# Modeling the Impact of Baryons on Subhalo Populations with Machine Learning

Ethan O. Nadler<sup>1</sup><sup>(1)</sup>, Yao-Yuan Mao<sup>2</sup><sup>(1)</sup>, Risa H. Wechsler<sup>1,3</sup><sup>(1)</sup>, Shea Garrison-Kimmel<sup>4</sup><sup>(1)</sup>, and Andrew Wetzel<sup>4,5,6</sup><sup>(1)</sup>

<sup>1</sup> Kavli Institute for Particle Astrophysics and Cosmology and Department of Physics, Stanford University, Stanford, CA 94305, USA; enadler@stanford.edu Department of Physics and Astronomy and Pittsburgh Particle Physics, Astrophysics and Cosmology Center (PITT PACC),

University of Pittsburgh, Pittsburgh, PA 15260, USA

<sup>3</sup> SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

<sup>4</sup> TAPIR, Mailcode 350-17, California Institute of Technology, Pasadena, CA 91125, USA

The Observatories of the Carnegie Institution for Science, Pasadena, CA 91101, USA <sup>6</sup>Department of Physics, University of California, Davis, CA 95616, USA

Received 2017 December 12; revised 2018 April 12; accepted 2018 May 1; published 2018 June 1

#### Abstract

We identify subhalos in dark matter-only (DMO) zoom-in simulations that are likely to be disrupted due to baryonic effects by using a random forest classifier trained on two hydrodynamic simulations of Milky Way (MW)-mass host halos from the Latte suite of the Feedback in Realistic Environments (FIRE) project. We train our classifier using five properties of each disrupted and surviving subhalo: pericentric distance and scale factor at first pericentric passage after accretion and scale factor, virial mass, and maximum circular velocity at accretion. Our five-property classifier identifies disrupted subhalos in the FIRE simulations with an 85% out-of-bag classification score. We predict surviving subhalo populations in DMO simulations of the FIRE host halos, finding excellent agreement with the hydrodynamic results; in particular, our classifier outperforms DMO zoom-in simulations that include the gravitational potential of the central galactic disk in each hydrodynamic simulation, indicating that it captures both the dynamical effects of a central disk and additional baryonic physics. We also predict surviving subhalo populations for a suite of DMO zoom-in simulations of MW-mass host halos, finding that baryons impact each system consistently and that the predicted amount of subhalo disruption is larger than the host-to-host scatter among the subhalo populations. Although the small size and specific baryonic physics prescription of our training set limits the generality of our results, our work suggests that machine-learning classification algorithms trained on hydrodynamic zoom-in simulations can efficiently predict realistic subhalo populations.

Key words: dark matter – galaxies: abundances – galaxies: halos – methods: numerical

### 1. Introduction

The ACDM cosmological model provides a remarkably successful framework in which the observed large-scale distribution of galaxies can be understood in terms of the underlying distribution of dark matter halos. However, there are several outstanding "small-scale" problems associated with ACDM cosmology (see Bullock & Boylan-Kolchin 2017 and Del Popolo & Le Delliou 2017 for recent reviews). For example, dark matter-only (DMO) simulations predict large numbers of low-mass subhalos that contribute to an ever-rising low-mass end of the subhalo mass function. If these low-mass subhalos exist and host galaxies, we should observe many more dwarf satellites than currently detected around the Milky Way (MW) or the Andromeda Galaxy (M31); this is often dubbed the missing-satellites problem (Klypin et al. 1999; Moore et al. 1999). Meanwhile, the "too big to fail" (TBTF; Boylan-Kolchin et al. 2011) problem arises because the number of subhalos with high maximum circular velocities  $(V_{\rm max} \gtrsim 15 \,{\rm km \, s^{-1}})$  found in DMO simulations of MW-mass systems substantially exceeds the number of such subhalos inferred to exist around the MW and M31. Equivalently, observational estimates for the masses of the subhalos that host the dwarf satellites of the MW and M31 fall below the masses predicted by DMO simulations (Boylan-Kolchin et al. 2012).

While these small-scale problems present challenges to the ACDM paradigm, a number of promising astrophysical solutions to each problem have been proposed. For example, it is now understood that cosmic reionization suppresses star formation in low-mass subhalos, while supernova (SN) feedback can suppress star formation in more massive subhalos, potentially resolving the missing-satellites problem (Bullock et al. 2000; Somerville 2002). Proposed solutions to the TBTF problem build on these ideas by invoking stellar feedback to soften central density cusps and deplete subhalos of dark matter (Governato et al. 2012; Pontzen & Governato 2012), along with enhanced subhalo disruption via tidal stripping or disk shocking, to destroy many of the high- $V_{\text{max}}$ subhalos found in DMO simulations. Several authors have suggested that these mechanisms can yield subhalo populations in agreement with those inferred observationally for the MW and M31 (Zolotov et al. 2012; Brooks et al. 2013; Brook & Di Cintio 2015; Wetzel et al. 2016; Sawala et al. 2017).

Indeed, recent high-resolution hydrodynamic simulations that self-consistently resolve star formation, stellar feedback, and the formation of central galactic disks indicate that the missing-satellites and TBTF problems can largely be mitigated for the subhalo populations of MW-mass host halos. For example, Wetzel et al. (2016) and Garrison-Kimmel et al. (2017, hereafter GK17) studied the subhalo populations of two MW-mass host halos from the Latte simulation suite of the Feedback in Realistic Environments (FIRE) project (Hopkins et al. 2014) using the "zoom-in" simulation technique (Katz & White 1993; Oñorbe et al. 2015). These authors found that the total number of subhalos in each simulation is reduced by about a factor of two relative to corresponding DMO simulations with identical initial conditions, and they also found significantly fewer subhalos with high circular velocities in the hydrodynamic runs. Moreover, the subhalo populations in both of these systems are consistent with a variety of observational

probes for the MW and M31, which suggests that the missingsatellites and TBTF problems can be resolved in these particular simulations (Wetzel et al. 2016). Zhu et al. (2016) reached similar conclusions by comparing hydrodynamic zoom-in simulations of MW-mass host halos from the Aquarius Project (Springel et al. 2008) to DMO simulations of the same hosts.

These results rely on a limited number of high-resolution simulations of MW-mass host halos; unfortunately, studying a large, diverse sample of subhalo populations in hydrodynamic zoom-in simulations is currently infeasible. While many authors have justifiably focused on the subhalo populations of MW-mass host halos, since these are particularly relevant to the original TBTF problem, it is important to assess whether analogous TBTF problems arise for the subhalo populations of more massive host halos. In addition, understanding whether the TBTF problem is consistently mitigated in a range of simulations with different baryonic physics implementations is necessary in order to make robust conclusions. Quantifying the impact of baryonic physics on subhalo populations more generally will be important in order to interpret results from large-scale surveys, including the Dark Energy Spectroscopic Instrument (DESI Collaboration et al. 2016) and the Large Synoptic Survey Telescope (LSST Science Collaboration et al. 2009), and from targeted searches for satellites of MWlike galaxies outside the Local Group, such as the Satellites Around Galactic Analogs Survey (SAGA; Geha et al. 2017).

Thus, models that can incorporate a variety of hydrodynamic simulations to predict realistic subhalo populations directly from DMO simulations are worth exploring. As a first step toward such a model, we present a machine-learning classification algorithm to identify subhalos in DMO zoom-in simulations of MW-mass host halos that are likely to be disrupted due to baryonic effects in hydrodynamic resimulations. In particular, we train a random forest classifier on disrupted and surviving subhalos from the FIRE zoom-in simulations presented in GK17, and we use the classifier to predict surviving subhalo populations in DMO zoom-in simulations. Our aim is to explore whether this algorithm can capture the effects of baryons in existing hydrodynamic simulations and how the particular baryonic physics in these simulations alters subhalo populations in independent DMO simulations. Rather than providing a detailed comparison of different classification algorithms, we show that a simple random forest classifier predicts subhalo populations in excellent agreement with hydrodynamic results when applied to DMO simulations of the FIRE host halos. This technique is efficient, since a trained classifier can immediately predict surviving subhalo populations from relatively inexpensive DMO simulations. We view classification as a promising technique for predicting subhalo disruption because classifiers will become more robust as the number of high-resolution hydrodynamic simulations to train on increases. In particular, classification algorithms can be trained on a variety of zoom-in simulations to capture the impact of baryons on subhalo populations for a range of host halo masses, central galaxy types, formation histories, and subgrid physics prescriptions.

In addition to the practical utility of our results for predicting realistic subhalo populations, our work provides insights into subhalo disruption in hydrodynamic simulations and relates to the small-scale challenges described above. For example, our random forest classifier determines how strongly various subhalo properties correlate with disruption, which indicates the importance of different disruption mechanisms, including tidal effects and stellar feedback, given the specific baryonic physics prescription in these simulations. To explore the relative importance of these disruption mechanisms, we compare the surviving subhalo populations that we predict from DMO simulations of the FIRE host halos to the DMOplus-disk simulations presented in GK17, which are designed to capture the dynamical effects of the central galactic disk that develops in each hydrodynamic simulation. In particular, by performing DMO zoom-in simulations of two systems with analytic disk potentials tuned to match the galactic disks that develop in the corresponding hydrodynamic simulations, GK17 found subhalo populations in good agreement with the hydrodynamic results, particularly in the innermost regions ( $r \lesssim 100$  kpc). This result suggests that, for MW-mass halos with a central galactic disk, the tidal effects of the disk are largely responsible for disrupting both the low- $V_{\text{max}}$  subhalos relevant to the missing-satellites problem and the high- $V_{max}$ subhalos relevant to the TBTF problem. Our machine-learning predictions are consistent with the DMO-plus-disk simulations at low  $V_{\rm max}$ , but we find enhanced disruption for subhalos with  $V_{\rm max} \gtrsim 15$  km s<sup>-1</sup>, and our results match the FIRE simulations more closely for such subhalos. Interestingly, several authors have suggested that baryonic physics efficiently creates cored subhalo density profiles in this regime (Chan et al. 2015; Tollet et al. 2016; Fitts et al. 2017). We therefore argue that baryonic effects within subhalos, such as stellar feedback, can help to relieve the tension between the subhalo populations predicted by DMO simulations and those inferred from observations of the Local Group.

Our work also has broader implications for studying the galaxy-halo connection. For example, by using our classifier to predict surviving subhalo populations for the suite of DMO zoom-in simulations of MW-mass host halos from Mao et al. (2015), we find that the average amount of subhalo disruption due to baryonic effects is larger than the host-to-host scatter among various subhalo populations. Thus, models that utilize subhalo statistics from these simulations should account for enhanced subhalo disruption when marginalizing over the effects of baryonic physics. Several semianalytic models (e.g., Lu et al. 2016, 2017) incorporate subhalo velocity functions predicted by DMO zoom-in simulations of MW-mass host halos in order to constrain the properties of the MW satellite galaxies and their host halos, and it is plausible that the physical insights provided by these models could change when more realistic subhalo populations are used as input.

This paper is organized as follows. In Section 2, we describe the FIRE simulations that we use to train our random forest classifier, as well as the DMO and DMO-plus-disk simulations presented in GK17 to which we compare our results. In Section 3, we describe our training and cross-validation methods, and we test our classifier by predicting disrupted subhalos in two FIRE zoom-in simulations. We present our main results in Section 4. In Section 4.1, we predict surviving subhalo populations in DMO simulations of the FIRE host halos, and we present velocity functions and radial distributions for our predicted subhalo populations; in Section 4.2, we predict surviving subhalo populations for the suite of DMO zoom-in simulations from Mao et al. (2015), and we discuss the implications for satellite searches. We address avenues for future work and summarize our conclusions in Section 5. We adopt cosmological parameters consistent with each simulation that we analyze. In particular, we use h = 0.702,  $\Omega_{\rm m} = 0.272$ ,  $\Omega_{\rm b} = 0.0455$ , and  $\Omega_{\Lambda} = 0.728$  for our analysis of the FIRE simulations and h = 0.7,  $\Omega_{\rm m} = 0.286$ ,  $\Omega_{\rm b} = 0.047$ , and  $\Omega_{\Lambda} = 0.714$  for our analysis of the MW zoom-in simulation suite. Note that we express distances in physical kpc and velocities in km s<sup>-1</sup>.

### 2. Simulation Data

We train our random forest classifier using subhalos from the hydrodynamic zoom-in simulations presented in GK17. These authors studied the subhalo populations of two MW-mass host halos, referred to as m12i  $(M_{\rm vir} = 1.1 \times 10^{12} \ M_{\odot})$  and m12f  $(M_{\rm vir} = 1.6 \times 10^{12} \ M_{\odot})$ , which were simulated as part of the Latte suite from the FIRE project (Hopkins et al. 2014). These simulations were performed using the FIRE-2 code (Hopkins et al. 2017), which includes the same radiative heating and cooling, star formation, and stellar feedback prescriptions as the original FIRE-1 code in addition to several numerical improvements. The simulations were run in the same cosmological volume (side length 60  $h^{-1}$  Mpc) as the AGORA project (Kim et al. 2014); the m12i and m12f zoom-in simulation regions each contain a single host halo at redshift z = 0 that has no MW-mass neighbors within 3 Mpc. The m12i simulation was originally presented in Wetzel et al. (2016); m12f, which was simulated using the same parameters and pipeline, was first presented in GK17. The baryonic mass resolution in these simulations is  $\sim$ 7000  $M_{\odot}$ , while the dark matter particle mass is  $3.5 \times 10^4 M_{\odot}$ , corresponding to a subhalo mass resolution of  $\sim 3 \times 10^6 M_{\odot}$ . We refer the reader to GK17 and Hopkins et al. (2017) for details on the initial conditions, gravitational-force softenings, and models for radiative heating/cooling, star formation, and stellar feedback in these simulations. Halo catalogs were created using AHF (Knollmann & Knebe 2009), and merger trees were generated using the consistent-trees merger code (Behroozi et al. 2013b).

We will compare our results to the m12i and m12f subhalo populations from three sets of simulations: the hydrodynamic FIRE simulations described above, DMO simulations that were run with identical initial conditions, and the dark matter–plus– disk potential (DISK) simulations presented in GK17. The DISK simulations are identical to the corresponding DMO simulations, but they include gravitational potentials designed to capture the effects of the central disks in the hydrodynamic simulations. In particular, a disk potential is added to each DMO zoom-in simulation at z = 3, and its parameters and evolution are tuned to match the central disk that develops in the corresponding FIRE simulation. We refer the reader to GK17 for a detailed description of the DISK simulations.

Figure 1 in GK17 illustrates the dark matter substructure in m12i for each type of simulation. The visual differences between the FIRE and DMO subhalo populations qualitatively show that baryonic physics in the FIRE simulations lowers both the total number of surviving subhalos and the number of high- $V_{\text{max}}$  subhalos that contribute to the TBTF problem. This figure also shows that the DISK simulation captures the majority of the subhalo disruption in m12i, particularly in the innermost regions ( $r \leq 100 \text{ kpc}$ ), which implies that the central disk is largely responsible for the subhalo disruption in the corresponding hydrodynamic simulation. We have verified the quantitative results in GK17 by calculating velocity functions

and radial distributions for the m12i and m12f subhalo populations in the FIRE, DISK, and DMO simulations. Note that, as in GK17, we scale all subhalo masses by a factor of  $1 - f_b$  and all subhalo circular velocities by a factor of  $\sqrt{1 - f_b}$  in our post-processing of the DMO and DISK halo catalogs, where  $f_b = \Omega_b/\Omega_m \simeq 0.17$  is the cosmic baryon fraction. The mass correction accounts for the fact that the baryonic mass in the hydrodynamic simulations is included in the dark matter particles in the DMO simulations, and the circular velocity correction is an approximate way to account for reduced subhalo densities due to stellar feedback, similar to the prescription in Zolotov et al. (2012). Neither of these corrections affect our results.

To study disrupted subhalos in the FIRE simulations, we select subhalos that disappear from the m12i and m12f halo catalogs after z = 3. We choose this redshift in order to match the initial redshift of the DISK simulations in GK17; note that there are very few subhalos disrupted before z = 3 that pass our subsequent minimum circular velocity cuts. We restrict our analysis to first-order subhalos (i.e., we exclude subhalos of subhalos); thus, for a disrupted subhalo to be included in our catalog, it must contribute to the host halo at z = 0. Operationally, each disrupted subhalo must have a descendant ID equal to the ID of a main-branch progenitor of the final host halo. Meanwhile, we define surviving subhalos as those that remain in the halo catalog at z = 0 and have a parent ID that is equal to the host ID, which similarly excludes higher-order subhalos.

To ensure that we study well-resolved subhalos, we restrict both disrupted and surviving subhalos to those with peak circular velocity  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  in our fiducial model, where  $V_{\text{peak}}$  is defined as the largest maximum circular velocity a subhalo attains along its entire main branch. This is a conservative choice; for example, GK17 presented velocity functions using the cut  $V_{\text{max}} > 5 \text{ km s}^{-1}$ , where  $V_{\text{max}}$  is the maximum circular velocity at z = 0. However, this cut ensures that we train our algorithm on subhalos that are consistent with those we will classify in a lower-resolution zoom-in simulation suite. By choosing a  $V_{\text{peak}}$  threshold rather than a  $V_{\text{max}}$ threshold, we also avoid biasing our subhalo selection with a redshift-dependent cut, since  $V_{\text{peak}}$ —unlike  $V_{\text{max}}$ —is not defined at a particular redshift. The  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  cut results in a combined total of 566 surviving subhalos and 872 disrupted subhalos from m12i and m12f, which we combine to form our fiducial training set. In the Appendix, we examine the impact of different training sets and minimum circular velocity cuts, and we present the results using the  $V_{\text{max}}$  cut employed in GK17 for comparison.

## 3. Random Forest Classification

# 3.1. Overview

We use the random forest algorithm from the package Scikit-Learn (Pedregosa et al. 2011) to classify disrupted and surviving subhalos. We refer the reader to the Scikit-Learn documentation for a detailed description of the algorithm, but we outline the most important aspects here. A random forest is a collection of decision trees, each of which is tuned to classify objects based on their input properties. Each tree in the forest is trained on a random sample of the training data with replacement, using a random subset of the input features at each split in the learning process, with the goal of

predicting the classes of the objects in the training set as accurately as possible according to some metric. For example, the default Scikit-Learn implementation minimizes the Gini impurity of the classifier's prediction. The random forest prediction for a given object is the majority vote of the tuned decision trees, while the classification probability is equal to the fraction of trees that predict a certain class. In this work, we label subhalos as either surviving until z = 0 or disrupted at some earlier time; thus, our random forest objects are subhalos, and our decision trees vote for whether each subhalo is disrupted or survives until z = 0. Note that our model does not explicitly include enhanced mass stripping due to baryonic effects, since we simply label subhalos as disrupted or surviving.

We train our classifier using the disrupted and surviving subhalos from m12i and m12f described above. We train on subhalo properties that depend on the entire history of each subhalo to avoid biasing the classifier by using properties defined at specific redshifts—for example, at z = 0 for surviving subhalos or at the final available redshift for disrupted subhalos. In particular, since we aim to classify subhalos in DMO halo catalogs that have survived to z = 0 but are likely to be disrupted in hydrodynamic resimulations, training our classifier with only present-day properties results in too many surviving subhalos because of the systematic evolution of subhalo properties over time.

Thus, we train on the following properties: pericentric distance and scale factor at first pericentric passage after accretion (dperi, aperi), and scale factor, virial mass, and maximum circular velocity at accretion ( $a_{acc}$ ,  $M_{acc}$ ,  $V_{acc}$ ). In principle, we could train the classifier on additional subhalo properties at pericenter or accretion; these properties could also include information about the host halo, such as subhalo scale radius in units of the host halo's scale radius. Indeed, random forests are well-suited to classifying objects using a large number of features because of the randomized nature of the training process, so we could even use *every* available subhalo property at pericenter and accretion to train the classifier. However, we will show that our five-property classifier performs very well, so we adopt this model to simplify our analysis and avoid overfitting the training data. In addition, we checked whether including the present-day properties  $V_{\text{max}}$  and  $M_{\rm vir}$  improves our classifier, finding that these properties are much less informative than features defined at pericenter or accretion. We discuss the correlations among the training features below, and we explore the feature selection in more detail in the Appendix.

We calculate the aforementioned subhalo features from the merger trees as follows. We define accretion as the last snapshot, working backward in time from z = 0 (for surviving subhalos) or the redshift of disruption (for disrupted subhalos), at which a subhalo's host ID is equal to the main halo's ID. Physically, this occurs when a subhalo enters the virial radius of the host halo for the final time.<sup>7</sup> We then take  $a_{acc}$ ,  $M_{acc}$ , and  $V_{acc}$  as the scale factor, virial mass, and maximum circular velocity at the time of accretion for each subhalo. We define pericenter as the first snapshot after accretion at which a subhalo reaches a local minimum in its three-dimensional distance from the center of the host halo. We inspected



Figure 1. Normalized joint and marginal distributions of pericentric distance and scale factor at first pericentric passage after accretion for surviving (blue) and disrupted (red) subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  in the m12i and m12f FIRE simulations. We select disrupted subhalos starting at a = 0.25 (z = 3).

individual subhalo orbits and determined that selecting the distance from the center of the host halo at the first snapshot after accretion at which a subhalo's separation from the host increases provides an accurate estimate of  $d_{peri}$ .<sup>8</sup> For subhalos that do not reach a local minimum in their separation from the host halo after accretion, we define  $d_{peri}$  as the instantaneous distance from the center of the host. In particular, for surviving subhalos on infalling orbits that have not experienced a pericentric passage by z = 0, we define  $d_{peri}$  as the distance from the host at z = 0. Analogously, for destroyed subhalos on infalling orbits that have not reached pericenter by the time of disruption, we define  $d_{peri}$  as the distance from the host at the time of disruption.

### 3.2. Choice of Subhalo Features

We choose the subhalo properties listed above because we expect them to correlate with subhalo disruption. Several of these properties are motivated by the results in GK17, which show that most of the subhalo disruption in m12i and m12f is caused by the central galactic disk in each simulation. For example, Figure 1 shows the joint and marginal distributions of  $d_{\text{peri}}$  and  $a_{\text{peri}}$  for disrupted and surviving subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  in m12i and m12f. Disrupted subhalos, shown in red, tend to have closer pericentric passages that occur at earlier times—or smaller values of  $a_{\text{peri}}$ —than their surviving counterparts, which are shown in blue. The  $d_{\text{peri}}$  distributions make sense physically; subhalos that pass close to the center of the host experience significant tidal forces due to the galactic disk and are therefore more likely to disrupt.<sup>9</sup> Next,

<sup>&</sup>lt;sup>7</sup> Note that a subhalo could have been contained within the host halo's virial radius at an earlier time and later reaccreted; we select the final accretion event for each subhalo.

 $<sup>\</sup>frac{8}{8}$  Given a spacing of ~25 Myr between halo catalog snapshots and a generous subhalo orbital velocity of ~300 km s<sup>-1</sup> at pericenter, the uncertainty in  $d_{peri}$  is only ~8 kpc.

<sup>&</sup>lt;sup>9</sup> GK17 found that the amount of disruption is largely insensitive to the shape and mass of the central disk, so subhalo disruption in these simulations is at least partly due to disk shocking rather than tidal stripping.



**Figure 2.** Normalized joint and marginal distributions of maximum circular velocity and virial mass at accretion for surviving (blue) and disrupted (red) subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  in the m12i and m12f FIRE simulations. We select disrupted subhalos starting at a = 0.25 (z = 3); note that  $V_{\text{acc}} < V_{\text{peak}}$  for subhalos that are stripped prior to infall (e.g., see Behroozi et al. 2014).

consider the  $a_{peri}$  dependence: subhalos that reach pericenter earlier have relatively low masses at pericenter and tend to experience more pericentric passages, both of which contribute to enhanced disruption. Although  $a_{peri}$  and  $a_{acc}$  are somewhat degenerate properties, we find that including  $a_{acc}$  improves our results, likely because subhalos accreted at higher redshifts are tidally stripped for longer periods of time, making them more susceptible to disruption.

Figure 2 illustrates the  $V_{\rm acc}$  and  $M_{\rm acc}$  distributions for disrupted and surviving subhalos in m12i and m12f. Interestingly, even though these features mainly contain information about internal rather than orbital subhalo properties, they are useful for identifying disrupted subhalos; as we show below, these properties account for 16% of the total feature importance score for our fiducial five-property classifier. At the low-mass end of the subhalo population, subhalos with lower values of  $V_{\rm acc}$  are more likely to be disrupted. In particular, the survival of low-mass subhalos at fixed  $M_{\rm acc}$  is dictated by tidal effects that preferentially disrupt lowerconcentration subhalos, i.e., subhalos with smaller values of  $V_{\rm acc}$  at fixed  $M_{\rm acc}$ . However, at the high-mass end of the subhalo population, subhalos with larger values of  $V_{\rm acc}$  are more likely to be disrupted. This behavior suggests that baryonic mechanisms, in addition to the tidal effects of the central disk, contribute to subhalo disruption in the FIRE simulations. Specifically, it is plausible that  $V_{\rm acc}$  and  $M_{\rm acc}$ encode information about stellar feedback, which can soften central density cusps. In particular, we expect high-mass subhalos with larger values of  $V_{\rm acc}$  to host more massive galaxies and experience more significant baryonic feedback, i.e., high-mass subhalos with larger values of  $V_{\rm acc}$  are more likely to be disrupted. Thus, even though  $M_{\rm acc}$  and  $V_{\rm acc}$  are highly correlated, it is useful to train on both properties because subhalo concentration determines  $V_{\rm acc}$  at fixed  $M_{\rm acc}$  and provides physical information about whether a subhalo subject to given tidal forces is disrupted. The advantage of random forest classification is that it captures these complex relationships between subhalo properties and subhalo disruption.

## 3.3. Training and Validation

To train our classifier, we use the GridSearchCV function to search the space of random forest hyperparameters and select the ones that yield the highest out-of-bag (OOB) classification score averaged over ten cross-validation folds of the training data.<sup>10</sup> These hyperparameters include the number of trees in the forest, the depth of each tree, the maximum number of features used by each tree, and the loss function. We train the classifier using a randomly selected 75% of the disrupted and surviving subhalos from our fiducial training set, with replacement. The number of folds and the ratio of the test-training split do not affect our results. The raw percentage of subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$ from the hydrodynamic m12i and m12f simulations that are identified correctly by our classifier is 95%. We cross-validate this result by computing the OOB classification score, which is defined as the percentage of subhalos from the training data that the random forest classifies correctly when each tree does not vote on subhalos in its own training set. The optimal OOB score for our fiducial five-property classifier is 85%, and we find that at least 20 trees are needed to achieve this OOB score. Our classifier therefore identifies subhalos accurately, although the gap between the overall classification accuracy and the OOB scores suggests that we mildly overfit the training data. In particular, the raw accuracy is higher than the OOB score because decision trees are allowed to vote on subhalos within their respective training sets when classifying all subhalos. To illustrate the relative importance of each subhalo feature, Table 1 shows the OOB score along with the percentage of correct and incorrect predictions for subhalos in the test set, which is the set of all subhalos that are not included in the training set, for five different classifiers. We calculate these scores for each classifier by using the hyperparameters determined by GridSearchCV and averaging the results over 100 test-training splits. Each row of Table 1 lists the results for a classifier trained using an additional subhalo feature; as we add training features, the OOB score and the total classification accuracy generally improve. Note that there are more disrupted subhalos than surviving subhalos in our fiducial training set, so the raw classification accuracy for each set of features is higher than the mean classification accuracy inferred from Table 1. Thus, while the classification accuracy for surviving subhalos decreases when  $M_{\rm acc}$  and  $V_{\rm acc}$  are added, the increase in classification accuracy for disrupted subhalos outweighs this effect. We emphasize, however, that  $d_{peri}$ ,  $a_{peri}$ , and  $a_{acc}$  contain most of the information about subhalo disruption in m12i and m12f.

Next, we examine our classifier's receiver operating characteristic (ROC) curve, which illustrates the rate of true-versus false-positive classifications for subhalos in the test set. The ROC curve for our five-property classifier is shown in Figure 3. The red and black lines illustrate perfect (100% true-positive rate) and random (true-positive rate equal to false-positive rate) classifiers. We quantify our classifier's performance by calculating the area under the ROC curve (AUC), which confirms that the random forest classifies subhalos in the

<sup>&</sup>lt;sup>10</sup> In *n*-fold cross-validation, the training set is divided into *n* subsets of equal size; n - 1 of these subsets are used for training, the remaining subset is used for cross-validation, and this procedure is repeated once for each possible cross-validation subset.



**Figure 3.** True- vs. false-positive classification rate for our fiducial fiveproperty random forest classifier, which is trained on subhalos from the m12i and m12f FIRE simulations with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$ . These classification rates apply to subhalos that are not included in the training set. The AUC is equal to 1 for a perfect classifier (red), 0.93 for our random forest classifier (blue), and 0.5 for a random classifier (black).

Table 1Performance Metrics for Five Different Random Forest Classifiers Trained onDisrupted and Surviving Subhalos from the m12i and m12f FIRE Simulationswith  $V_{peak} > 10 \ {\rm km \ s^{-1}}$ 

Training Features	OOB Score	Classification Accur- acy (Disrupted)	Classification Accur- acy (Surviving)
d <sub>peri</sub>	72%	$80\%\pm3\%$	$58\%\pm4\%$
$d_{\rm peri}, a_{\rm peri}$	82%	$88\%\pm2\%$	$72\%\pm3\%$
$d_{\rm peri}, a_{\rm peri}, a_{\rm acc}$	85%	$87\%\pm2\%$	$82\%\pm4\%$
$d_{ m peri}, a_{ m peri}, a_{ m acc}, M_{ m acc}$	85%	$88\%\pm2\%$	$81\%\pm3\%$
	85%	89% ± 2%	$80\%\pm4\%$

Note. The first column lists the subhalo features used to train each classifier. The second column lists the OOB classification score, which is the percentage of subhalos in the training data identified correctly when each tree does not vote on subhalos in its own training set. The third and fourth columns list the percentage of disrupted and surviving subhalos in the test set that are identified correctly by each classifier, averaged over 100 test-training splits. The test set is the collection of subhalos from the m12i and m12f FIRE simulations with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  that are not included in the training set. We also indicate the standard deviation of each classification accuracy. Note that the ratio of disrupted to surviving subhalos in our fiducial halo catalog is roughly 3:2.

FIRE simulations accurately: its AUC is 0.93, while a random classifier has an AUC equal to 0.5, and a perfect classifier has an AUC equal to 1. Note that Figure 3 shows the ROC curve for a particular test-training split, but the scatter in the ROC curves for different test-training splits is small.

### 3.4. Importance of Subhalo Features

The random forest algorithm determines the feature importance of the various subhalo properties included in the training process. The feature importance indicates the relative importance of each



Figure 4. Feature importance scores for the five subhalo properties used to classify disrupted and surviving subhalos in the m12i and m12f FIRE simulations. The colored bars above each property indicate the feature importance scores averaged over 100 test-training splits when that property is added to the training features. Thus, the columns correspond to the five different classifiers in Table 1. For a given classifier, each property's score indicates its relative importance for classifying disrupted and surviving subhalos. Here  $d_{peri}$  and  $a_{peri}$  are the pericentric distance and scale factor at first pericentric passage after accretion, and  $a_{acc}$ ,  $M_{acc}$ , and  $V_{acc}$  are the scale factor, virial mass, and maximum circular velocity at accretion.

property for predicting whether a given subhalo is disrupted or survives until z = 0. In particular, a property's feature importance score is the suitably normalized change in the OOB classification score when the property is randomly shuffled among the subhalos in the training set. Thus, the property with the highest feature importance score is the most important for classifying disrupted and surviving subhalos in the m12i and m12f simulations. Figure 4 illustrates the mean feature importance scores for each classifier listed in Table 1; for a given classifier, the scores are averaged over 100 test-training splits, and the same hyperparameters are used for each realization. For our fiducial five-property classifier, which corresponds to the fifth column of Figure 4, we find mean feature importance scores of 0.28 for  $d_{peri}$ , 0.21 for  $a_{peri}$ , 0.35 for  $a_{acc}$ , 0.08 for  $M_{\rm acc}$ , and 0.08 for  $V_{\rm acc}$ . The variance in the feature importance scores for different test-training splits is small, and the scores depend very weakly on the random forest hyperparameters.

Figure 4 shows that pericentric distance is an important property for determining whether a given subhalo is disrupted; subhalos with close pericentric passages are more likely to be destroyed. The scale factors at accretion and at first pericentric passage after accretion are also important features. In particular, subhalos that accrete and reach pericenter earlier are preferentially disrupted. The fact that  $a_{\rm acc}$  has the highest feature importance score suggests that the number of pericentric passages, rather than the distance and scale factor associated with each individual passage, is most strongly correlated with subhalo disruption. However, we note that interpreting the feature importance scores for  $d_{peri}$  and  $a_{peri}$  is complicated by the fact that we defined these properties as the instantaneous distance and scale factor at the final available snapshot for subhalos on infalling orbits that have not reached their true pericenter. The true pericenters for such subhalos occur at smaller values of  $d_{peri}$  and larger values of  $a_{peri}$  than we have assigned here; in a more detailed analysis, we would need to calculate these features by fitting individual subhalo orbits. However, the fraction of disrupted (surviving) subhalos in our

fiducial training set that have not reached their true pericenter by the time of disruption (z = 0) is only 17% (20%), so the feature importance for  $d_{peri}$  and  $a_{peri}$  is reasonably accurate.

### 3.5. Model Limitations

Finally, we note that our classification method, like any other model, has its limitations. In particular,

- our classifier is only trained on two zoom-in simulations of MW-mass host halos with a specific baryonic physics prescription, and thus it is not clear how well our algorithm will perform on subhalo populations associated with higher- or lower-mass host halos;
- 2. neither of the hosts that we train on experience a recent major merger, so our classifier might not apply to halos with significantly different formation histories;
- 3. both hosts form a central galactic disk that is responsible for most of the subhalo disruption, so our classifier mainly captures the dynamical effects of a central disk.

We discuss these limitations in more detail and comment on how they might affect our results in the following section.

### 4. Results

We now present our main results. In Section 4.1, we use our classifier to identify subhalos in DMO simulations of m12i and m12f that are likely to be disrupted in hydrodynamic resimulations. We analyze our predicted surviving subhalo populations by comparing the velocity functions and radial distributions to those from the FIRE, DISK, and DMO simulations in GK17. In Section 4.2, we predict surviving subhalo populations for the suite of DMO zoom-in simulations of MW-mass host halos from Mao et al. (2015), and we study the resulting velocity functions, radial distributions, and implications for satellite searches.

# 4.1. Predictions for DMO Simulations of the FIRE Halos

## 4.1.1. Subhalo Feature Distributions

There are about twice as many surviving subhalos at z = 0 in the DMO simulations of m12i and m12f as in the corresponding hydrodynamic simulations. As we have discussed, we expect many of these subhalos to be disrupted due to baryonic effects, including stellar feedback, enhanced tidal stripping, and disk shocking, and our random forest classifier can identify such subhalos based on their internal and orbital properties. In particular, to identify subhalos in the m12i and m12f DMO simulations that are likely to be disrupted by baryonic effects, we select subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  at z = 0, and we use our trained classifier to predict whether these subhalos should have been destroyed at some earlier time using their values of  $d_{\text{peri}}$ ,  $a_{\text{peri}}$ ,  $a_{\text{acc}}$ ,  $M_{\text{acc}}$ , and  $V_{\text{acc}}$ . Note that this method does not require matching subhalos between DMO and hydrodynamic simulations.

Figure 5 shows the joint and marginal distributions of  $d_{\text{peri}}$  and  $a_{\text{peri}}$  for surviving subhalos from the m12i and m12f DMO simulations predicted by our random forest classifier. The random forest predicts a surviving subhalo population in  $d_{\text{peri}}-a_{\text{peri}}$  space that agrees well with the hydrodynamic data; we also find good agreement in the spaces defined by the other subhalo features. Of course, since our classifier is trained on subhalos from the m12i and m12f FIRE simulations, we expect it to perform particularly well on the corresponding DMO simulations, which have



Figure 5. Normalized joint and marginal distributions of pericentric distance and scale factor at first pericentric passage after accretion for surviving subhalos in the m12i and m12f FIRE simulations (blue); surviving subhalos from the corresponding DMO simulations are shown in green. The unfilled contour and purple histograms show the most probable surviving subhalo population from the m12i and m12f DMO simulations predicted by our random forest classifier.

identical initial conditions. Nevertheless, these results are encouraging: even though there is no galactic disk or stellar feedback in the DMO simulations, our classifier efficiently predicts subhalo populations that are in good agreement with the hydrodynamic results. In particular, once the classifier has been trained on the hydrodynamic simulations, it can immediately predict surviving subhalo populations from DMO halo catalogs. Simulations that include baryonic effects by hand, such as the DISK simulations presented in GK17, are complementary to our approach, since they provide more direct physical modeling at the expense of increased computational costs.

In general, at least three mechanisms contribute to enhanced subhalo disruption in the m12i and m12f hydrodynamic simulations relative to the DMO simulations: tidal effects due to the central galactic disk, stellar feedback, and characteristic changes in subhalo orbits due to the presence of baryons. The results from the DISK simulations in GK17 indicate that the central disk is the main source of subhalo disruption in these simulations, but the frequency of disruption events might be enhanced by stellar feedback, which can soften central density cusps (Governato et al. 2012; Pontzen & Governato 2012; Zolotov et al. 2012; Di Cintio et al. 2014); as noted above, we multiply all circular velocities in the DMO and DISK simulations by a factor of  $\sqrt{1 - f_b}$  to approximate this effect. Meanwhile, Zhu et al. (2017) analyzed the orbital properties of subhalos in hydrodynamic and DMO zoom-in simulations of an MW-mass host halo from the Aquarius Project and found that the distributions of subhalos in different orbital families change when baryons are included. It is difficult to assess the importance of the characteristic differences in internal and orbital subhalo properties between hydrodynamic and DMO simulations in general; however, the fact that we

predict subhalo feature distributions starting from DMO halo catalogs that agree with hydrodynamic results suggests that these effects are relatively unimportant.

### 4.1.2. Subhalo Counts

Having shown that we can predict the feature distributions of surviving subhalos from DMO simulations of m12i and m12f, we turn to our predictions for the number of surviving subhalos as a function of various properties. In Figure 6, we present our predictions for the m12i and m12f velocity functions; the top panels show the velocity functions evaluated using  $V_{\text{max}}$ , and the bottom panels show the velocity functions evaluated using  $V_{\text{peak}}$ . The blue lines show the most probable surviving subhalo populations predicted by our random forest algorithm for each host halo; we also plot the FIRE, DISK, and DMO results for comparison. We restrict the velocity functions to subhalos within 300 kpc of the center of their respective host at z = 0, since this roughly corresponds to the virial radii of m12i and m12f. Similarly, Figure 7 shows the distribution of tangential and radial orbital velocities for subhalos within 300 kpc of their respective host at z = 0, and Figure 8 shows the radial distribution of surviving subhalos at z = 0 within each host halo. In Figures 6–8, we only include subhalos with  $V_{\text{peak}} > 10$ km  $s^{-1}$  to match the cut used in our fiducial training set. The bottom panels in these figures show the number of surviving subhalos predicted by the most probable realization of our random forest classifier divided by the number of subhalos found in each hydrodynamic simulation. We also plot the Poisson error associated with the random forest predictions as shaded areas in each figure. In the Appendix, we show that the intrinsic scatter in the random forest predictions is small.

There are several interesting aspects of Figures 6–8 that are worth exploring. Our random forest algorithm predicts that the amount of substructure in each host is significantly reduced relative to the DMO simulations, bringing the velocity functions and radial distributions into good agreement with the FIRE results. The random forest predictions for the velocity functions are comparable to the DISK simulations at low velocities, which indicates that the effects of the disk are largely encoded in the subhalo properties that we use to train our classifier, at least for subhalos with low values of  $V_{\text{max}}$  or  $V_{\text{peak}}$ . However, the random forest predicts more subhalo disruption than the DISK simulations for  $V_{\text{max}} \gtrsim 15$  or  $V_{\text{peak}} \gtrsim 20 \,\mathrm{km \, s^{-1}}$  and generally matches the FIRE results more closely in these regimes. The minor discrepancies for  $V_{\text{max}} \gtrsim 15 \,\mathrm{km \, s^{-1}}$  are likely caused by enhanced mass stripping due to baryonic effects, which would shift the predictions toward smaller velocities at high  $V_{\text{max}}$ .

Our predicted radial distributions are also generally comparable to the DISK simulations; however, for 30 kpc  $\leq r \leq 100$  kpc, where the disk should be particularly effective at disrupting subhalos, our classifier predicts more subhalo disruption than the m12i DISK simulation and matches the FIRE results more closely for both hosts. Finally, Figure 7 shows that our classifier predicts a substantial reduction in the number of subhalos with low tangential velocities, even though it is not explicitly trained on orbital velocities. Our predicted tangential and radial velocity distributions are similar to the DISK results for m12i, while we slightly overpredict the number of high- $V_{tan}$  and high- $V_{rad}$ subhalos for m12f. Comparing our predictions to the DISK simulations is a particularly useful way to assess whether our classifier captures baryonic physics beyond the dynamical effects of a central galactic disk, since the DISK simulations do not modify internal subhalo properties. Thus, Figures 6–8 suggest that our classifier captures both the tidal effects of a disk and additional baryonic processes that contribute to subhalo disruption.

Our random forest classifier predicts that many subhalos with large values of  $V_{\text{max}}$  and  $V_{\text{peak}}$  should be disrupted, while these subhalos are not necessarily destroyed in the DISK simulations (see Figures 6 and 14). These subhalos either orbit at large radii, so that they are not significantly affected by the disk, or they are too tightly bound to be disrupted by the disk alone. We find that 45% (84%) of the disrupted subhalos from m12i and m12f with  $V_{\text{peak}} > 20 \text{ km s}^{-1}$  have pericentric passages within 50 kpc (100 kpc) of their respective hosts. The disk does not seem to be the main factor that contributes to the destruction of the remaining subhalos, though a combination of stellar feedback and tidal forces could lead to their disruption. Interestingly, the region of the  $V_{\text{max}}$  and  $V_{\text{peak}}$  functions where we predict enhanced subhalo disruption relative to the DISK simulations ( $V_{\text{max}} \gtrsim 15$  and  $V_{\text{peak}} \gtrsim 20 \,\text{km s}^{-1}$ ) corresponds to the regime where baryonic physics can efficiently create cored subhalo density profiles (Chan et al. 2015; Tollet et al. 2016; Fitts et al. 2017). It is also intriguing that our classifier predicts both the  $V_{\text{max}}$  and  $V_{\text{peak}}$  functions accurately, even though it does not account for enhanced mass stripping beyond the  $\sqrt{1-f_b}$  circular velocity correction, which does not reproduce the hydrodynamic results on its own (for example, compare the "Raw DMO" and "DMO" curves in GK17). Since  $V_{\text{max}}$  and  $V_{\text{peak}}$  are proxies for satellite luminosity, our method can therefore be extended to predict satellite galaxy populations associated with MW-mass host halos (see Figure 11); in addition, it can be used to constrain the cumulative mass functions of dark and luminous substructures relevant to gravitational-lensing analyses. Clearly, a more diverse training sample is required in order to make robust predictions regarding the populations of satellite galaxies around the MW and MW analogs from the SAGA survey. Nonetheless, Figures 6 and 8 show that classification algorithms can predict subhalo populations in good agreement with hydrodynamic simulations, providing an efficient way to explore the range of possible satellite galaxy populations associated with a particular host halo.

One could argue that the efficiency of our approach is outweighed by the fact that we must train our classifier on computationally expensive hydrodynamic simulations in order to predict surviving subhalo populations for corresponding DMO simulations. However, as we demonstrate in the following section, our method can be used to predict surviving subhalo populations when hydrodynamic simulations are unavailable. Of course, the surviving subhalo populations we predict in this paper are specific to the FIRE simulations that we use to train our classifier. Nonetheless, even though the generality of our results is limited by the small size of our training set, our work suggests that random forest classification can be used to predict realistic subhalo populations given a sufficiently diverse sample of hydrodynamic training simulations. In addition, we emphasize that our classifier is trained on simulations that yield satellite populations that are consistent with the observed mass functions and velocity dispersion functions for satellites of the MW and M31.

## 4.2. Predictions for a Suite of DMO Zoom-in Simulations

### 4.2.1. Subhalo Counts

We now use our classifier to identify subhalos from a suite of independent DMO zoom-in simulations that are likely to be disrupted in hydrodynamic resimulations. In particular, we



**Figure 6.** Velocity functions for subhalos hosted by m12i (left) and m12f (right), predicted from DMO simulations of these hosts by our random forest classifier (blue). The top panels show velocity functions evaluated using the maximum circular velocity at z = 0, and the bottom panels show velocity functions evaluated using the peak circular velocity  $V_{\text{peak}}$ . Our classifier is trained on subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  from both hydrodynamic simulations. The FIRE (red), DISK (dot-dashed), and DMO (dashed) results are shown for comparison; recall that we scale circular velocities in the DISK and DMO halo catalogs by a factor of  $\sqrt{1 - f_b}$ . Dotted lines show the DMO results scaled by a factor of 1/2 for comparison. We restrict these velocity functions to subhalos within 300 kpc of their respective hosts at z = 0. The bottom panels show the random forest and  $N_{\text{FIRE}}$ , where  $N_{\text{pred}}$  is the number of surviving subhalos predicted by the random forest and  $N_{\text{FIRE}}$  is the number of subhalos in each FIRE simulation. Shaded areas show the standard deviation about the most probable random forest prediction for 1000 draws from a Poisson distribution with a mean value of  $N_{\text{pred}}$  at each value of  $V_{\text{max}}$  or  $V_{\text{peak}}$ .

predict surviving subhalo populations for the 45 zoom-in simulations of MW-mass host halos from Mao et al. (2015). We refer the reader to Mao et al. (2015) for a detailed description of the simulations, but we briefly highlight the most important aspects for this work. The host halos lie in the mass range  $M_{\rm vir} = 10^{12\pm0.03} M_{\odot}$  and have a variety of formation histories; we plot the mass accretion histories for these hosts in

Figure 9. Note that m12i and m12f have formation histories that are consistent with these host halos, so we expect our model to perform well on this simulation suite. Of course, our model would not accurately predict subhalo disruption for hosts with significantly different formation histories due to the limited size of our training set. The zoom-in simulations were run at a lower resolution than the DMO simulations of m12i



Figure 7. Distributions of tangential orbital velocities (top) and radial orbital velocities (bottom) for subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  hosted by m12i (left) and m12f (right) at z = 0, predicted from DMO simulations of these hosts by our random forest classifier (blue). The classifier is trained on subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  from both FIRE simulations. We restrict these distributions to subhalos within 300 kpc of their respective hosts at z = 0. The various curves and panels are described in Figure 6.

and m12f; the dark matter particle mass is  $3 \times 10^5 M_{\odot}$ , and Mao et al. (2015) estimated that  $V_{\rm max} \sim 9 \,\rm km \, s^{-1}$  is a conservative lower limit for the subhalo circular velocity resolution. Halo catalogs and merger trees were generated using the ROCKSTAR halo finder and the consistenttrees merger code (Behroozi et al. 2013a, 2013b). Again, we scale all subhalo masses by  $1 - f_b$  and all subhalo circular velocities by  $\sqrt{1 - f_b}$  in our post-processing of the halo catalogs. As noted at the end of Section 1, the cosmological parameters for these simulations are slightly different than

those used in the FIRE simulations, and we adjust the parameters in our analysis accordingly.

In Figure 10, we plot the maximum circular velocity functions and radial distributions for the subhalo populations from this simulation suite, along with those predicted by the most probable realization of our random forest classifier for each simulation. We also plot the results from the m12i and m12f FIRE simulations, along with the mean DMO curves scaled by a constant factor, for comparison. In particular, we scale the mean DMO curves by a factor of 2/3 so that the average number of



**Figure 8.** Radial distributions of subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  hosted by m12i (left) and m12f (right) at z = 0, predicted from DMO simulations of these hosts by our random forest classifier (blue). The classifier is trained on subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  from both FIRE simulations, and the various curves and panels are described in Figure 6. The scaled DMO curve overpredicts the number of surviving subhalos at small radii by an order of magnitude, highlighting the enhanced subhalo disruption in the inner regions of the hydrodynamic simulation due to the central disk.



Figure 9. Mass accretion histories for the suite of DMO zoom-in simulations of MW-mass host halos presented in Mao et al. (2015). The black line shows the mean mass accretion history for the 45 hosts, and the shaded area shows the associated  $\pm 1\sigma$  standard deviation. Mass accretion histories for the m12i and m12f FIRE simulations are shown in orange and green, respectively.

subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  and r < 300 kpc matches the mean random forest prediction. The random forest predictions were generated using the method described above. We classify subhalos in each zoom-in simulation using the features  $d_{\text{peri}}$ ,  $a_{\text{peri}}$ ,  $a_{\text{acc}}$ ,  $M_{\text{acc}}$ , and  $V_{\text{acc}}$ , and we restrict the velocity functions to subhalos within 300 kpc of their respective host at z = 0. We plot the most probable realization of the random forest prediction for each host. The intrinsic scatter in our random forest predictions is small.

Figure 10 shows that the reduction in the total number of subhalos predicted by our random forest classifier is larger than the host-to-host scatter for the subhalo populations in these zoomin simulations. In particular, our classifier predicts that the total number of subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  and r < 300 kpc is reduced by a factor of 2/3, while the  $1\sigma$  host-to-host scatter corresponds to an 87% reduction at most. This suggests that subhalo disruption due to baryonic effects, such as stellar feedback and the tidal influence of a central galactic disk, should not be neglected in semianalytic models that use the subhalo populations predicted by these DMO simulations as input. In particular, for MW-mass host halos that contain a central galactic disk similar to those found in the m12i and m12f FIRE simulations, the reduction in substructure due to the disk and other baryonic processes is larger than the scatter in subhalo abundance from host to host, so the impact of baryonic physics cannot be accounted for simply by marginalizing over the subhalo populations of host halos with a range of formation histories.

While the average amount of subhalo disruption is larger than the host-to-host scatter among the subhalo populations in these simulations, the impact of baryons on individual subhalo populations is largely consistent. In particular, our classifier predicts that the hosts with the most subhalos tend to have the largest number of surviving subhalos once baryonic effects are taken into account. Moreover, the number of DMO subhalos and predicted number of surviving subhalos above different  $V_{\rm max}$  thresholds and within various hostcentric radii are highly correlated for this simulation suite. For example, the Spearman rank correlation coefficient between the number of surviving subhalos with  $V_{\text{peak}} > 10 \,\text{km s}^{-1}$  and  $r < 300 \,\text{kpc}$  predicted by the DMO simulations and our classifier is 0.74. This implies that the shapes of the velocity functions and radial distributions are not strongly affected by baryonic physics; indeed, the scaled DMO curves in Figure 10 are very similar to the random forest predictions, except at small radii, where subhalos are preferentially disrupted in the training data. The fractional amount of subhalo disruption is also consistent among the



**Figure 10.** Velocity functions (top) and radial distributions (bottom) for the suite of 45 zoom-in simulations of MW-mass host halos presented in Mao et al. (2015). The thick lines show the mean number of subhalos predicted by the DMO simulations (black) and our random forest classifier (blue), which is trained on the m12i and m12f FIRE simulations; the shaded areas show the  $\pm 1\sigma$  standard deviation of these predictions. The thin lines show the DMO result and the most probable random forest prediction for each host. The thick dotted lines show the mean DMO velocity function and radial distribution scaled by a factor of 2/3 for visual comparison, and the orange and green lines show the results for m12i and m12f, respectively. Note that the scaled DMO line in the top panel is mostly obscured by the random forest vs. the number of such subhalos in the corresponding DMO simulations. The thick dotted line in the inset shows the constant fraction of surviving subhalos corresponding to the scaled DMO curves, and the thin dash-dotted line shows a 1: 1 relationship for comparison. The vertical line at  $V_{max} = 9 \text{ km s}^{-1}$  in the top panel represents a conservative resolution limit for these simulations.

zoom-in simulations. In particular, the number of predicted surviving subhalos with  $V_{\text{max}} > 10 \text{ km s}^{-1}$  and r < 300 kpcfor all 45 hosts is given by  $N_{\text{pred}}/N_{\text{DMO}} = 0.65 \pm 0.09$ . To illustrate these results, the inset in the bottom panel of Figure 10 shows the number of predicted surviving subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$  and r < 300 kpc for each host versus the corresponding number of subhalos in each DMO simulation. The inset shows that the random forest predictions are consistent with an overall scaling of the DMO subhalo

populations. Thus, subhalo disruption due to baryonic effects can be parameterized rather simply for these host halos in the context of our disruption model. We leave a detailed exploration of such a parameterization to future work informed by a wider range of hydrodynamic simulations, but we note that a simple one-parameter rescaling would not be sufficient to model subhalo disruption in detail; for example, Figure 10 shows that the shape of the mean radial subhalo distribution is somewhat altered by baryonic physics. Finally, we note that



**Figure 11.** Luminosity functions for the DMO zoom-in simulations presented in Mao et al. (2015; black), inferred using the  $V_{\text{peak}}$ -luminosity abundancematching relation tuned to the *r*-band luminosity function from the GAMA survey (Loveday et al. 2015). We do not apply scatter in the  $V_{\text{peak}}-M_{r,o}$ relation to highlight the host-to-host variability. The blue line shows our mean prediction for the luminosity function of surviving satellites, and shaded areas show  $\pm 1\sigma$  and  $\pm 2\sigma$  standard deviations. We also plot luminosity functions for m12i (orange) and m12f (green). Here  $M_{r,o}$  is the observed *r*-band luminosity, and the vertical line corresponds to the completeness limit of the SAGA survey.

our random forest classifier predicts that these zoom-in simulations typically contain more high- $V_{\text{max}}$  subhalos than m12i or m12f and more subhalos at small radii than m12i. Determining whether these differences represent statistical fluctuations or systematic differences between the FIRE simulations and this simulation suite would require a larger sample of hydrodynamic results for comparison.

## 4.2.2. Implications for Satellite Searches

Our model, when applied to MW-size zoom-in simulations, suggests that MW-mass host halos are somewhat less likely to host bright satellite galaxies such as the Magellanic Clouds and that they have more extended radial satellite profiles than those inferred from DMO simulations. At face value, both of these predictions seem to be in tension with observations of MW satellites (e.g., see Kim et al. 2017). However, the MW itself could be an outlier, so here we also examine our model's predictions for the satellite populations of MW analogs.

To estimate the impact of baryonic physics on the luminosity functions of MW analogs, which can be compared to the results of the SAGA survey, we use the V<sub>peak</sub>-luminosity abundancematching relation tuned to the r-band luminosity function from the GAMA galaxy survey (Loveday et al. 2015); we refer the reader to Geha et al. (2017) for details on the abundance-matching procedure. Figure 11 shows the resulting luminosity functions for the MW zoom-in suite, along with the luminosity functions for these hosts predicted by our random forest classifier. We neglect the scatter in the  $V_{\text{peak}}$ -luminosity relation for this simple estimate because the host-to-host scatter among the zoom-in simulations is larger than the intrinsic scatter in the luminosity function introduced by abundance matching. Our classifier predicts a significant reduction in the number of bright satellites associated with MW analogs; the number of satellites with observed r-band magnitudes  $M_{r,o} < -12.3$  inferred from the DMO simulations is  $3.0 \pm 1.6$ , while our random forest predicts that only  $1.5 \pm 1.3$  such satellites exist. Although these estimates of surviving satellite populations are simplistic, it will be interesting to compare predictions informed by hydrodynamic simulations to observational results as the number of systems with high completeness limits improves.

### 5. Conclusions and Discussion

To conclude, we summarize our main results as follows.

- 1. We train a random forest classifier on disrupted and surviving subhalos in two hydrodynamic zoom-in simulations of MW-mass host halos from the FIRE project using five properties of each subhalo:  $d_{peri}$ ,  $a_{peri}$ ,  $a_{acc}$ ,  $M_{acc}$ , and  $V_{acc}$ .
- 2. Our classifier identifies subhalos in the FIRE simulations with an 85% OOB classification score and predicts surviving subhalo populations from DMO simulations of these hosts that are in excellent agreement with the hydrodynamic results, often outperforming the DMO-plus-disk simulations presented in Garrison-Kimmel et al. (2017).
- We argue that our classifier captures the effects of the central galactic disks that develop in the FIRE simulations, in addition to other baryonic disruption mechanisms such as stellar feedback.
- 4. We use our classifier to predict surviving subhalo populations for the suite of DMO zoom-in simulations of MW-mass host halos presented in Mao et al. (2015), finding that the average amount of subhalo disruption is larger than the host-to-host scatter; however, the baryonic impact on each subhalo population is largely consistent, with  $N_{\rm pred}/N_{\rm DMO} = 0.65 \pm 0.09$  for subhalos with  $V_{\rm peak} > 10 \ {\rm km \ s^{-1}}$  and  $r < 300 \ {\rm kpc}$ .

We refer the reader to the end of Section 3 for a summary of the limitations of our classification method.

There are several interesting avenues for future work. For example, since we find that the average amount of subhalo disruption due to baryonic physics is larger than the host-to-host scatter among the suite of zoom-in simulations analyzed above, this characteristic reduction in the number of subhalos should be taken into account when marginalizing over the effects of baryonic physics for MW-mass host halos that contain a central galactic disk. Thus, it is plausible that the reduced number of surviving subhalos will change the conclusions drawn from semianalytic models that use the subhalo populations predicted by such simulations (e.g., Lu et al. 2016, 2017).

Another potential application of our results concerns the radial segregation of dark matter subhalos with respect to various subhalo properties. Subhalo segregation, as studied by van den Bosch et al. (2016), directly depends on the subhalo populations predicted by DMO simulations. Since subhalo populations that have been altered by baryonic effects systematically differ from those predicted by DMO simulations, subhalo segregation could be affected by baryonic physics, and our classifier provides an efficient method for predicting surviving subhalo populations in order to explore this possibility. Meanwhile, resolving the disruption of individual subhalos in detail is an important challenge for current simulations; for example, van den Bosch (2017) estimated that 80% of all subhalo disruption in the Bolshoi simulation is numerical, rather than physical. The Latte simulations have  $\sim 4000 \times$  smaller dark matter particle mass and  $\sim$ 35× smaller dark matter force softening than Bolshoi, so these effects are likely much less severe, but it is nevertheless worth exploring whether artificial disruption persists in high-resolution

hydrodynamic simulations and how these numerical artifacts might influence our results.

The algorithm presented in this paper is extremely simple, using only five subhalo properties as training features. Nevertheless, these properties encode the majority of the information about subhalo disruption in the ml2i and ml2f FIRE simulations, yielding a classifier that predicts surviving subhalo populations from DMO simulations that are in excellent agreement with hydrodynamic results. Of course, as the number of hydrodynamic training simulations grows, it will be worthwhile to explore more sophisticated classification algorithms and study the feature selection in more detail. It will be interesting to assess how well a classifier can perform in principle, since there are characteristic differences between DMO and hydrodynamic simulations, including changes in subhalo orbits due to the presence of barvons (Zhu et al. 2017), that our simple model cannot capture. Our results hint that these characteristic differences are relatively unimportant, but further tests should be performed using a larger sample of training simulations.

As more high-resolution zoom-in simulations become available, it will become feasible to train classifiers on increasingly diverse data sets, allowing for more robust predictions. Once a classifier has been trained on a wide variety of hydrodynamic simulations, it can predict a range of surviving subhalo populations associated with different central galaxy types and halo formation histories directly from DMO simulations. It is worth exploring whether these predictions can be used as input for neural networks in order to generate large samples of mock halo catalogs, perhaps eliminating the need for certain types of simulations entirely.

Machine-learning algorithms have the potential to identify large samples of realistic subhalo populations that can be used as input for models that populate subhalos with galaxies. Comparing the surviving subhalo populations predicted by such algorithms for host halos on different mass scales could provide insight into the original TBTF problem for MW-mass systems and analogous problems for host halos of different masses. Moreover, comparing the results of classification algorithms that are trained on hydrodynamic simulations with different implementations of baryonic physics would be a promising step toward parameterizing the impact of baryons on the abundance and properties of dark matter subhalos.

We have made our code and trained classifier publicly available at github.com/ollienad/subhalo\_randomforest; please contact the authors with data requests. We thank Frank van den Bosch and Andrew Hearin for useful discussions. This research was supported in part by NSF grant AST-1517148. Y-YM is supported by the Samuel P. Langley PITT PACC Postdoctoral Fellowship. Support for SG-K was provided by NASA through Einstein Postdoctoral Fellowship grant number PF5-160136 awarded by the *Chandra* X-ray Center, which is operated by the Smithsonian Astrophysical Observatory for NASA under contract NAS8-03060. AW was supported by a Caltech-Carnegie Fellowship, in part through the Moore Center for Theoretical Cosmology and Physics at Caltech, and by NASA through grants HST-GO-14734 and HST-AR-15057 from STScI.

This research made use of computational resources at SLAC National Accelerator Laboratory, a U.S. Department of Energy Office; the authors are thankful for the support of the SLAC computational team. This research was supported in part by the National Science Foundation under grant No. NSF PHY17-48958 through the Kavli Institute for Theoretical Physics program "The Galaxy-Halo Connection Across Cosmic Time." This research made use of the Python Programming Language, along with many community-developed or maintained software packages, including IPython (Pérez & Granger 2007), Jupyter (jupyter.org), Matplotlib (Hunter 2007), NumPy (van der Walt et al. 2011), Pandas (McKinney 2010), Scikit-Learn (Pedregosa et al. 2011), SciPy (Jones et al. 2001), and Seaborn (seaborn. pydata.org). This research made extensive use of the arXiv and NASA's Astrophysics Data System for bibliographic information.

# Appendix Scatter, Feature Selection, and Resolution

We perform several tests to check the robustness of our results. First, we examine the scatter in the random forest



Figure 12. Velocity functions for m12i (left) and m12f (right) predicted by the most probable realization of our random forest classifier when trained only on subhalos from m12i (orange) or m12f (green) with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$ . Blue lines show 200 realizations of the prediction for our fiducial classifier, which is trained on subhalos from both hosts, and red lines show the FIRE results. While there is a difference between the total number of surviving subhalos predicted by classifiers trained only on m12i or m12f, the scatter about the most probable prediction for our fiducial classifier is small.



**Figure 13.** Velocity functions and radial distributions of subhalos in m12i (left) and m12f (right) predicted by the most probable realization of random forest classifiers trained on subhalos from both hydrodynamic simulations with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$ . The classifiers use the features  $d_{\text{peri}}$  (blue);  $d_{\text{peri}}$  and  $a_{\text{peri}}$  (red);  $d_{\text{peri}}$ ,  $a_{\text{peri}}$ ,  $a_{\text{peri}}$ ,  $a_{\text{acc}}$ , and  $M_{\text{acc}}$  (green);  $d_{\text{peri}}$ ,  $a_{\text{acc}}$ , and  $M_{\text{acc}}$  (orange); and  $d_{\text{peri}}$ ,  $a_{\text{peri}}$ ,  $a_{\text{acc}}$ , and  $V_{\text{acc}}$  (cyan), corresponding to the rows of Table 1 and the columns of Figure 4. The solid red lines show the FIRE results.

predictions for the m12i and m12f  $V_{max}$  functions and whether the details of the training data affect our results. Figure 12 shows the  $V_{max}$  functions from 200 realizations of our fiducial classifier, along with the most probable realization of classifiers trained only on subhalos from m12i or m12f with  $V_{peak} > 10 \,\mathrm{km \, s^{-1}}$ . The scatter about the most probable prediction for our fiducial classifier is small; in particular, the intrinsic scatter of the random forest prediction is comparable to or smaller than the Poisson noise over the entire velocity function for each host. Thus, even though the prediction for the total number of surviving subhalos is different for classifiers trained on m12i or m12f separately, this uncertainty does not propagate to our fiducial classifier.

Next, we explore the choice of training features. In particular, we test how adding subhalo features affects our results for the  $V_{\text{max}}$  functions and radial distributions of the

surviving subhalo populations predicted from DMO simulations of m12i and m12f. In Figure 13, we plot the most probable  $V_{\text{max}}$  functions and radial distributions predicted by five classifiers that each use an additional training feature, corresponding to the rows of Table 1 and the columns of Figure 4. As we add subhalo features, the predicted distributions approach the FIRE results. Interestingly,  $d_{\text{peri}}$  (or  $a_{\text{peri}}$  or  $a_{\text{acc}}$ ) alone provides most of the information needed to match the total number of surviving subhalos with  $V_{\text{peak}} > 10 \text{ km s}^{-1}$ and r < 300 kpc, but adding additional features improves the predictions at large  $V_{\text{max}}$  and small radii.

Finally, we study how our results depend on the resolution limits used for the training data. In Figure 14, we show the m12i and m12f velocity functions and radial distributions predicted by a classifier trained on subhalos from both hosts with  $V_{\text{peak}} > 5 \text{ km s}^{-1}$ , which is less restrictive than the



**Figure 14.** Velocity functions (top) and radial distributions (bottom) for subhalos hosted by m12i (left) and m12f (right), predicted by the most probable realization of our random forest classifier trained on subhalos from m12i and m12f with  $V_{\text{peak}} > 5 \text{ km s}^{-1}$  (blue). The FIRE (red), DISK (dot-dashed), and DMO (dashed) results are shown for comparison. We restrict these plots to subhalos within 300 kpc of their respective hosts at z = 0 and with  $V_{\text{max}} > 5 \text{ km s}^{-1}$ , where  $V_{\text{max}}$  is the maximum circular velocity evaluated at z = 0, to allow for a direct comparison with the results in Garrison-Kimmel et al. (2017).

 $V_{\text{peak}} > 10 \text{ km s}^{-1}$  cut used in our primary analysis. We plot the results for subhalos with  $V_{\text{max}} > 5 \text{ km s}^{-1}$ , where  $V_{\text{max}}$  is the maximum circular velocity evaluated at z = 0, which allows for a direct comparison with the results in GK17. Our conclusions are unaffected by changing the minimum circular velocity. In fact, our predictions match the hydrodynamic results even more closely than before in the low- $V_{\text{max}}$  regime, since this less restrictive cut significantly increases the number of subhalos at the low- $V_{\text{max}}$  end of the training set. Thus, our classifier can be applied to simulations with a range of resolution thresholds if appropriate cuts are applied to the training data.

# **ORCID** iDs

Ethan O. Nadler <sup>(i)</sup> https://orcid.org/0000-0002-1182-3825 Yao-Yuan Mao <sup>(i)</sup> https://orcid.org/0000-0002-1200-0820 Risa H. Wechsler (1) https://orcid.org/0000-0003-2229-011X Shea Garrison-Kimmel (1) https://orcid.org/0000-0002-4655-8128

Andrew Wetzel () https://orcid.org/0000-0003-0603-8942

#### References

- Behroozi, P. S., Wechsler, R. H., Lu, Y., et al. 2014, ApJ, 787, 156
- Behroozi, P. S., Wechsler, R. H., & Wu, H.-Y. 2013a, ApJ, 762, 109
- Behroozi, P. S., Wechsler, R. H., Wu, H.-Y., et al. 2013b, ApJ, 763, 18
- Boylan-Kolchin, M., Bullock, J. S., & Kaplinghat, M. 2011, MNRAS, 415, L40
- Boylan-Kolchin, M., Bullock, J. S., & Kaplinghat, M. 2012, MNRAS, 422, 1203
- Brook, C. B., & Di Cintio, A. 2015, MNRAS, 450, 3920
- Brooks, A. M., Kuhlen, M., Zolotov, A., & Hooper, D. 2013, ApJ, 765, 22
- Bullock, J. S., & Boylan-Kolchin, M. 2017, ARA&A, 55, 343
- Bullock, J. S., Kravtsov, A. V., & Weinberg, D. H. 2000, ApJ, 539, 517

- Chan, T. K., Kereš, D., Oñorbe, J., et al. 2015, MNRAS, 454, 2981
- Del Popolo, A., & Le Delliou, M. 2017, Galax, 5, 17
- DESI Collaboration, Aghamousa, A., & Aguilar, J. 2016, eprint (arXiv:1611. 00036)
- Di Cintio, A., Brook, C. B., Dutton, A. A., et al. 2014, MNRAS, 441, 2986
- Fitts, A., Boylan-Kolchin, M., Elbert, O. D., et al. 2017, MNRAS, 471, 3547
- Garrison-Kimmel, S., Wetzel, A., Bullock, J. S., et al. 2017, MNRAS, 471, 1709
- Geha, M., Wechsler, R. H., Mao, Y.-Y., et al. 2017, ApJ, 847, 4
- Governato, F., Zolotov, A., Pontzen, A., et al. 2012, MNRAS, 422, 1231
- Hopkins, P. F., Kereš, D., Oñorbe, J., et al. 2014, MNRAS, 445, 581
- Hopkins, P. F., Wetzel, A., Keres, D., et al. 2017, eprint (arXiv:1702.06148) Hunter, J. D. 2007, CSE, 9, 90
- Jones, E., Oliphant, T., Peterson, P., et al. 2001, SciPy: Open source scientific tools for Python, scipy.org
- Katz, N., & White, S. D. M. 1993, ApJ, 412, 455
- Kim, J.-h., Abel, T., Agertz, O., et al. 2014, ApJS, 210, 14
- Kim, S. Y., Peter, A. H. G., & Hargis, J. R. 2017, eprint (arXiv:1711.06267)
- Klypin, A., Kravtsov, A. V., Valenzuela, O., & Prada, F. 1999, ApJ, 522, 82
- Knollmann, S. R., & Knebe, A. 2009, ApJS, 182, 608 Loveday, J., Norberg, P., Baldry, I. K., et al. 2015, MNRAS, 451, 1540
- LSST Science Collaboration, Abell, P. A., & Allison, J. 2009, eprint (arXiv:0912.0201)
- Lu, Y., Benson, A., Mao, Y.-Y., et al. 2016, ApJ, 830, 59

- Lu, Y., Benson, A., Wetzel, A., et al. 2017, ApJ, 846, 66
- Mao, Y.-Y., Williamson, M., & Wechsler, R. H. 2015, ApJ, 810, 21
- McKinney, W. 2010, in Proc. 9th Python in Science Conf., ed. S. van der Walt & J. Millman, 51
- Moore, B., Ghigna, S., Governato, F., et al. 1999, ApJL, 524, L19
- Oñorbe, J., Boylan-Kolchin, M., Bullock, J. S., et al. 2015, MNRAS, 454, 2092
- Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, Journal of Machine Learning Research, 12, 2825
- Pérez, F., & Granger, B. E. 2007, CSE, 9, 21
- Pontzen, A., & Governato, F. 2012, MNRAS, 421, 3464
- Sawala, T., Pihajoki, P., Johansson, P. H., et al. 2017, MNRAS, 467, 4383
- Somerville, R. S. 2002, ApJL, 572, L23
- Springel, V., Wang, J., Vogelsberger, M., et al. 2008, MNRAS, 391, 1685
- Tollet, E., Macciò, A. V., Dutton, A. A., et al. 2016, MNRAS, 456, 3542
- van den Bosch, F. C. 2017, MNRAS, 468, 885
- van den Bosch, F. C., Jiang, F., Campbell, D., & Behroozi, P. 2016, MNRAS, 455, 158
- van der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, CSE, 13, 22
- Wetzel, A. R., Hopkins, P. F., Kim, J.-h., et al. 2016, ApJL, 827, L23
- Zhu, Q., Hernquist, L., Marinacci, F., Springel, V., & Li, Y. 2017, MNRAS, 466, 3876
- Zhu, Q., Marinacci, F., Maji, M., et al. 2016, MNRAS, 458, 1559
- Zolotov, A., Brooks, A. M., Willman, B., et al. 2012, ApJ, 761, 71