo plastic rearrangement events at low temperatures.

III. CLASSIFICATION IN THE FLUID PHASE

We find that we *can* learn to classify extreme diffusive events even far above the onset temperature using local structure. Hyperplanes are characterized by a normal vector and a bias; the direction of the normal corresponds to the linear combination of features that has been learned, and the bias is an offset that best separates the training set given that direction. We expect the *direction* to encode the key physical features governing rearrangements, whereas we expect the *bias* may be strongly dependent on details such as the choice of time window or the temperature of the training set. For instance, with our fixed-threshold definition of a rearrangement the total number of rearranging particles increases as T increases, so even if the same underlying structural variable controls rearrangements, the optimal bias of the hyperplane will shift to maximize the soft margin in the training set data. Because of this, we want to remove the influence of the bias on our later results. In Fig. 1 we show the training accuracy of fluidity as a function of the bias, and during our transfer learning approach later we will select values that maximize our classification accuracy not on the training but on a low-temperature test set.

We believe it is noteworthy that at such high temperatures, *any* structural features predictive of dynamics can be found. We find that even a softness classifier, i.e., a classifier trained on activated dynamics, has some ability to classify diffusive events in the fluid phase: as shown in the inset, the accuracy of the high- T training sets is almost as good as the classifiers trained at those temperatures. The optimal bias that needs




FIG. 1. Fluidity classifies rare events at high and low temperatures. The points show the fivefold cross-validation accuracy of linear SVMs trained on extreme diffusive samples at $T = 1.0, 1.2, 1.4, 1.8, 2.0$ (dark blue to light red) as a function of the classifier’s bias. Each classifier achieves near-peak accuracy for small values of the bias. In contrast, the inset shows the test classification accuracy of the “softness” classifier trained on activated dynamics at $T = 0.45$ and applied to the extremes of diffusive events at different temperatures, for which very different values of the bias optimize performance.

to be chosen is quite different, but the *direction* in feature space learned is quite similar. This finding encourages us to more explicitly frame a transfer learning task from the high-temperature to the low-temperature regime. Concretely, we apply the fluid-phase classifiers, trained at temperatures ranging from $T = 1$ to $T = 2$, to labeled data from $T = 0.47$. As shown in Fig. 2, even though we have trained on data well above T_0 , we find that our classifiers maintain substantial classification accuracy. Again, the optimal bias varies strongly




FIG. 2. A transfer learning approach connects extreme diffusive events above T_0 with activated dynamics below T_0 . The main plot shows the test accuracy of linear SVMs, trained on extreme diffusive samples at $T = 1.0, 1.2, 1.4, 1.8, 2.0$ (dark blue to light red), as applied to a test set of activated dynamics at $T = 0.47$ as a function of the classifier’s bias relative to the optimal bias for that choice of temperature. The inset shows the self part of the Van Hove correlation function for large particles at a timescale of 10 LJ time units. (Black dots are for $T = 0.45$.)




FIG. 3. The probability of rearrangement conditioned on fluidity reveals energy barriers below the onset temperature. (a) The log probability of rearrangement conditioned on $\mathcal{F}^{2.0}$ vs inverse temperature. Point colors correspond to different bins of fluidity, as indicated in the legend. (b) The same features for rearrangements conditioned on S . In all cases, dotted lines are Kramers form fits.

with training and testing temperature, but the *direction* in feature space is extremely highly correlated.

To highlight how surprising this is, in the inset we show the self part of the Van Hove function characterizing single-particle displacements, choosing as a timescale the same window we used for p_{hop} . At high temperatures this distribution is essentially Gaussian and involves a substantial number of particles moving *many* times their own size; at low temperatures this distribution is nontrivial and has an exponential tail corresponding to hopping motions whose size is less than the single-particle diameter.

IV. INTERPRETABILITY OF FLUIDITY AND SOFTNESS

Using $T > T_0$ classifiers, we are able to obtain reasonable accuracy for training sets (which are, by definition, constructed from atypical particles at the various training temperatures), but remarkably, we find that fluidity has the same kind of physical interpretability as softness. We define a rearrangement as a particle having an instantaneous value of $p_{\text{hop}} > p_c$ and fit the probability of rearranging P_R to a Kramers form [35]: $P_R = \frac{1}{T} \exp[\Sigma(\mathcal{F}^T)] \exp[-\Delta E(\mathcal{F}^T)/T]$. Just as for softness, we show in Fig. 3 that fluidity partitions




FIG. 4. Collapse of inferred landscape features from different training temperatures. The energy barrier as a function of the distance to the optimized classifier, ΔE vs x_α , as inferred from Kramers fits to $P_R(S)$ shows little variation across a wide range of classifier training temperatures. The inset showing the entropic contribution similarly collapses in this representation.

the overall system dynamics into a collection of barrier-hopping processes characterized by an energy barrier scale ΔE and an entropic contribution Σ . We also find that our prediction of the onset temperature itself, whether from the intersection of the Kramers form fits or more qualitatively from where the data collapse, is the same across our softness and fluidity classifiers, suggesting that a consistent physical interpretation is being learned.

The identification of a scalar value—fluidity—that encodes the energy barrier characterizing an activated process by training a classifier on diffusive events is striking. Given the crossover nature of the onset temperature, perhaps this qualitative result could have been expected for training temperatures close to T_0 , but it holds even when training far above T_0 , as shown in Fig. 3(a). How do the energy barriers learned by these classifiers compare to the energy barriers learned by classifiers trained on supercooled data, i.e., to those from softness? A direct answer to this question is complicated by two aspects of the training and testing procedure.

The first is that there is no reason to think that the hyperplane *bias* should be held constant when moving from one task to another. This is implicit in the relatively large shifts in bias needed in the inset of Fig. 1 and in the test accuracy for suboptimal choices of bias in Fig. 2. The second issue is related to the fact that we study systems across such a wide temperature range that the distribution of the structural features changes substantially (a similar issue arose in the context of applying classifiers to systems at different *densities* [36]). To account for these, we compare the physical interpretations of the different classifiers by defining $x_\alpha = (\bar{w}_\alpha \cdot \bar{F} - b_\alpha^{\text{best}})/\sigma_\alpha$. That is, we adjust the bias to the optimal value when the classifier is applied to a common ($T = 0.47$) training set and rescale the feature vector by the standard deviation of the distribution of fluidity (or softness) at the training temperature. With this choice, in Fig. 4 we show that the learned aspects of the landscape associated with particle structure, including both the energy barrier and entropic contribution, are almost *identical*. We note that fitting the data only in

TABLE I. Optimal bias b of the hyperplane during training, the standard deviation of fluidity (or softness for $T = 0.45$) at the training temperatures, and the optimal bias b^{best} that maximizes test accuracy at $T = 0.47$. The last column displays the projection of our classifiers onto the softness classifier.

T	b	σ_T	b^{best}	$\hat{w}_S \cdot \hat{w}_\alpha$
0.45	0.0096	1.1188	0.55710	1.0000
1.0	-0.0056	0.7653	-0.2385	0.8246
1.2	-0.0171	0.7028	-0.4644	0.8071
1.4	-0.0099	0.6760	-0.4644	0.6643
1.8	-0.0262	0.6286	-0.5774	0.7104
2.0	-0.0341	0.6286	-0.6151	0.7198

the regime unambiguously below the onset temperature, i.e., the points for which $T < 0.8$, does not qualitatively change these results. We speculate that there may be some correlation between training at higher temperatures and a hint of a slight curvature in the data but do not yet have sufficient data to confirm this.

Given these results, Table I reports several values that contribute to the formation of the results: the bias (i.e., the bias that achieves the highest accuracy during training), the standard deviation of fluidity and softness, and the optimal choice of bias when the classifier is applied to a test set at $T = 0.47$. We also report, as a simple measure of the similarity of the classifiers, the dot product between the normal vector describing each classifier and that of the softness classifier.

V. DISCUSSION

Taken together, our results establish a surprising connection between the structural features that control activated events at low temperatures and those apparently responsible for the tails of the distribution of diffusive events above the onset temperature. Although many approaches have considered the link between local structural arrangements and dynamical arrest in the supercooled regime [37], much of this knowledge is set aside when studying the liquid phase. Our finding that structure is relevant even above the temperature of the liquid-gas critical point (roughly $T = 1.2$ in this model [38]) suggests that further pursuing this avenue of research may prove fruitful. The connections between structure and dynamics across temperatures that we find may be a consequence of the only modestly growing structural length scales over the temperature range studied, but we again emphasize that the qualitative character of the dynamics changes significantly over these same temperatures.

A natural hypothesis might be that our classification accuracy stems from an ability to identify fluid-phase particles that do not diffuse very much: perhaps we are identifying rare particles that consistently sample a similar, high-barrier part of the energy landscape and do not truly distinguishing both immobile and highly mobile particles. We show in the Appendix that this hypothesis fails and that using both tails of the diffusive-motion distribution is crucial to our results. We comment that our main finding—that one can take a classifier built on fluid-phase data without barrier-hopping dynamics, apply it to data in a dynamically heterogeneous phase, and

infer the existence of energy barriers there—is reminiscent of the results reported in Ref. [26]. That work considered a biologically inspired model with highly unusual glassy dynamics [39–41] meant to mimic the behavior of dense cellular materials. There, it was speculated that it was the anomalous sub-Arrhenius behavior of the model that was responsible for the success of the transfer learning task; the results presented here suggest an alternative explanation may be needed.

Our work highlights what we believe continue to be crucial unanswered questions: *why* do these machine learning methodologies learn simple structural order parameters that correspond to local energy barriers in disordered phases of matter? What aspects of the training lead to this result? And to what extent can we use this result to uncover new, relevant descriptions for the physics of amorphous solids? We note that there is some indication that the specific methodology used here and earlier—linear SVMs—may not be crucial to recover this physical interpretation; Ref. [23] hinted at a similar result using a Graph Neural Network-inspired linear-regression-based model. We believe it will be crucial to compare different machine learning techniques as applied to predicting glassy dynamics [21,42] not only along dimensions of predictive capacity, generalizability, efficiency, and training cost but *also* in terms of their physical interpretability.

Data used in this paper are available from Zenodo [43].

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. DMR-2143815. We thank I. Nemenman, S. Ridout, A. Swain, A. Liu, and E. Weeks for stimulating discussions.

APPENDIX

1. Fluidity distribution

Similar to the distribution of softness shown in Ref. [19], the distribution of fluidity, when measured at different test temperatures, remains approximately Gaussian. This is shown in Fig. 5 for two training temperatures ($T = 1.2$ and $T = 2.0$). The mean of the distribution behaves monotonically as the test temperature changes.

2. Determination of the rearrangement threshold for fluidity

The dependence of softness on various choices for training set thresholds was explored in Ref. [19], and our choices are consistent with widely employed values. In the case of fluidity, as noted above, we selected threshold values to have comparably sized training sets given the length of our simulations (i.e., to generate of the order of 10^4 elements of a balanced training set). We first emphasize that the qualitative outcome of our analysis is not crucially dependent on the choice of threshold. In Figs. 6 and 7 we show (at a fixed representative temperature above the onset temperature, $T = 1.2$) that our classification accuracy for the test set is only weakly dependent on the precise choice of p_l and p_u : sensibly, more extreme events are slightly easier to classify, but we are far enough into the tail of rare events that the effect is a small one.




FIG. 5. The distribution of fluidity at different test temperatures for two considered training temperatures. Training at both (a) $T = 1.2$ and (b) $T = 2.0$, the distribution of fluidity is approximately Gaussian. The mean is a monotonic function of the test temperature.

We note, however, that our ability to classify and to successfully transfer learn *is* dependent on the selection of appropriate structural examples that correspond to both extreme




FIG. 6. Classification accuracy above the onset temperature as a function of the misclassification penalty C using different thresholds for defining extreme events. Three different choices for the fraction of the distribution of events to use in the training set give broadly consistent accuracies for the test set, with no systematic changes in the optimal hyperparameter.

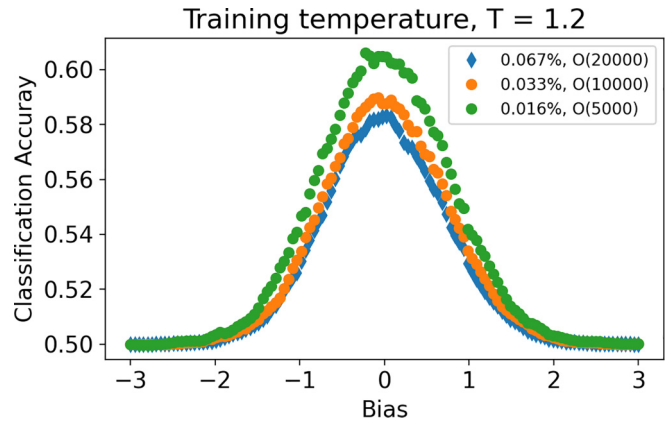


FIG. 7. Classification accuracy above the onset temperature as a function of bias for a fixed training hyperparameter. Using a fixed misclassification hyperparameter of $C = 10^{-2}$, we again see that the accuracy is only weakly dependent on the chosen event thresholds.

ends of the p_{hop} distribution. As discussed in the main text, one natural hypothesis is that local structures corresponding to the lower tail of events capture sufficiently deep minima that influence even particles in the fluid phase and that our ability to classify stems entirely from these rare, relatively immobile particles. To test this, we construct a classifier using local structures from the lower tail of the p_{hop} distribution as our nonrearranging particles and construct the other class from *randomly selected* particles with $p_{\text{hop}} > p_l$. We train such a classifier at $T = 1.2$ and evaluate it on the test set of $T = 0.47$ (as in Fig. 2 in the main text). We show in Fig. 8 that the accuracy of this “immobile vs randomly selected” classifier is substantially lower than the accuracy of the “immobile vs highly mobile” fluidity classifier.

3. Probability of rearrangements as a function of x_α

In this section, we show data for the probability of rearrangement as a function of x_α against $1/T$ for different training temperatures; these data lead directly to the results




FIG. 8. Test accuracy against bias of a low mobility vs random particle classifier. The plot shows that the test accuracy of the classifier (blue lower points) is lower than the classifier built using both extreme ends of the p_{hop} distribution (orange upper points).




FIG. 9. The probability of rearrangement conditioned on x_α against $1/T$. Dotted lines are fits to the Kramers form for different bins of x_α .

reported in Fig. 4 in the main text. The probability of rearrangement conditioned on x_α is the fraction of local structure with value x_α that has a corresponding p_{hop} value greater than p_c , where p_c is the rearrangement cutoff used at low T .

Figure 9, mirroring Fig. 3, shows the probability of rearrangement as a function of $1/T$ for the softness and the fluidity classifiers. In Fig. 4, we report the inferred energy barriers and “entropic barriers” associated with each of these fits.

-
- [1] L. Berthier, Dynamic heterogeneity in amorphous materials, *Physics* **4**, 42 (2011).
- [2] C. A. Angell, Formation of glasses from liquids and biopolymers, *Science* **267**, 1924 (1995).
- [3] S. Sastry, P. G. Debenedetti, and F. H. Stillinger, Signatures of distinct dynamical regimes in the energy landscape of a glass-forming liquid, *Nature (London)* **393**, 554 (1998).
- [4] P. G. Debenedetti and F. H. Stillinger, Supercooled liquids and the glass transition, *Nature (London)* **410**, 259 (2001).
- [5] A. Banerjee, M. K. Nandi, S. Sastry, and S. M. Bhattacharyya, Determination of onset temperature from the entropy for fragile to strong liquids, *J. Chem. Phys.* **147**, 024504 (2017).
- [6] G. Folena, S. Franz, and F. Ricci-Tersenghi, Rethinking Mean-Field Glassy Dynamics and Its Relation with the Energy Landscape: The Surprising Case of the Spherical Mixed p -Spin Model, *Phys. Rev. X* **10**, 031045 (2020).
- [7] A. S. Keys, L. O. Hedges, J. P. Garrahan, S. C. Glotzer, and D. Chandler, Excitations Are Localized and Relaxation Is Hierarchical in Glass-Forming Liquids, *Phys. Rev. X* **1**, 021013 (2011).
- [8] M. D. Ediger, Spatially heterogeneous dynamics in supercooled liquids, *Annu. Rev. Phys. Chem.* **51**, 99 (2000).
- [9] L. Berthier, G. Biroli, J.-P. Bouchaud, L. Cipelletti, and W. V. Saarloos, *Dynamical Heterogeneities in Glasses, Colloids, and Granular Media*, International Series of Monographs on Physics, Vol. 150 (Oxford University Press, Oxford, 2011).
- [10] M. D. Ediger, C. A. Angell, and S. R. Nagel, Supercooled liquids and glasses, *J. Phys. Chem.* **100**, 13200 (1996).

- [11] U. R. Pedersen, T. B. Schröder, and J. C. Dyre, Phase Diagram of Kob-Andersen-Type Binary Lennard-Jones Mixtures, *Phys. Rev. Lett.* **120**, 165501 (2018).
- [12] L. Berthier and R. L. Jack, Structure and dynamics of glass formers: Predictability at large length scales, *Phys. Rev. E* **76**, 041509 (2007).
- [13] T. Kawasaki, T. Araki, and H. Tanaka, Correlation between Dynamic Heterogeneity and Medium-Range Order in Two-Dimensional Glass-Forming Liquids, *Phys. Rev. Lett.* **99**, 215701 (2007).
- [14] C. Patrick Royall, S. R. Williams, T. Ohtsuka, and H. Tanaka, Direct observation of a local structural mechanism for dynamic arrest, *Nat. Mater.* **7**, 556 (2008).
- [15] R. L. Jack, A. J. Dunleavy, and C. P. Royall, Information-Theoretic Measurements of Coupling between Structure and Dynamics in Glass Formers, *Phys. Rev. Lett.* **113**, 095703 (2014).
- [16] M. L. Manning and A. J. Liu, Vibrational Modes Identify Soft Spots in a Sheared Disordered Packing, *Phys. Rev. Lett.* **107**, 108302 (2011).
- [17] A. Widmer-Cooper, H. Perry, P. Harrowell, and D. R. Reichman, Irreversible reorganization in a supercooled liquid originates from localized soft modes, *Nat. Phys.* **4**, 711 (2008).
- [18] E. D. Cubuk, S. S. Schoenholz, J. M. Rieser, B. D. Malone, J. Rottler, D. J. Durian, E. Kaxiras, and A. J. Liu, Identifying Structural Flow Defects in Disordered Solids Using Machine-Learning Methods, *Phys. Rev. Lett.* **114**, 108001 (2015).
- [19] S. S. Schoenholz, E. D. Cubuk, D. M. Sussman, E. Kaxiras, and A. J. Liu, A structural approach to relaxation in glassy liquids, *Nat. Phys.* **12**, 469 (2016).
- [20] D. M. Sussman, S. S. Schoenholz, E. D. Cubuk, and A. J. Liu, Disconnecting structure and dynamics in glassy thin films, *Proc. Natl. Acad. Sci. USA* **114**, 10601 (2017).
- [21] V. Bapst, T. Keck, A. Grabska-Barwińska, C. Donner, E. D. Cubuk, S. S. Schoenholz, A. Obika, A. W. R. Nelson, T. Back, D. Hassabis, and P. Kohli, Unveiling the predictive power of static structure in glassy systems, *Nat. Phys.* **16**, 448 (2020).
- [22] E. Boattini, S. Marín-Aguilar, S. Mitra, G. Foffi, F. Smallenburg, and L. Filion, Autonomously revealing hidden local structures in supercooled liquids, *Nat. Commun.* **11**, 5479 (2020).
- [23] E. Boattini, F. Smallenburg, and L. Filion, Averaging Local Structure to Predict the Dynamic Propensity in Supercooled Liquids, *Phys. Rev. Lett.* **127**, 088007 (2021).
- [24] E. D. Cubuk, A. J. Liu, E. Kaxiras, and S. S. Schoenholz, Unifying framework for strong and fragile liquids via machine learning: A study of liquid silica, [arXiv:2008.09681](https://arxiv.org/abs/2008.09681).
- [25] G. Jung, G. Biroli, and L. Berthier, Predicting Dynamic Heterogeneity in Glass-Forming Liquids by Physics-Inspired Machine Learning, *Phys. Rev. Lett.* **130**, 238202 (2023).
- [26] I. Tah, T. A. Sharp, A. J. Liu, and D. M. Sussman, Quantifying the link between local structure and cellular rearrangements using information in models of biological tissues, *Soft Matter* **17**, 10242 (2021).
- [27] W. Kob and H. C. Andersen, Testing mode-coupling theory for a supercooled binary Lennard-Jones mixture: The van Hove correlation function, *Phys. Rev. E* **51**, 4626 (1995).
- [28] D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications* (Elsevier, Amsterdam, 2023).
- [29] J. A. Anderson, J. Glaser, and S. C. Glotzer, HOOMD-blue: A Python package for high-performance molecular dynamics and hard particle Monte Carlo simulations, *Comput. Mater. Sci.* **173**, 109363 (2020).
- [30] J. Behler and M. Parrinello, Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces, *Phys. Rev. Lett.* **98**, 146401 (2007).
- [31] C. W. Hsu, C. C. Chang, and C. J. Lin, A practical guide to support vector classification, 2003.
- [32] R. Candelier, O. Dauchot, and G. Biroli, Building Blocks of Dynamical Heterogeneities in Dense Granular Media, *Phys. Rev. Lett.* **102**, 088001 (2009).
- [33] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay, Scikit-learn: Machine learning in python, *J. Mach. Learn. Res.* **12**, 2825 (2011).
- [34] L. Bocquet, A. Colin, and A. Ajdari, Kinetic Theory of Plastic Flow in Soft Glassy Materials, *Phys. Rev. Lett.* **103**, 036001 (2009).
- [35] P. Hänggi, P. Talkner, and M. Borkovec, Reaction-rate theory: Fifty years after Kramers, *Rev. Mod. Phys.* **62**, 251 (1990).
- [36] I. Tah, S. A. Ridout, and A. J. Liu, Fragility in glassy liquids: A structural approach based on machine learning, *J. Chem. Phys.* **157**, 124501 (2022).
- [37] C. P. Royall and S. R. Williams, The role of local structure in dynamical arrest, *Phys. Rep.* **560**, 1 (2015).
- [38] V. Testard, L. Berthier, and W. Kob, Intermittent dynamics and logarithmic domain growth during the spinodal decomposition of a glass-forming liquid, *J. Chem. Phys.* **140**, 164502 (2014).
- [39] D. M. Sussman, M. Paoluzzi, M. C. Marchetti, and M. L. Manning, Anomalous glassy dynamics in simple models of dense biological tissue, *Europhys. Lett.* **121**, 36001 (2018).
- [40] Y.-W. Li, L. L. Y. Wei, M. Paoluzzi, and M. P. Ciamarra, Softness, anomalous dynamics, and fractal-like energy landscape in model cell tissues, *Phys. Rev. E* **103**, 022607 (2021).
- [41] D. E. P. Pinto, D. M. Sussman, M. M. T. da Gama, and N. A. M. Araújo, Hierarchical structure of the energy landscape in the Voronoi model of dense tissue, *Phys. Rev. Res.* **4**, 023187 (2022).
- [42] R. M. Alkemade, E. Boattini, L. Filion, and F. Smallenburg, Comparing machine learning techniques for predicting glassy dynamics, *J. Chem. Phys.* **156**, 204503 (2022).
- [43] T. M. Obadiya and D. M. Sussman, Dataset for "using fluid structures to encode predictions of glassy dynamics", Zenodo (2023), doi: [10.5281/zenodo.7469766](https://doi.org/10.5281/zenodo.7469766).