



The Role of Mass and Environment on Satellite Distributions around Milky Way Analogs in the ROMULUS25 Simulation

Jordan Van Nest¹ , Ferah Munshi² , Charlotte Christensen³ , Alyson M. Brooks^{4,5} , Michael Tremmel^{6,7} , and Thomas R. Quinn⁸

¹ Homer L. Dodge Department of Physics & Astronomy, University of Oklahoma, 440 W. Brooks St., Norman, OK 73019, USA; jdvannest@ou.edu

² Department of Physics & Astronomy, George Mason University, 4400 University Dr., MSN: 3F3 Fairfax, VA 22030-4444, USA

³ Department of Physics, Grinnell College, 1116 8th Ave., Grinnell, IA 50112, USA

⁴ Department of Physics & Astronomy, Rutgers, The State University of New Jersey, 136 Frelinghuysen Road, Piscataway, NJ 08854, USA

⁵ Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Ave., New York, NY 10010, USA

⁶ Astronomy Department, Yale University, P.O Box 208120, New Haven, CT 06520, USA

⁷ Physics Department, University College Cork, T12 K8AF Cork, Ireland

⁸ Astronomy Department, University of Washington, Box 351580, Seattle, WA 98195-1580, USA

Received 2023 April 4; revised 2023 August 29; accepted 2023 September 8; published 2023 October 12

Abstract

We study satellite counts and quenched fractions for satellites of Milky Way analogs in ROMULUS25, a large-volume cosmological hydrodynamic simulation. Depending on the definition of a Milky Way analog, we have between 66 and 97 Milky Way analogs in ROMULUS25, a 25 Mpc per-side uniform volume simulation. We use these analogs to quantify the effect of environment and host properties on satellite populations. We find that the number of satellites hosted by a Milky Way analog increases predominantly with host stellar mass, while environment, as measured by the distance to a Milky Way-mass or larger halo, may have a notable impact in high isolation. Similarly, we find that the satellite quenched fraction for our analogs also increases with host stellar mass, and potentially in higher-density environments. These results are robust for analogs within 3 Mpc of another Milky Way-mass or larger halo, the environmental parameter space where the bulk of our sample resides. We place these results in the context of observations through comparisons to the Exploration of Local VolumE Satellites and Satellites Around Galactic Analogs surveys. Our results are robust to changes in Milky Way analog selection criteria, including those that mimic observations. Finally, as our samples naturally include Milky Way-Andromeda pairs, we examine quenched fractions in pairs versus isolated systems. We find potential evidence, though not conclusive, that pairs, defined as being within 1 Mpc of another Milky Way-mass or larger halo, may have higher satellite quenched fractions.

Unified Astronomy Thesaurus concepts: [Galaxy evolution \(594\)](#); [Galaxy quenching \(2040\)](#); [Dwarf galaxies \(416\)](#)

1. Introduction

The satellites of the Milky Way and its neighbors in the Local Group, thanks to their proximity, have often served as our basis of understanding satellite and dwarf galaxy formation and evolution. In the past decades there has been an explosion in our understanding of satellites around our own Milky Way (Mateo 1998; Koposov et al. 2008; Simon 2019; Drlica-Wagner et al. 2020, and references within) and Andromeda (Ibata et al. 2014; Martin et al. 2016; McConnachie et al. 2018, and references within). Further, in the age of ultrafaint galaxy detection, the low surface brightness end of the Milky Way’s satellite distribution continues to grow (e.g., Drlica-Wagner et al. 2015; Koposov et al. 2015; Simon 2019). As we continue to discover fainter objects nearby, the question of the Milky Way’s uniqueness becomes an important one. Applying what we learn locally to the Universe at large would not be appropriate if the Local Group could be considered “atypical.”

To test for any potential discrepancy, surveys such as the “Satellites Around Galactic Analogs” (SAGA; Geha et al. 2017; Mao et al. 2021) and “Exploration of Local VolumE Satellites” (ELVES; Carlsten et al. 2020, 2021a, 2021b, 2022)

study the satellite distributions of galaxies similar to our own, placing the Milky Way in a broader, cosmological context. The SAGA survey is an ongoing effort to compile spectroscopically complete satellite luminosity functions of 100 Milky Way analogs with distances between 20 and 40 Mpc, providing vastly improved statistics for the bright end of these satellite distributions (down to $M_R = -12.3$). In complement to the SAGA survey’s probing of distant Milky Way-like systems, the ELVES survey seeks to map the satellite distributions of the hosts within the Local Volume fully (<12 Mpc) down to $M_V = -9$.

Working in tandem, SAGA and ELVES will provide a better understanding of both what a “typical” Milky Way-like halo will look like and what influences an environment like the Local Volume can impart. The SAGA survey has found that the luminosity function of the Milky Way is consistent with their observations of other systems, but that the host-to-host scatter in the number of satellites is large (Mao et al. 2021). SAGA also finds that the total number of satellites in a system correlates with the host’s K -band luminosity. Similar to SAGA, the ELVES survey finds that satellite abundance correlates with host mass and that the Milky Way is typical for its mass. However, Carlsten et al. (2021b) find that the observed luminosity functions of local hosts are typically “flatter” than predicted by the cosmological model; the stellar-to-halo mass relation tends to underpredict bright satellites and overpredict faint ones, a result found also by Geha et al. (2017). These



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

results highlight the power of a larger sample of galaxies and their satellites to provide context for understanding satellite dwarf galaxies.

One of the most interesting discrepancies to be highlighted so far is that the quenched fraction of Local Group (Milky Way and Andromeda) satellites is not in agreement with SAGA's results: the SAGA sample exhibits lower quenched fractions than those found in the Local Group. On the other hand, the ELVES survey finds higher quenched fractions among the Local Volume than in the SAGA sample, though still not as high as the Local Group. Although Mao et al. (2021) carefully attempt to quantify incompleteness in the SAGA survey, it remains an open question whether SAGA may be missing faint, red, or low surface brightness satellites which would be predominantly quenched (Carlsten et al. 2022; Font et al. 2022), or whether the Local Group is a true outlier in terms of quenched satellite fraction.

In general, various simulations of Milky Way-like galaxies tend to find good agreement in their resulting quenched satellite fractions, lying somewhere between the Local Group and Local Volume fractions (Akins et al. 2021; Engler et al. 2021; Karunakaran et al. 2021; Samuel et al. 2022). These simulations generally find that galaxies that infall into a host Milky Way with stellar masses above $M_* \sim 10^8 M_\odot$ are better able to retain their gas and continue star forming for extended periods. On the other hand, galaxies with stellar masses below $10^8 M_\odot$ instead tend to experience ram pressure stripping that strips gas and quenches their star formation (SF), and the quenching timescales can often be quite short (<2 Gyr; Wetzel et al. 2015; Simpson et al. 2018; Simons et al. 2020; Akins et al. 2021). These results lead to high predicted quenched satellite fractions as luminosity decreases.

On the theoretical front, many analyses use zoom-in simulations of a handful of Milky Way analogs (e.g., Akins et al. 2021; Samuel et al. 2022), though Font et al. (2022) use the ARTEMIS suite of 24 cosmological Milky Way-mass zooms to interpret the ELVES and SAGA results. Font et al. (2022) find that applying a surface brightness limit to the ARTEMIS satellites can bring the quenched fractions and radial distributions into line with the SAGA results, suggesting that SAGA is missing faint surface brightness galaxies. Fainter surface brightnesses correlate with more quenching at a fixed luminosity in ARTEMIS, and thus bias the SAGA results if true. On the other hand, Engler et al. (2023) found that a surface brightness cut could not bring the TNG50 satellite quenched fractions fully into agreement with SAGA, though it did bring the simulation and observational results more into line. Engler et al. (2023) were able to use TNG50, a 50 Mpc on-a-side uniform cosmological volume, to study a larger sample of Milky Way analogs and look for statistical trends. In this work, we use ROMULUS25, a 25 Mpc on-a-side uniform cosmological volume with comparable resolution to TNG50, to study similar trends. We particularly focus on the questions of how host mass and large-scale environment impact both satellite counts and quenched fractions for our simulated Milky Way analogs.

The paper is outlined as follows. We begin in Section 2 by describing the ROMULUS25 simulation. In Section 3 we outline our various methods for identifying Milky Way analogs, as well as their satellites. In Section 4 we present our primary results, focusing on the general size of the satellite populations

and their quenched fractions. We then discuss and summarize our results in Sections 5 and 6, respectively.

2. Simulation

For this work, we use the ROMULUS25 simulation (Tremmel et al. 2017). ROMULUS25 was run with CHANGA (Menon et al. 2015), which includes standard physics modules previously used in GASOLINE (Wadsley et al. 2004, 2008, 2017) such as a cosmic UV background (Haardt & Madau 2012) including self-shielding (Pontzen et al. 2008), SF, “blastwave” supernova (SN) feedback (Stinson et al. 2006), and low-temperature metal cooling (Bromm et al. 2001). CHANGA implements an updated smoothed particle hydrodynamics (SPH) routine that uses a geometric mean density in the SPH force expression, allowing for the accurate simulation of shearing flows with Kelvin-Helmholtz instabilities (Wadsley et al. 2017). Finally, a time-dependent artificial viscosity and an on-the-fly time-step adjustment (Saitoh & Makino 2009) system allow for more realistic treatment of weak and strong shocks (Wadsley et al. 2017).

ROMULUS25 assumes a Λ cold dark matter (Λ CDM) model with cosmological parameter values following results from Planck ($\Omega_0 = 0.3086$, $\Lambda = 0.6914$, $h = 0.6777$, and $\sigma_8 = 0.8288$; Planck Collaboration et al. 2016). The simulation has a Plummer equivalent force softening of 250 pc (a spline softening of 350 pc is used, which converges to a Newtonian force at 700 pc). Unlike many similar cosmological runs, the dark matter particles were *oversampled* relative to gas particles, such that the simulation was run with initially 3.375 times more dark matter particles than gas. This increased dark matter resolution allows for the ability to track the dynamics of supermassive black holes within galaxies (Tremmel et al. 2015). The result is a dark matter particle mass of $3.39 \times 10^5 M_\odot$ and gas particle mass of $2.12 \times 10^5 M_\odot$. These relatively low dark matter particle masses decrease numerical effects resulting from two-body relaxation and energy equipartition, which occur when particles have significantly different masses, both of which can affect the structure of simulated galaxies (e.g., Ludlow et al. 2019). ROMULUS25 has been shown to reproduce important galaxy and supermassive black hole scaling relations (Tremmel et al. 2017; Ricarte et al. 2019; Sharma et al. 2022a, 2022b).

2.1. Star Formation and Gas Cooling

Gas cooling at low temperatures is regulated by metal abundance as in Guedes et al. (2011), as well as SPHs that include both thermal and metal diffusion as described in Shen et al. (2010) and Governato et al. (2015; thermal and metal diffusion coefficients set to 0.3, see Tremmel et al. 2017, 2019 for an in-depth discussion). SF and associated feedback from SNe are crucial processes that require subgrid models in cosmological simulations like ROMULUS25. Following Stinson et al. (2006), SF is regulated with parameters that encode SF efficiency in dense gas, couple SN energy to the interstellar medium (ISM), and specify the physical conditions required for SF. These parameters were calibrated using several dozen zoom-in simulations of dwarf to Milky Way-mass galaxies (Tremmel et al. 2017) and are as follows.

1. The normalization of the SF efficiency, $c_{\text{SF}} = 0.15$, and formation timescale, $\Delta t = 10^6$ yr, are both used to calculate the probability p of creating a star particle from

a gas particle that has a dynamical time t_{dyn}

$$p = \frac{m_{\text{gas}}}{m_{\text{star}}}(1 - e^{-c_{\text{SF}}\Delta t/t_{\text{dyn}}}). \quad (1)$$

2. The fraction of SN energy coupled to the ISM, $\epsilon_{\text{SN}} = 0.75$.
3. Minimum density, $n_* = 0.2 \text{ cm}^{-3}$, and maximum temperature, $T_* = 10^4 \text{ K}$, thresholds beyond which cold gas is allowed to form stars.

Star particles form with a mass of $6 \times 10^4 M_{\odot}$, or 30% the initial gas particle mass. ROMULUS25 assumes a Kroupa initial mass function (Kroupa 2001) with associated metal yields and SN rates. Feedback from SNe uses the “blastwave” implementation (Stinson et al. 2006), with thermal energy injection and a cooling shutoff period approximating the “blastwave” phase of SN ejecta when cooling is inefficient.

2.2. Halo Identification

Amiga Halo Finder (AHF; Knollmann & Knebe 2009) was applied to ROMULUS25 to identify dark matter halos, subhalos, and the baryonic content within. AHF uses a spherical top-hat collapse technique (Bryan & Norman 1998) to calculate each halo’s virial radius (R_{vir}) and mass (M_{vir}). Halos are considered resolved if their virial mass is at least $3 \times 10^9 M_{\odot}$ at $z = 0$. This corresponds to a dark matter particle count of $\sim 10^4$, and a stellar mass of at least $10^7 M_{\odot}$ (star particle count of ~ 150). Following Munshi et al. (2013), stellar masses were scaled by a factor of 0.5 as a photometric correction for a better comparison to values inferred from typical observational techniques, and all magnitudes use the Vega zero-point.

3. Analog and Satellite Identification

There is no concrete definition of what constitutes a Milky Way analog; observational surveys like SAGA and ELVES make sample cuts using K -band magnitudes as proxies for stellar mass, while simulations have access to more exact values for halo properties such as stellar mass and virial radius. In this work, we select samples of Milky Way analogs according to three different criteria sets in order to test if the selection criteria can influence the resultant satellite distribution. Our samples are defined as follows.

1. A *general M_{vir} restriction*: any halo where $10^{11.5} < M_{\text{vir}}/M_{\odot} < 10^{12.5}$.
2. A *general M_* restriction*: any galaxy where $10^{10} < M_*/M_{\odot} < 10^{11}$. This corresponds to the host stellar mass range outlined in Section 2.1.2 of SAGA II (Mao et al. 2021). A stellar mass of $10^{10} M_{\odot}$ also corresponds to the lower limit of the ELVES survey (Carlsten et al. 2022).
3. An *M_K +environmental restriction*: any galaxy where $-24.6 < M_K < -23$. Additionally, no neighbor within 300 kpc can have $M_{K,\text{neighbor}} < M_{K,\text{MW}} - 1.6$. This corresponds to the K -band magnitude cut and environmental restrictions from SAGA II (Mao et al. 2021).

We also explore two different ways to identify a satellite galaxy. First, we consider galaxies within the host’s virial radius down to a stellar mass of $10^7 M_{\odot}$, the resolution limit for ROMULUS25. This corresponds to a magnitude limit of $M_R \approx -12.6$. We note that the magnitude limit for the SAGA survey is $M_R = -12.3$ (though they have four satellites below

Table 1
A Summary of Our Samples of Milky Way Analogs and Satellites

(1) Analog Criteria	(2) Analog Radius	(3) N_{MW}	(4) N_{Sats}	(5) $\text{Max}(N_{\text{Sat}})$
M_{vir}	...	67	138	8
M_*	R_{vir}	97	210	13
$M_K + \text{Env.}$...	77	148	13
M_{vir}	...	66	125	6
M_*	300 kpc	90	171	7
$M_K + \text{Env.}$...	77	137	6
	(SAGA II)			

Note. (1) The criteria for identifying Milky Way analogs; (2) the virial radius of the Milky Way analog for the purpose of identifying satellites; (3) the total number of Milky Way analogs; (4) the total number of satellites with $M_* > 10^7 M_{\odot}$; and (5) the largest number of satellites hosted by a single Milky Way analog.

this limit; see Mao et al. 2021), so our samples do not probe the lowest-mass regions of the SAGA or ELVES sample spaces. In addition, we perform a selection where satellites are identified by being within 300 kpc of a Milky Way analog, rather than the analog’s virial radius, as a more direct comparison to the SAGA and ELVES surveys. We note, however, that these surveys use 2D projected distances while in this work we use true 3D distances. In the event that a satellite is hosted by multiple analogs, it is ascribed to the most massive host. Any satellites that fall into the criteria of a Milky Way analog are not included in the satellite distribution. As a final step, any analogs that host a “satellite” more massive than themselves are removed from consideration. This cut is partially responsible for the slight variation in the number of Milky Way analogs under the same criteria when switching between the R_{vir} and 300 kpc satellite identifications. Our sample of Milky Way analogs and satellites are summarized for each criteria set in Table 1.

Figure 1 shows the three Milky Way analog samples that we focus on in this work: M_{vir} and M_* with R_{vir} and M_K with 300 kpc. While the samples largely overlap, we find that none of them are simple subsets of the others. As they approach the boundaries of the selection cuts, the samples diverge from one another. For example, the stellar mass sample probes virial masses below the virial mass cut, and vice versa. This is the result of natural scatter within the stellar–halo mass relation, which was shown in Tremmel et al. (2017) to match observations (Moster et al. 2013; Kravtsov et al. 2018). Within the overlapping regions of the criteria, there are galaxies considered Milky Way analogs in some samples but not others. This occurs as a result of the environmental criteria in the SAGA sample, which could remove analogs that are still within the K -band magnitude limits.

In Figure 2, we compare the normalized distributions of hosts and satellites from our largest sample, M_* with simulated R_{vir} (in order to encapsulate the full magnitude range of our samples), to data from SAGA II and ELVES (Mao et al. 2021; Carlsten et al. 2022). We note that the ELVES satellites are weighted according to their likelihood estimates (P_{sat} in Table 9 of Carlsten et al. 2022), so each satellite adds its likelihood as a count rather than 1. In panel (a), we see that our hosts’ span in K -magnitude space matches well with the ELVES sample, while the SAGA II sample (by definition) resides in $-24.6 < M_K < -23$. The peaks of the host distributions are

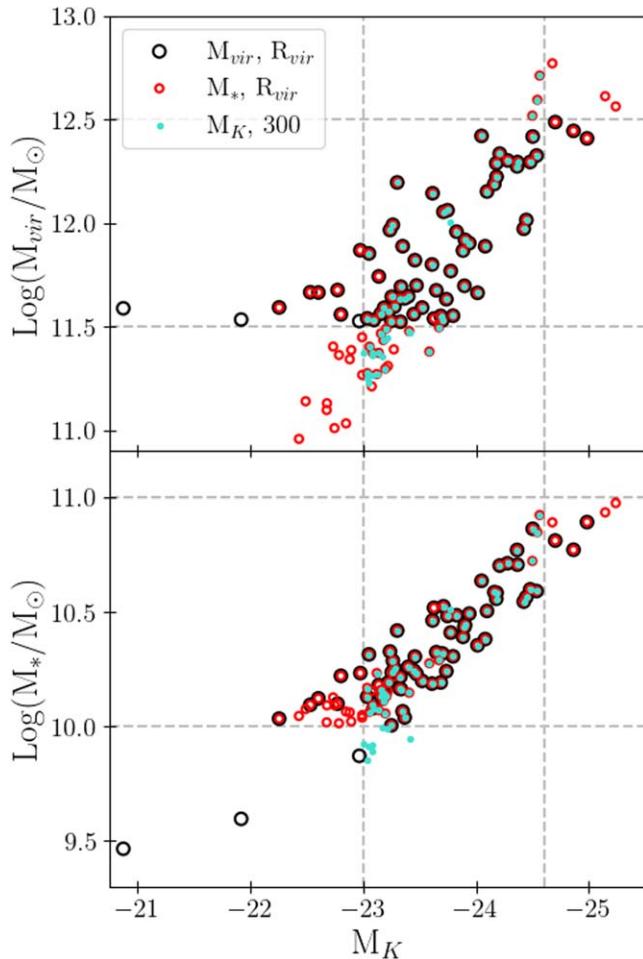


Figure 1. Virial and stellar masses plotted against K -band magnitudes for three of our Milky Way analog samples. The dotted lines denote the mass and magnitude cuts used in our samples. The samples diverge at the different boundaries, and even within the boundaries there are analogs that exist in some definitions but not others.

in good agreement as well, though we note our peak is at a slightly dimmer magnitude than the observational data. In panel (b), we see that our satellite distribution is in very good agreement with the SAGA II data, though we have an interesting lack of satellites at $M_K \approx -17$. The ELVES data probe much dimmer satellites (due to the difference in observational limits), but when only considering satellites brighter than $M_K = -12$, the ELVES sample is still more concentrated at low-mass satellites when compared to SAGA II and ROMULUS25. This is consistent with ELVES finding steeper luminosity functions (fewer high-mass satellites and more low-mass ones) in their sample when compared to SAGA, and might also contribute to the different quenched fractions found by the two surveys (see Section 4.2 for discussion).

4. Results

Figure 3 shows the V -band satellite luminosity function for our sample of Milky Way analogs alongside data from the Milky Way and several Milky Way-like systems. The outer gray region outlines the space occupied by our M_* within R_{vir} sample, while the black line and inner dark gray region indicate the mean and standard deviation. The Milky Way and M31

data are taken from Geha et al. (2017). The NGC 4258 and NGC 4631 data were taken from Carlsten et al. (2021b), the M94 data from Smercina et al. (2018), and the M101 data from Bennet et al. (2019). We find that our sample of Milky Way analogs is in good agreement with these observations. We note that the space occupied by our sample remains largely unaffected when changing the Milky Way analog criteria.

4.1. Host Effects on Satellite Accumulation

In order to study how the physical properties of our Milky Way analogs affect their satellite populations, we separated our sample according to their mass and environment. Figure 4 shows the average number of satellites hosted by the Milky Way analogs where the analogs are binned according to their stellar mass and minimum distance to a Milky Way-sized or larger halo, hereafter $D_{\text{MW+}}$. In calculating $D_{\text{MW+}}$, we consider the closest galaxy outside the system of the analog (i.e., not a satellite) that exceeds the minimum criteria of Milky Way analog under the given criteria. The text in each bin details N (the number of analogs in that bin) and σ (the standard deviation of the number of satellites hosted by analogs in that bin). A plot is shown for both our M_{vir} with simulated R_{vir} (left) and SAGA II comparison (right) samples. In all of our samples, the number of hosted satellites appears to increase with host mass, and potentially with decreasing $D_{\text{MW+}}$. However, this latter trend cannot be verified by eye as the box size of ROMULUS25 yields a lack of data in the upper regions of this plot (i.e., highly isolated hosts), so the apparent trend is not statistically significant.

While these macroscopic trends are present across all of our simulated samples, there are some notable differences in the distributions. We see that while the M_{vir} definition includes analogs at a lower stellar mass, the number of analogs below $M_* = 10^{10} M_\odot$ is much larger in the SAGA II sample. Additionally, in the higher-mass bins there is fluctuation in both the number of analogs and hosted satellites due to the changing of the satellite selection radius from R_{vir} to 300 kpc.

In an effort to quantify the “by-eye” trends seen in Figure 4, we looked at the specific frequency of the number of satellites hosted by our Milky Way analogs, S_N , normalized to their mass and environment. We use the following specific frequency equations adapted from Harris & van den Bergh (1981):

$$S_{N,\text{env}} = N_{\text{sat}} \times 10^{0.4(D-1.5)}, \quad (2)$$

$$S_{N,\text{mass}} = N_{\text{sat}} \times 10^{0.4(M-10.3)}. \quad (3)$$

Here, N_{sat} is the number of satellites hosted by the Milky Way analog, D is $D_{\text{MW+}}$ in units of Mpc, and M is $\log(M_*/M_\odot)$. The normalization values of 1.5 Mpc and 10.3 were chosen to be roughly the averages of the M_* with simulated R_{vir} sample.

Figure 5 shows the specific frequencies normalized to mass and environment for our M_{vir} and M_* with simulated R_{vir} sample, as well as our SAGA II comparison sample. In looking at the trend with mass, the S_N values consistently increase with the stellar masses of the Milky Way analogs. These results, which are present in all of our Milky Way analog samples, indicate that stellar mass exerts a large influence on satellite accumulation. The SAGA and ELVES surveys both observe this trend of satellite abundance increasing with host mass, though the trends they find are slightly weaker than ours (see Section 5.1 for discussion). Further, a study of seven nearby Milky Way-like systems with the Hyper Suprime-Cam on the

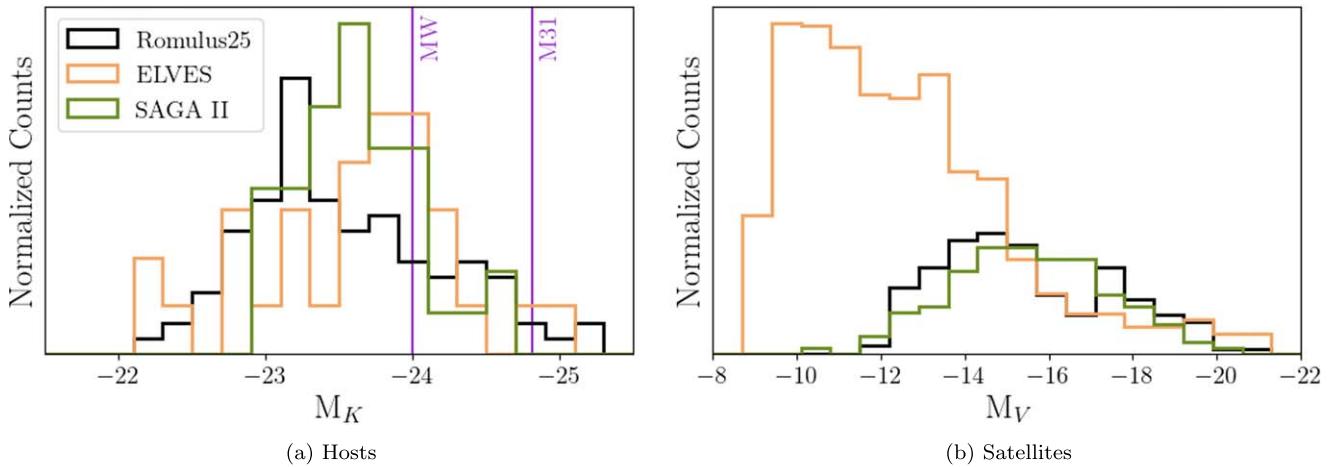


Figure 2. Normalized histograms of (a) hosts in K -band magnitude and (b) satellites in V -band magnitude for our M_* with simulated R_{vir} sample. We make direct comparisons to SAGA II (Mao et al. 2021) and ELVES (Carlsten et al. 2022) data, with the Milky Way and M31 values taken from the latter. The ELVES satellites are weighted according to their likelihood measurements. For a fair comparison, the satellite distributions in (b) are all normalized to the samples’ number of satellites brighter than $M_V = -14$, the approximate completeness limit for ROMULUS25.

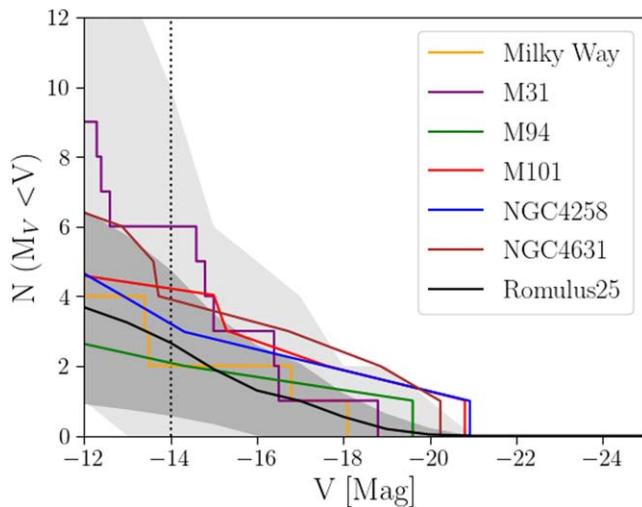


Figure 3. The V -band satellite luminosity function for our Milky Way analog sample under the M_* with simulated R_{vir} criteria. The black line and dark gray region represent the mean and single standard deviation of our sample, respectively, while the outer light gray region encompasses our entire sample. We compare to the Milky Way and M31 (Geha et al. 2017), M94 (Smercina et al. 2018), M101 (Bennet et al. 2019), and NGC 4258 and NGC 4631 (Carlsten et al. 2021b). The dotted vertical line marks the approximate completeness limit for ROMULUS25. Our sample is in good agreement with these observations.

Subaru telescope observes this trend as well (Nashimoto et al. 2022). The trend of satellite abundance with host mass was also found by Font et al. (2021) using the ARTEMIS suite of zoom-in simulations (Font et al. 2020), and by Engler et al. (2021) using the TNG50 simulation.

In looking at the trend with environment, Figure 5(b), we see some interesting behavior. The S_N values increase somewhat linearly until $D_{\text{MW+}} \approx 3.5$ Mpc, where future points go either to zero or extreme outliers. This would suggest that N_{sat} increases as hosts become more isolated, but we note that a majority of our hosts ($\sim 60\%-70\%$) have $D_{\text{MW+}} < 2$ Mpc, so beyond this distance our samples get increasingly small, resulting in the large error bars and stochasticity of the higher- $D_{\text{MW+}}$ points. Thus, we see no definitive trend of satellite accumulation with environment, though one might

become present with a larger sample of more isolated hosts. However, we do not believe that we can fully rule out an environmental impact on satellite accumulation through a measurement of specific frequency. Figure 6 shows the average number of satellites hosted by our analogs when binned by their $D_{\text{MW+}}$ measurement. We see that if we split our analogs into subsamples of “pairs” and “isolated” based on having $D_{\text{MW+}} < 1$ Mpc or > 1 Mpc, the average numbers of hosted satellites, plotted as red diamonds, are notably different between the subsamples. This difference, which is present in all of our samples except $M_K + \text{Env.}$ with simulated R_{vir} , is driven primarily by the low number of satellites hosted by analogs in the $\sim 3\text{--}5$ Mpc bin. Although this is where our sample tapers off in $D_{\text{MW+}}$ space, we can see that the number of analogs in this bin is certainly nonnegligible, and may be hinting at a strong environmental impact on satellite accumulation in more extreme isolation. The stochasticity of our trends in Figure 5(b) compared to the more direct information from Figure 6 leads us to believe that normalizing specific frequency to such an extremely variable parameter (in this case, large-scale environment) does not yield a reliable measurement.

4.2. Host Effects on Satellite Quenching

In addition to studying the number of satellites hosted by our analogs, we also analyzed the quenched fraction of the satellites. When studying the quenched fraction (f_q), we only consider satellites with a stellar mass of at least $10^8 M_\odot$, as ROMULUS25 may be subject to numerical overquenching below this mass (Wright et al. 2021). A galaxy is considered quenched if its instantaneous specific star formation rate (ssSFR) is below 10^{-11} yr^{-1} . The instantaneous ssSFR is calculated using the expected SFR from gas particles meeting the temperature and density thresholds for SF given in Section 2.1. In Figure 7, we show our quenched fractions as a function of host K -band magnitude for the $M_K + \text{Env.}$ with 300 kpc satellite selection (our SAGA II comparison sample), and compare our results to data from the SAGA and ELVES surveys (Mao et al. 2021; Carlsten et al. 2022). For a direct comparison, we only consider SAGA and ELVES satellites with stellar masses above $10^8 M_\odot$. We note, however, that the SAGA and ELVES surveys’ methods of determining

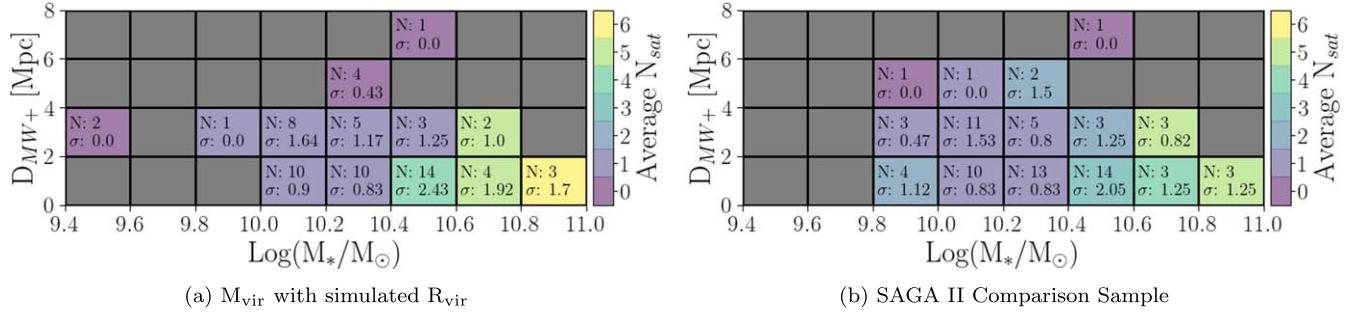


Figure 4. The average number of satellites hosted by Milky Way analogs as a function of stellar mass and environment (distance to a Milky Way–sized or larger halo). The text in each box indicates the number of Milky Way analogs in that parameter space, as well as the standard deviation among the number of satellites. The left plot shows the M_{vir} with simulated R_{vir} sample, while the right plot shows the sample most analogous to the SAGA Survey. In both cases, the number of satellites appears to increase as stellar mass increases.

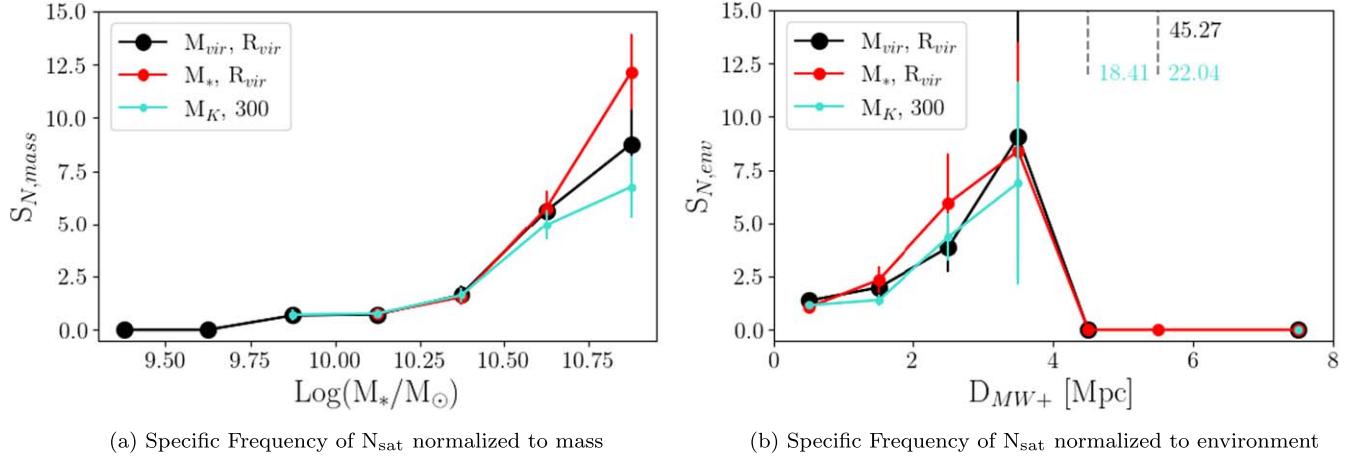


Figure 5. The specific frequencies of the number of satellites hosted by Milky Way analogs normalized to their (a) stellar mass and (b) distance to a Milky Way–sized or larger halo, D_{MW+} . The plots show the results for the M_{vir} (black) and M_* (red) with simulated R_{vir} and SAGA II (blue) samples. Error bars represent the standard error within each bin (σ/\sqrt{N}). With the exception of some large outliers, the $S_{N,env}$ values do not show statistically significant trends. However, the $S_{N,mass}$ values show a clear positive trend.

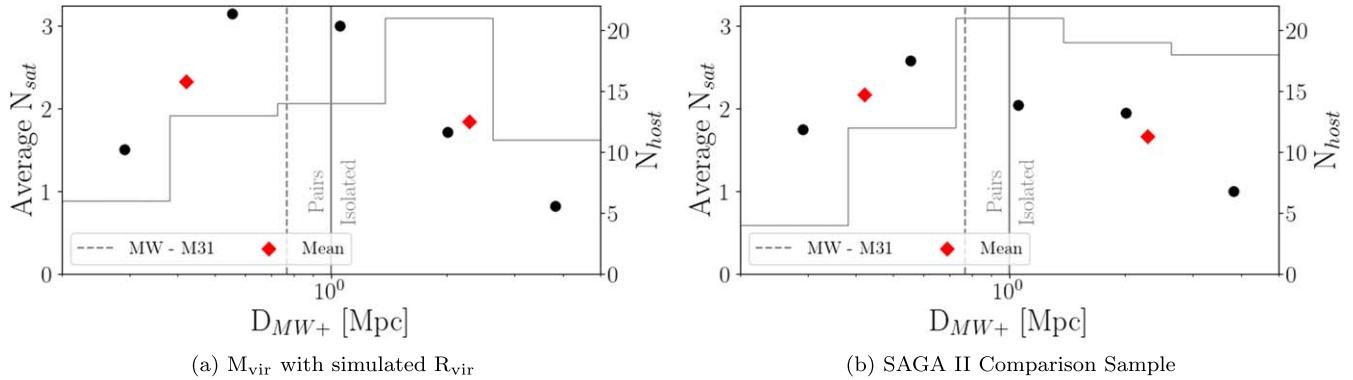


Figure 6. Points show the average number of satellites hosted by Milky Way analogs binned by D_{MW+} for our (a) M_{vir} with simulated R_{vir} and (b) SAGA II comparison sample. Hosts whose D_{MW+} measurements are within 1 Mpc are considered to be in “pairs,” while the rest are considered “isolated,” and the means of these two subsamples are plotted as red diamonds. The histogram data shows the number of hosts in each bin, and the Milky Way–M31 distance is plotted as a dashed vertical line for reference. In both samples the isolated subset hosts fewer average satellites than the pairs, and this is driven predominantly by the most isolated bin.

quenching are different than ours: SAGA considers a satellite quenched if it lacks strong H α emission (equivalent width (EW) of H α $< 2 \text{ \AA}$) and ELVES considers a satellite quenched if it exhibits an early-type morphology, i.e., not exhibiting clear star-forming structures such as blue clumps or dust lanes (this correlates with color as well; see Carlsten et al. 2021a for an in-depth discussion). Our sSFR quenched definition was shown (see Sharma et al. 2022a) to yield a good match to galaxies

identified observationally as quenched using $\text{EW}[\text{H}\alpha] < 2 \text{ \AA}$ and $D_n 4000 > 0.6 + 0.1 \log_{10} M_*$ (as in Geha et al. 2012).

While all three samples show the quenched fractions increasing with host brightness, our simulated sample exhibits slightly larger quenched fractions than the observational surveys, with the exception of the lowest-mass bin where the difference becomes significant (see Section 5.2 for discussion). The SAGA and ELVES data are in very good agreement up to

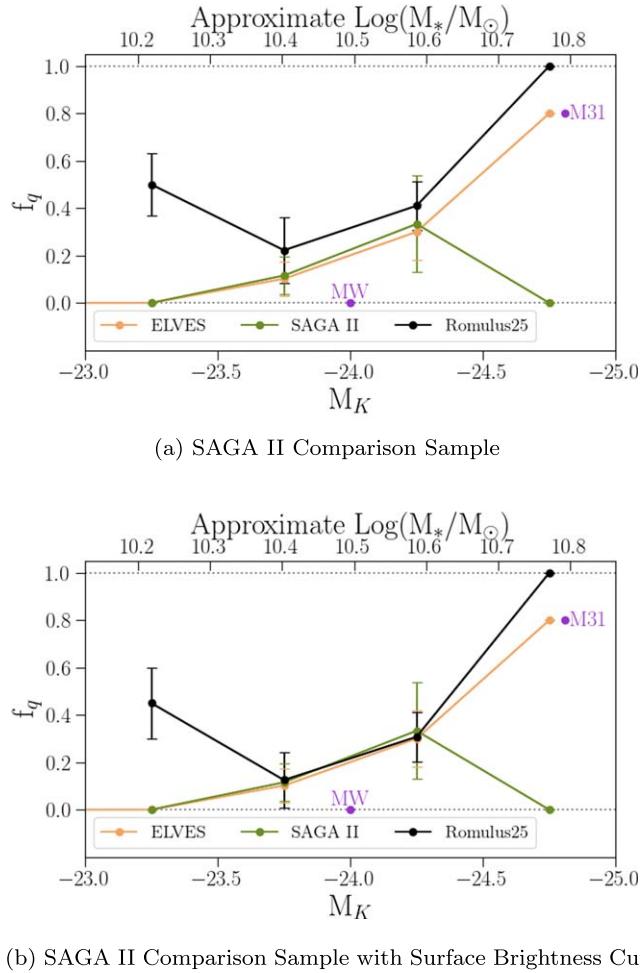


Figure 7. (a) Quenched fraction plotted against K -band magnitude for the $M_K + \text{Env}$. sample from ROMULUS25 compared to SAGA II (Mao et al. 2021) and ELVES (Carlsten et al. 2022) data. (b) The same sample from ROMULUS25 with addition criteria of requiring satellites to have $\mu_{\text{eff},r} < 25 \text{ mag arcsecond}^{-2}$. As a direct comparison, the SAGA and ELVES data plotted here only contain satellites with stellar masses above $10^8 M_{\odot}$. Error bars represent the standard error within each bin. Approximate stellar mass values are taken from a linear fit between M_K and $\log(M_*/M_{\odot})$ for our Milky Way analogs. The Milky Way and M31 values are taken from ELVES, and also only consider satellites with stellar masses above $10^8 M_{\odot}$. While all three samples show the quenched fractions increasing with host brightness, our ROMULUS25 sample exhibits slightly larger quenched fractions (particularly in the faintest bin) unless low surface brightness galaxies are removed. The SAGA and ELVES data are also in good agreement up to the brightest magnitude bin where the sample sizes are small.

the brightest magnitude bin, where the sample sizes are only one host for ELVES (M31) and two hosts for SAGA (NGC 5792 and NGC 7541). This agreement within the high-mass satellite subset is interesting, as the SAGA and ELVES quenched fractions are quite different when considering their full samples. Carlsten et al. (2022) find that the quenched fractions of the Local Volume are significantly higher than the SAGA sample (their Figures 11 and 12), particularly in the low-mass satellite regime. In Figure 2(a), we see that the ELVES survey contains a much larger number of faint satellites when compared to SAGA, but also that ELVES hosts (along with those of ROMULUS25) probe fainter magnitudes.

In studying the ARTEMIS simulations, Font et al. (2022) found that the SAGA detection methods may be preferentially selecting star-forming or recently quenched satellites near their

completeness limit, missing a notable population of quenched dwarfs. This detection bias could explain the difference between SAGA and ELVES low-mass satellites, both the abundances and quenched fractions. Following Font et al. (2022), in Figure 2(b) we apply an additional cut to our SAGA II comparison sample by requiring satellites to have $\mu_{\text{eff},r} < 25 \text{ mag arcsecond}^{-2}$. As in the ARTEMIS simulations, we find that this cut lowers the resultant quenched fractions, and brings our results (particularly the middle bins) into excellent agreement with SAGA and ELVES.

To quantify the trend of quenched fraction with mass seen in Figure 7, and to search for a trend with environment, we again used the specific frequency in Equation (2) with N_{sat} replaced by f_q . Figure 8(a) shows our quenched fraction specific frequencies for the M_{vir} and M_* with simulated R_{vir} samples, as well as our SAGA II comparison sample. We find that, as with the number of hosted satellites, we see a trend of S_N with host mass, indicating that larger hosts are expected to yield higher quenched fractions. However, we note that this trend is not as strong as the one seen in Figure 5(a). Interestingly, when applying the satellite surface brightness criteria in Figure 8(b), we see that our trend of quenched fraction with host mass is effectively erased. As the high-mass end of Figure 8(b) is strongly affected by this surface brightness cut, it seems that the preferentially quenched satellites below this threshold are more common in higher-mass hosts, which is consistent with Figure 8(a) implying a larger number of quenched galaxies in this regime. Our results agree with Engler et al. (2023) who, using the TNG50 run from the IllustrisTNG simulations, found that massive hosts exhibit systematically larger satellite quenched fractions. Further, Engler et al. (2023) found no difference in quenching between isolated and paired analogs when considering satellites within 300 kpc of their host (see Section 5.3 for discussion).

To look for a trend with environment, we examined the quenched fractions of our systems plotted against $D_{\text{MW+}}$. Figure 9 bins our analogs in $D_{\text{MW+}}$ space, and shows the quenched fraction of all satellites hosted by analogs in each bin. Again, pairs are identified as having $D_{\text{MW+}} < 1 \text{ Mpc}$. The figure shows data for our M_* and $M_K + \text{Env}$. samples (both R_{vir} and 300 kpc), where the color of each point represents the number of satellites in each bin. We find that the average quenched fraction is higher among pairs than isolated analogs, though the magnitude of this difference is not ubiquitous across our samples (see Section 5.3 for discussion). In fact, while not shown in the plot, the M_{vir} criteria exhibit no notable difference in the average quenched fractions of paired and isolated analogs. We also find that in the switch from 300 kpc to R_{vir} when identifying satellites, hosts typically have either the same or lower quenched fractions and a higher satellite count. This indicates that restricting satellites to within 300 kpc for this host range is more likely to exclude satellites, and that the satellites beyond 300 kpc are predominantly star forming; though this is still only when considering satellites with $M_* > 10^8 M_{\odot}$. In applying the surface brightness cut in subplots (c) and (d) we find that, while the resultant averaged quenched fractions are lower, the difference between the isolated group and pairs notably increases in all samples.

Similar to Figure 6, this difference between our “pairs” and “isolated” subsamples is largely driven by the most isolated bin, where the quenched fractions are extremely low. This, again, could be alluding to a strong environmental effect on

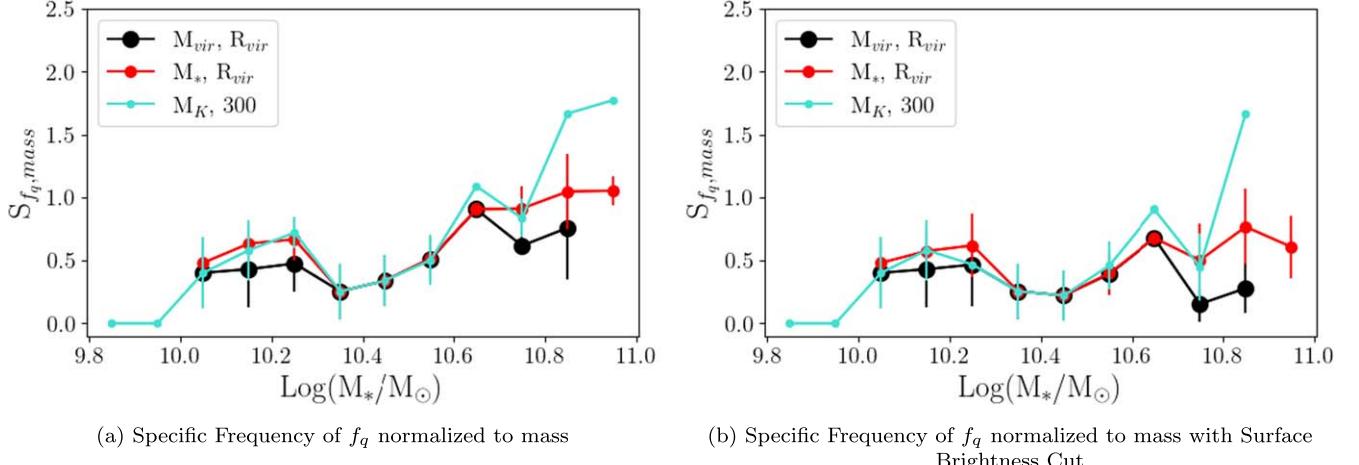


Figure 8. The specific frequencies of the quenched fraction of satellites hosted by Milky Way analogs normalized to their stellar mass. The plots show the results for the M_{vir} (black) and M_* (red) with simulated R_{vir} and SAGA II (blue) samples. Subplot (b) applies the surface brightness cut to satellites ($\mu_{eff,r} < 25$ mag arcsecond $^{-2}$). Error bars represent the standard error within each bin. As with N_{sat} , there is a notable positive trend with Milky Way analog mass. However, applying the surface brightness cut to our satellites effectively removes the trend with mass.

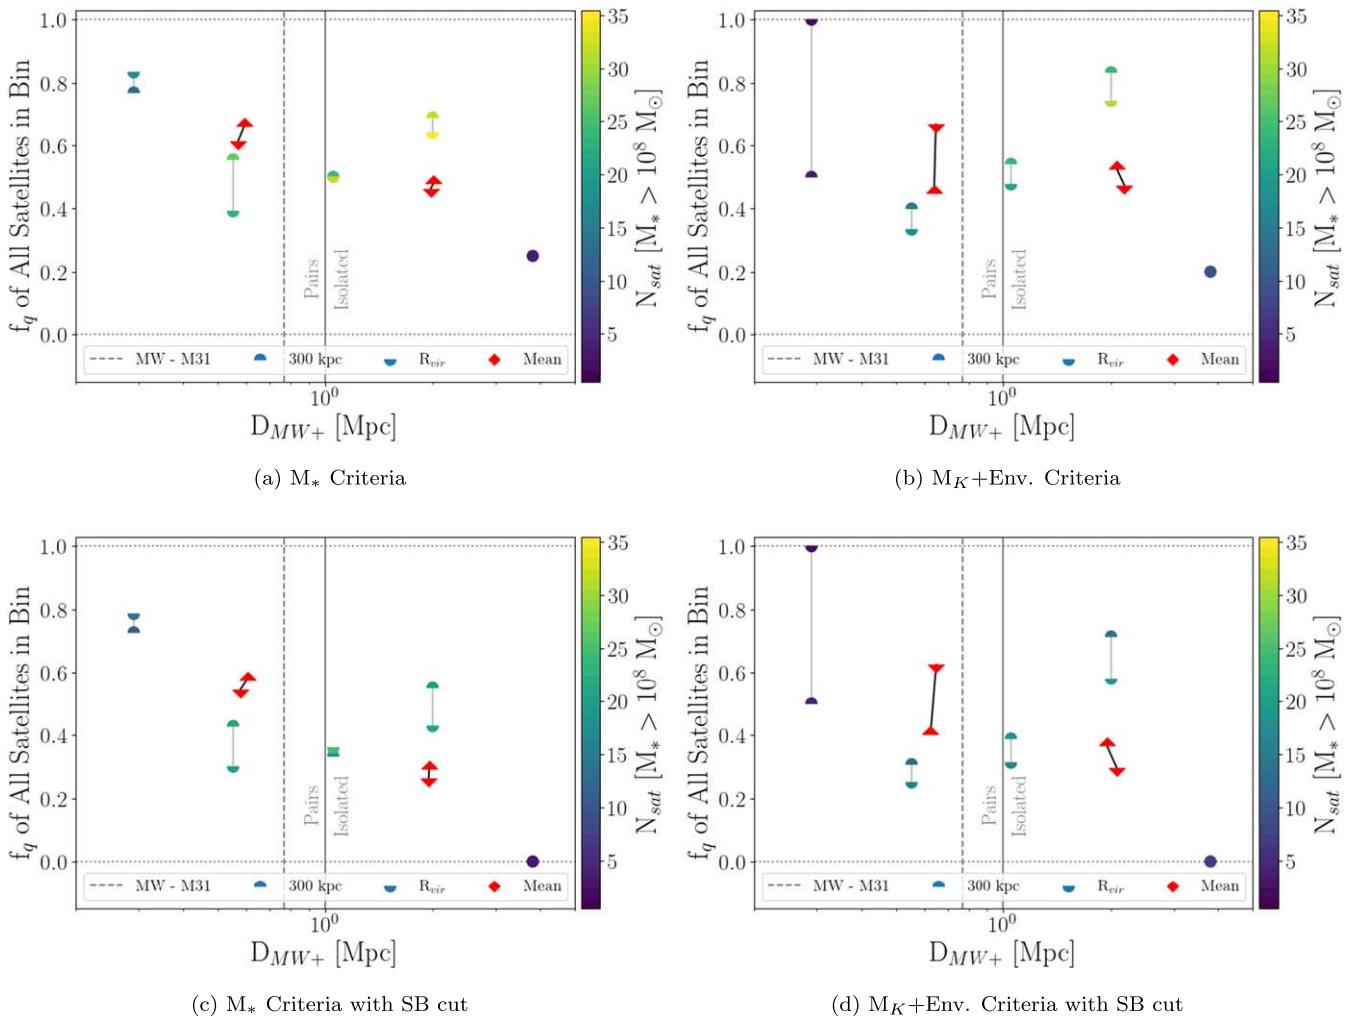


Figure 9. Quenched fraction plotted against the distance to the closest Milky Way halo or larger for the (a) M_* and (b) $M_K + \text{Env.}$ criteria. Analogs are binned by D_{MW+} and the quenched fractions are taken from an aggregate of all satellites in each bin. The upper-hemisphere points represent analogs from the sample with a Milky Way radius of 300 kpc, while lower-hemisphere points represent the sample using the simulated virial radius with vertical lines bridging the two. The colors of each point represent the number of hosted satellites with stellar mass greater than $10^8 M_\odot$. The samples are separated into “Pairs” and “Isolated” by whether the closest Milky Way or larger halo is within 1 Mpc, and the means of each sample are denoted by the red diamonds. For reference, the Milky Way–M31 distance is plotted with a vertical dashed line. Subplots (c) and (d) require satellites to have $\mu_{eff,r} < 25$ mag arcsecond $^{-2}$. The pairs exhibit a slightly higher mean quenched fraction, and changing the satellite selection radius from 300 kpc to R_{vir} typically takes the quenched fraction to an equivalent or lower value. Applying the surface brightness cut to our satellites strengthens the disparity between paired and isolated samples.

satellite quenching in the highly isolated regime that we do not quite capture in ROMULUS25. This result is at odds with the aforementioned study of TNG50 by Engler et al. (2023), as well as the study of FIRE-2 by Samuel et al. (2022), both of whom find no difference in quenching between isolated and paired analogs. However, the simulations studied in Samuel et al. (2022) are zoom-in simulations, and thus would not capture the large-scale isolation in which we begin to see differences between paired and isolated hosts.

While this work does not consider signs of conformity within satellites, we did perform some cursory analysis on how the quenched fractions and number of hosted satellites relate to the instantaneous SFRs of our analogs. We find that the quenched fractions exhibit no notable trend with host SFR. In looking at the average number of hosted satellites, though, we do see a correlation of more populated systems having a higher host SFR. However, we believe this is just a reflection of the number satellites increasing with host mass, as scaling to sSFR largely removes this correlation.

5. Discussion

In Section 3, we discussed the various methods by which we identified Milky Way analogs and satellites. While shifting between these definitions has no effect on our conclusions, there are subtle impacts worth noting.

5.1. Satellites within R_{vir} versus 300 kpc

In Figures 4 and 5 we showed host stellar mass to be a driving factor in satellite accumulation, but this trend is less prominent when using our SAGA II comparison sample. This appears to be the result of using 300 kpc to identify satellites, not the selection on K -band magnitude, as our $M_K + \text{Env.}$ with R_{vir} sample actually exhibits the strongest trend. In fact, identifying satellites via a 300 kpc selection rather than R_{vir} reduces the strength of the mass trend in all criteria (though the trend is still prominent). The weakening of the trends is the result of analogs in the high-mass regime (where the trends manifest), which have virial radii larger than 300 kpc and exclude satellites in this shift to 300 kpc. This shift is in agreement with the ARTEMIS simulations (Font et al. 2021), in which satellite abundance trends strongly with host mass, but the trend is weakened when SAGA observation selection criteria are applied ($M_{r,\text{sat}} < -12$, $\mu_{\text{eff},r} < 25 \text{ mag arcsecond}^{-2}$, and within 300 kpc of the host).

When considering quenched fractions, our choice of satellite selection radius also seems to have a noticeable effect on our $M_K + \text{Env.}$ sample. In Figure 9(b), the switch from R_{vir} to 300 kpc typically raises the quenched fraction while lowering the number of satellites (with the notable exception of the least isolated bin). Thus, within the context of satellites with $M_* > 10^8 M_\odot$, it seems applying a satellite cut of 300 kpc to the K -band magnitude analog selection primarily removes star-forming satellites from massive hosts, and biasing the global quenched fraction high.

Since the 300 kpc selection results in a more centrally located satellite population, it is likely that these satellites had an earlier infall time and underwent more ram pressure stripping when compared to satellites near or beyond 300 kpc from the host. This effect is present in Figure 8 as well, wherein the SAGA II comparison sample exhibits the strongest trend of quenched fraction with host mass. These results are consistent

with those from the TNG50 simulation (Engler et al. 2023), another large-volume, uniform-resolution simulation with comparable resolution to ROMULUS25.

5.2. Quenched Fraction Discrepancy

The shift from R_{vir} to 300 kpc, however, does not explain why our quenched fractions are higher than those of SAGA and ELVES (Figure 7). Our SAGA II comparison sample uses 300 kpc as a selection radius, and our results indicate that if SAGA and ELVES had access to the virial radii of their hosts, their quenched fractions would be lower. Donnari et al. (2021) find that the adopted definition of quenching and using 2D projected distances can both affect the resultant quenched fractions. Notably, the quenched fractions of ROMULUS25 are in better agreement with the observations when satellites with $\mu_{\text{eff},r} < 25 \text{ mag arcsecond}^{-2}$ are removed, in agreement with Font et al. (2022). The exception is the faintest bin, where a key factor may be the resolution of ROMULUS25. The lower resolution of the volume is unable to resolve a multiphase ISM, i.e., there is no extremely dense gas (Tremmel et al. 2019, 2020, and references within). Thus, all of the gas is “puffy” and overly susceptible to ram pressure stripping and quenching.

Dickey et al. (2021) found that large-scale cosmological simulations overquench isolated galaxies below $M_* = 10^9 M_\odot$ when compared to the Sloan Digital Sky Survey. The authors attribute this to overefficient feedback, which is typically tuned to recreate quenched fractions found in the Local Volume. In looking at ROMULUS25, Sharma et al. (2022a), also found that isolated dwarfs exhibit a higher quiescent fraction when compared to observations, but that this can be entirely attributed to the presence of massive black holes and their feedback. Although we are not studying isolated dwarfs in this work, it is still likely that these feedback properties are influencing our results. We note, however, that there are only six satellites in our SAGA II comparison sample with black holes and $M_* > 10^8 M_\odot$, so this does not notably affect our results.

5.3. Isolated versus Paired Hosts

Figure 9 suggests that pairs exhibit higher quenched fractions than more isolated analogs, but there are some caveats preventing us from making a more robust statement about the environment’s effect on the quenched fraction. First, we are only considering satellites with $M_* > 10^8 M_\odot$. Within our SAGA II sample, this is only $\sim 56\%$ of our total satellite population and they are hosted by $\sim 72\%$ of our Milky Way analogs with a nonzero satellite count (or $\sim 49\%$ of all Milky Way analogs), so a large section of our population is being removed. Second, our simulation box size prevents us from having a large sample of highly isolated analogs; only $\sim 14\%$ of our SAGA II sample analogs have $D_{\text{MW}} > 3 \text{ Mpc}$. Finally, by ignoring low-mass satellites, we are looking at the quenched fractions of several systems with few satellites (only one or two satellites). Around 43% of the high-mass satellite-hosting analogs in our SAGA II sample contain only one satellite above our resolution limit, so their quenched fractions can only occupy the extremes of 0 and 1, and in Figures 8 and 9 these systems are being considered equally alongside systems with as many as eight high-mass satellites (though the binning in Figure 9 should alleviate this issue). These combined effects yield a sample that is lacking low-mass satellites (and thus the analogs’ full satellite distributions) as well as highly isolated

hosts, making it difficult for us to extrapolate our results to the Universe at large.

Recently, Engler et al. (2023; TNG50) and Samuel et al. (2022; FIRE-2) found no difference in the satellite quenched fractions of paired and isolated hosts in their simulations. Further, Garrison-Kimmel et al. (2019; FIRE-2) found that the satellites of isolated Milky Way-mass galaxies have nearly identical SF histories to satellites of Milky Way analogs in Local Group-like pairs. However, these results were only considering satellites within 300 kpc of the host. In looking further out to 300–1000 kpc, Engler et al. (2023) find that paired, Local Group-like hosts exhibit significantly larger quenched fractions than their isolated counterparts.

6. Conclusions

Using the ROMULUS25 simulation, we have created various samples of Milky Way analogs along with their satellite distributions. We explored the role of host mass and environment on satellite numbers and quenched fractions. Our results can be summarized as follows.

1. When testing various criteria for defining a Milky Way analog, from more theoretically motivated (M_{vir}) to more observationally motivated (M_* and SAGA-like), we find that the resultant samples do not fully overlap. Within the overlapping regions, galaxies may also be defined as analogs in one sample but not in another due to environmental criteria (see Table 1 and Figure 1).
2. The number of satellites hosted by a Milky Way analog increases predominantly with host stellar mass, while environment may have a significant impact in high isolation ($D_{\text{MW+}} > 3 \text{ Mpc}$; see Figures 4–6).
3. The quenched fraction (for satellites with $M_* > 10^8 M_\odot$) of our analogs increases with host mass (see Figures 7(a) and 8(b)), but applying a surface brightness cut to satellites can erase this trend (see Figure 8(d)).
4. Being in a pair may yield higher satellite quenched fractions, but it is hard to draw statistically robust results given the small volume of ROMULUS25 and the fact that we can only study satellites down to $M_* = 10^8 M_\odot$ to avoid numerical overquenching (see Figure 9).

We find that the distributions of both the Milky Way and M31 are well explained by our sample, with M31 being at the highly populated edge of our sample space. This is in agreement with the SAGA and ELVES surveys, where ELVES found the Local Volume to be slightly more populated and exhibiting a steeper luminosity function when compared to the full SAGA sample. Additionally, we are in agreement with ELVES in finding that the stellar mass of a Milky Way analog seems to be the dominant factor in both the number of hosted satellites and the number of quenched satellites. Interestingly, in our study of quenching, we find that the SAGA and ELVES results are in good agreement for satellites with $M_* > 10^8 M_\odot$, suggesting that their discrepancy in quenched fraction comes from lower-mass satellites, which we are unable to probe here due to numerical effects that artificially quench simulated galaxies. However, our results support the notion put forward in Font et al. (2022) that SAGA is missing a large population of low surface brightness satellites near its detection limit that are preferentially quenched.

Acknowledgments

J.D.V. is supported by the Homer L. Dodge Fellowship from the University of Oklahoma. Long before the University of Oklahoma was established, the land on which the University now resides was the traditional home of the “Hasinai” Caddo Nation and Kirikiris Wichita & Affiliated Tribes; more information can be found here: <http://www.ou.edu/cas/nas/land-acknowledgement-statement>. C.C. was supported by the NSF under CAREER grant AST-1848107. A.M.B. was partially supported by NSF grant AST-1813871. M.T. was supported by an NSF Astronomy and Astrophysics Postdoctoral Fellowship under award AST-2001810. The Romulus simulations are part of the Blue Waters Sustained Petascale Computing project, which is supported by the National Science Foundation (awards OCI-0725070 and ACI-1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana–Champaign and its National Center for Supercomputing Applications. This work is also part 12 of a Petascale Computing Resource Allocations allocation support by the National Science Foundation (award number OAC-1613674). This work also used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant No. ACI-1548562.

Data Availability

The data for this work were generated from a proprietary branch of the CHANGA N-Body+SPH code (Menon et al. 2015). The public repository for CHANGA is available on github <https://github.com/N-BodyShop/changa>. Analysis was conducted using the publicly available softwares pynbody (Pontzen et al. 2013; <https://github.com/pynbody/pynbody>) and TANGOS (Pontzen & Tremmel 2018; <https://github.com/pynbody/tangos>). These results were generated from the ROMULUS25 cosmological simulation. The raw output from this simulation can be accessed upon request from Michael Tremmel (mtremmel@ucc.ie), along with the TANGOS database files that were generated from these outputs and directly used for this analysis.

ORCID iDs

Jordan Van Nest  <https://orcid.org/0000-0003-3789-3722>
Ferah Munshi  <https://orcid.org/0000-0002-9581-0297>
Charlotte Christensen  <https://orcid.org/0000-0001-6779-3429>
Alyson M. Brooks  <https://orcid.org/0000-0002-0372-3736>
Michael Tremmel  <https://orcid.org/0000-0002-4353-0306>

References

Akins, H. B., Christensen, C. R., Brooks, A. M., et al. 2021, *ApJ*, **909**, 139
Bennet, P., Sand, D. J., Crnojević, D., et al. 2019, *ApJ*, **885**, 153
Bromm, V., Ferrara, A., Coppi, P. S., & Larson, R. B. 2001, *MNRAS*, **328**, 969
Bryan, G. L., & Norman, M. L. 1998, *ApJ*, **495**, 80
Carlsten, S. G., Greco, J. P., Beaton, R. L., & Greene, J. E. 2020, *ApJ*, **891**, 144
Carlsten, S. G., Greene, J. E., Beaton, R. L., Danieli, S., & Greco, J. P. 2022, *ApJ*, **933**, 47
Carlsten, S. G., Greene, J. E., Greco, J. P., Beaton, R. L., & Kado-Fong, E. 2021a, *ApJ*, **922**, 267
Carlsten, S. G., Greene, J. E., Peter, A. H. G., Beaton, R. L., & Greco, J. P. 2021b, *ApJ*, **908**, 109
Dickey, C. M., Starkenburg, T. K., Geha, M., et al. 2021, *ApJ*, **915**, 53
Donnari, M., Pillepich, A., Nelson, D., et al. 2021, *MNRAS*, **506**, 4760
Drlica-Wagner, A., Bechtol, K., Mau, S., et al. 2020, *ApJ*, **893**, 47
Drlica-Wagner, A., Bechtol, K., Rykoff, E. S., et al. 2015, *ApJ*, **813**, 109

Engler, C., Pillepich, A., Joshi, G. D., et al. 2023, *MNRAS*, **522**, 5946

Engler, C., Pillepich, A., Pasquali, A., et al. 2021, *MNRAS*, **507**, 4211

Font, A. S., McCarthy, I. G., & Belokurov, V. 2021, *MNRAS*, **505**, 783

Font, A. S., McCarthy, I. G., Belokurov, V., Brown, S. T., & Stafford, S. G. 2022, *MNRAS*, **511**, 1544

Font, A. S., McCarthy, I. G., Poole-Mckenzie, R., et al. 2020, *MNRAS*, **498**, 1765

Garrison-Kimmel, S., Wetzel, A., Hopkins, P. F., et al. 2019, *MNRAS*, **489**, 4574

Geha, M., Blanton, M. R., Yan, R., & Tinker, J. L. 2012, *ApJ*, **757**, 85

Geha, M., Wechsler, R. H., Mao, Y.-Y., et al. 2017, *ApJ*, **847**, 4

Governato, F., Weisz, D., Pontzen, A., et al. 2015, *MNRAS*, **448**, 792

Guedes, J., Callegari, S., Madau, P., & Mayer, L. 2011, *ApJ*, **742**, 76

Haardt, F., & Madau, P. 2012, *ApJ*, **746**, 125

Harris, W. E., & van den Bergh, S. 1981, *AJ*, **86**, 1627

Ibata, R. A., Lewis, G. F., McConnachie, A. W., et al. 2014, *ApJ*, **780**, 128

Karunakaran, A., Spekkens, K., Oman, K. A., et al. 2021, *ApJL*, **916**, L19

Knollmann, S. R., & Knebe, A. 2009, *ApJS*, **182**, 608

Koposov, S., Belokurov, V., Evans, N. W., et al. 2008, *ApJ*, **686**, 279

Koposov, S. E., Belokurov, V., Torrealba, G., & Evans, N. W. 2015, *ApJ*, **805**, 130

Kravtsov, A. V., Vikhlinin, A. A., & Meshcheryakov, A. V. 2018, *AstL*, **44**, 8

Kroupa, P. 2001, *MNRAS*, **322**, 231

Ludlow, A. D., Schaye, J., Schaller, M., & Richings, J. 2019, *MNRAS Lett.*, **488**, L123

Mao, Y.-Y., Geha, M., Wechsler, R. H., et al. 2021, *ApJ*, **907**, 85

Martin, N. F., Ibata, R. A., Lewis, G. F., et al. 2016, *ApJ*, **833**, 167

Mateo, M. L. 1998, *ARA&A*, **36**, 435

McConnachie, A. W., Ibata, R., Martin, N., et al. 2018, *ApJ*, **868**, 55

Menon, H., Wesolowski, L., Zheng, G., et al. 2015, *ComAC*, **2**, 1

Moster, B. P., Naab, T., & White, S. D. M. 2013, *MNRAS*, **428**, 3121

Munshi, F., Governato, F., Brooks, A. M., et al. 2013, *ApJ*, **766**, 56

Nashimoto, M., Tanaka, M., Chiba, M., et al. 2022, *ApJ*, **936**, 38

Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, *A&A*, **594**, A13

Pontzen, A., Governato, F., Pettini, M., et al. 2008, *MNRAS*, **390**, 1349

Pontzen, A., Roškar, R., Stinson, G., & Woods, R. 2013, *pynbody: N-Body/SPH Analysis for Python*, Astrophysics Source Code Library, ascl:[1305.002](https://ascl.net/1305.002)

Pontzen, A., & Tremmel, M. 2018, *ApJS*, **237**, 23

Ricarte, A., Tremmel, M., Natarajan, P., & Quinn, T. 2019, *MNRAS*, **489**, 802

Saitoh, T. R., & Makino, J. 2009, *ApJL*, **697**, L99

Samuel, J., Wetzel, A., Santistevan, I., et al. 2022, *MNRAS*, **514**, 5276

Sharma, R. S., Brooks, A. M., Tremmel, M., Bellovary, J., & Quinn, T. R. 2022a, arXiv:[2211.05275](https://arxiv.org/abs/2211.05275)

Sharma, R. S., Brooks, A. M., Tremmel, M., et al. 2022b, *ApJ*, **936**, 82

Shen, S., Wadsley, J., & Stinson, G. 2010, *MNRAS*, **407**, 1581

Simon, J. D. 2019, *ARA&A*, **57**, 375

Simons, R. C., Peebles, M. S., Tumlinson, J., et al. 2020, *ApJ*, **905**, 167

Simpson, C. M., Grand, R. J. J., Gómez, F. A., et al. 2018, *MNRAS*, **478**, 548

Smercina, A., Bell, E. F., Price, P. A., et al. 2018, *ApJ*, **863**, 152

Stinson, G., Seth, A., Katz, N., et al. 2006, *MNRAS*, **373**, 1074

Tremmel, M., Governato, F., Volonteri, M., & Quinn, T. R. 2015, *MNRAS*, **451**, 1868

Tremmel, M., Karcher, M., Governato, F., et al. 2017, *MNRAS*, **470**, 1121

Tremmel, M., Quinn, T. R., Ricarte, A., et al. 2019, *MNRAS*, **483**, 3336

Tremmel, M., Wright, A. C., Brooks, A. M., et al. 2020, *MNRAS*, **497**, 2786

Wadsley, J. W., Keller, B. W., & Quinn, T. R. 2017, *MNRAS*, **471**, 2357

Wadsley, J. W., Stadel, J., & Quinn, T. 2004, *NewA*, **9**, 137

Wadsley, J. W., Veeravalli, G., & Couchman, H. M. P. 2008, *MNRAS*, **387**, 427

Wetzel, A. R., Tollerud, E. J., & Weisz, D. R. 2015, *ApJL*, **808**, L27

Wright, A. C., Tremmel, M., Brooks, A. M., et al. 2021, *MNRAS*, **502**, 5370