



HAL
open science

Finding defects in glasses through machine learning

Simone Ciarella, Dmytro Khomenko, Ludovic Berthier, Felix Mocanu, David Reichman, Camille Scalliet, Francesco Zamponi

► **To cite this version:**

Simone Ciarella, Dmytro Khomenko, Ludovic Berthier, Felix Mocanu, David Reichman, et al.. Finding defects in glasses through machine learning. *Nature Communications*, 2023, 14 (1), pp.4229. 10.1038/s41467-023-39948-7. hal-04165248

HAL Id: hal-04165248

<https://cnrs.hal.science/hal-04165248>

Submitted on 15 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Finding defects in glasses through machine learning

Simone Ciarella,^{1,*} Dmytro Khomenko,^{2,3,*} Ludovic Berthier,^{4,5} Felix C. Mocanu,¹ David R. Reichman,² Camille Scalliet,⁶ and Francesco Zamponi¹

¹*Laboratoire de Physique de l'École Normale Supérieure, ENS, Université PSL, CNRS, Sorbonne Université, Université de Paris, 75005 Paris, France*

²*Department of Chemistry, Columbia University, 3000 Broadway, New York, NY 10027, USA*

³*Dipartimento di Fisica, Sapienza Università di Roma, P.le A. Moro 2, I-00185, Rome, Italy*

⁴*Yusuf Hamied Department of Chemistry, University of Cambridge, Lensfield Road, Cambridge CB2 1EW, United Kingdom*

⁵*Laboratoire Charles Coulomb (L2C), Université de Montpellier, CNRS, 34095 Montpellier, France*

⁶*DAMTP, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge CB3 0WA, United Kingdom*

(Dated: May 17, 2023)

Structural defects control the kinetic, thermodynamic and mechanical properties of glasses. For instance, rare quantum tunneling two-level systems (TLS) govern the physics of glasses at very low temperature. Due to their extremely low density, it is very hard to directly identify them in computer simulations. We introduce a machine learning approach to efficiently explore the potential energy landscape of glass models and identify desired classes of defects. We focus in particular on TLS and we design an algorithm that is able to rapidly predict the quantum splitting between any two amorphous configurations produced by classical simulations. This in turn allows us to shift the computational effort towards the collection and identification of a larger number of TLS, rather than the useless characterization of non-tunneling defects which are much more abundant. Finally, we interpret our machine learning model to understand how TLS are identified and characterized, thus giving direct physical insight into their microscopic nature.

I. INTRODUCTION

When a glass-forming liquid is cooled rapidly, its viscosity increases dramatically and it eventually transforms into an amorphous solid, called a glass, whose physical properties are profoundly different from those of ordered crystalline solids [1]. At even lower temperature, around 1 K, the specific heat of a disordered solid is much larger than that of its crystalline counterpart as it scales linearly rather than cubically with temperature. Similarly, the temperature evolution of the thermal conductivity in glasses is quadratic, rather than cubic [2–11]. A theoretical framework rationalizing such anomalous behavior was provided by Anderson, Halperin and Varma [12] and by Phillips [13, 14]. They argued that the energy landscape of amorphous solids contains many nearly-degenerate minima, connected by the localized motion of a few atoms, that can act as tunneling defects, called two-level systems (TLS) [15–19]. Since then, localized structural defects have been understood to play a crucial role in many other glass properties [20]. Understanding the microscopic origin of such localized defects and how to control their density and physical properties is a major goal not only to improve our fundamental understanding of amorphous solids, but also for technological applications, such as optimizing the performance of certain quantum devices [21, 22].

The development of particle-swap computer algorithms [23, 24] has allowed the creation of computer glasses at unprecedentedly low temperatures. Combined with potential energy landscape exploration algorithms [25–34], this provides a powerful method to investigate the nature of defects in materials prepared under conditions comparable to experimental studies [35, 36]. These tools have enabled direct numerical observation of TLS [35, 37], confirming the experimental result [7, 8, 10, 11, 38] that the density of tunneling defects is strongly depleted as the kinetic stability of a glass increases. Similar results have been obtained for a different kind of defect, namely soft vibrational modes [39]. The direct detection of TLS revealed some of their microscopic features, namely that fewer atoms participate in the TLS of stable glasses [35]. It was also shown that TLS are not in a one-to-one correspondence with soft harmonic [35, 36, 40] or localized [36, 41] modes.

The main limitation of the direct landscape exploration method is its large computational cost, making it hard to construct the large library of defects needed for a robust statistical analysis of their physical properties. After accumulating a large number of inherent structures (IS), one must run a computationally expensive algorithm to find the minimum energy path connecting pairs of IS in order to determine if the pair forms a proper defect (e.g., to form a TLS, a defect must have a quantum energy splitting within thermal energy at 1 K). The very large number of IS pairs detected makes it impossible to characterize all of them. In previous works, some *ad hoc* filtering rules were introduced in order to identify candidate TLS and focus computational effort on them [35–37]

* These authors contributed equally.
email: simone.ciarella@ens.fr, dmytro.khomenko@uniroma1.it




Figure 1. **Numerical search for two-level systems.** (a) Exploring the potential energy landscape: different glass samples define different metabasins in the rough landscape. (Inset) Each glass metabasin is explored via molecular dynamics simulations (black arrow) during which frequent energy minimization (dashed arrows) generates a large number of inherent structures (IS). Previous works restricted the search for candidate defects only to pairs of IS explored consecutively in the dynamics. (b) Our machine-learning approach instead considers all IS pairs, irrespective of the dynamical exploration, and rapidly provides a robust prediction for their properties, such as the quantum splitting. The candidates selected by the ML model are then analyzed via a minimum energy path-finding protocol (NEB algorithm) and their properties are computed exactly and compared with the ML prediction.

but the success rate of such filters is poor. A considerable computational effort, which consisted in sampling $\sim 10^8$ IS, resulted in the direct detection of about 60 TLS. It is then obvious that most of the computational effort has been wasted in the study of pairs that form defects which do not tunnel at low temperatures. Looking for TLS is akin to finding the proverbial needle in a haystack.

In this paper we demonstrate the relevance of machine learning techniques to predict with enhanced accuracy whether a pair of inherent structures forms a defect of a certain type. Recently, machine learning (ML) was shown to be extremely effective in using structural indicators to predict structural, dynamical or mechanical properties of glassy solids [20, 42–50]. Here, we use supervised learning to streamline the identification of defects. We focus in particular on the classical energy barrier and the quantum splitting associated with defects, which are relevant to identify TLS. Our study has two goals: (i) develop a faster way to identify TLS compared to the standard approach described below [35, 37] in order to collect a statistically significant number of tunneling defects; (ii) determine the structural and dynamical features characterizing TLS as well as their evolution with glass preparation. To address (i) we show that our machine learning model can be trained in a few minutes using a small quantity of data, after which the model is able to identify candidate TLS with high speed and accuracy. To address (ii) we determine which static features are the most important for the model prediction. We show that TLS are not necessarily pairs of IS explored consecutively in the dynamics. We conclude by explaining how the ML model distinguishes TLS from non-TLS and how it is able to identify glasses prepared at different temperatures. While here we mostly focus on TLS, which are the rarest defects in glasses, our method should easily apply to other problems, such as supercooled liquid

dynamics, plasticity or devitrification of glassy solids.

II. RESULTS

In the following, we focus for concreteness on the problem of detecting rare tunneling TLS. We also demonstrate that our method can successfully predict the classical energy barrier between two energy minima, with applications to the efficient detection of other kinds of defects.

A. Machine learning approach

The standard procedure [32, 33, 35] to identify TLS is sketched in Fig. 1(a). The following steps aim at identifying potential candidates for TLS (see the Methods section for details):

1. Equilibrate the system at the preparation temperature T_f . Glasses with lower T_f have an increased glass stability.
2. Run molecular dynamics to sample configurations along a dynamical trajectory at the exploration temperature $T < T_f$.
3. Perform energy minimization from the sampled configurations to produce a time series of energy minima, or inherent structures (IS).
4. Analyze the number of transitions recorded between pairs of IS in the dynamical exploration of step 2, and select the pairs of IS that are explored consecutively.

Step 4 was necessary because it is computationally impossible to analyze all pairs of IS, as the number of pairs scales quadratically with the number of minima. The filter defined in step 4 was physically motivated by the fact that TLS tend to originate from IS that are not too distant in order to have a reasonable tunneling probability. As such it is likely that those pairs of IS get explored one after the other during the exploration dynamics in step 2. Overall, given N_{IS} inherent structures, this procedure selects for $\mathcal{O}(N_{IS})$ pairs to be analyzed.

Once potential candidates are selected, the procedure continues as follows:

5. For each selected pair of IS, look for the minimum energy path and the classical barrier between them by running a minimum energy path finding algorithm, such as the nudge elastic band (NEB) algorithm [51–53]. This provides the value of the potential energy $V(\xi)$ along the minimum energy path between the pair, where $0 \leq \xi \leq 1$ is the reaction coordinate.
6. Select pairs whose energy profile $V(\xi)$ has the form of a double well (DW), i.e. exclude paths with multiple wells.
7. Solve the one-dimensional Schrödinger equation

$$-\frac{\hbar^2}{2md^2\epsilon}\partial_\xi^2\Psi(\xi) + V(\xi)\Psi(\xi) = \mathcal{E}\Psi(\xi), \quad (1)$$

where ξ is a normalized distance along the reaction path $\xi = x/d$ and energy is normalized by a Lennard-Jones energy scale ϵ , the effective mass m and the distance d are calculated as in Ref. [35]. We obtain the quantum splitting (QS) $E_{qs} = \mathcal{E}_2 - \mathcal{E}_1$ from the first two energy levels \mathcal{E}_1 and \mathcal{E}_2 .

The quantum splitting is the most relevant parameter for TLS because when $E_{qs} \sim T$ the system can transition from one state to the other via quantum tunneling [13]. In particular, since we choose to report the data in units that correspond to Argon [35], a double well excitation will be an active TLS at $T=1$ K when $E_{qs} < 0.0015\epsilon$, where ϵ sets the energy scale of the pair interactions in the simulated model.

Overall, since at low temperature the landscape exploration dynamics is slow, one would like to spend most of the computational time on steps 2-3 to construct a large library of pairs of IS. A first problem is that when the library of IS grows larger it takes a lot of time to perform steps 5-7. Moreover, the main bottleneck lies in the fact that most of the pairs that go through the full procedure turn out not to be TLS. The large computational time dedicated to steps 5-7 is thus wasted. Furthermore, many pairs of IS can be close but not sampled consecutively during the dynamics, owing to the high-dimensional nature of the potential energy landscape.

We now introduce our machine learning (ML) approach to the problem, whose main advantage is to consider all pairs of IS as TLS candidates. As shown below,

our approach can detect TLS which are otherwise excluded in step 4. We distinguish two phases: training and deployment. Our supervised training approach, detailed in the Methods section and sketched in Fig. 2, takes just a few hours of training on a single CPU. It requires an initial dataset of $\mathcal{O}(10^4)$ full NEB calculations, whose collection is the most time-consuming part of the training phase. Once training is complete, the ML model can be deployed to identify new TLS. Its workflow is similar to the standard one, with some major improvements. It proceeds with the following steps:

1. - 3. The first 3 steps are similar to the standard procedure to obtain a collection of inherent structures from a dynamical exploration.
4. Apply the ML model to all possible pairs of IS to predict which pairs form a DW potential.
5. Apply the ML model to predict the quantum splitting (QS) for all predicted DW and filter out the configurations that are not predicted to be TLS by the ML model.
6. - 8. For the pairs predicted to be TLS by the ML model only, run the NEB algorithm and select the pairs that form a DW potential. Solve the one-dimensional Schrödinger equation in order to obtain the exact value of the quantum splitting.

In the Methods section we provide details on how steps 1-3 are performed: glass preparation, exploration of the potential energy landscape via molecular dynamics simulations and minimization procedure, as well as NEB computation. We also explain how it is possible to use steps 4-5 as a single shot or as an iterative training approach, see Methods Sec. IV H.

Importantly, the well-trained ML model has two significant advantages over the standard approach. First, $\mathcal{O}(N_{IS}^2)$ pairs of IS are scanned to identify TLS, compared to a much smaller number $\mathcal{O}(N_{IS})$ in the standard procedure. Second, if a pair of IS passes step 5 and goes through the full procedure it is very likely to be a real TLS. As a consequence, by using the ML approach one can spend more time doing steps 2-3 to produce new IS, since fewer pairs pass step 5. At the same time, for any given number of IS, the ML approach can analyze all possible pairs and is therefore able to identify many more TLS, as we demonstrate below.

B. Quality of the machine learning prediction

In Refs. [35, 36], the authors analyze a library of 14202, 23535 and 117370 pairs of inherent structures for a continuously polydisperse system of soft repulsive particles, equilibrated at reduced temperatures $T_f = 0.062$, 0.07, and 0.092, respectively. The standard approach described in Sec. II A leads to the identification of 61, 291 and 1008 TLS for the three temperatures, respectively.

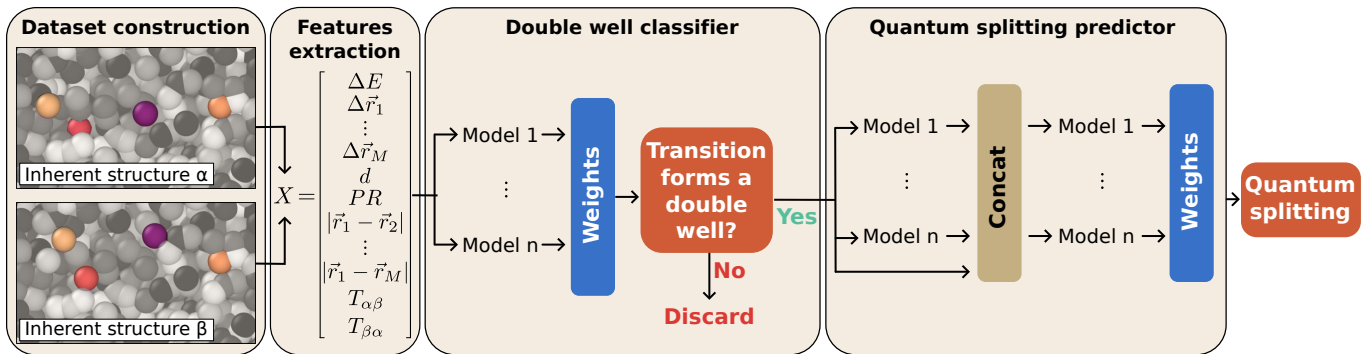


Figure 2. **Flowchart of the machine learning model.** The dataset is constructed by comparing all the pairs of inherent structures (IS), focusing on the M atoms that displace the most between two IS (highlighted by colors: bright, resp. dark, indicate small, resp. large particle radius). Specific features are extracted to construct the input vector X . We then train a classifier to predict whether a pair of inherent structures forms a double well potential (DW) or not. The DW are finally processed using a multi-layer stacking strategy to predict the quantum splitting of the double well potential. Our pipeline analyses a given pair of IS in about $\sim 10^{-4}$ s.

Note that this approach uses pairs of IS that are selected by the dynamical information contained in the transition matrix between pairs of IS [35]. This was done to filter out all non-DW potentials. For all pairs in this small subset, the quantum splitting was then calculated.

Instead, the ML approach starts by independently evaluating the relevant information contained in each IS and constructs all possible combinations, even for pairs that are not dynamically connected in the landscape exploration. Following the steps discussed in Sec. II A the model is then able to predict which of all pairs form a DW, as well as the value of their quantum splitting, very accurately. From a quantitative perspective, this means that the same dynamical trajectories now contain many more TLS candidates in the ML approach compared to the standard approach.

In this section we briefly describe the flowchart of the model summarized in Fig. 2. A detailed description of the machine learning model is provided in the Methods section. First, for all the available IS, we evaluate a set of static quantities that we use to construct the input features for each pair of IS. By convention, we label the IS with $\alpha = 1, \dots, N_{IS}$ by increasing potential energy $E_1 < \dots < E_{N_{IS}}$ where E_α is the potential energy of IS_α . We use the convention that $\alpha < \beta$. The input feature for a pair $\alpha\beta$ of IS consists in the potential energy difference $\Delta E = E_\beta - E_\alpha$ between the two minima, the displacements $\Delta \vec{r}_i$ of the M particle which displace the most between the two IS (labeled with $i = 1, \dots, M$ by decreasing displacement), as well as the relative positions between those M particles, the total displacement of particles d , and participation ratio PR , all computed by comparing the two IS. We also use as input feature the number of transitions recorded in the dynamical exploration $T_{\alpha\beta}$ (from the lowest energy IS to the highest, in our convention) and $T_{\beta\alpha}$ (from highest to lowest IS). See Methods Sec. IV E for more details. We then apply in series two *model ensembles*. Model ensembling consists

in averaging the output of different ML models to achieve better predictions compared to each of them separately. The first ensemble is trained to classify DW (Methods, Sec. IV F), which is a necessary condition for TLS. A DW is defined when the minimum energy path between the two IS resembles a quartic potential, as sketched in Fig. 1(b). For the pairs that are predicted to be DW, a second model ensemble (Methods, Sec. IV G) is used to predict the quantum splitting (QS), which determines if the pair is a TLS or not.

To showcase the performance of the ML model we report in Fig. 3(a)-(c) the exact QS calculated from the NEB procedure, compared with the value predicted by the model. We see that the data concentrates around the diagonal, indicating good correlation between true and predicted values. The Pearson correlation reported in the figure provides a quantitative measure for the correlation. The train/test split is performed by randomly selecting 10% of the pairs to be used only for the evaluation. We have trained three independent models to work at the three different glass preparation temperatures. As explained in the Methods (Sec. IV E), the model needs the information about the $M \ll N$ particles that displace the most between the two IS only to achieve the excellent precision demonstrated in Fig. 3. In the SI (Fig. S2) we show that the optimal value is $M = 3$, i.e. information on only three particles is needed for the model to identify TLS, confirming the low participation ratio in TLS. Furthermore, the models have been trained using the smallest number of samples, randomly selected from all the IS pairs available, that allow the model to reach its top performance. We have also performed an analysis of the optimal training time. Details on these points are provided in the SI (Sec. S1). The performances presented in Fig. 3 are achieved by training the model for ~ 10 hours of single CPU time, but we also show in the SI that it is possible to already achieve $> 90\%$ of this performance by training the ensemble for only 10 minutes.




Figure 3. **Machine learning prediction for the quantum splitting and classical energy barrier between pairs of inherent structures (IS).** (a)-(c) Quantum splitting and (d)-(f) energy barrier predicted by the ML model compared to the exact value. The model was not trained on these IS pairs. Glass stability decreases from left to right: glasses are equilibrated at (a, d) $T_f = 0.062$, (b, e) $T_f = 0.07$, and (c, f) $T_f = 0.092$. The ML model was trained on 7000, 10000 and 30000 samples respectively, using information on the $M = 3$ particles with largest displacements. All models were trained for ~ 10 hours of single CPU time.

The ML approach introduced here is also easily generalizable to target other quantities related to state-to-state transition, such as excitations and higher energy effects. We modified the quantum splitting predictor to instead predict the classical energy barrier between two IS states. If the minimal energy path between two IS forms a DW, we define the classical energy barrier as the energy difference between the saddle point and the lowest energy minimum. In Fig. 3(d)-(f) we report the value of the energy barrier predicted by the ML model compared to the exact value calculated from the NEB procedure. We use the same hyperparameters and features as for the quantum splitting predictor. Such a high performance demonstrates that our ML approach can predict other types of transitions between states, associated to distinct kinds of defects.

C. Capturing elusive TLS with machine learning

We now use ML in order to speed up the TLS search. This highly efficient method allows us to collect a library of TLS of unprecedented size, generated from numerical simulations with the same interaction potential as in Ref. [35], see Methods for its definition. First of all, we reprocess the data produced to obtain the results presented in Ref. [35] with our new ML method based on iterative training (Methods, Sec. IV H), obtaining new information about the connection between TLS and dy-

namics. Next, we perform ML-guided exploration to collect as many TLS as possible. This sizable library of TLS allows us to perform for the first time a detailed statistical analysis of TLS and compare their distribution to the distribution of double wells. We perform this analysis for glasses of three different stabilities. Finally, we discuss the microscopic features of TLS not only by looking at their statistics, but also by analyzing what the ML model has learned, and how it expresses its predictions.

Prior to this paper, it was not possible to evaluate all the IS pairs collected in Ref. [35]. For this reason, the authors introduced a filtering rule based on the assumption that high transition rates during the dynamic landscape exploration is a good indicator that the minimum transition path between two IS forms a double well. Therefore, Ref. [35] discarded all pairs $\alpha\beta$ of IS such that the number of jumps $T_{\alpha\beta}$ (from low to high energy IS) and $T_{\beta\alpha}$ (high to low) during the MD exploration is smaller than four, i.e. $\min(T_{\alpha\beta}, T_{\beta\alpha}) < 4$. This reduced the number of pairs to 14202, 21109 and 117339 for glasses prepared at $T_f = 0.062$, 0.07 and 0.092 respectively. In order to have comparable data at the three temperatures, for $T_f = 0.092$ we only consider a subset of glasses corresponding to 30920 IS pairs.

The results of the TLS search are summarized in the red columns of Tab. I. Overall, the standard procedure reaches a rate of TLS found per NEB calculation of $4 \cdot 10^{-3}$, $13 \cdot 10^{-3}$, and $8 \cdot 10^{-3}$ with increasing T_f . We compare these numbers with those obtained with our it-




Figure 4. **The machine-learning-driven exploration identifies an unprecedentedly large number of two-level systems.** (a) We compare the model predictions to the calculated values at the end of our iterative training, for stable glasses $T_f = 0.062$. The confusion matrix is color coded in the background and reported on the bottom right of panel (a). Horizontal dashed lines report the percentage of TLS that are predicted below that value of quantum splitting, showing that more than 95% of the TLS are within twice the TLS threshold of 0.0015. (b) Cumulative distribution of quantum splitting $n(E_{qs})$ divided by E_{qs} . (c) Histograms of the number of TLS and DW per glass, at the three preparation temperatures $T_f = 0.092$, 0.07, and 0.062 from top to bottom. We have considered 5, 30, and 237 glasses, respectively. Results are reported for Ar units.

erative training procedure applied to the same data, see green columns of Tab. I. We immediately notice two major improvements. First, the overall number of TLS that we find from the same data set is more than twice larger. Second, the ratio of TLS per NEB calculation is more than 15 times larger, corresponding to $62 \cdot 10^{-3}$, $211 \cdot 10^{-3}$, and $194 \cdot 10^{-3}$ with increasing T_f .

We conclude that the iterative ML approach is much more efficient than the standard procedure, and also that TLS do not necessarily have a large dynamical transition rate, since the dynamical-filtering approach discards more than half of them.

T_f	Standard procedure (Ref. [35])		Iterative training (this paper)	
	# TLS	# NEBs	# TLS	# NEBs*
0.062 (Ar)	61	14202	156	2500
0.062 (NiP)	28	14202	59	2000
0.07 (Ar)	291	21109	1057	5000
0.07 (NiP)	46	21109	152	4000
0.092 (Ar)	245	30920	776	4000
0.092 (NiP)	28	30920	123	6000

* This number does **not** include the data that we use to pre-train the model.

Table I. **Comparison of the standard procedure with our iterative training approach.** Analysis of data collected in Ref. [35]. We report results for different glass stabilities, decreasing from top to bottom, using Argon units (Ar) [30] and NiP metallic glass parameters (NiP) [27]. The standard procedure finds less than half of the TLS found with the ML and is computationally much more expensive.

D. Differences between DW and TLS

With our ML-driven exploration of the energy landscape we can focus the numerical effort on DW and favorable TLS candidates, while processing a larger number and/or longer exploration trajectories. This allows us to consider a larger set of independent glasses of the same type as those treated in Ref. [35], which is particularly relevant for ultrastable glasses generated at the lowest temperature $T_f = 0.062$. While in Ref. [35] the collection of 61 TLS required more than 14000 NEB calculations, we are able to identify 872 TLS running 11 iterations of iterative training using only a total of 5500 NEB calculations in addition to the ~ 6000 used for pretraining. In the next section we analyze these results to discuss the nature of TLS.

The database of glasses that we analyze with iterative training contains 5 times more IS than in Ref. [35], and we find up to 15 times more TLS by running around half of the NEB calculations. We report in Fig. 4 the results from this extended TLS search. In Fig. 4(a) we report predicted and true values for the quantum splitting, with a background color coding for the confusion matrix. The threshold is set by the fact that TLS have $E_{qs} < 0.0015$. The number of data points in each quadrant is reported in the inset. The horizontal dashed lines highlight the percentage of true TLS detected by running the NEB algorithm for all points with a predicted quantum splitting below the line. Due to false negatives, it is better to also consider transitions slightly above the TLS threshold. We see that all TLS are identified by considering only the pairs that are predicted to be within twice the quantum splitting threshold of TLS. All TLS thus are safely detected by running 2484 NEB calculations, out of 4147 DW in total. In Fig. 4(b) we report the cumu-




Figure 5. **Determining important features to predict quantum splittings and classify double well potentials.** Importance of the different features for the ML models (a) predicting the quantum splitting (QS), and (b) classifying double wells, both at $T_f = 0.062$. The features include the single particle displacements $\Delta\vec{r}_i$ (which decrease with increasing label i), their relative position compared to the most displacing particle $|\vec{r}_1 - \vec{r}_i|$ and the number of recorded transitions from the lowest to highest energy IS in the dynamic exploration $T_{\alpha\beta}$ (and $T_{\beta\alpha}$ from high to low-energy IS). The features are ordered from top to bottom by decreasing importance. Each point corresponds to a single IS pair, with a color coding for the feature value (red: high, green: low). The points are spread vertically for readability. The feature impact on the model output is given on the horizontal axis: large SHAP values predict large QS values, low SHAP values predict low QS (more likely to be a TLS).

relative density of TLS quantum splitting $n(E_{qs})$, which according to the TLS model scales as $n(E_{qs}) \sim n_0 E_{qs}$ at low E_{qs} [12, 14]. We show indeed that $n(E_{qs})/E_{qs}$ converges to a plateau n_0 for small E_{qs} . We have recorded $n_0 = 0.67, 4.47$ and 25.14 in units of $\epsilon^{-1}\sigma^{-3}$. This is approximately $10^{23}, 10^{24}$ and 10^{25} $\text{eV}^{-1}\text{cm}^{-3}$ in Ar units. In Refs. [35, 37], we discuss the comparison between numerical and experimental TLS densities. The ML approach allows us to collect significantly better statistics compared to Ref. [35], confirming that the TLS density n_0 decreases by several orders of magnitude from hyperquenched to ultrastable glasses. Lastly, in Fig. 4(c) we report the histograms of the number of TLS and DW per glass at the three temperatures. We see that when the glasses are ultrastable ($T_f = 0.062$) most of the glasses have very few TLS. Conversely, poorly annealed ($T_f = 0.092$) glasses show a very unbalanced distribution, with a few glasses that contain most of the DW and TLS.

E. Interpretation of the ML model

The ML model contains precious information about the distinctive structure of TLS. First, the present and previous works [7, 8, 10, 11, 35, 37, 38] find that the density of TLS decreases upon increasing glass stability, which in our simulations is controlled by the preparation temperature T_f . Thus, one may also expect temperature-dependent TLS features. In the SI we show that when the ML model is trained on glasses prepared at T_{train} and deployed on glasses prepared at $T_{\text{prediction}} \neq T_{\text{train}}$ there is only a minor performance drop and the model is able to perform reasonably well. This implies that the model captures distinctive signatures of TLS that do not depend

strongly on the preparation temperature. Yet, we also show in the SI that it is very easy to train another ML model to predict the temperature itself and eventually add it to the pipeline.

Overall, the ML model is not only able to capture the different microscopic features of TLS, but it can also suggest what the specific influence of each feature is. To interpret this information we calculate Shapley values [54] for each input feature and report them in Fig. 5 for the quantum splitting predictor (a) and the double well classifier (b). The features are ranked from the most important (top) to the less important (bottom). We first discuss the quantum splitting predictor Fig. 5(a). The horizontal axis reports their impact (SHAP value) on the model output: large positive SHAP values predict on average a high value of the quantum splitting (QS). The data points are colored following the value of the feature itself. The most important feature is the classical energy splitting ΔE corresponding to the potential energy difference between the two IS. In our model, a large value of classical splitting ΔE (red) implies a large QS, i.e. non-TLS transitions. The second most important feature is the largest single particle displacement $\Delta\vec{r}_1$, which has to be larger than a threshold corresponding to 0.3σ in order to predict a TLS (low SHAP, hence low QS). The total displacement d is the third most important and shows a similar effect. All the remaining features have a less clear and much smaller effect on the model prediction and they only collaborate collectively to the final QS prediction. Details on features definition are provided in the Methods. In the SI we show that it is possible to obtain very good performance even when removing some of the features with the largest Shapley values, which means that the ML interpretation is not unique.




Figure 6. **Microscopic features of two-level systems and double well potentials in ultrastable glasses.** Probability distributions of (a) number of recorded transitions between the two inherent structures $T_{\alpha\beta} + T_{\beta\alpha}$, (b) classical energy splitting ΔE , (c) largest particle displacement $\Delta\bar{r}_1$, (d) total displacement d and (e) WKB off-diagonal element Δ_0 . We color coded in red the regions of parameters where we do not expect to find TLS, which instead concentrate in the green regions. The points in the yellow regions could be any of TLS, DW or non-DW. The green regions could serve as an alternative way to rapidly identify TLS. Data for stable glasses created at $T_f = 0.062$.

According to this Shapley analysis we explain the ML prediction in the following way: the energy difference ΔE between two IS is the main predictor for the quantum splitting, and it has to be small for TLS. Then, the largest particle displacement $\Delta\bar{r}_1$ is necessary to understand if the two IS are similar and what is their stability (we show in the SI that $\Delta\bar{r}_1$ is the most important feature to identify the glass stability). Then the total displacement d complements this information and gives local information about the displacements of the other particles. Lastly, all the other inputs provide fine tuning to refine the final prediction and are discussed in more detail in the SI. Interestingly, in the SI we also show that even without the two most important features, the ML approach can still identify TLS candidates reasonably well.

F. Microscopic features of TLS

We have shown that by following a ML-driven approach it is possible to collect a significant library of TLS for any preparation temperature. However it may be useful to discuss alternative strategies to rapidly identify TLS. In general, since TLS are extremely rare defects [7, 8, 10, 11, 38] a filtering rule is necessary in order to reduce the number of possible candidates. In particular, Ref. [35–37] proposed to use the number of transitions recorded between two IS during the MD exploration, and to exclude pairs that are not explored consecutively. This is based on the assumption that DW (and

consequently TLS) correspond to IS pairs that are close to each other and should thus be visited consecutively in the dynamics (non-zero number of recorded transitions).

Instead, we prove in Fig. 6 that a filter based on dynamical information only is a poor predictor. In Fig. 6(a) we report the distribution of the number dynamical transitions recorded between two inherent structures α and β , $T_{\alpha\beta} + T_{\beta\alpha}$ (both from low to high, and high to low-energy IS). We report three curves measured for TLS, DW, and all pairs, measured at $T_f = 0.062$. While the slowly decaying tail of TLS and DW suggests that they often exhibit a large transition rate, actually most TLS and DW are formed by IS with no recorded transitions between them in the dynamic exploration.

Our interpretation is that even though the transition from one IS to the other is favorable, the landscape has such a large dimensionality ($3N$) that even very favorable transitions may never take place in a finite exploration time. This issue can become more severe when the trajectories are shorter, for example if the exploration is performed in parallel. We confirm this observation with the results reported in Tab. I, where we have used our iterative training approach to re-analyze the data of Ref. [35], including pairs with no recorded transition, and found many more TLS.

We conclude that even though the number of recorded transitions is the most important feature to predict which pair forms a double well, as seen in Fig. 5(b), a filter based solely on them still misses many pairs of interest and therefore is not the most efficient.

In Fig. 6(b) we focus on the distribution of the classical splitting ΔE , or energy difference between the two IS. When ΔE is large, the transition path between IS rarely forms a DW, or a TLS (red region). On the other hand, there are many pairs with a very small ΔE which are not necessarily more likely to be DW or TLS, hence the yellow region (could be any of TLS, DW, or non-DW). Ultimately we find a ‘sweet spot’ (green region), where TLS are more frequent. The ML model also captures this feature, as seen from the SHAP parameter of ΔE in Fig. 5(a). The next most important feature according to the ML model is the largest particle displacement $\Delta\bar{r}_1$, reported in Fig. 6(c). When it is larger than $\sim 0.8\sigma$ we rarely find TLS and DW, but we do not find them also when $\Delta\bar{r}_0 < 0.3\sigma$. The second row in Fig. 5(a) confirms that the ML model has discovered this feature. In Fig. 6(d) we report the total displacement d . If $d > 0.9\sigma$ the pair is so different that it is not likely to be a TLS or DW, while this probability increases for smaller d . In Fig. 6(e) we report the distribution of off-diagonal elements Δ_0 , measured using the WKB approximation as explained in Ref. [35]. We find that the distribution obtained from TLS and DW scales as $1/\Delta_0$, in good agreement with the standard TLS model [12].

Finally, if one is interested in identifying TLS in a ‘quick and dirty’ way, we propose to use the number of recorded transitions to filter DW from non-DW, and then to select a sweet spot for the classical energy splitting and the displacements for selecting optimal TLS candidates.

III. DISCUSSION

In this paper we have introduced a machine-learning approach to explore complex energy landscapes, with the goal of efficiently locating double wells (DW) and two-level systems (TLS). We demonstrate that it is possible to use ML to rapidly estimate the quantum splitting of a pair of inherent structures (IS) and accurately predict if a DW is a TLS or not. We also show that our ML approach can be used to predict very accurately the energy barrier between pairs of IS, which would be useful to analyze supercooled liquid dynamics, or for the response to a mechanical deformation. Overall, this approach allows us to collect a library of defects of unprecedented size that would be impossible to obtain without the support of ML.

The ML model uses as input the information calculated from a pair of inherent structures. After just a few minutes of supervised training it is able to infer with high accuracy the quantum splitting of any new pair of inherent structures. We establish that our ML model based on model ensembling and gradient boosting is fast and precise. Its efficiency allows us to introduce an iterative training procedure, where we perform a small batch of predictions and then retrain the model.

After performing statistical analysis over the unprecedented number of TLS collected with our method, we

have discovered that many DW and TLS are not consecutively visited during the dynamical exploration. We reanalyzed the data collected for the study of Ref. [35] and found that more than half of the TLS were missed, because the corresponding IS were not visited consecutively and the pair was consequently discarded. Our ML approach not only finds more than twice the number of TLS from the same data, but it also requires significantly fewer calculations. We conclude that ML significantly improves the existing approaches. The ML method allows us to propose a ‘quick and dirty’ way to predict TLS: a) use $T_{\alpha\beta}, T_{\beta\alpha}$ for predicting DW; b) for those which are DW, use the classical energy splitting between the two IS to predict which are TLS.

We also discuss the microscopic nature of DW and TLS. We perform a Shapley analysis to dissect the ML model and understand what it learns, and we compare this with the extended statistics of TLS that we are able to collect. We find that the quantum splitting is mostly related to the classical energy splitting and the displacements of the particles. Overall, the Shapley analysis suggests that TLS are characterized by one particle that displaces between 0.3 and 0.9 of its size, while the total displacement and the energy difference between the two states remains small. The local structure around the particle is not as important, nor is the number of times we actually see this transition during the exploration dynamics.

Lastly we investigate the effect that glass stability (equivalent to the preparation temperature in our simulations) has on double wells and TLS. The ML model learns that at higher temperatures all the pairs are characterized by more collective rearrangements, but TLS are similar for any preparation temperature.

Ultimately, since our ML approach is extremely efficient in exploring the energy landscape and is easy to generalize to target any type of state-to-state transition (as we show for the energy barriers), we hope that our method will be used in the future to analyze not only TLS, but also many other examples of phenomena related to specific transitions between states in complex classical and quantum settings.

IV. METHODS

A. Glass-forming model

We study a three-dimensional polydisperse mixture of $N = 1500$ particles of equal mass m . The particle diameters σ_i are drawn from the normalized distribution $P(0.73 < \sigma < 1.62) \propto 1/\sigma^3$. Two particles i and j separated by a distance r_{ij} interact via the repulsive pair potential

$$v(r_{ij})/\epsilon = (\sigma_{ij}/r_{ij})^{12} + c_0 + c_2(r_{ij}/\sigma_{ij})^2 + c_4(r_{ij}/\sigma_{ij})^4, \quad (2)$$

only if $r_{ij} \leq 1.25\sigma_{ij}$, with the non-additive interaction $\sigma_{ij} = 0.5(\sigma_i + \sigma_j)(1 - 0.2|\sigma_i - \sigma_j|)$. The polynomial coef-

ficients c_0, c_2, c_4 ensure continuity of v and its first two derivatives at the interaction cutoff. We study the system at number density $\rho = 1$ in a cubic box with periodic boundary conditions. We express energies and lengths in units of ϵ and the average diameter σ , respectively. Times measured during molecular dynamics (MD) simulations are expressed in units of $\sqrt{\epsilon/m\sigma^2}$.

B. Glass sample preparation

We fully equilibrate $n_g = 5, 50, 200$ independent configurations of the liquid at preparation temperatures $T_f = 0.092, 0.07, 0.062$, respectively. We do so employing the hybrid MD/particle-swap Monte Carlo algorithm described in Ref. [24]. The algorithm alternates between blocks of $5N$ attempts of particle-swap Monte Carlo moves and short MD runs of duration $t_{MD} = 0.1$ to thermalize the liquid efficiently. Glassy samples are then created by rapidly cooling the equilibrium configurations to $T = 0.04$ using regular MD with a Berendsen thermostat [55]. The preparation temperature T_f is thus Tool's fictive temperature [56] and characterizes the degree of stability of the glass: glasses prepared at a lower T_f are more stable. We compare these T_f with characteristic temperature scales. The mode-coupling crossover temperature is $T_d = 0.1$, and the extrapolated experimental glass transition temperature, where the relaxation time is 12 decades larger than at the onset of glassy dynamics, is $T_g = 0.067$ [24]. The lower $T_f = 0.062$ glasses are ultra-stable, while the higher $T_f = 0.092$ are hyperquenched.

C. Energy landscape exploration and transition matrix $T_{\alpha\beta}$

We use classical molecular dynamics (MD) to explore the potential energy landscape of the glasses. We run MD simulations at $T = 0.04$ in the NVE ensemble with a time step $dt = 0.01$. Configurations along the MD trajectory are quenched to the closest potential energy minimum, or inherent structure (IS), every $\tau_{\text{quench}} = 0.2, 0.1, 0.05$ (for glasses prepared at $T_f = 0.062, 0.07, 0.092$, respectively) using the conjugate gradient method. The quench period τ_{quench} is chosen such that two consecutive quenches typically reach different IS, separated by one energy barrier. For each n_g glass sample we perform 100, 100, 200 MD trajectories starting from different initial velocities, each lasting 40000, 100000, 10000 time steps (low to high T_f). For $T_f = 0.092$ we used a subset of the data obtained in [35].

The transition matrix elements $T_{\alpha\beta}$ count how many times IS_α and IS_β are visited consecutively, i.e. IS_α is reached at time t , and IS_β at time $t + \tau_{\text{quench}}$. Overall, $T_{\alpha\beta}$ is a number that counts the number of transitions observed from IS_α to IS_β , with no physical dimensions. Since this quantity depends on the specific trajectories collected, it varies for different quenching rates and sim-

ulation times. Here, we demonstrate that one advantage of our ML approach over a brute-force approach, as employed in Ref. [35], is to consider all IS pairs, not only those visited consecutively in the MD trajectory (characterized by $T_{\alpha\beta} > 0$), as potential TLS. This expands massively the pool of candidates, while ensuring that computational effort is targeted to IS pairs that are likely to be TLS.

To analyze the transition between two IS we compute the multi-dimensional minimum energy path separating them. This is done by the nudged elastic band (NEB) method [51, 52] implemented in the LAMMPS package. We use 40 images to interpolate the minimal energy path, that are connected by springs of constant $\kappa = 1.0 \epsilon\sigma^{-2}$, and use the FIRE algorithm in the minimization procedure [51, 53]. The NEB algorithm outputs a one-dimensional potential energy profile $V(\xi)$ defined for the reduced coordinate ξ , between the two minima.

D. Quantum splitting computation

We extrapolate the potential obtained from the NEB, defined only between the two minima, to obtain a full double well potential. We used a linear extrapolation of the NEB reaction path. Let us denote \mathbf{r}_1 and \mathbf{r}_2 the coordinates of the particles in the first two images of the system along the reaction path (\mathbf{r}_1 is an energy minimum). We extrapolate the potential V starting from \mathbf{r}_1 and measuring the potential energy of the configuration moving in the direction $\mathbf{r}_1 - \mathbf{r}_2$. We perform a similar extrapolation at the other minimum.

Once the classical potential $V(\xi)$ is obtained by extrapolation as discussed above, we solve numerically the Schrödinger Eq. (1) using a finite difference method. In general, the Laplacian term should take into account curvature effects along the reaction coordinates, as discussed in Ref. [35]. For simplicity, we neglect these effects and use the standard second derivative along the reaction coordinate.

E. Dataset and features construction for ML approach

The first step of our machine learning approach is the evaluation of a set of static quantities for all the available IS. This set consists of: potential energy, particle positions, averaged bond-orientational order parameters [57] determined via a Voronoi tessellation, that we denote $\{q_2, q_3 \dots q_{12}\}$, and finally particle radii. The cost of this operation scales as the number of available states N_{IS} , but we use these quantities to calculate the features of $\sim N_{IS}^2$ pairs. A detailed analysis (see Sec. S2 in the SI) shows that the bond-orientational parameters and the particle sizes are not very useful for the ML model. Since their calculation is slower than all the other features, we

do not include them in the final version of the ML approach.

To construct the input features for each pair of IS we combine the information of the two states evaluating the following:

- Energy splitting ΔE : energy difference between the two IS.
- Displacements $\Delta\vec{r}_i$: displacement vector of particle i between the two configurations. When used in this context index i increases with decreasing displacement $\Delta\vec{r}_1 > \Delta\vec{r}_2 > \dots$
- Total displacement d : total distance between the two IS defined as $d^2 = \sum_i |\Delta\vec{r}_i|^2$.
- Participation ratio PR : defined as $PR = (d^2)^2 / (\sum_i |\Delta\vec{r}_i|^4)$.
- Distance from the displacement center $|\vec{r}_1 - \vec{r}_i|$: we measure the average distance of particle i from the center of displacement \vec{r}_1 , identified as the position of the particle that moves the most. This quantity identifies the typical size of the region of particles that rearrange.
- Transition matrix $T_{\alpha\beta}$ (resp. $T_{\beta\alpha}$): number of times the dynamics explores consecutively the lowest (resp. highest) then the highest (resp. lowest) energy minimum.

The crucial step of the feature construction is that we can reduce the number of features by considering only the M particles whose displacement is the largest between pairs of IS. We make this assumption because we expect that the low temperature dynamics is characterized by localized rearrangements involving only a small fraction of the particles [7, 8, 10, 11, 35, 38]. In Sec. S1 of the SI, we confirm this assumption by showing that the ML model achieves optimal performances even when M is very small. So, the choice of $M \ll N$ makes the ML model computationally effective without any performance drop.

F. Double well classifier

A necessary condition in order for a pair of IS to be a TLS is that the transition between the pair forms a double well (DW) potential. A DW is defined when the minimum energy path between the two IS resembles a quartic potential, as sketched in Fig. 1(b). The final goal of the ML model is to predict the quantum splitting (QS) of the pair and identify pairs with low QS. The first obstacle in the identification of TLS is that DW represent only a small subgroup of all IS pairs. For instance in Ref. [35], only $\sim 0.5\%$ of all the IS pairs are DW at the lowest temperature. It is then mandatory to filter out pairs that are not likely to be a DW.

In the machine learning field there are usually many different models that can be trained to achieve similar performances, with complexity ranging from polynomial regression to deep neural networks. Here, we perform model ensembling and use ensembles both for DW classification and QS prediction. Model ensembling consists in averaging the output of different ML models to achieve better predictions compared to each of them separately. In practice, we use the publicly available AutoGluon library [58]. In this approach, we train in a few minutes a single-stack ensemble that is able to classify DW with $> 95\%$ accuracy. In the SI, sec. S1, we justify this choice of ML model and provide details on performances and hyperparameters.

In particular, we get the best results using ensembles of gradient boosting models [59, 60], which have proven to be the optimal choice in estimating barrier heights of chemical reactions computed with similar methods [61]. The gradient boosting approach predicts the probability $p(y_i) = G(x_i)$ that a pair x_i is a DW, where $y_i = 1$ if the pair is a DW and 0 otherwise. It is based on a series of $m = 1, \dots, n_{WL}$ weak learners W_m that minimize the log-loss score

$$\frac{1}{n} \sum_i y_i \log [p(y_i)] - (1 - y_i) \log [1 - p(y_i)], \quad (3)$$

where n is the number of IS pairs. In this approach, each of the weak learners W_m attempts to improve over the result of its predecessor by predicting the residual $h_{m-1}(x_i) = y_i - W_{m-1}(x_i)$. The final prediction thus becomes

$$p(y_i) = \sum_{m=1}^{n_{WL}} W_m(x_i | W_{l < m}). \quad (4)$$

This approach usually outperforms random forests [62], where the prediction is just the average over the weak learners $p(y_i) = 1/n_{WL} \sum_{m=1}^{n_{WL}} W_m(x_i)$. We then build a stack of gradient boosting models such that the final prediction is given by $p(y_i) = 1/n_{GB} \sum_k^{n_{GB}} c_k G_k(x_i) / \sum_k c_k$, which is the weighted average over the n_{GB} models in the ensemble with learnable weights c_k . Overall, our DW classifier turns to be very accurate and rapid, achieving $> 95\%$ accuracy after only 10 minutes of training. As such, it is convenient to use it to filter out the pairs that do not require the attention of the QS predictor because they cannot be TLS anyway.

G. Quantum splitting predictor

We want to predict the quantum splitting of a pair of IS for which the features discussed in Sec. IV E have been computed. We need this prediction to be very precise, because we know that a pair can be considered a TLS when $E_{qs} < 0.0015\epsilon$, but E_{qs} can vary significantly so errors may be large. In the SI (see Sec. S1) we show that models such as deep neural networks and regression

Collection of a IS pair	50 s
NEB+Schrödinger for a IS pair	10^3 s
Predict if a IS pair is a DW	10^{-5} – 10^{-4} s
Predict E_{qs} for a IS pair	10^{-5} – 10^{-4} s
Train DW classifier	10^3 – 10^4 s
Train QS predictor	10^3 – 10^4 s

Table II. Computational time needed to perform our ML approach, on an Intel i9-9980HK CPU.

are not stable or powerful enough to achieve satisfying results. We thus follow the same strategy introduced for DW classification, by using model ensembling [58] and gradient boosting [59, 60, 62, 63]. Compared to the DW predictor, each model G_k in the ensemble now performs a regression task by predicting $E_{qs,k} = G_k(x_i)$. We then construct a multi-layer stack (schematized in Fig. 2), where the prediction of the first stack $E_{qs}^{(0)} = 1/n_{GB} \sum_k^{n_{GB}} c_k G_k(x_i) / \sum_k c_k$ is concatenated to x_i and used as input for the following stack. At the same time, we also perform k -fold bagging [64], which consists in splitting the data in k subsets used to train k copies of each model with different data. This has shown to be particularly effective in improving the prediction for small datasets [58].

In order to train the model we first collect a set of E_{qs} examples. The size of this training set is discussed in the SI (see Fig. S2) where we find that the minimum number is around 10^4 . We can use some of the data already collected in previous work in Ref. [35] for the training. Moreover, since we are interested in estimating with more precision the lowest values of E_{qs} we train the model to minimize the following loss function

$$\mathcal{L} = \frac{\sum_{i=1}^n w_i (E_{qs,true} - E_{qs,predicted})^2}{n \sum_{i=1}^n w_i}, \quad (5)$$

which is a weighted mean-squared error. The weights correspond to $w_i = 1/E_{qs,true}$ in order to give more importance to low E_{qs} values. We thus train our model to provide a very accurate prediction of the value E_{qs} for any given pair. Once the model is trained it takes only $\sim 10^{-4}$ s to predict the QS of a new pair (compared to 1 minute to run the standard procedure). If we predict a value $E_{qs} < 0.0015\epsilon$, then we have identified a TLS much faster.

H. Iterative training procedure

We finally introduce an approach to optimally employ our ML model to process new data: the iterative training procedure. To produce the results reported in Fig. 3 we trained the model once using a subset of the already available data. This is a natural way to proceed when the

goal is to process new data that are very similar to the training set, and the training set is itself large enough. However, since the goal of the proposed ML model is to ultimately drive the landscape exploration and collect new samples, the single-training approach may encounter two types of problems. First, at the beginning there may be not enough data and, second, the findings of the model do not provide any additional feedback.

To solve both problems we introduce the iterative training procedure. The idea of iterative training is to use the predictive power of ML to create and expand its own training set, consequently enhancing its performance by iteratively retraining over the new data. Compared to standard active learning methods, iterative training does not focus on new samples with the highest model uncertainty, but instead it iterates the predictions on the samples below the threshold of interest. Details on the method and parameters are discussed in Sec. S3 of the SI. In practice, we start from a training set of $K_0 \sim 10^3 - 10^4$ randomly selected pairs to have an initial idea of the relation between input and output. We then use the ML approach outlined in Fig. 2 to predict the $K_i = 500$ pairs with the lowest QS. For these TLS candidates, we perform the full procedure to calculate the true QS and determine whether the pair is a DW or a TLS. In the SI (Fig. S6), we report the result of this procedure when we process a new set of trajectories from the same poly-disperse soft sphere potential as in Ref. [35]. In general the first few iterations of iterative training have a poor performance. In fact we find that $> 70\%$ of the first K_i pairs are actually non-DW. After collecting additional K_i measurements, we retrain the model. We report in Tab. II the average time for each step of the ML procedure. The retraining can be done in ~ 10 min, after which the model is able to predict the next K_i pairs with lowest QS. Overall, to process N_{ISpairs} we estimate that the computational time of the iterative approach is $t_i = [K_0 \cdot 10^2 + N_{\text{iter}} (K_i \cdot 10^2 + 10^3 + N_{\text{ISpairs}} \cdot 10^{-5})]$ s. If $N_{\text{ISpairs}} > 10^9$ it is possible to significantly reduce t_i by permanently discarding the worst pairs, but this is not needed here. We iterate this procedure N_{iter} times, until the last batch of K_i candidates contains less than 1% of the total number of TLS. We believe that continuing this iterative procedure would lead to the identification of even more TLS/DW, but this is out of the scope of this paper.

Data availability

The data that support the findings of this study are available upon request from the corresponding author.

Code availability

The codes that support the findings of this study are available upon request from the corresponding author.

- [1] P. G. Debenedetti and F. H. Stillinger, Supercooled liquids and the glass transition, *Nature* **410**, 259 (2001).
- [2] R. C. Zeller and R. O. Pohl, Thermal conductivity and specific heat of noncrystalline solids, *Physical Review B* **4**, 2029 (1971).
- [3] M. T. Laponen, R. C. Dynes, V. Narayanamurti, and J. P. Garno, Measurements of the time-dependent specific heat of amorphous materials, *Physical Review B* **25**, 1161 (1982).
- [4] J. F. Berret and M. Meißner, How universal are the low temperature acoustic properties of glasses?, *Zeitschrift für Physik B Condensed Matter* **70**, 65 (1988).
- [5] A.-M. Boiron, P. Tamarat, B. Lounis, R. Brown, and M. Orrit, Are the spectral trails of single molecules consistent with the standard two-level system model of glasses at low temperatures?, *Chemical Physics* **247**, 119 (1999).
- [6] A. L. Burin, J. M. Leveritt, G. Ruyters, C. Schötz, M. Bazrafshan, P. Fassel, M. von Schickfus, A. Fleischmann, and C. Enss, Low-temperature dipolar echoes in amorphous dielectrics: Significance of relaxation and decoherence free two-level systems, *EPL (Europhysics Letters)* **104**, 57006 (2013).
- [7] D. R. Queen, X. Liu, J. Karel, T. H. Metcalf, and F. Hellman, Excess specific heat in evaporated amorphous silicon, *Physical Review Letters* **110**, 135901 (2013).
- [8] T. Pérez-Castaneda, R. J. Jiménez-Riobóo, and M. A. Ramos, Two-level systems and boson peak remain stable in 110-million-year-old amber glass, *Physical Review Letters* **112**, 165901 (2014).
- [9] T. Pérez-Castaneda, C. Rodríguez-Tinoco, J. Rodríguez-Viejo, and M. A. Ramos, Suppression of tunneling two-level systems in ultrastable glasses of indomethacin, *Proceedings of the National Academy of Sciences* **111**, 11275 LP (2014).
- [10] X. Liu, D. R. Queen, T. H. Metcalf, J. E. Karel, and F. Hellman, Hydrogen-free amorphous silicon with no tunneling states, *Physical Review Letters* **113**, 25503 (2014).
- [11] D. R. Queen, X. Liu, J. Karel, H. C. Jaks, T. H. Metcalf, and F. Hellman, Two-level systems in evaporated amorphous silicon, *Journal of Non-Crystalline Solids* **426**, 19 (2015).
- [12] P. W. Anderson, B. I. Halperin, and C. M. Varma, Anomalous low-temperature thermal properties of glasses and spin glasses, *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics* **25**, 1 (1972).
- [13] W. A. Phillips, Tunneling states in amorphous solids, *Journal of Low Temperature Physics* **7**, 351 (1972).
- [14] W. A. Phillips, Two-level states in glasses, *Reports on Progress in Physics* **50**, 1657 (1987).
- [15] A. J. Leggett and D. C. Vural, “tunneling two-level systems” model of the low-temperature properties of glasses: Are “smoking-gun” tests possible?, *The Journal of Physical Chemistry B* **117**, 12966 (2013).
- [16] D. Zhou, Universal ratio of t1s-phonon coupling constants in low-temperature amorphous solids, *Journal of Physics: Condensed Matter* **32**, 55704 (2019).
- [17] H. M. Carruzzo and C. C. Yu, Why phonon scattering in glasses is universally small at low temperatures, *Physical Review Letters* **124**, 75902 (2020).
- [18] C. Artiaco, F. Balducci, and A. Scardicchio, Signatures of many-body localization in the dynamics of two-level systems in glasses, *Physical Review B* **103**, 214205 (2021).
- [19] W. Ji, Toward understanding the depletion of two-level systems in ultrastable glasses, *arXiv preprint arXiv:2112.10105* (2021).
- [20] D. Richard, M. Ozawa, S. Patinet, E. Stanifer, B. Shang, S. A. Ridout, B. Xu, G. Zhang, P. K. Morse, J.-L. Barrat, L. Berthier, M. L. Falk, P. Guan, A. J. Liu, K. Martens, S. Sastry, D. Vandembroucq, E. Lerner, and M. L. Manning, Predicting plasticity in disordered solids from structural indicators, *Phys. Rev. Materials* **4**, 113609 (2020).
- [21] J. M. Martinis, K. B. Cooper, R. McDermott, M. Steffen, M. Ansmann, K. D. Osborn, K. Cicak, S. Oh, D. P. Pappas, R. W. Simmonds, and C. C. Yu, Decoherence in josephson qubits from dielectric loss, *Physical Review Letters* **95**, 210503 (2005).
- [22] K. Bagchi and M. Ediger, Controlling structure and properties of vapor-deposited glasses of organic semiconductors: Recent advances and challenges, *The Journal of Physical Chemistry Letters* **11**, 6935 (2020).
- [23] A. Ninarello, L. Berthier, and D. Coslovich, Models and algorithms for the next generation of glass transition studies, *Physical Review X* **7**, 021039 (2017).
- [24] L. Berthier, E. Flenner, C. J. Fullerton, C. Scalliet, and M. Singh, Efficient swap algorithms for molecular dynamics simulations of equilibrium supercooled liquids, *Journal of Statistical Mechanics* **2019**, 64004.
- [25] F. H. Stillinger and T. A. Weber, Hidden structure in liquids, *Physical Review A* **25**, 978 (1982).
- [26] T. A. Weber and F. H. Stillinger, Interactions, local order, and atomic-rearrangement kinetics in amorphous nickel-phosphorous alloys, *Physical Review B* **32**, 5402 (1985).
- [27] A. Heuer and R. J. Silbey, Microscopic description of tunneling systems in a structural model glass, *Phys. Rev. Lett.* **70**, 3911 (1993).
- [28] D. Dab, A. Heuer, and R. J. Silbey, Low temperature properties of glasses: a preliminary study of double well potentials microscopic structure, *Journal of Luminescence* **64**, 95 (1995).
- [29] A. Heuer and R. J. Silbey, Collective dynamics in glasses and its relation to the low-temperature anomalies, *Physical Review B* **53**, 609 (1996).
- [30] F. Demichelis, G. Viliani, and G. Ruocco, Properties of the double well potential and relaxation processes in a model glass, *PhysChemComm* **2**, 20 (1999).
- [31] J. Reinisch and A. Heuer, Local properties of the potential-energy landscape of a model glass: Understanding the low-temperature anomalies, *Physical Review B* **70**, 64201 (2004).
- [32] J. Reinisch and A. Heuer, What is moving in silica at 1 k? a computer study of the low-temperature anomalies, *Physical Review Letters* **95**, 155502 (2005).
- [33] A. Heuer, Exploring the potential energy landscape of glass-forming systems: from inherent structures via metabasins to macroscopic transport, *Journal of Physics: Condensed Matter* **20**, 373101 (2008).
- [34] T. Damart and D. Rodney, Atomistic study of two-level systems in amorphous silica, *Physical Review B* **97**,

- 14201 (2018).
- [35] D. Khomenko, C. Scalliet, L. Berthier, D. R. Reichman, and F. Zamponi, Depletion of two-level systems in ultra-stable computer-generated glasses, *Physical Review Letters* **124**, 225901 (2020).
- [36] D. Khomenko, D. R. Reichman, and F. Zamponi, Relationship between two-level systems and quasilocalized normal modes in glasses, *Physical Review Materials* **5**, 55602 (2021).
- [37] F. C. Mocanu, L. Berthier, S. Ciarella, D. Khomenko, D. R. Reichman, C. Scalliet, and F. Zamponi, Microscopic observation of two-level systems in a metallic glass model, arXiv preprint arXiv:2209.09579 (2022).
- [38] M. D. Ediger, Vapor-deposited glasses provide clearer view of two-level systems, *Proceedings of the National Academy of Sciences* **111**, 11232 LP (2014).
- [39] E. Lerner and E. Bouchbinder, Low-energy quasilocalized excitations in structural glasses, *The Journal of chemical physics* **155**, 200901 (2021).
- [40] L. Wang, A. Ninarello, P. Guan, L. Berthier, G. Szamel, and E. Flenner, Low-frequency vibrational modes of stable glasses, *Nature Communications* **10**, 26 (2019).
- [41] B. B. Laird and H. R. Schober, Localized low-frequency vibrational modes in a simple model glass, *Physical Review Letters* **66**, 636 (1991).
- [42] S. S. Schoenholz, E. D. Cubuk, D. M. Sussman, E. Kaxiras, and A. J. Liu, A structural approach to relaxation in glassy liquids, *Nature Physics* **12**, 469 (2016).
- [43] D. M. Sussman, M. Paoluzzi, M. C. Marchetti, and M. L. Manning, Anomalous glassy dynamics in simple models of dense biological tissue, *EPL (Europhysics Letters)* **121**, 36001 (2018).
- [44] E. Boattini, M. Ram, F. Smallenburg, and L. Filion, Neural-network-based order parameters for classification of binary hard-sphere crystal structures, *Molecular Physics* **116**, 3066 (2018).
- [45] 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, *Nature Physics* **16**, 448 (2020).
- [46] J. Paret, R. L. Jack, and D. Coslovich, Assessing the structural heterogeneity of supercooled liquids through community inference, *The Journal of Chemical Physics* **152**, 144502 (2020).
- [47] E. Boattini, F. Smallenburg, and L. Filion, Averaging local structure to predict the dynamic propensity in supercooled liquids, *Physical Review Letters* **127**, 88007 (2021).
- [48] S. A. Ridout, J. W. Rocks, and A. J. Liu, Correlation of plastic events with local structure in jammed packings across spatial dimensions, *Proceedings of the National Academy of Sciences* **119**, e2119006119 (2022).
- [49] I. Tah, S. A. Ridout, and A. J. Liu, Fragility in glassy liquids: A structural approach based on machine learning, *The Journal of Chemical Physics* **157**, 124501 (2022).
- [50] G. Jung, G. Biroli, and L. Berthier, Predicting dynamic heterogeneity in glass-forming liquids by physics-informed machine learning, arXiv preprint arXiv:2210.16623 (2022).
- [51] G. Henkelman, B. P. Uberuaga, and H. Jónsson, A climbing image nudged elastic band method for finding saddle points and minimum energy paths, *The Journal of Chemical Physics* **113**, 9901 (2000).
- [52] W. E, W. Ren, and E. Vanden-Eijnden, String method for the study of rare events, *Physical Review B* **66**, 52301 (2002).
- [53] J. Guérolé, W. G. Nöhring, A. Vaid, F. Houllé, Z. Xie, A. Prakash, and E. Bitzek, Assessment and optimization of the fast inertial relaxation engine (fire) for energy minimization in atomistic simulations and its implementation in lammmps, *Computational Materials Science* **175**, 109584 (2020).
- [54] S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee, From local explanations to global understanding with explainable ai for trees, *Nature Machine Intelligence* **2**, 56 (2020).
- [55] H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren, A. DiNola, and J. R. Haak, Molecular dynamics with coupling to an external bath, *The Journal of Chemical Physics* **81**, 3684 (1984).
- [56] A. Q. Tool, Relation between inelastic deformability and thermal expansion of glass in its annealing range, *Journal of the American Ceramic Society* **29**, 240 (1946).
- [57] W. Lechner and C. Dellago, Accurate determination of crystal structures based on averaged local bond order parameters, *The Journal of Chemical Physics* **129**, 114707 (2008).
- [58] N. Erickson, J. Mueller, A. Shirkov, H. Zhang, P. Larroy, M. Li, and A. Smola, Autogluon-tabular: Robust and accurate automl for structured data, arXiv preprint arXiv:2003.06505 (2020).
- [59] L. Prokhorenkova, G. Gusev, A. Vorobev, A. V. Dorogush, and A. Gulin, Catboost: unbiased boosting with categorical features, *Advances in neural information processing systems* **31** (2018).
- [60] H. Zhang, S. Si, and C.-J. Hsieh, Gpu-acceleration for large-scale tree boosting, arXiv preprint arXiv:1706.08359 (2017).
- [61] X. García-Andrade, P. G. Tahoces, J. Pérez-Ríos, and E. M. N. nez, Barrier height prediction by machine learning correction of semiempirical calculations, *The Journal of Physical Chemistry A* **127**, 2274 (2023).
- [62] C. Bentéjac, A. Csörgő, and G. Martínez-Munoz, A comparative analysis of gradient boosting algorithms, *Artificial Intelligence Review* **54**, 1937 (2021).
- [63] X. García-Andrade, P. G. Tahoces, J. Pérez-Ríos, and E. M. Núñez, Barrier height prediction by machine learning correction of semiempirical calculations, arXiv preprint arXiv:2208.02289 (2022).
- [64] B. Parmanto, P. W. Munro, and H. R. Doyle, Reducing variance of committee prediction with resampling techniques, *Connection Science* **8**, 405 (1996).
- [65] P. Geurts, D. Ernst, and L. Wehenkel, Extremely randomized trees, *Machine Learning* **63**, 3 (2006).

Acknowledgements

We thank E. Flenner, G. Folena, M. Ozawa, J. Sethna, S. Elliott, G. Ruocco and W. Schirmacher for useful discussions. This work was granted access to the HPC resources of MesoPSL financed by the Region Ile de France and the project Equip@Meso (reference ANR-10-EQPX-29-01) of the programme Investissements d’Avenir super-

vised by the Agence Nationale pour la Recherche. This project received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program, Grant No. 723955 – GlassUniversality (FZ), and from the Simons Foundation (#454933, LB, #454955, FZ, #454951 DR) and by a Visiting Professorship from the Leverhulme Trust (VP1-2019-029, LB). CS acknowledges support from the Herchel Smith Fund and Sidney Sussex College, University of Cambridge.

Author contributions

SC developed the machine learning model and approach; DK, CS performed the simulations; DK, LB, CS, DRR, FZ developed the simulation model; SC, DK, FCM, FZ analyzed the data; SC, DK, LB, FCM, DRR, CS, FZ wrote the paper; LB, CS, DRR, FZ provided resources and supervision.

Competing interests

The authors declare no competing interests.

SUPPLEMENTARY INFORMATION

S1. WHICH MACHINE LEARNING APPROACH SHOULD WE USE?

A. Comparison between ensembled and non-ensembled models

In the main manuscript we discuss how to transform the problem of computing the quantum splitting of two inherent structures into a supervised learning regression problem. Still, many machine learning models can in principle answer this question. We opted for the AutoGluon library [58] which is based on model ensembling to find an optimal solution. Model ensembling consists in averaging the output of different ML models and create an ensemble-averaged model that outperforms each one of its components. In particular, the AutoGluon library performs model ensembling in the form of stacking (stacking different layers of models that are applied in series) and bagging (training multiple instances of the same model on different subsets of data and then combining their prediction).

In Fig. S1(a) we report the most performant ensembling after 20 hours of training, at $T_f = 0.062$, $M = 5$ and $N_{\text{samples}} = 10^4$. The R2-score reported in the figure shows that the most performant model is a bagged ensemble of *CatBoost*, a gradient boosting method [59], which then corresponds to the final *WeightedEnsemble* for this training instance. Notice that in the main paper, we always use the best ensemble model. In addition to *CatBoost*, very good performances are also usually achieved by *LightGBM*, another gradient boosting method [60], and *ExtraTrees*, which is an extremely randomized tree ensemble [65]. Overall, gradient boosting methods typically achieve the highest R2-score. Thus, in the main text when we report results obtained with the ‘best models’, we refer to ensembles of gradient boosting methods assembled with Autogluon.

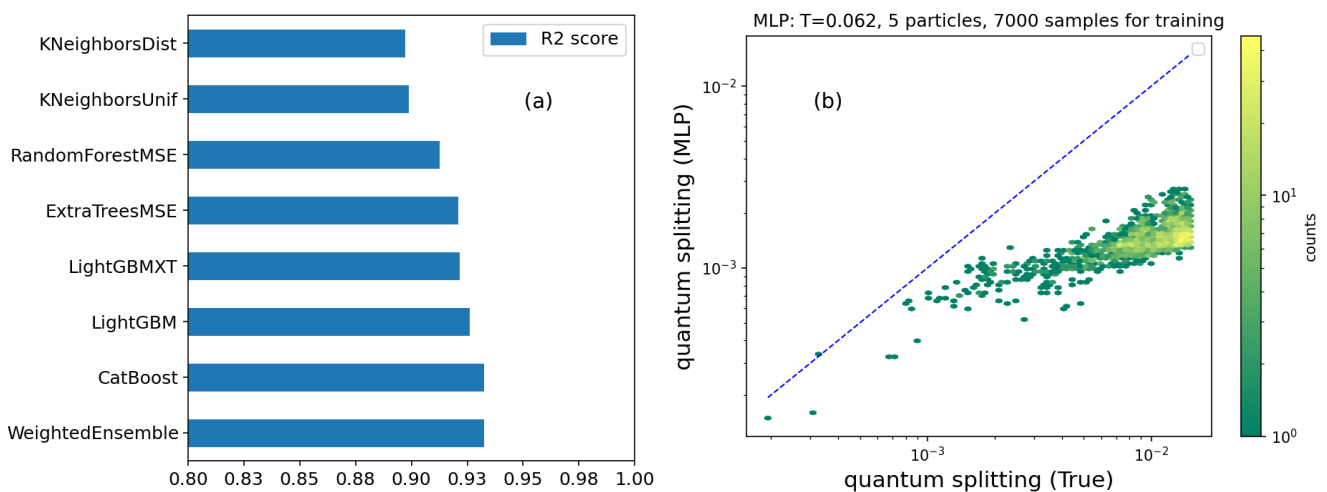


Figure S1. Comparison of different ML model ensembles to predict the quantum splitting at $T_f = 0.062$ using $M = 5$ and $N_{\text{samples}} = 7 \cdot 10^4$. (a) R2-score of the test set obtained for an ensemble of models. The full collection was trained for 20 hours of CPU time. In the main manuscript we use the *WeightedEnsemble*. (b) An optimal Neural Network (MLP with 10 hidden layers of size $1.1 \cdot N_{\text{features}}$) does not produce good predictions ($R2 < 0$).

We motivate our choice of complex ensemble model constructed with AutoGluon by showing that simpler models do not perform at the same level. For the system at $T_f = 0.062$ and in the optimal condition corresponding to $M = 5$ and $N_{\text{samples}} = 7000$ the simplest model to at least capture some correlation was a Multi-layer-perceptron (MLP) with 10 hidden layers of size $1.1 \cdot N_{\text{features}}$, using the same input and output structure as in the main text. A first disadvantage of this approach is that several additional hyperparameters have to be selected, such as learning rate, number and size of hidden layers, activation function, etc. We varied these parameters and report in Fig. S1(b) results for the model with the best performance. This optimal MLP was trained until the iteration per epoch was consistently below 10^{-10} , corresponding to ~ 10 hours. Still, the results in Fig. S1(b) demonstrate that the MLP performance is much worse than that of the ensemble models used in the main text. We believe that with sufficient training time and optimal hyperparameters, the MLP can perform better. However, our model ensembling approach is much faster and does not require any hyperparameter optimization, so we decided to not invest too many computational resources in the MLP. In particular, we believe that the model in Fig. S1(b) is overfocusing on the low quantum splitting pairs, that weight more in the loss function defined in Eq. (3). While in this case the predictions could be scaled by a

constant to achieve significant improvement, the problem with such an approach is that it is not reliable, and this scaling constant is not known a priori adding an additional layer of complication. For this reason we prefer to use more complex machine-learning models.

B. Parameters to predict the quantum splitting

A significant advantage of the model ensemble approach is that most of the hyperparameters are automatically optimized by the ensembling. Still, some degrees of freedom have to be fixed. In particular we need to fix the number of particles to be considered (M), how many samples to use (N_{samples}) and for how long to train the ensemble. In the next sections we motivate the choices reported in the main text.




Figure S2. Optimizing the Quantum splitting predictor. Effect of (a) the number of particles M , (b) number of samples N_{samples} and (c) training time on performance, while other parameters assume their optimal value. We report the R2-score after training on the data of Ref. [35]. Overall we conclude that optimal performances can be achieved for $M = 3$, $\sim 10^4$ samples in just 10 minutes of training.

- **M : input particles** – In Fig. S2(a) we report the effect of M , the number of particles considered in the ML procedure. A large M hinders the performance of the ML model because it has to process a larger input vector. The peak performance is achieved at $M = 3$, which is a relatively low number of particles to define a TLS. This confirms that the participation rate in TLS is low [35].
- **N_{samples} : number of samples** – To evaluate the optimal value of N_{samples} we report in Fig. S2(b) the variation of the R2-score upon using more samples for training. Notice that we keep the other parameters in the optimal condition for each temperature. The shape of the curves let us conclude that a good number of training samples is 7000, 10000, 30000 over a total of 14202, 23535, 117370 for $T_f = 0.062, 0.07, 0.092$ respectively. Notice that the results depend on the specific quality of the samples provided for training. It is possible to achieve better predictions by selecting better initial samples, more uniformly distributed, which are different from the random choice we make for simplicity. Still, the values reported consistently produce good results, independently from the initial sample composition.
- **Training time** – In Fig. S2(c) we report the effect of the training time on performance. The R2-score approaches the plateau in 10 minutes of training on a single CPU and reaches the final value after 10^3 minutes (~ 16 hours). Overall, we find that in iterative training, where we retrain the model several times, a good balance between performance and speed can be reached by training for ~ 10 minutes. Instead, in a more classic single training approach we suggest to train for $> 10^3$ minutes. Performing the training in parallel can further reduce the training time.

C. Parameters to identify double wells

The ML model has to process a large number of unfiltered minima obtained during the exploration simulation. While we can anticipate that only a small fraction of inherent structure pairs will be a TLS, we also know that the minimum energy path connecting any pair of IS will not form a DW and as it is likely to contain intermediate minima. To exclude those non-DW pairs, we train a classifier using as input the same quantities that we measure to predict




Figure S3. Accuracy of the double well classifier as a function of (a) M , (b) training time, and (c) number of samples. We report the accuracy after training at the lowest temperature $T_f = 0.062$ using the data from Ref. [35].

the QS, as described in the main text. In Fig. S3 we report the results of a hyperparameter scan at $T_f = 0.062$. These results are similar to those obtained for the QS prediction suggesting that at $T_f = 0.062$ we can use (a) $M = 3$, (b) 10²s of training and (c) $\sim 10^4$ samples in order to achieve $> 95\%$ accuracy, measured over a validation set. We conclude that we can use the same hyperparameters for the DW classifier and the QS predictor.

S2. FEATURES

A. Removing irrelevant features

After finding the optimal parameters to train the ML model, we can extract the specific effect and overall importance of each input parameter, as well as the score drop when each of them is removed. Since each additional feature corresponds to additional computational time and memory, we investigate which features are crucial, and which ones can be excluded to make the ML pipeline faster and more efficient without affecting performance.




Figure S4. Importance of different features on the quantum splitting predictor. (a) Performance drop when each specific feature is removed, normalized to the first most important. Different colors correspond to different glass preparation temperatures. (b) Shapley values calculated at $T_f = 0.062$. In general, at any temperature the most important feature is the energy difference (ΔE) of the two IS, which is followed by the value of the displacement of the M particles. The Shapley values on the right also report the effect of the specific features: the splitting ΔE has to be small (green) in order to predict low QS (i.e. negative SHAP), while instead the displacements have to be large (red) in order to point towards low QS.

In Fig. S4(a) we report the feature importance (in log scale) after a single iteration of training containing all the data from Ref. [35]. Different colors refer to glasses prepared at different temperatures. Independently from the temperature, the most important features are the energy difference ΔE between the two IS, followed by the total displacement of particles $\sum_i \Delta \vec{r}_i$. The third most important information is the positions of the particles that displaced the most $\sum_i |\vec{r}_1 - \vec{r}_i|$. After them, we find that the arrangement of the first shell of neighbors represented by the q

parameters and the particle sizes σ_i are insignificant (and so is their variation between the two IS of the pair Δ -reported in Fig. S4), since their importance is more than two orders of magnitude smaller than the one of the energy and displacements. Since the q parameters are also the slowest to compute, we decided to not include them in the final pipeline reported in the main manuscript.

Next, to quantify the role of those inputs, we calculate their Shapley values [54] shown in Fig. S4(b) for $T_f = 0.062$. The color codes for the value taken by the specific feature (red when the value is high, green when low) and reports its scaled impact on the model output. The x axis reports the Shapley value, where SHAP > 0 implies that the specific feature is pushing towards predicting a high QS, while SHAP < 0 indicates that the feature promotes low QS values. While no input feature alone can predict a TLS (because |SHAP| is small) there are some visible trends: (i) the energy asymmetry is the most important feature and should be small in order to predict a low QS, (ii) the particle displacements have to be larger than a threshold in order to predict low QS.

According to these results we rationalize the ML prediction. The energy difference between two IS is the main predictor for the quantum splitting, which is never too large for DW, then the displacements are necessary to understand if the two IS are similar and what their stability is (see the temperature classification section in the SI). The displacement nucleus size complements the information contained in the displacements and gives local information about the participation ratio. Finally, all the other features do not provide any improvement. Since the bond order parameters are computationally expensive to calculate, we do not include them in the final ML approach we propose in the main text. The performance reported in the main text confirms that our choice is justified.

B. Reducing the number of features

We justify the choice of features discussed in sec. IVa and Fig. 2 by presenting the performance of our ML model as a function of feature number. In Tab. S1 we rank the features according to their Shapley values (see Fig. 6). We report (blue line) in Fig. S5 the accuracy of the DW classifier (a) and the QS predictor (b) as a function of number of features, following the Shapley ranking of Tab. S1. The DW classifier reaches its best accuracy with six features. In the initial study Ref. [35, 37], the matrix \mathbf{T} was the only information used to analyze the pair of IS. Here, the classifier already reaches $\sim 90\%$ accuracy using the transition matrix only (two features). We add (orange line) the performances when we exclude $\Delta E, T_{\alpha\beta}$ and $T_{\beta\alpha}$, which are the most important features overall. Surprisingly, the DW classifier still reaches its maximum accuracy with two features ($\Delta\bar{r}_3$ and $|\bar{r}_1 - \bar{r}_3|$) as highlighted in the inset. The QS predictor shows instead very good predictions even using a single feature (ΔE in the blue curve, or $\Delta\bar{r}_1$ in the orange).

Shapley ranking	1 st	2 nd	3 rd	4 th	5 th	6 th	7 th	8 th	9 th	10 th
DW classification	$T_{\alpha\beta}$	$T_{\beta\alpha}$	$\Delta\bar{r}_3$	$ \bar{r}_1 - \bar{r}_3 $	$ \bar{r}_1 - \bar{r}_2 $	d	$\Delta\bar{r}_1$	ΔE	PR	$\Delta\bar{r}_2$
QS prediction	ΔE	$\Delta\bar{r}_1$	d	$\Delta\bar{r}_2$	$ \bar{r}_1 - \bar{r}_2 $	$ \bar{r}_1 - \bar{r}_3 $	$\Delta\bar{r}_3$	$T_{\alpha\beta}$	$T_{\beta\alpha}$	PR

Table S1. Ranking the features used in the main manuscript by their Shapley values.

S3. ITERATIVE TRAINING

The standard supervised learning approach consists in collecting a significant amount of data, then using them to train the ML model and finally use the model predictions. While overall very powerful, this scheme is not optimal when the goal of the model is to drive the exploration in a space much larger than the available data. This is the case of our main study where we employ the ML model to identify IS pairs with low QS. In order to make our approach more efficient and generalizable we developed the *iterative training* procedure.

We start the iterative training procedure from a sample of K_0 pairs for which we calculate the QS. We empirically find that $K_0 \sim 5000$ achieves a good balance between precision and time. This sample has to be balanced between DW and non-DW, but there is no need to include TLS in the initial sample. It is possible to start from a smaller K_0 , at the cost of a performance drop in the first few iterations of training. From the initial sample we perform a first training that takes 10 minutes for the DW classifier and 10 minutes for the QS predictor. We then use the model to predict the $K_i = 500$ IS pairs with the lowest QS. We report the cumulative number of TLS, DW, and non-DW at each iteration in Fig. S6(a). From Fig. S6(b) we see that during the first iteration most of the $K_i = 500$ best candidates are actually non-DW, so the model is performing poorly. This is the reason why we suggest to perform only $K_i = 500$ NEBs for each iterations, otherwise the first iteration would lead to wasting time to perform calculations over non-DW. On




Figure S5. Effect of the number of features. Accuracy of the DW classifier (a) and Pearson correlation score of the QS predictor (b), as a function of the number of features used. Along the blue curves the features are ranked by their importance using their Shapley values, so the first features to be used are the most important. For the orange curve we exclude ΔE , $T_{\alpha\beta}$ and $T_{\beta\alpha}$ to evaluate the performances of the ML model without the features with the best Shapley values.




Figure S6. Performance of the iterative training procedure at $T_f = 0.062$ for the data collected in the main manuscript. Starting from $K_0 = 5000$ samples we perform iterations of $K_i = 500$ predictions for 11 iterations. In (a) we report the cumulative number of double wells (DW), two-level systems (TLS) and non double wells (Non-DW). In (b) we break down this result by color coding the confusion matrix of the ML model at each iteration of iterative training.

the other hand, in the first iteration 72 TLS are already found by running 500 NEBs. This is to be compared with Ref. [35] in which 61 TLS were found by running > 14000 NEBs. After the first retraining (during iteration 2) the performance of the ML model is already excellent and less than 30% of the 500 best pairs are non-DW, while 134 are newly found TLS.

We also used our iterative training to reprocess the data of Ref. [35], and we report the results in Table 1 (main text). We show in Fig. S7 a detailed analysis of the procedure at $T_f = 0.062$. While the standard approach was able to identify 61 TLS running > 14000 NEB+Schrödinger calculations, iterative training finds 156 TLS, by running only 2500 NEBs. This confirms that more than half of the total TLS were hidden among the pairs discarded by Ref. [35]. In details, Fig. S7(a) shows the cumulative number of DW and TLS that we find, while Fig. S7(b) reports the confusion matrix for the different steps of iterative training.

Overall, these results demonstrate that the ML approach is not only faster than a manual filtering rule based on the transition matrix, but also much more effective. The pool of IS pairs excluded from the analysis in the original approach effectively contains a significant number of TLS. In conclusion, a ML driven exploration is not only more efficient than manual filtering, but also necessary to capture the correct statistics.




Figure S7. Reprocessing the data of Ref. [35] at $T_f = 0.062$ with iterative training. In (a) we report the cumulative number of double wells (DW), two-level systems (TLS) and non double wells (Non-DW). While Ref. [35] (red dashed line) identified 61 TLS running > 14000 NEBs, iterative training finds 156 TLS from 2500 NEBs. In (b) we report the confusion matrix of the 5 steps of iterative training that we run.

S4. TEMPERATURE AND STABILITY

Our ML approach is able to predict the quantum splitting of a pair of IS from static information. It is known that TLS have different features depending on their preparation temperature, or glass stability [35]. Here we address the following questions: (i) how difficult is it for a machine to distinguish data corresponding to different temperatures, (ii) what are the most important features for this task, and (iii) does our ML approach learn temperature-independent features?

A. Temperature classification

To understand how TLS and IS features evolve with temperature, we trained a multi-layer-perceptron (MLP) to classify the temperature corresponding to a specific pair. We find that it is possible to rapidly train a classifier reaching accuracy $> 95\%$. Depending on the value of M , the input layer is composed by N_{features} neurons, where the features are explained in the main text. After the input layer, there are n hidden fully connected layers, all of the same size s . The activation function for each neuron-neuron connection is a ReLU function. The last layer is composed of 3 neurons that represent the probability that a given input belongs to one of the three classes: $T_f = 0.062$ or $T_f = 0.07$ or $T_f = 0.092$. The performance of the MLP are evaluated by measuring the accuracy, which is the percentage of the corrected predictions. In Fig. S8(a) we report the effect of increasing the size of the hidden layers, while in Fig. S8(b) we report the effect of a larger number of hidden layers.

Our results show that a MLP with 2 hidden layers of 60 neurons has a 95% accuracy after less than 1h of supervised training. The answer to question (i), is that one can identify the temperature at which a pair of IS was obtained.

B. Static signatures of glass stability

We answer question (ii) by measuring the Shapley (SHAP) values from the temperature classifier. In Fig. S9 (a) we report the SHAP value of the most important input features for the MLP prediction that considers only the $M = 5$ particles that displaced the most. The smallest ($i = 5$) and the largest ($i = 1$) particle displacements emerge as the most important input, noted $\Delta\bar{r}_5$ and $\Delta\bar{r}_1$, respectively. Their average effect on the model prediction is shown in Fig.S9(b,c). The particle that displaced the most ($i = 1$) shows a large displacement at low preparation temperature, while the particle that displaced the least ($i = M$) exhibits a large displacement at high temperature. Our results suggest that higher temperatures are characterized by more collective rearrangements, leading to large $\Delta\bar{r}_5$. In glasses prepared at low temperature instead, transitions are characterized by a particle displacing significantly more than the others.




Figure S8. Temperature classifier accuracy as a function of number of hidden layers (a) and their size (b). We report results for $M = 5$. Results suggest that already with 2 hidden layers of 60 neurons it is possible to achieve an accuracy above 95%.




Figure S9. Microscopic differences originating from different temperature preparation/glass stability, quantified using the Shapley values for the temperature classifier. (a) Summary of the SHAP values for the 9 most important input features, measured for a subset of 1000 pairs processed by the T -classifier. Predicted glass preparation temperature as a function of (b) $\Delta\bar{r}_5$, and (c) $\Delta\bar{r}_1$, with all other inputs taking their average value.

C. Transferability and crossvalidation

We address question (iii) by testing how our model performs when trained at T_{train} and deployed to predict the quantum splitting of pairs at $T_{\text{predict}} \neq T_{\text{train}}$. In Fig. S10 each column corresponds to a training temperature: $T_{\text{train}} = 0.092$ (left), 0.07 (middle) and 0.062 (right). Then, computed quantum splittings are compared with ML prediction made on samples obtained at T_{predict} , decreasing from top to bottom. We see in Fig. S10 that the predictive power of a model trained at a different temperature is slightly lower compared to the model trained at the same temperature (Fig. 3 of main). Still, the transferability of the model is good, especially when the model is trained at $T_{\text{train}} > T_{\text{predict}}$. This situation is the most useful, since it is easier to obtain data at higher T . We conclude that if not enough data is available at low temperature, it is possible to rely on model transferability. This relies on the fact that TLS share some general temperature-independent features.

In summary, we have seen that IS and TLS have different microscopic features when generated from different stabilities. It is then optimal to train the ML model at a fixed temperature only. If no easier way to measure T /stability are available, it is possible to use ML to classify the stability of each IS by adding another block to the workflow, as reported in the main manuscript. Alternatively, at the cost of a finite accuracy drop it is even possible to transfer the model predictions at different temperatures and thus train where data collection is fast and easy and apply it where it is not.

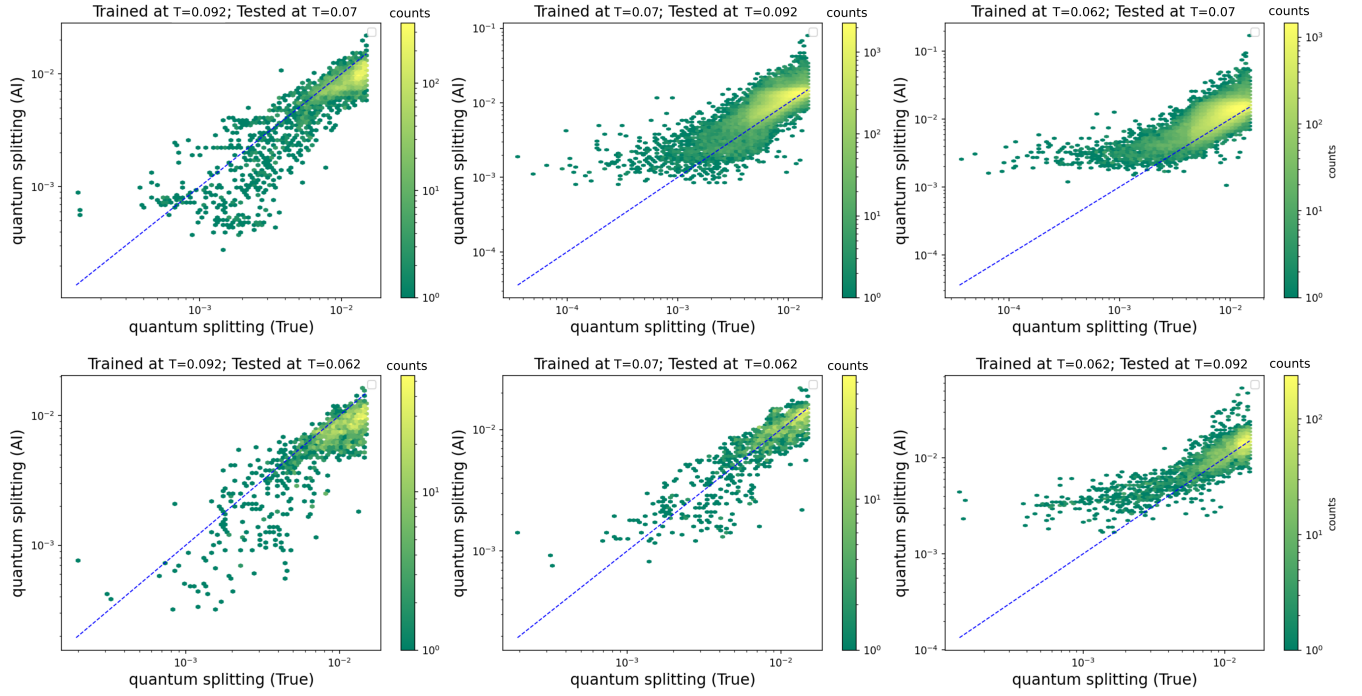


Figure S10. Transferability and crossvalidation of the machine learning model. Exact quantum splitting against ML prediction. The quantum splitting predictor is trained at T_{train} and deployed at $T_{\text{prediction}} \neq T_{\text{train}}$. From left to right: $T_{\text{train}} = 0.092, 0.07, 0.062$. From top to bottom, decreasing $T_{\text{prediction}}$. Performances are not on par with Fig. 3, but good in particular when $T_{\text{train}} > T_{\text{predict}}$, which is of practical interest.