The Astrophysical Journal Letters, 993:L55 (7pp), 2025 November 10 https://doi.org/10.3847/2041-8213/ae1584

The First RELHIC? Cloud-9 is a Starless Gas Cloud*

Gagandeep S. Anand[1], Alejandro Benítez-Llambay[2], Rachael Beaton[1], Andrew J. Fox[3], Julio F. Navarro[4], and Elena D'Onghia[5]
[1] Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA
[2] Dipartimento di Fisica G. Occhialini, Universitá degli Studi di Milano Bicocca, Piazza della Scienza, 3 I-20126 Milano MI, Italy
[3] AURA for ESA, Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA
[4] Department of Physics and Astronomy, University of Victoria, Victoria, BC V8P 5C2, Canada
[5] Department of Astronomy, University of Wisconsin, Madison, WI 53706, USA
Received 2025 August 27; revised 2025 October 6; accepted 2025 October 21; published 2025 November 10

Abstract

Five-hundred-meter Aperture Spherical Telescope observations have recently identified a compact H I cloud (hereafter Cloud-9) in the vicinity of the spiral galaxy M94. This identification has been confirmed independently by Very Large Array and Green Bank Telescope observations. Cloud-9 has the same recession velocity as M94, and is therefore at a similar distance ($\sim$4.4 Mpc). It is compact ($\sim 1'$ radius, or $\sim$1.4 kpc), dynamically cold ($W_{50} = 12\,\mathrm{km\,s^{-1}}$), nonrotating, and fairly massive, with an H I mass of $\sim 10^6\,M_\odot$. Here we present deep Hubble Space Telescope/Advanced Camera for Surveys imaging designed to search for a luminous stellar counterpart. We visually rule out the presence of any dwarf galaxy with stellar mass exceeding $10^{3.5}\,M_\odot$. A more robust color–magnitude diagram-based analysis conservatively rules out a $10^4\,M_\odot$ stellar counterpart with $99.5^{+0.5}_{-8.2}\%$ confidence. The nondetection of a luminous component reinforces the interpretation that this system is a reionization-limited H I cloud (RELHIC); i.e., a starless dark matter halo filled with hydrostatic gas in thermal equilibrium with the cosmic ultraviolet background. Our results make Cloud-9 the leading RELHIC candidate of any known compact H I cloud. This provides strong support for a cornerstone prediction of the Lambda cold dark matter model, namely the existence of gas-filled starless dark matter halos on subgalactic mass scales, and constrains the present-day threshold halo mass for galaxy formation.

Unified Astronomy Thesaurus concepts: Galaxies (573); Hubble Space Telescope (761); Cosmology (343)

1. Introduction

A fundamental prediction of the Lambda cold dark matter (ΛCDM) model is the formation of dark matter halos over a vast range of scales, with a mass function that rises steeply toward low masses (W. H. Press & P. Schechter 1974; A. Jenkins et al. 2001). The abundance of halos far exceeds that of known galaxies, implying that not all halos are able to host luminous galaxies. This has been interpreted to mean that galaxies only form in halos that exceed a "critical" mass, $M_{\rm crit}(z)$, a result supported by numerical simulations and theoretical modeling (e.g., G. R. Blumenthal et al. 1984; M. Hoeft et al. 2006; T. Okamoto & C. S. Frenk 2009; T. Sawala et al. 2016a; A. Benitez-Llambay & C. Frenk 2020, and references therein). After reionization, this threshold is set by the balance between gravity, gas cooling, and photoheating by the cosmic ultraviolet background (UVB). At the present epoch, $M_{\rm crit} \approx 10^{9.7}\,M_\odot$ (e.g., A. Benitez-Llambay & C. Frenk 2020; O. Nebrin et al. 2023).

Halos below $M_{\rm crit}$ do not retain much photoheated gas and, unless they were able to form stars before reionization, remain starless. By contrast, halos above $M_{\rm crit}$ are massive enough to overcome thermal pressure, enabling galaxy formation. Halos just below $M_{\rm crit}$ occupy an intermediate regime: they are not massive enough to form stars today, but are massive enough to retain some of their gas. Of these, the most massive retain enough baryons in their central regions for hydrogen recombination to occur. These massive, H I-rich, starless systems were termed reionization-limited H I clouds (RELHICs) by A. Benítez-Llambay et al. (2017).

Cosmological hydrodynamical simulations suggest that, despite lacking stars, RELHICs should host compact, nearly spherical cores of neutral hydrogen (A. Fattahi et al. 2016; T. Sawala et al. 2016b; A. Benítez-Llambay et al. 2017). The gas is in hydrostatic equilibrium with the dark matter and in thermal equilibrium with the UVB. These properties make RELHICs detectable through their narrow 21 cm emission line, with characteristic velocity widths of $W_{50} \lesssim 20\,\mathrm{km\,s^{-1}}$. These simulations also indicate that RELHICs are susceptible to ram pressure stripping, and, therefore, are unlikely to be found very close to a massive galaxy.

Because their gas is pressure-supported, RELHICs offer a unique window into the underlying dark matter distribution, bypassing the complexities of galaxy formation that often prevent robust inferences on the nature of dark matter (see, e.g., J. S. Bullock & M. Boylan-Kolchin 2017 for a review). Detecting RELHICs would be a breakthrough, providing clear evidence for bound, collapsed dark matter halos on subgalactic scales, and offering constraints on the nature of dark matter in a previously unexplored regime.

A few RELHIC candidates have recently been identified with the Five-hundred-meter Aperture Spherical Telescope (FAST), which is well suited for surveying large areas of the northern sky in H I. An example is FAST J0139+4328

* Based on observations made with the NASA/ESA Hubble Space Telescope, obtained at the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS5-26555. These observations are associated with program 17712.

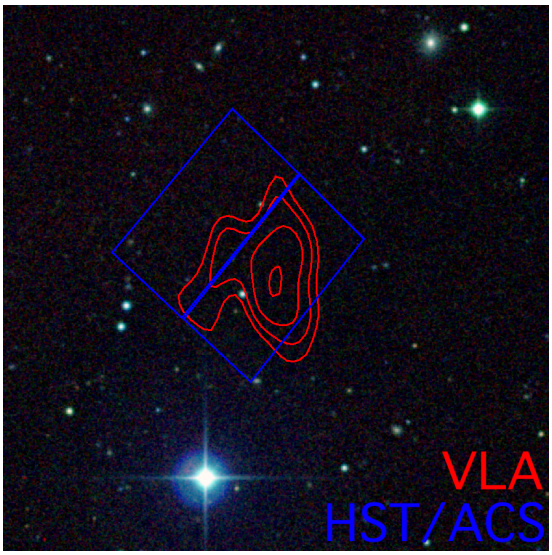

Figure 1. Digitized Sky Survey image covering a $10' \times 10'$ region around Cloud-9. The VLA H I contours (A. Benítez-Llambay et al. 2024) are shown in red, and the footprint of our HST/ACS observations is shown in blue.

(J.-L. Xu et al. 2023), which has no detected luminous counterpart yet contains an H I mass of $\approx 8.3 \times 10^7\,M_\odot$. However, the double-peaked shape of its H I profile indicates that the gas is rotationally supported, suggesting that the gas has collapsed into a disk, and that J0139+4328 may actually contain a faint galaxy. Indeed, the current upper limit on the stellar mass of any luminous counterpart is only $\approx 7 \times 10^5\,M_\odot$. This limit should be improved by deeper imaging before it can be claimed that J0139+4328 is actually "dark."

Other notable candidates include some of the ALFALFA ultracompact high-velocity clouds (HVCs) cataloged by E. A. K. Adams et al. (2013). However, most of the candidate clouds are likely associated with the Milky Way disk, and, therefore, not RELHICs, which could not survive the ram pressure stripping due to Galactic gas. In addition, those clouds not associated with the Milky Way are either too extended, too irregular, or exhibit line widths too broad to be consistent with RELHICs (A. Benítez-Llambay et al. 2017). The scarcity of robust candidates is perhaps unsurprising: the relatively shallow depth of ALFALFA limits sensitivity to the low fluxes predicted for RELHICs, while the large beam of the Arecibo radio telescope prevents them from being spatially resolved.

Among all candidate detections to date, one stands out as a particularly compelling: Cloud-9, recently identified in the vicinity of M94 with FAST by R. Zhou et al. (2023). Unlike previous dark H I cloud candidates, Cloud-9 exhibits properties consistent with RELHICs: its recessional velocity ($304\,\mathrm{km\,s^{-1}}$) is positive and similar to that of M94; its 21 cm line profile has a width of $W_{50} \lesssim 20\,\mathrm{km\,s^{-1}}$, and its total H I flux is consistent with that expected for a RELHIC at the M94 distance (A. Benitez-Llambay & J. F. Navarro 2023; R. Zhou et al. 2023; A. Benítez-Llambay et al. 2024; A. Karunakaran & K. Spekkens 2024). The small angular separation from M94 ($\sim$70 kpc in projection) and comparable recessional velocity suggest a physical association, providing a lower limit on their relative distance.

Higher-resolution Very Large Array (VLA) observations (see Figure 1) further support the idea that Cloud-9 is at M94's distance: the column density isocontours—which appeared round in FAST data—appear slightly distorted in the VLA image, possibly indicating ram pressure interactions with the gaseous halo of M94 (A. Benítez-Llambay et al. 2024).

Detailed analysis of Cloud-9 indicates that, at the assumed distance of M94, the system requires a substantial amount of gravitational mass beyond its H I content to remain in hydrostatic equilibrium (A. Benitez-Llambay & J. F. Navarro 2023; A. Benítez-Llambay et al. 2024). Interpreting this additional mass as cold dark matter yields a tight constraint on the total halo mass of $\sim 5 \times 10^9\,M_\odot$, remarkably close to $M_{\rm crit}$ at the present day. Moreover, a comparison with mock RELHIC models shows that Cloud-9's properties are broadly consistent with a ΛCDM RELHIC, with a total neutral hydrogen mass of $M_{\rm HI} \sim 10^6\,M_\odot$ (A. Benitez-Llambay & J. F. Navarro 2023). Together, these results make Cloud-9 the firmest RELHIC candidate known to date.

A key challenge to interpreting Cloud-9 as a RELHIC lies in the loose upper limit on its stellar content available at present. Its H I mass and radial distribution closely resemble those of Leo T, a dwarf galaxy in the outskirts of the Milky Way (M. J. Irwin et al. 2007; A. Benítez-Llambay et al. 2024). Notably, Leo T's stellar mass ($\sim 10^5\,M_\odot$) is comparable to the current upper limit for a luminous counterpart to Cloud-9, based on imaging from the DESI Legacy Survey (R. Zhou et al. 2023). This makes it difficult to exclude the presence of a Leo T–like stellar counterpart. Confirming the RELHIC nature of Cloud-9 therefore requires deeper observations capable of significantly lowering this upper limit.

In this Letter, we present deep Hubble Space Telescope (HST) Advanced Camera for Surveys (ACS) imaging of Cloud-9, obtained to search explicitly for a stellar counterpart. In Section 2 we describe the HST observations and data reduction. In Section 3 we place limits on the stellar mass of a stellar counterpart to Cloud-9. We discuss our results in Section 4 and use Section 5 to provide a brief summary.

2. HST Observations and Reductions

We obtained observations with the HST/ACS Wide-Field Channel (WFC) as part of program GO-17712 in Cycle 32 (A. Benitez-Llambay et al. 2024). The observations were taken in four visits of two orbits each between 2025 February 17 and 19, and were designed to target a depth of 4 mag below the tip of the red giant branch (TRGB) at the distance of M94 ($d_{\rm M94} = 4.41$ Mpc; G. S. Anand et al. 2021).

The nominal position of Cloud-9 (based on the VLA H I maximum) was placed at the center of the WFC1 chip (see Figure 1). Each visit employed a small, 5 pixel dither to allow rejection of cosmic rays and other detector artifacts. A loss of guide stars occurred near the end of the final F814W exposure during the first visit, truncating that frame, but otherwise the dataset was unaffected. In total, the observations comprise 9036 s in the F606W filter ("wide-V"), and 8749.4 s in the F814W filter (approximately *I* band).

Because our data were obtained in four separate visits, we first aligned all underlying exposures using the charge-transfer-corrected FLC files. For the F606W images, we applied the tweakreg procedure from the DrizzlePac software suite (R. J. Avila et al. 2015) to align the individual exposures, and then combined them into a deep "drizzled" image with AstroDrizzle (also within the DrizzlePac package). We repeated this process independently for the

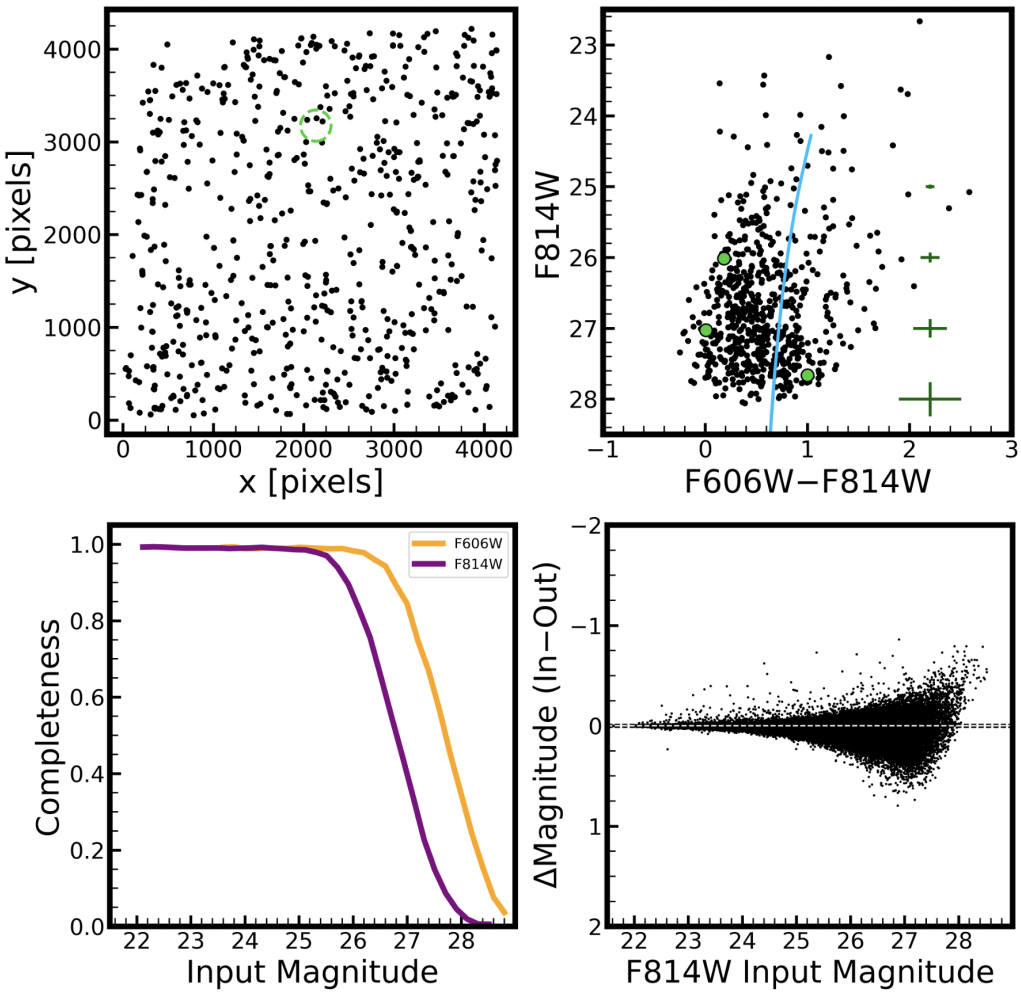

Figure 2. Top left: the spatial distribution of sources fulfilling our quality criteria. The green dashed circle indicates the effective radius (8.″4 or 180 pc projected) of a Leo T analog at the distance of M94, and is centered at the location of Cloud-9. Top right: CMD of the field. The three sources within the green circle in the top left panel are shown in green. We also show an RGB isochrone for an old, metal-poor stellar population (10 Gyr, [Fe/H] = −2.0 dex) drawn from MIST, and placed at the distance of M94. The green crosses on the right-hand side indicate typical photometric uncertainties for various F814W magnitudes at the color of the RGB isochrone. Bottom left: photometric completeness distributions for both of our filters, as determined by the artificial star experiments. Bottom right: difference between input and output magnitudes from the artificial star experiments for the primary CMD filter (F814W).

F814W data. We then mutually aligned the two drizzled images, and used `tweakback` to back-propagate this improved alignment to the original `FLC` files. Using these adjusted `FLC` exposures, we again created drizzled images of the fields in the separate filters, ensuring consistent alignment across the full dataset.

We performed point-spread function (PSF) photometry with the `DOLPHOT` stellar photometry package (A. E. Dolphin 2000a; A. Dolphin 2016), adopting the reduction parameters recommended by the PHAT program (B. F. Williams et al. 2014, 2023) for ACS/WFC imaging. These parameters have been extensively validated across a wide range of extragalactic fields. The F814W drizzled image served as the undistorted reference frame, while the PSF photometry was carried out on the individual `FLC` images to avoid issues introduced by resampling of the PSF in the drizzled images. All magnitudes are reported in the Vega photometric system.

The base photometric output was culled using quality cuts—we tested several iterations from the literature (B. F. Williams et al. 2014; G. S. Anand et al. 2021; J. T. Warfield et al. 2023; C. E. Murray et al. 2024), and found that the following relatively strict cuts provided the most effective rejection of artifacts, such as spurious sources found in the diffraction spikes of bright foreground stars: $\texttt{Crowd} < 0.5$ (both filters), $|\texttt{Sharp}| < 0.1$ (both filters), $\texttt{Type} \leqslant 1$, $\texttt{S/N} \geqslant 5$ (both filters), and $\texttt{Flag} \leqslant 2$ (both filters). The final catalog contains 667 sources.

To assess photometric bias, completeness, and measurement uncertainties, we conducted artificial star experiments with `DOLPHOT`. We injected and recovered approximately 500,000 artificial stars one at a time using the same reduction setup. The same quality cuts were applied to the resulting artificial star catalog. The bottom panels of Figure 2 show our photometric completeness distributions in both filters, as well as the differences between input and output magnitudes for artificial stars in our primary filter (F814W).

The color–magnitude diagram (CMD) of the full ACS field, together with the spatial distribution of detected sources, is shown in Figure 2. This CMD closely resembles those derived from other "background" fields, such as the Hubble eXtreme Deep Field (I. S. Jang & M. G. Lee 2017). In the top left panel, a green dashed circle marks the location

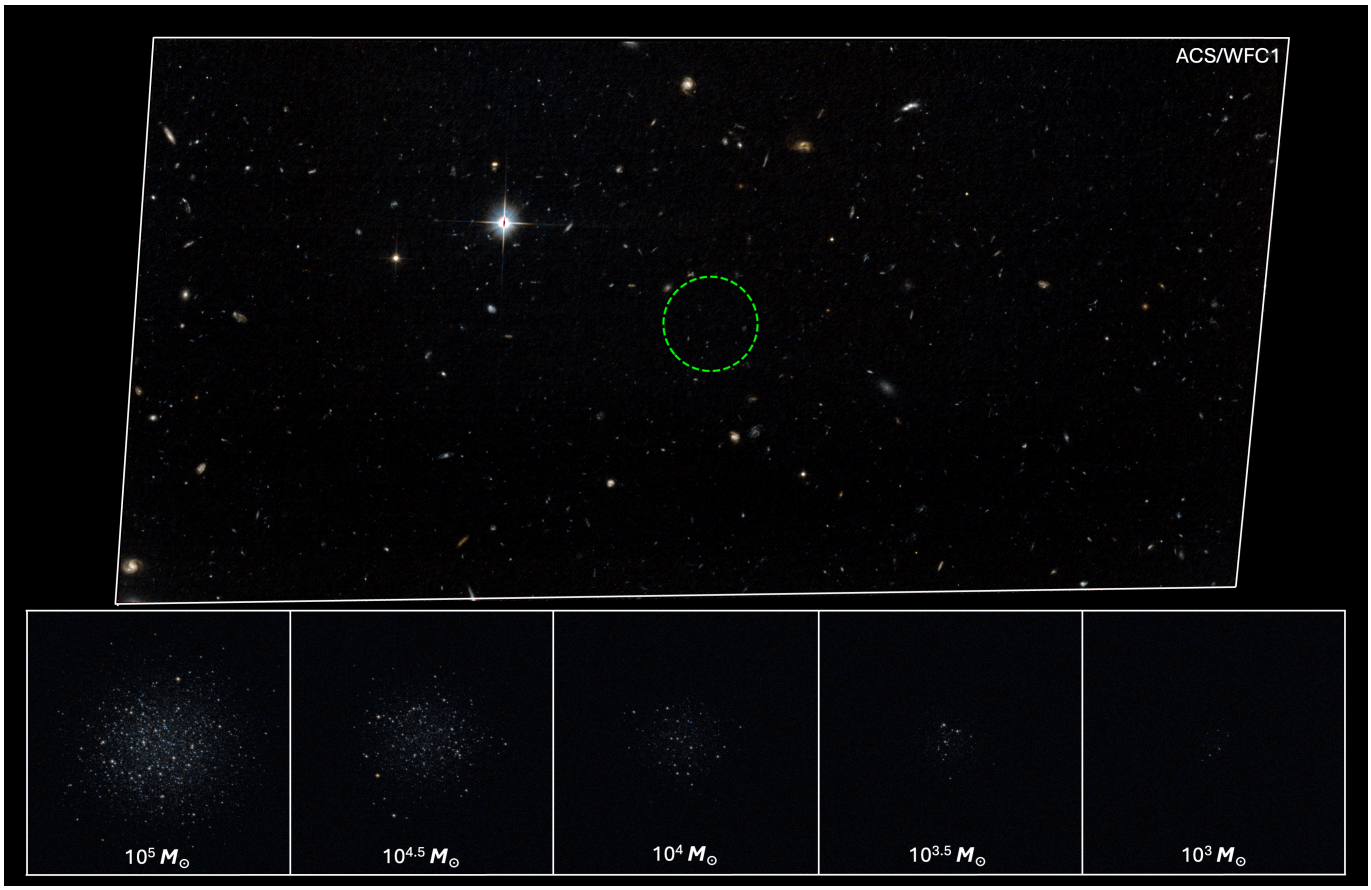

Figure 3. Top: a color composite of our ACS/WFC1 imaging ($\sim 202'' \times 100''$). The green circle ($r = 8\farcs4$) marks the VLA H I maximum column density, with a radius corresponding to the effective radius of a Leo T analog at the distance of M94. Bottom: simulated dwarf galaxies spanning a range of stellar masses, generated under the same observing conditions and shown at the same physical scale as the ACS data. A Leo T analog ($M_\star = 10^5\,M_\odot$) would be readily detected, but no stellar counterpart is visible down to at least $M_\star \sim 10^{3.5} M_\odot$.

of the VLA H I maximum, with a radius of $8\farcs4$, corresponding to the effective radius of Leo T (180 pc) at the distance of M94. In the top right panel, green symbols highlight the (only) three sources that fall within this region, alongside a 10 Gyr, metal-poor ([Fe/H] $= -2.0$) red giant branch (RGB) isochrone at this distance, drawn from MIST (J. Choi et al. 2016; A. Dotter 2016). The green crosses on the right indicate the observed photometric uncertainties for stars spanning a range of magnitudes at the color of the isochrone.

3. Searching for a Stellar Counterpart

3.1. Visually Ruling Out a Leo T Analog

With our deep ACS imaging and photometry in hand, we examined the field for a stellar counterpart to Cloud-9. Using the `ArtPop` software package (J. P. Greco & S. Danieli 2022), we generated simulated images of dwarf galaxies at the distance of M94 under our ACS/WFC observing conditions, including the telescope parameters, filter set, exposure times, photometric zero-points, and PSFs.

Following the stellar mass–size relation for nearby dwarf satellites of Milky Way analogs (including M94) from S. G. Carlsten et al. (2021), we adopted as a starting point a Leo T analog with $M_\star = 10^5\,M_\odot$ and size $r_{\rm eff} = 180$ pc (M. J. Irwin et al. 2007; D. R. Weisz et al. 2012), and decreased the stellar mass by factors of $\sim$3 down to $M_\star = 10^3\,M_\odot$ (with $r_{\rm eff} = 57$ pc). We assumed an old, metal-poor stellar population of age 10 Gyr and [Fe/H] $= -2.0$, noting that the precise metallicity has little impact on the results. We also note that the adoption of an old stellar population is the most conservative, as a stellar counterpart with any fraction of its stars at a young or intermediate age would produce even more numerous and bright stars that would be visible in our HST observations.

The top panel of Figure 3 shows our ACS/WFC1 imaging, centered on the VLA H I centroid. The bottom panel shows the artificially observed dwarf galaxies, shown at the same physical scale as the ACS data. It is clear that a Leo T analog with $M_\star \sim 10^5\,M_\odot$ would be readily visible if present. Such a counterpart would produce dozens of resolved stars above our magnitude limit (see the HST CMD in D. R. Weisz et al. 2012). In contrast, no comparable resolved structure is present in the image, with nondetections extending down to $M_\star \sim 10^{3.5}\,M_\odot$ or lower (though we caution that stochasticity in the number of bright stars is significant at these low stellar masses). While these simulations do not capture all the intricate detector effects associated with ACS—such as charge-transfer efficiency losses or bias striping—these effects are largely corrected by the ACS pipeline. Importantly, the simulations are intended primarily as a visual tool, illustrating what a hypothetical stellar counterpart to Cloud-9 would look like if present in our data. In the next subsection, we introduce a more quantitative framework for our analysis.

3.2. Stellar Population Simulations

Given the absence of a visible stellar counterpart to Cloud-9 in our HST imaging, we next quantify the limits on its underlying stellar mass. A system with $M_\star = 10^5\,M_\odot$ would be unmistakable, but we also considered a possible $M_\star = 10^4\,M_\odot$ counterpart. While many star formation history codes (e.g., A. E. Dolphin 2002; B. A. Jacobs et al. 2011; L. N. Makarova et al. 2023) can fit for a total stellar mass in their procedure, they are not designed to determine upper limits for stellar mass content in the apparent absence of stars. Instead, we used `ArtPop` (J. P. Greco & S. Danieli 2022) to simulate stellar (Vega) magnitudes for a given input stellar mass in our ACS filters, incorporating the observational uncertainties and completeness measured from our artificial star experiments. We adopted the same stellar population parameters (10 Gyr old population with Fe/H $= -2.0$ at the distance of M94) and generated 10,000 CMD realizations to capture stochastic variations at this relatively low stellar mass. To mimic the conditions of our HST observations, we applied our full observational selection functions—photometric bias, completeness, and uncertainties—to each simulated CMD.

Turning to the observed CMD, we detect three nominal sources within the $8\farcs4$ region centered on Cloud-9. To account for the $\sim 9''$ uncertainty in the position of the H I centroid, we shift the center of our nominal region within this positional uncertainty $\sim$200,000 times, yielding a mean of 3.5 ± 1 sources associated with Cloud-9. However, because of the large population of unresolved background galaxies (or substructures within them), these sources may not be genuine stars. To quantify the background source population, we use the opposing (WFC2) chip. We place $\sim$200,000 same-sized regions across the chip, shifting their centers to cover the full WFC2 area (avoiding chip edges). This yields a mean of 3.7 ± 2 background sources per region. Since we observe 3.5 ± 1 sources in the Cloud-9 region, this corresponds to an overdensity of -0.2 ± 2.2 sources associated with Cloud-9, consistent with no excess at the location of the VLA H I maximum.

Figure 4 shows the number of stars recovered in 10,000 simulation iterations, performed separately for a range of input stellar masses. For a population with $M_\star = 10^4\,M_\odot$, a source was recovered in 99.5% of cases, allowing us to rule out such a population with that confidence level. To incorporate the uncertainty in the derived overdensity (-0.2 ± 2.2 sources) we proceed as follows. On the high end, the maximum overdensity corresponds to 2.0 sources; in the same simulations, 8.7% of realizations yield two or fewer stars. On the low end, a source population of -2.4 sources is nonphysical, and thus excluded with 100% confidence. Taken together, this results in a $99.5^{+0.5}_{-8.2}\%$ confidence limit against a $M_\star = 10^4\,M_\odot$ stellar counterpart. Finally, we stress that these limits remain conservative: a genuine stellar counterpart would produce additional resolved or semiresolved stars that fail our CMD-selection criteria. No such population is observed, further strengthening the case against a stellar counterpart to Cloud-9.

We can also frame the question in terms of maximum stellar mass consistent with our observations. By running simulations across a wide range of stellar masses, we find that a population with $M_\star = 10^{3.49} M_\odot$ best matches the data at the nominal center: after including observational biases, incompleteness, and photometric errors, such a system yields on average one or fewer RGB stars, consistent with our observations. We therefore adopt $M_\star = 10^{3.5} M_\odot$ as the baseline upper limit for any stellar counterpart to Cloud-9.

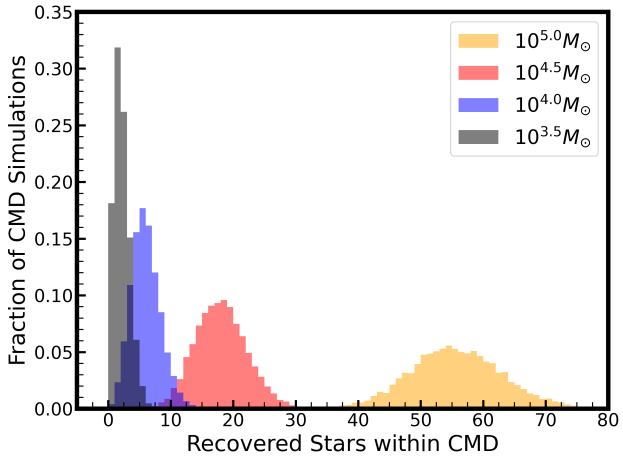

Figure 4. Histograms of 10,000 simulations performed for each of a range of input stellar masses (shown in the legend). The x-axis values show the remaining number of visible stars in the simulated CMDs after applying our observational systematics.

Finally, we note that ram pressure interactions can in principle induce a spatial offset between the stellar and gaseous components of low-mass galaxies. However, the H I morphology of Cloud-9 is relatively regular with only mild asymmetry, suggesting that any such effect would be small. To test this possibility directly, we searched for clusters of point sources within a radius of 1 kpc ($47''$) from the VLA H I maximum, comparable to the extent of the compact H I core (for reference, the top-to-bottom angular size of the WFC1 chip shown in Figure 3 is $\sim 100''$). In this region we identified a cluster of nine sources. However, given that the search covers a large fraction of the WFC1 chip, such local maxima are expected from the underlying background galaxy clustering. A control search on the opposing WFC2 chip from its center outward yields two nonoverlapping clusters of nine sources, confirming that the WFC1 cluster is not a statistically significant overdensity.

4. Discussion

The absence of a faint stellar counterpart associated with Cloud-9 is striking and reinforces its identification as a likely RELHIC. While isolated ultrafaint dwarfs exist with stellar masses comparable to, or below, the upper limit from our analysis, none host such a substantial neutral hydrogen reservoir. Adopting a total H I mass of $M_{\rm HI} \approx 1.4 \times 10^6\,M_\odot$ from Green Bank Telescope observations (A. Karunakaran & K. Spekkens 2024), together with our stellar mass upper limit, implies an H I-to-stellar mass ratio of $M_{\rm HI}/M_\star \gtrsim 443$, exceeding by orders of magnitude the ratios typical of dwarf galaxies. For comparison, Leo T, KK 153 (J.-L. Xu et al. 2025), and Leo P (K. B. W. McQuinn et al. 2015) have similar H I masses but stellar counterparts yielding $M_{\rm HI}/M_\star \lesssim 4$, consistent with typical values for gas-rich dwarfs (see, e.g., F. Lelli 2022; I. D. Karachentsev et al. 2024). Figure 5 places Cloud-9 in this context, showing it relative to Leo P, Leo T, KK 153, newly detected dwarfs with FAST (I. D. Karachentsev et al. 2024), and the ALFALFA Sloan Digital Sky Survey catalog (A. Durbala et al. 2020).

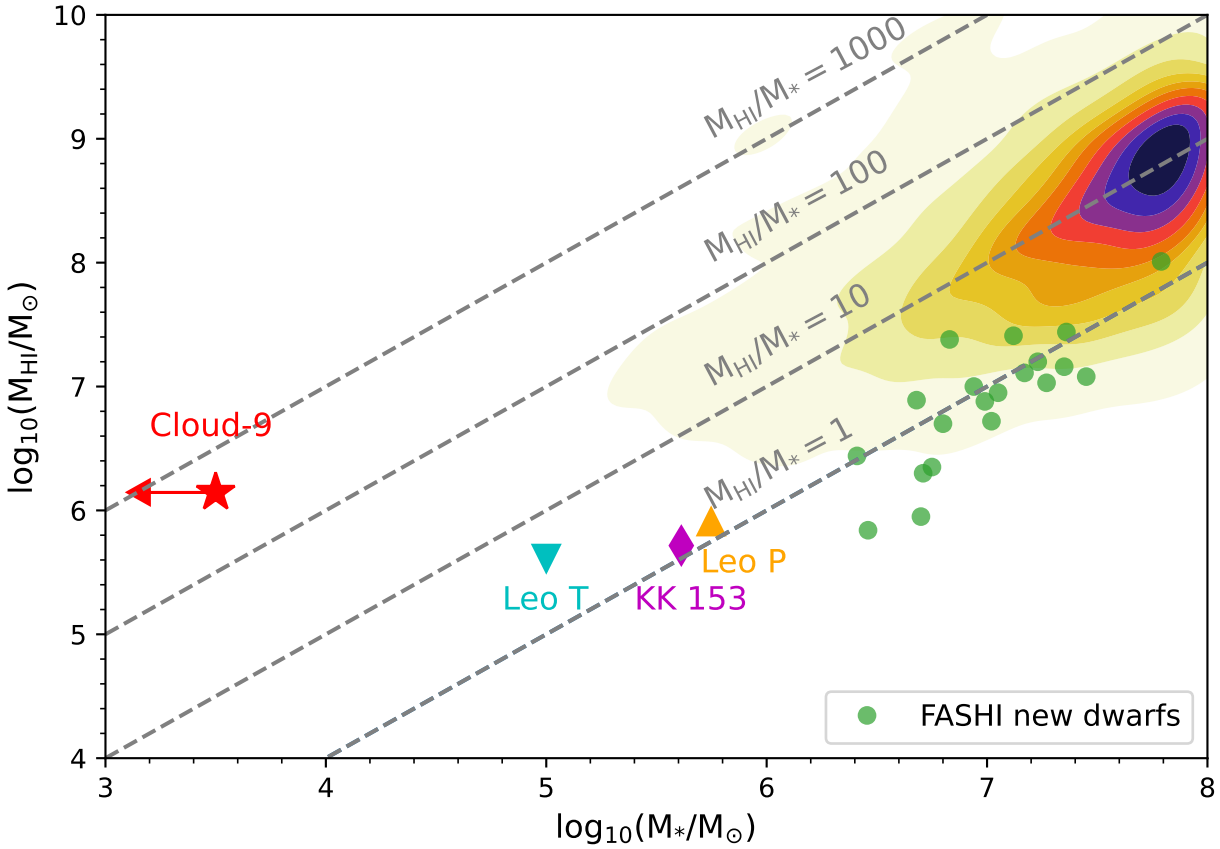

Figure 5. H I–stellar mass relation for the ALFALFA Sloan Digital Sky Survey catalog (isocontours; A. Durbala et al. 2020). Newly identified dwarf galaxies in the local volume detected with the FAST All Sky HI Survey (FASHI) are shown as green points (I. D. Karachentsev et al. 2024). For comparison, Leo T, KK 153, Leo P, and Cloud-9 are also shown. The stellar mass upper limit for Cloud-9 (red star symbol) underscores its extreme gas richness relative to its possible stellar mass, if interpreted as an ultrafaint galaxy.

One can consider alternative interpretations for Cloud-9. One possibility is that it is a foreground HVC of the Milky Way that happens to be projected close on the sky to M94. This scenario is disfavored by Cloud-9's recessional velocity, which closely matches M94's velocity, as well as the absence of foreground HVCs with comparable velocities in this region of the sky. The closest HVCs to M94 are complexes M, C, and K (B. P. Wakker & H. van Woerden 1991; T. Westmeier 2018), but they all have negative velocities ($v_{\rm LSR} \approx -200$ to $-100\,{\rm km\,s^{-1}}$), whereas M94 and Cloud-9 are located at $v_{\rm LSR} \approx +300\,{\rm km\,s^{-1}}$. A second (related) possibility is that Cloud-9 could be associated with the Magellanic Stream, at a distance of ≈50–100 kpc. Such an explanation has been proposed for other compact H I HVCs observed at high latitude (J. T. van Loon et al. 2006). However, the stream lies on the opposite side of the sky to M94, with its closest point (the Leading Arm) around 70° away, rendering this explanation unlikely for Cloud-9. A third possibility is that Cloud-9 is a neutral gas cloud in pressure equilibrium with M94's hot circumgalactic medium. For a cloud at $\approx 10^4$ K, confinement by ambient gas at $\sim 10^6$ K requires a density contrast of ∼100. Cloud-9's mean density, $\sim 2 \times 10^{-2}\,{\rm cm^{-3}}$, would then imply a circumgalactic density of $\sim 2 \times 10^{-4}\,{\rm cm^{-3}}$, which is a plausible value. Thus, static pressure confinement cannot be excluded on simple pressure-balance grounds. However, such an equilibrium configuration is unlikely to be long-lived: Cloud-9's perturbed H I morphology suggests ongoing ram pressure interactions, and unless we are observing it at a very particular moment, a purely non-self-gravitating, pressure-confined cloud of this size would be disrupted within a few tens of Myr (e.g., R. I. Klein et al. 1994).

These arguments lead us to favor the RELHIC or ultrafaint dwarf interpretation over these alternatives. The nondetection of a stellar counterpart, the extreme H I-to-stellar mass ratio, and the consistency with the predicted mass threshold strongly support the RELHIC interpretation. Therefore, we conclude that our HST observations make Cloud-9 the most compelling RELHIC candidate to date, and a rare case highlighting the physics that govern the onset of galaxy formation.

5. Summary and Future Work

Our main conclusions may be summarized as follows:

1. Visual inspection of our deep HST/ACS imaging reveals no obvious stellar counterpart within our HST/ACS imaging. Using `ArtPop` to generate resolved galaxies following the local stellar mass–size relation with simulated HST observing conditions, we can visually rule out a faint galaxy counterpart down to $M_\star \approx 10^{3.5}\,M_\odot$.
2. Quantitatively, 10,000 CMD realizations of faint systems, mimicking the HST observations and applying identical selection cuts, indicate that a $10^4\,M_\odot$ stellar counterpart can be ruled out at 99.5% confidence within a projected radius of 180 pc centered on Cloud-9. This radius matches the size of a Leo T–like ($M_\star \sim 10^5\,M_\odot$) galaxy at the distance of M94.

3. A similar analysis using simulated CMDs yields an average upper stellar mass limit of $10^{3.5}\,M_{\odot}$, in excellent agreement with our visual comparison between the HST/ACS imaging field and simulated observations.
4. Extending the search radius to account for a potential offset of up to $\sim$1 kpc between the VLA peak emission and any stellar counterpart results in no statistically significant detection above the background level, as determined by analyzing sources on the opposing (WFC2) chip.

We suggest that future work on Cloud-9 should focus on three complementary fronts. First, deeper optical observations, particularly with JWST, could further lower the upper limit on its stellar mass. Second, numerical simulations should explore whether the slightly perturbed morphology of Cloud-9 is consistent with RELHICs experiencing ram pressure stripping. Third, deep Hα imaging could probe the expected ringlike emission in its outer regions (C. Sykes et al. 2019). Together with existing radio data, these observations would enable joint constraints on the intensity of the local UVB—which may be influenced by Cloud-9's proximity to M94—while also providing an opportunity to map the dark matter distribution beyond Cloud-9's core.

In the ΛCDM framework, the existence of a critical halo mass scale for galaxy formation naturally predicts galaxies spanning orders of magnitude in stellar mass at roughly fixed halo mass. This threshold marks a sharp transition at which galaxy formation becomes increasingly inefficient (A. Benitez-Llambay & C. Frenk 2020), yielding outcomes that range from halos entirely devoid of stars to those able to form faint dwarfs, depending sensitively on their mass assembly histories. Even if Cloud-9 were to host an undetected, extremely faint stellar component, our HST observations, together with FAST and VLA data, remain fully consistent with these theoretical expectations. Cloud-9 thus appears to be the first known system that clearly signals this predicted transition, likely placing it among the rare RELHICs that inhabit the boundary between failed and successful galaxy formation. Regardless of its ultimate nature, Cloud-9 is unlike any dark, gas-rich source detected to date.

Acknowledgments

We thank the anonymous referee, as well as Jacco van Loon and Yakov Faerman, for useful comments that improved the quality of our work. A.B.L. acknowledges support by the Italian Ministry for Universities (MUR) program "Dipartimenti di Eccellenza 2023-2027" within the Centro Bicocca di Cosmologia Quantitativa (BiCoQ), and support by UNIMIB's Fondo Di Ateneo Quota Competitiva (project 2024-ATEQC-0050). J.F.N. acknowledges the hospitality of the Max-Planck Institute for Astrophysics, the Donostia International Physics Centre, and Durham University during the completion of this manuscript. Support for program GO-17712 was provided by NASA through a grant from the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS5-26555.

The HST/ACS data are available at this MAST DOI: 10.17909/10gr-vg66.

Facility: HST (ACS).

*Software:*DOLPHOT (A. E. Dolphin 2000b; A. Dolphin 2016), DrizzlePac (R. J. Avila et al. 2015).

ORCID iDs

Gagandeep S. Anand https://orcid.org/0000-0002-5259-2314
Alejandro Benítez-Llambay https://orcid.org/0000-0001-8261-2796
Rachael Beaton https://orcid.org/0000-0002-1691-8217
Andrew J. Fox https://orcid.org/0000-0003-0724-4115
Julio F. Navarro https://orcid.org/0000-0003-3862-5076
Elena D'Onghia https://orcid.org/0000-0003-2676-8344

References

Adams, E. A. K., Giovanelli, R., & Haynes, M. P. 2013, ApJ, 768, 77
Anand, G. S., Rizzi, L., Tully, R. B., et al. 2021, AJ, 162, 80
Avila, R. J., Hack, W., Cara, M., et al. 2015, in ASP Conf. Ser. 495, Astronomical Data Analysis Software and Systems XXIV (ADASS XXIV), ed. A. R. Taylor & E. Rosolowsky (San Francisco, CA: ASP), 281
Benitez-Llambay, A., Beaton, R. L., Anand, G. S., et al. 2024, HST Proposal. Cycle 32, ID. #17712
Benítez-Llambay, A., Dutta, R., Fumagalli, M., & Navarro, J. F. 2024, ApJ, 973, 61
Benitez-Llambay, A., & Frenk, C. 2020, MNRAS, 498, 4887
Benitez-Llambay, A., & Navarro, J. F. 2023, ApJ, 956, 1
Benítez-Llambay, A., Navarro, J. F., Frenk, C. S., et al. 2017, MNRAS, 465, 3913
Blumenthal, G. R., Faber, S. M., Primack, J. R., & Rees, M. J. 1984, Natur, 311, 517
Bullock, J. S., & Boylan-Kolchin, M. 2017, ARA&A, 55, 343
Carlsten, S. G., Greene, J. E., Greco, J. P., Beaton, R. L., & Kado-Fong, E. 2021, ApJ, 922, 267
Choi, J., Dotter, A., Conroy, C., et al. 2016, ApJ, 823, 102
Dolphin, A., 2016 DOLPHOT: Stellar photometry, Astrophysics Source Code Library, ascl:1608.013
Dolphin, A. E. 2000a, PASP, 112, 1397
Dolphin, A. E. 2000b, PASP, 112, 1383
Dolphin, A. E. 2002, MNRAS, 332, 91
Dotter, A. 2016, ApJS, 222, 8
Durbala, A., Finn, R. A., Crone Odekon, M., et al. 2020, AJ, 160, 271
Fattahi, A., Navarro, J. F., Sawala, T., et al. 2016, MNRAS, 457, 844
Greco, J. P., & Danieli, S. 2022, ApJ, 941, 26
Hoeft, M., Yepes, G., Gottlöber, S., & Springel, V. 2006, MNRAS, 371, 401
Irwin, M. J., Belokurov, V., Evans, N. W., et al. 2007, ApJL, 656, L13
Jacobs, B. A., Tully, R. B., Rizzi, L., et al. 2011, AJ, 141, 106
Jang, I. S., & Lee, M. G. 2017, ApJ, 836, 74
Jenkins, A., Frenk, C. S., White, S. D. M., et al. 2001, MNRAS, 321, 372
Karachentsev, I. D., Karachentseva, V. E., Kaisin, S. S., & Zhang, C.-P. 2024, A&A, 684, L24
Karunakaran, A., & Spekkens, K. 2024, RNAAS, 8, 24
Klein, R. I., McKee, C. F., & Colella, P. 1994, ApJ, 420, 213
Lelli, F. 2022, NatAs, 6, 35
Makarova, L. N., Tully, R. B., Anand, G. S., et al. 2023, ApJ, 943, 139
McQuinn, K. B. W., Skillman, E. D., Dolphin, A., et al. 2015, ApJ, 812, 158
Murray, C. E., Lindberg, C. W., Yanchulova Merica-Jones, P., et al. 2024, ApJS, 275, 5
Nebrin, O., Giri, S. K., & Mellema, G. 2023, MNRAS, 524, 2290
Okamoto, T., & Frenk, C. S. 2009, MNRAS, 399, L174
Press, W. H., & Schechter, P. 1974, ApJ, 187, 425
Sawala, T., Frenk, C. S., Fattahi, A., et al. 2016a, MNRAS, 456, 85
Sawala, T., Frenk, C. S., Fattahi, A., et al. 2016b, MNRAS, 457, 1931
Sykes, C., Fumagalli, M., Cooke, R., Theuns, T., & Benítez-Llambay, A. 2019, MNRAS, 487, 609
van Loon, J. T., Stanimirović, S., Evans, A., & Muller, E. 2006, MNRAS, 365, 1277
Wakker, B. P., & van Woerden, H. 1991, A&A, 250, 509
Warfield, J. T., Richstein, H., Kallivayalil, N., et al. 2023, RNAAS, 7, 23
Weisz, D. R., Zucker, D. B., Dolphin, A. E., et al. 2012, ApJ, 748, 88
Westmeier, T. 2018, MNRAS, 474, 289
Williams, B. F., Durbin, M., Lang, D., et al. 2023, ApJS, 268, 48
Williams, B. F., Lang, D., Dalcanton, J. J., et al. 2014, ApJS, 215, 9
Xu, J.-L., Zhu, M., Yu, N., et al. 2023, ApJL, 944, L40
Xu, J.-L., Zhu, M., Yu, N.-P., et al. 2025, ApJL, 982, L36
Zhou, R., Zhu, M., Yang, Y., et al. 2023, ApJ, 952, 130