

OH Megamasers in H I Surveys: Forecasts and a Machine-learning Approach to Separating Disks from Mergers

Hayley Roberts ¹, Jeremy Darling ¹, and Andrew J. Baker ²

¹ Center for Astrophysics and Space Astronomy, Department of Astrophysical and Planetary Science, University of Colorado, 389 UCB, Boulder, CO 80309-0389, USA; hayley.roberts@colorado.edu

Abstract

OH megamasers (OHMs) are rare, luminous masers found in gas-rich major galaxy mergers. In untargeted neutral hydrogen (H I) emission-line surveys, spectroscopic redshifts are necessary to differentiate the $\lambda_{\rm rest}=18\,{\rm cm}$ masing lines produced by OHMs from H I 21 cm lines. Next-generation H I surveys will detect an unprecedented number of galaxies, most of which will not have spectroscopic redshifts. We present predictions for the numbers of OHMs that will be detected and the potential "contamination" they will impose on H I surveys. We examine the Looking at the Distant Universe with the MeerKAT Array (LADUMA), a single-pointing deep-field survey reaching redshift $z_{\rm H~I}=1.45$, as well as potential future surveys with the Square Kilometre Array (SKA) that would observe large portions of the sky out to redshift $z_{\rm H~I}=1.37$. We predict that LADUMA will potentially double the number of known OHMs, creating an expected contamination of 1.0% of the survey's H I sample. Future SKA H I surveys are expected to see up to 7.2% OH contamination. To mitigate this contamination, we present methods to distinguish H I and OHM host populations without spectroscopic redshifts using near- to mid-IR photometry and a k-Nearest Neighbors algorithm. Using our methods, nearly 99% of OHMs out to redshift $z_{\rm OH}\sim 1.0$ can be correctly identified. At redshifts out to $z_{\rm OH}\sim 2.0$, 97% of OHMs can be identified. The discovery of these high-redshift OHMs will be valuable for understanding the connection between extreme star formation and galaxy evolution.

Unified Astronomy Thesaurus concepts: Galaxy mergers (608); Spiral galaxies (1560); Radio spectroscopy (1359); Hydroxyl masers (771); Megamasers (1023); Galaxy evolution (594); H I line emission (690); Surveys (1671)

1. Introduction

OH megamasers (OHMs) are luminous 18 cm masers found in (ultra)luminous infrared galaxies ([U]LIRGs) produced predominantly by major galaxy mergers. The dominant masing lines occur at 1667 and 1665 MHz in the ground state of OH. These masing lines have isotropic luminosities between $10^1-10^4L_{\odot}$ and line widths ranging from 10 to 1000 km s⁻¹ due to Doppler broadening (Darling 2005). This rare phenomenon has only been discovered in roughly 110 galaxies out to redshift z=0.264 (Darling & Giovanelli 2002a).

OHMs are associated with high molecular gas density, $n({\rm H_2}) \sim 10^4 \, {\rm cm}^{-3}$, and strong far-IR radiation, making them markers for some of the most extreme star formation observed in our local universe (Darling 2007; Lockett & Elitzur 2008). These galaxies can provide significant information about how extreme star formation relates to galaxy evolution, particularly once discovered at higher redshifts. As products of gas-rich major mergers, OHMs can also provide an independent measure of the galaxy merger rate at a specific evolutionary stage. Further, the masing lines present in OHMs can be utilized as Zeeman magnetometers, providing in situ measurements of magnetic fields in nearby galaxies (Robishaw et al. 2008; McBride et al. 2014). OHMs are useful astronomical tools for understanding many aspects of galaxy evolution.

However, in untargeted emission-line surveys for neutral hydrogen (H I), an OH emission line at $z_{\rm OH}$ can "spoof" a 21 cm H I line at $z_{\rm HI}$ if $\nu_{\rm H~I}/(1+z_{\rm H~I})=\nu_{\rm OH}/(1+z_{\rm OH})$ (Briggs 1998), where $\nu_{\rm H~I}=1420.4$ MHz and $\nu_{\rm OH}=1667.4$ MHz. For example, rest-frame H I corresponds to OH at a redshift of $z_{\rm OH}=0.174$, while H I at redshift $z_{\rm H~I}=0.1$ would correspond to OH at a

redshift of $z_{\rm OH} = 0.291$. These two lines have similar line widths in their respective environments: H I in spiral galaxies and OH in major galaxy mergers. Distinguishing between these lines often requires knowledge of a galaxy's spectroscopic redshift to determine the rest wavelength for an observed emission line (e.g., Hess et al. 2021). For galaxies that do not have spectroscopic redshifts, disentangling H I from OH is not straightforward. This ambiguity particularly becomes a problem for high-redshift untargeted line surveys, which will not have prior spectroscopic redshifts for many of their detected galaxies.

Despite the versatility and capacity of OHMs to serve as tools for studying galaxy evolution, line confusion in H I surveys can serve as a source of contamination for the survey's main goals. Nonetheless, next-generation H I surveys, such as those with the Square Kilometre Array (SKA) and its precursors, will be able to detect many new OHMs at redshifts never before reached. Finding these unique galaxies in the pool of H I disk detections will enable an exciting new era in OHM science.

LADUMA (Looking at the Distant Universe with MeerKAT Array; Blyth et al. 2016) is a survey with the MeerKAT radio interferometer, a precursor instrument for the SKA, that will be susceptible to OH/H I confusion. LADUMA will be the deepest neutral hydrogen survey to date and is expected to detect H I out to redshifts $z_{\rm HI}$ =1.45. LADUMA's main science goals are related to studying the neutral atomic gas content of galaxies, meaning that OHM detections will contaminate its H I samples. At low redshift, spectroscopic redshifts are generally known, so the contamination rate should be small. At greater distances where fewer spectroscopic redshifts are currently known and OHM prevalence increases due to the

Department of Physics and Astronomy, Rutgers, The State University of New Jersey, 136 Frelinghuysen Road, Piscataway, NJ 08854-8019, USA Received 2020 November 23; revised 2021 February 22; accepted 2021 February 22; published 2021 April 14

elevated merger rate, the contamination rate threatens to be higher.

This paper presents predictions for the numbers of OH megamasers that will be detected by LADUMA in Section 2 and by other H I surveys in Section 3. We then present methods for distinguishing OH from H I in LADUMA and other untargeted line surveys using near- to mid-IR photometry in Section 4. We discuss these results and their limitations in Section 5 and summarize our conclusions in Section 6.

Throughout this work, we assume a flat Λ CDM cosmology with $H_0 = 70 \text{ km s}^{-1} \text{ Mpc}^{-1}$, $\Omega_m = 0.3$, and $\Omega_{\Lambda} = 0.7$.

2. OH Megamasers in LADUMA

In this section, we present our predictions for the number of OHMs that will be detected in LADUMA. A few inputs are needed to enable these predictions: an OH luminosity function, the volume of the survey, and the galaxy merger rate. The following subsections will cover each of these. We then make use of this information to predict the number of OH megamasers and the resultant OHM contamination rate in other H I surveys.

2.1. The OH Luminosity Function

To create a prediction for the number of OHMs that will be found in the LADUMA survey, we need to integrate the OH luminosity function (OHLF) over the volume and luminosity limits of LADUMA. The OHLF presented in Darling & Giovanelli (2002b) was constructed from OHMs detected by Arecibo and is valid for $2.2 < \log(L_{\rm OH}/L_{\odot}) < 3.8$ and $0.1 < z_{\rm OH} < 0.23$. An obvious source of uncertainty is the extrapolation of the OHLF to cover redshifts out to $z_{\rm OH} = 1.876$ and luminosities as low as $1L_{\odot}$ for the LADUMA survey.

The OHLF is defined as the number of OHMs with luminosity $(L_{\rm OH})$ per unit comoving volume $({\rm Mpc}^3)$ per logarithmic interval in $L_{\rm OH}$ and can be parameterized as

$$\Phi(L_{\rm OH}) = b L_{\rm OH}^a. \tag{1}$$

The values of *a* and *b* presented in Darling & Giovanelli (2002b) were determined using an error-weighted least-squares fit. We use a Markov chain Monte Carlo (MCMC) method to refit the OHLF using the data presented in Darling & Giovanelli (2002b) to account for the correlations between fit parameters. We use emcee, a Python package for implementing MCMC (Foreman-Mackey et al. 2013). The refit OHLF is

$$\Phi(L_{\rm OH}) = (3.17 \pm 0.58) \times 10^{-6} (L_{\rm OH}/L_{\odot})^{-0.50 \pm 0.13} \text{ Mpc}^{-3} \text{ dex}^{-1}.$$
 (2)

The samples from the MCMC allow us to create a representative sample of possible OHLF parameters for predictions of the number of OH megamasers that will be observed. This is the OHLF that will be used throughout this work. We later compare how the difference between the two OHLFs changes the number of OHMs predicted to be detected.

2.2. Volume of the LADUMA Field

LADUMA will be observing an area encompassing the Extended Chandra Deep Field South (E-CDFS), covering 0.9 \deg^2 at z = 0. The field of view of MeerKAT will increase with redshift due to the array's larger primary beam at lower

frequencies. The total volume of the field can be calculated starting with the equation for comoving volume (Equation (28) from Hogg 1999) shown in Equation (3), where $E(z) = \sqrt{\Omega_M (1+z)^3 + \Omega_\Lambda}$, $D_H = c/H_0$, and D_A is the angular diameter distance:

$$dV_c = D_H \frac{(1+z)^2 D_A^2}{E(z)} d\Omega dz.$$
 (3)

To account for the increasing field of view at higher redshifts, we write $d\Omega$ in Equation (3) in terms of the redshift-dependent field of view of a telescope with primary beam diameter $1.22 \, c \, (1+z)/(\nu_0 \, D)$, where D is the diameter of a single dish (for MeerKAT, D=13.5 m) and ν_0 is the rest frequency of the line being observed. This diameter yields a redshift-dependent solid angle:

$$\Omega = \pi \left[\frac{1.22}{2} \frac{c (1+z)}{\nu_0 D} \right]^2. \tag{4}$$

Integrating Equation (3) over $d\Omega$ and substituting in Equation (4) yields the differential volume for the LADUMA field:

$$dV = \pi \left(\frac{1.22}{2}\right)^2 \left(\frac{D_A}{\nu_0 D}\right)^2 \frac{D_H (1+z)^4}{\sqrt{\Omega_m (1+z)^3 + \Omega_\Lambda}} dz.$$
 (5)

Integrating Equation (5) over the redshift limits of the LADUMA survey for detecting OH ($z_{OH} = 0.174-1.876$), we obtain a total volume of $V = 0.045 \, \text{Gpc}^3$.

2.3. Sensitivity of LADUMA

The luminosity sensitivity for integrating the OHLF changes with redshift and with MeerKAT band (either the UHF or L band). The sensitivity is calculated using Equation (6), where ΔS_{ν} is the interferometric sensitivity, SEFD is the system equivalent flux density as reported by MeerKAT,³ $\eta_{\rm corr}$ is the correlator efficiency, $N_{\rm ant}$ is the number of antennas, $N_{\rm pol}$ is the number of polarizations observed, $\Delta \nu$ is the bandwidth, and Δt is the integration time:

$$\Delta S_{\nu} = \frac{\text{SEFD}}{\eta_{\text{corr}} \sqrt{N_{\text{ant}} (N_{\text{ant}} - 1) N_{\text{pol}} \Delta \nu \Delta t}}.$$
 (6)

We assume that $\eta_{\rm corr}=1$, all $N_{\rm ant}=64$ antennas will be operating, and $N_{\rm pol}=2$. The assumed integration times are 333 hr for the L band and 3091 hr for the UHF band, with 3424 hr for the frequency range where the receivers overlap. We also assume a velocity width of $\Delta V=150\,{\rm km\,s^{-1}}$, the width of the average OH line at z=0, converted to hertz using $\Delta \nu/(\nu_{\rm OH}/(1+z))=\Delta V/c$. As LADUMA will be applying Briggs weighting to obtain a well-behaved synthesized beam, an extra noise penalty of 1.45 is included in sensitivity calculations. We also include the effects of primary beam attenuation away from the phase center, assuming two-dimensional Gaussian beams with half-power points defined as $0.5\times1.22\,\lambda/D$.

MeerKAT SEFD values can be accessed via the SARAO MeerKAT specifications page.

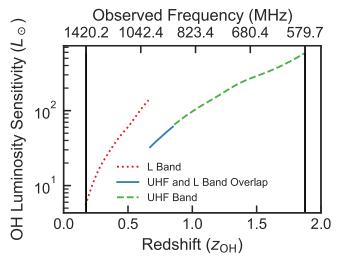


Figure 1. Isotropic 5σ luminosity sensitivity $L_{\rm OH}$ as a function of redshift and observed frequency. The vertical black lines show the low- and high-redshift limits of OH detection for LADUMA. The three different curves come from the two bands LADUMA will use, the UHF and L bands, and the frequency range where they overlap.

Using the sensitivity at each redshift, we calculate the 5σ luminosity limit as $L_{\rm min} = 4\pi D_L^2 \ 5 \ \Delta S_{\nu} \ \Delta \nu$, where D_L is the luminosity distance. Figure 1 shows how the luminosity sensitivity changes with redshift for LADUMA.

2.4. Dependence on the Major Merger Rate

An important element in the calculation of the number of OHMs is the galaxy merger rate because OHMs arise in merging galaxies. This consideration leads us to introduce a factor of $(1+z)^{\gamma}$ in Equation (7), where γ is determined by the galaxy merger rate (in the sense of merging events per comoving volume) and will be referred to as the merger rate evolution coefficient. Previous studies, based on both observations and simulations, have used different conventions for defining and parameterizing a "merger rate," and drawn different conclusions that depend on type of merger, redshift, mass, and many other factors (e.g., Lotz et al. 2011; Rodriguez-Gomez et al. 2015; Mundy et al. 2017; Mantha et al. 2018; Duncan et al. 2019; O'Leary et al. 2021, and references therein). The relevant merger rate for OHMs is one corresponding to gas-rich major mergers, whose evolution remains poorly constrained. Therefore, we select an intermediate estimate for this merger rate evolution coefficient of $\gamma \sim 2.2$ (Rodriguez-Gomez et al. 2015; Mundy et al. 2017), which will be assumed when not stated otherwise. In Section 2.5, we present OHM calculations for a conservatively large range of possible γ values, $0.0 \leqslant \gamma \leqslant 3.0$, which is slightly larger than the range suggested by recent studies (0.5 $\lesssim \gamma \lesssim$ 2.8; Mundy et al. 2017; Ferreira et al. 2020).

OHMs present a unique and independent way to measure the galaxy merger rate. Once we have a secure sample of OHMs, we will be able to provide an estimated γ for gas-rich major mergers. As more H I surveys take data, this method will be a robust way for tracing the cosmic history of major mergers.

2.5. MCMC Calculation of N(OH)

Using the MCMC fit samples for the OHLF discussed in Section 2.1, we can integrate over volume, luminosity, and merger rate to get a prediction of the number of OH

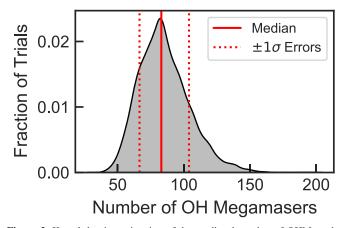


Figure 2. Kernel density estimation of the predicted number of OHMs to be found by LADUMA, assuming $\gamma = 2.2$, using 10,000 samples from the MCMC fit to the OHLF parameters. The expected number of OHM detections is 83^{+21}_{-17} .

megamasers to be detected at the 5σ level in LADUMA as shown in Equation (7):

$$N = \int_{L_{\min}}^{L_{\max}} \int_{0}^{V_{\text{total}}} \Phi(L_{\text{OH}}) (1+z)^{\gamma} dV d \log L_{\text{OH}}.$$
 (7)

Figure 2 shows a distribution of possible values for the number of OHMs using 10,000 samples from the MCMC fit to the OHLF, assuming that $\gamma = 2.2$. The median value is 82.99, and the associated 16th and 84th percentiles are 66.49 and 103.89. The 16th and 84th percentiles are commonly associated with -1σ and $+1\sigma$ limits (for a Gaussian distribution), respectively, and they will be referred to as such for the remainder of this paper. The mean of the distribution is 85.32. The number of OH megamasers we expect to detect with LADUMA is therefore 83^{+21}_{-17} . This total would nearly double the number of known OHMs.

As discussed in Section 2.4, the exact value of the merger rate evolution parameter γ is poorly constrained. Figure 3 presents the above calculation for values of γ ranging from 0.0 to 3.0. If γ is assumed to be 0.0 and the merger rate does not increase with redshift, then it is expected that LADUMA would detect 15^{+5}_{-3} OHMs.

We can compare the impact of the OHLF presented in this work to the OHLF in Darling & Giovanelli (2002b) by calculating the numbers of OHMs implied by the two. Figure 4 shows how the number of OHMs varies for the two OHLFs. When $\gamma = 2.2$, the MCMC approach adopted here implies a factor of 1.4 fewer detections. In general, the larger uncertainties from calculations with the OHLF in Darling & Giovanelli (2002b) are due to the larger uncertainties in that paper's OHLF parameters.

3. OH Contamination in Other H I Surveys

In this section, we generalize the above calculations for H I surveys underway or planned at other radio telescope arrays, particularly the Australian Square Kilometre Array Pathfinder (ASKAP), the Very Large Array (VLA), the APERture Tile In Focus (APERTIF) facility, and Phase I of the SKA. The VLA and ASKAP have H I surveys underway or planned. At the VLA, the COSMOS H I Large Extragalactic Survey (CHILES), a single-pointing survey, is currently analyzing data (Fernández et al. 2016). WALLABY (Widefield ASKAP *L*-band Legacy All-sky blind survey) is in the process

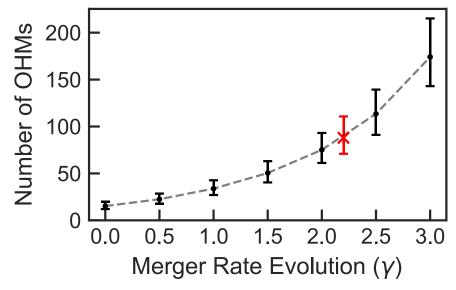


Figure 3. Number of OHM detections in LADUMA vs. merger rate evolution parameter γ . The fiducial $\gamma = 2.2$ is denoted by a red \times .

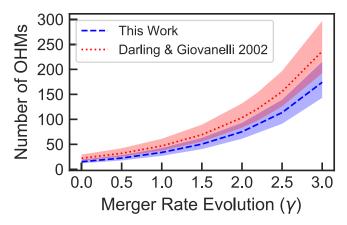


Figure 4. Comparison of OHM detection rate in LADUMA vs. merger rate evolution coefficient for the OHLFs presented in this work and in Darling & Giovanelli (2002b). The dashed and dotted lines show the median values, and the shaded areas represent $\pm 1\sigma$ uncertainties.

of observing pilot fields with ASKAP alongside DINGO (Deep Investigation of Neutral Gas Origins; Duffy et al. 2012). DINGO will have two tiers, Deep and Ultra Deep. For our analysis, we consider only DINGO-Deep. APERTIF is also currently collecting data on the Westerbork Synthesis Radio Telescope (WSRT) and will execute multiple surveys at different depths. We consider the Medium Deep Survey (MDS) as described in the APERTIF Survey Plan. The SKA is a future telescope that will come in two phases, with the first (SKA1) covering ~10% of the total collecting area of the second (SKA2; Abdalla et al. 2015). Due to the uncertainties in the schedule for SKA2, we consider only possible midfrequency SKA1 surveys for this analysis.

The calculation of numbers of OHMs for these H1 surveys uses Equation (7) above. For each survey, we calculate $L_{\rm min}$ as a function of redshift using the reported sensitivity and we calculate the volume using redshift ranges and sky coverage. All assumptions made about the survey or telescope for these calculations are presented in Table 1. The sensitivity column assumes a velocity width of $150\,{\rm km\,s^{-1}}$ as done for the LADUMA calculations. Concepts for the SKA1 surveys come from Staveley-Smith & Oosterloo (2014). Each of the three

Table 1 H I Survey Parameters for *N*(OH) Calculation

Survey	Redshift Range (z _{H I})	Sky Area (deg ²)	5σ Sensitivity (Jy)
LADUMA	0.0-1.45	0.90 ^a	48 ^b
CHILES	0.0 – 0.45	0.32 ^a	350
WALLABY	0.0-0.26	27,500	5,000
DINGO-Deep	0.0-0.26	150	156
APERTIF MDS	0.0-0.26	450	1250
SKA1 Medium wide	0.0-0.29	400	247 ^b
SKA1 Medium deep	0.0-0.50	20	52 ^b
SKA1 Deep	0.35 - 1.37	1 ^a	13 ^b
SKA1 All-sky	0.0-0.50	20,000	525 ^b

Notes.

Staveley-Smith & Oosterloo 2014).

fiducial surveys (medium wide, medium deep, and deep) assumes a total observing time of 1000 hr, while the all-sky commensal survey assumes an observing time of 10,000 hr. Staveley-Smith & Oosterloo (2014) note that angular resolutions finer than 10" are only accessible to the SKA for high column densities—a limitation that especially applies to the SKA1 Deep survey, whose nominal angular resolution is 2". We have not modeled the effects of resolution on the detectability of OHMs, but we note that at high resolutions, there could be a bias in favor of detecting (more compact) OHMs relative to (more extended) H I emitters.⁴

^a Single-pointing survey comoving volume calculations include an expanding field of view at higher redshifts and lower observing frequencies.

b LADUMA and SKA1 have published frequency-dependent sensitivities, which we employ for calculations in this paper. The values presented here are the mean sensitivities across the entire range of the observed frequencies.

References. CHILES (Fernández et al. 2013), WALLABY & DINGO-Deep (Duffy et al. 2012), APERTIF (Apertif Survey Plan), SKA1 (Braun et al. 2019;

This paper's predictions ignore the effects of radio frequency interference (RFI), which can vary for different sites and different array configurations. Here, too, it may in practice be systematically easier to recover more compact OHMs than more extended H I emitters in frequency ranges where RFI precludes the use of short-baseline data.

Survey	N(OH)	<i>N</i> (H I)	<i>N</i> (OH)/ <i>N</i> (H I)
LADUMA	$8.3^{+2.1}_{-1.7} \times 10^{1}$	8×10^3	$1.03^{+0.26}_{-0.20}\%$
CHILES	$5.5^{+1.1}_{-0.9} \times 10^{-1}$	3×10^2	$0.18^{+0.04}_{-0.03}\%$
WALLABY	$8.9^{+2.5}_{-1.8} \times 10^2$	6×10^5	$0.15^{+0.04}_{-0.03}\%$
DINGO-Deep	$7.5^{+2.3}_{-1.7} \times 10^{1}$	5×10^4	$0.15^{+0.04}_{-0.03}\%$
APERTIF MDS	$2.6^{+0.6}_{-0.5} \times 10^2$	3×10^5	$0.09^{+0.02}_{-0.01}\%$
SKA1 Medium wide	$2.5^{+0.7}_{-0.6} \times 10^2$	3.4×10^{4}	$0.73^{+0.21}_{-0.17}\%$
SKA1 Medium deep	$7.7^{+3.2}_{-2.3} \times 10^{1}$	2.5×10^{4}	$0.31^{+0.13}_{-0.09}\%$
SKA1 Deep	$1.9^{+0.6}_{-0.5} \times 10^2$	2.6×10^{3}	$7.20^{+2.32}_{-1.81}\%$
SKA1 All-sky	$4.1^{+1.0}_{-0.8} \times 10^3$	5.5×10^5	$0.75^{+0.18}_{-0.14}\%$

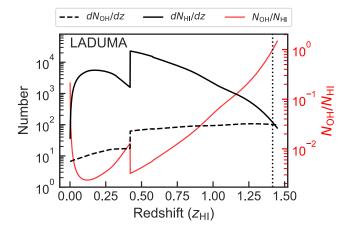
Note. N(OH) values assume merger rate evolution coefficient $\gamma = 2.2$.

For surveys featuring single pointings, we have calculated comoving volumes assuming that sky area increases $\propto (1 + z_{HI})^2$ due to the increasing size of the primary beam at lower frequencies. This calculation applies to CHILES, LADUMA, and a hypothetical SKA1 Deep survey and is noted in Table 1. For surveys covering larger sky areas through the use of multiple pointings across contiguous patches, sky area will be higher at $z_{\rm H~I} > 0$ than at $z_{\rm H~I} = 0$, but the change will be less dramatic because only the pointings that lie at the edges of the contiguous patches will contribute. Because this effect will be small in a fractional sense (smaller for larger sky areas) and will depend on the detailed distribution of patch sizes, we do not correct for it. We also choose to omit primary beam attenuation when predicting OHM contamination of other HI surveys, because we cite values for the numbers of HI detections that do not include this consideration (Staveley-Smith & Oosterloo 2014).

Table 2 presents the number of OHMs predicted to be detected in each survey for merger rate evolution coefficient $\gamma = 2.2$. Table 2 also presents the number of H I sources each survey expects to detect. The contamination column is the ratio of OHM detections to H I detections, which can be related to the fraction of an "H I sample" that will actually be OHMs mistaken for H I sources if spectroscopic redshifts are unavailable.

One noteworthy aspect from Table 2 is the much higher rate of contamination for LADUMA and the SKA1 Deep survey compared to the other surveys. These are distinctly different from the other H1 surveys due to the fact that they extend to significantly higher redshift. We therefore infer that H1 detections dominate at low redshifts (i.e., $z_{\rm H}$ 1 \lesssim 1) for all surveys. However, the OH detection density surpasses the H1 detection density at redshifts above $z_{\rm H}$ 1 —an effect that is pronounced for single-pointing surveys that have much larger relative fields of view at high versus low redshifts.

An earlier conclusion in the same vein was reached by Briggs (1998), who predicts that OHM contamination in H I surveys will increase with redshift. We explore how that contamination depends on redshift using both LADUMA and an expanded version of the SKA1 Deep survey for comparison. The SKA1 Deep survey presented in Staveley-Smith & Oosterloo (2014) only covers a frequency range of 600–1050 MHz that corresponds to only a portion of the full SKA1 midband (Braun et al. 2019). For the purpose of exploring OH contamination versus redshift, we assume a deep survey that exploits the full range of the midband and therefore covers a frequency range of 600–1420 MHz. Equation (7) is used to



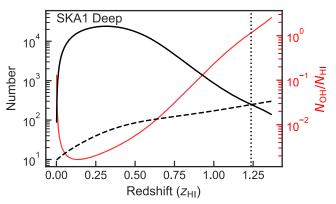


Figure 5. Projected numbers of OHMs and H I sources and the OHM fraction $(N_{OH}/N_{H\ I})$ vs. redshift $z_{H\ I}$ for LADUMA (above) and SKA1 Deep (below). The left axis shows the number of objects (plotted in black); the dashed curve indicates how the number of OH detections changes with redshift, and the solid curve indicates the same for H I detections. The right axis shows the ratio of OHM detections to H I detections (plotted in red). The vertical dotted lines indicate the redshifts where the numbers of OH and H I detections are equal. The discontinuities in the LADUMA curves originate from the overlap between the L and UHF bands. All calculations assume galaxy merger rate evolution parameter $\gamma=2.2$.

estimate how the number of OHMs varies with redshift. To estimate the number of H I detections per redshift interval, we use the following equation from Obreschkow et al. (2009):

$$\frac{dN/dz}{1 \text{ deg}^2} = 10^{c_1} z^{c_2} e^{-c_3 z},$$
 (8)

where the c_i are parameters specific to each H I survey. LADUMA's values are interpolated from Obreschkow et al. (2009) for each redshift using sensitivities calculated from Equation (6) as the limiting integrated flux and assuming line widths of $100 \, \mathrm{km \, s^{-1}}$, allowing us to determine the H I detection rate versus redshift. We follow a similar method for the SKA1 Deep survey using sensitivity values presented in Braun et al. (2019). We calculate how the number density for OH and H I varies with redshift, assuming redshift bins of dz = 0.01. Obreschkow et al. (2009) note that dN/dz will be $\geqslant 1\%$ underestimated for $z \leqslant z_c$ where z_c depends on the limiting flux. For both LADUMA and SKA1 Deep, on average, $z_c \sim 0.1$. Therefore, the numbers of H I detections are slightly underestimated for $z_{\mathrm{H \, I}} \leqslant 0.1$.

Figure 5 shows how the detection rate of OH and H I varies with redshift. H I detections dominate at low redshift for both

Table 3

Machine-learning Results for Distinguishing H I Emission Lines from OH

Megamasers

Mission	Features	OH Recall	OH Precision
WISE	W1, W1–W2, ν	0.985	0.974
WISE	W1–W2, W2–W3, ν	0.987	0.985
IRAC	[3.6], [3.6]–[4.5], ν	0.979	0.972
IRAC	[3.6]–[4.5], [5.8]–[8.0], ν	0.982	0.976

Note. The features column indicates what data were used to distinguish H I from OHM host galaxies. Each row includes the observed frequency (ν) of the line in question (OH or H I) to assist in separation of sources.

surveys. At higher redshifts, the OH detection rate and OH fraction grow significantly. This comparison also demonstrates how the OHM contamination rate depends on survey parameters, as discussed in Section 3. For LADUMA, OHMs will not outnumber the H I source at any redshift probed by the survey (i.e., for any $z_{\rm H~I} \le 1.45$); in comparison, SKA1 Deep's sensitivity as a function of frequency will yield a number of OH detections surpassing that of H I detections for $z_{\rm H~I} \ge 1.25$.

Briggs (1998) first presented this issue of OH contamination in untargeted H I surveys. Results from that paper demonstrated that by redshift $z_{\rm H~I} \sim 1.0$, OH detections would significantly outnumber HI detections for a much shallower survey $(5\sigma = 5 \text{ mJy})$ than LADUMA or SKA1 Deep. The predicted transition from a H I -dominated to an OH-dominated sample is qualitatively consistent with our findings, although the precise redshift at which this happens depends on survey parameters. In general, the deeper a survey observes over a given frequency range, the more H I emitters it will detect relative to OHMs. LADUMA and SKA1 Deep are both deeper than the hypothetical Briggs (1998) survey, allowing them to detect more HI sources at higher redshifts and thus to push out the projected redshift at which OH detections outnumber HI detections. The fact that this transition occurs at a lower redshift ($z_{\rm H\,I} = 1.25$) for SKA1 than for LADUMA owes to the fact that LADUMA's sensitivity improves at higher redshifts (due to its distribution of observing time), in contrast to the SKA1 Deep sensitivity.

4. Identifying OH Megamasers in Untargeted H I Surveys

Distinguishing an OH from an H I line is currently only done using the optical spectroscopic redshift of an object to determine an observed line's rest wavelength. Next-generation H I surveys will observe orders of magnitude more objects than previous surveys, as shown in Table 3, most of which will not have spectroscopic redshifts available. For that reason, we explore machine learning as a way to distinguish OH from H I emission lines using ancillary data.

4.1. Creating OH and H1 Models for Distinguishing Populations

4.1.1. Fitting OHM Host Galaxy SEDs

The limited number of OHMs creates serious limitations in understanding the OH population and how it differs from H I hosts. We therefore fit the spectral energy distributions (SEDs) of 111 OHMs using Multi-wavelength Analysis of Galaxy Physical Properties (MAGPHYS; Da Cunha et al. 2008), a software package that fits galaxy SEDs using physical

parameters of galaxies at the same redshifts and in the same photometric bands.

MAGPHYS fits SEDs from far-UV to far-IR, so we use photometry from that range for fitting OHM host SEDs. In total, we use photometry from eight sources: the Galaxy Evolution Explorer (Martin et al. 2005), the Sloan Digital Sky Survey (SDSS; Stoughton et al. 2002), the Two Micron All-Sky Survey (2MASS; Skrutskie et al. 2006), the Wide-field Infrared Survey Explorer (WISE; Wright et al. 2010), the Infrared Array Camera (IRAC) and Multiband Infrared Photometer for Spitzer (MIPS, both on Spitzer; Werner 2005), the Infrared Space Observatory (Kessler & Habing 1996), and the Infrared Astronomical Satellite (Beichman et al. 1988). In total, from these sources, we use up to 33 bands to fit OHM SEDs. We omit WISE band 1 (3.4 μ m) if Spitzer IRAC band 1 $(3.6 \mu m)$ exists for a given galaxy because the introduction of both causes poor fits and IRAC tends to have smaller uncertainties than WISE. The same is done for WISE and IRAC band 2 (4.6 and 4.5 μ m, respectively).

In Figure 6, we present examples of these SED fits. For comparison, we use the Atlas of Galaxy SEDs (Brown et al. 2014), which includes a total of seven OHM host galaxy SEDs. The fits we present are imperfect matches to the complete Brown et al. (2014) SEDs, and we emphasize that our SED fits were done with the limited scope of reproducing observations of OHMs in our particular wavelength regime of interest, UV to mid-IR. Our fits are limited in wavelength outside this range, particularly in the far-IR. Although far-IR photometry would provide very useful information, data are sparse and unavailable for our objects of interest, so we have chosen to omit far-IR photometry.

4.1.2. Emulating OHM Host Photometry

We use our SED fits to model OHM host galaxy observations for various missions and surveys as a function of redshift. We use PYPHOT,⁵ a package for calculating an object's photometry from its SED. It calculates the photometry using a given filter's transmission curve, $T(\lambda)$, by calculating the photon number flux:

$$N_{\text{tot}} = \frac{1}{hc} \int_{\lambda} f_{\lambda} \lambda T(\lambda) d\lambda, \tag{9}$$

where f_{λ} is the flux density as shown above in Figure 6.

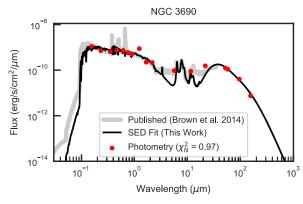
The greatest benefit of having SED fits is the ability to redshift them and mimic observations at higher redshifts. This scaling is done by adjusting the rest wavelength and the flux density by the inverse square of the luminosity distance and a redshift factor $(\propto D_L^{-2} \ (1+z)^{-1})$ and then re"observing" the OHM host. This method is used to create synthetic observations out to the desired redshifts.

4.1.3. Emulating H I Host Galaxy Photometry

For consistency, H I host galaxy photometry is created similarly. However, instead of fitting SEDs, we use 57 SEDs⁶

⁵ PYPHOT's documentation can be accessed at https://mfouesneau.github.io/docs/pyphot/.

⁶ This number was originally 58 SEDs; however, after some examination, NGC 7674 seems to behave much more like an OHM in the near- to mid-IR just like a spiral galaxy, despite looking like a classic spiral morphologically. We attempted to do follow-up observations to determine if it potentially possessed both emission lines; however the OH line cannot be observed due to RFI. We therefore removed this galaxy from our H I SED sample.



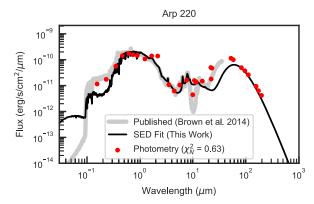


Figure 6. Examples of SED fits with MAGPHYS (black line). The published SED comes from Brown et al. (2014) and is shown as a thick gray line. Red dots denote the photometry used in a given fit.

published in the Atlas of Galaxy SEDs (Brown et al. 2014) that have previous HI detections, a population of mainly spiral galaxies. Because these sources are not drawn from a strictly HI-selected sample, they may not behave identically to samples from untargeted HI surveys, although we expect differences to be modest. These SEDs are also redshifted, and photometry is "measured" using PYPHOT.

4.2. Machine Learning to Distinguish OH from HI

To aid in determining if an emission line is an HI or OH detection, we use machine-learning algorithms to determine the likelihood of the line's classification. We employ a k-Nearest Neighbors (k-NN) algorithm that classifies objects based on a plurality vote of their neighbors' classes, where neighbors are determined within some parameter space (Goldberger et al. 2005), k-NN classification is a nonparametric method and a lazy learning algorithm. Lazy learning means that the algorithm itself does not make assumptions or generalizations based on the training data but instead uses those data to make direct decisions about the testing data. Algorithm "optimization" is purely done by our choices of the nearest-neighbor algorithm parameters used. The cost of using a lazy learning algorithm is the computation time in the testing phase. However, we are not testing on data sets large enough for slow speed to be problematic. This context makes the k-NN classification a robust and transparent method for our purposes.

The final classification parameters are WISE magnitudes and colors as well as the observed line frequency. Suess et al. (2016) demonstrate that WISE photometry can separate OHM and HIpopulations at low redshift. We also choose WISE because of its all-sky coverage, allowing it to be applicable to many different HIsurveys. Section 4.2.1 discusses the use of WISE magnitudes and colors to distinguish HI and OH populations, as well as the limitations of using WISE and its similarities to IRAC. Section 4.2.2 presents similar exercises using IRAC data, which have significant coverage over the LADUMA field but are otherwise less broadly applicable for other HI surveys.

4.2.1. OH and H1 Classification Using WISE

The analysis in Suess et al. (2016) is done with low-redshift (z < 0.1) objects and uses WISE bands W1, W2, W3, and W4 (3.4, 4.6, 12, 22 μ m). W3 and W4 are very insensitive compared to W1 and W2; thus, this method is limited by both object brightness and redshift. We focus on using machine

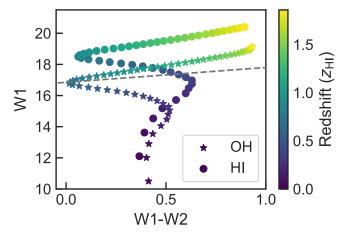


Figure 7. Predicted WISE magnitude vs. color for an OHM host galaxy (stars) and an H I galaxy (circles). H I redshift is denoted by color with redshift increasing with lightness. The corresponding OH points lie at the same observed frequency but are at a higher actual redshift. The gray dashed line represents the detection limit of WISE (the region above the line is undetectable by WISE band W1).

learning to sort using only W1 and W2 magnitudes, W1 - W2 color, and the observed line frequency.

One of the cuts from Suess et al. (2016) is done in the color-magnitude space of W1 versus W1 - W2 (or [3.4] versus [3.4] - [4.6]). We use this same parameter space for the k-NN algorithm. Examples of redshift evolution in this space are shown in Figure 7 for an OHM host and an H I source. These data are "measured" from their SEDs and show how OH -H I separability varies with redshift.

We use the 57 H I and 111 OHM host SEDs to test and train the k-NN algorithm. Each SED is redshifted to the maximum redshift detectable by LADUMA ($z_{\rm H}$ $_{\rm I}=1.45$ and $z_{\rm OH}=1.876$), with WISE photometry being "measured" roughly every dz=0.01. We then remove any data too faint to have a 5σ WISE detection.

k-NN algorithms require feature scaling or parameter normalization, because the algorithm is inherently built on the distances between a data point and its neighbors. Therefore, we normalize each parameter from 0 to 1. k-NN classification algorithms are dependent on a few parameters that can be optimized for a given case. Parameters that were varied and tested for our purposes include the number of neighbors that is included in the plurality vote on an object's classification (k), whether neighbors are weighted by their distances, and how

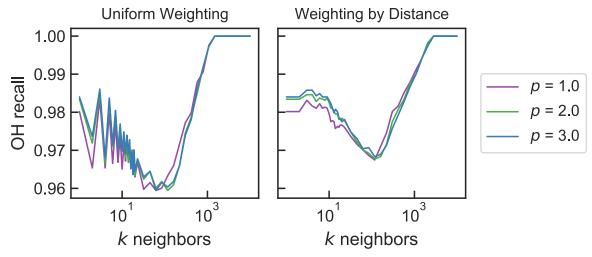


Figure 8. Parameter exploration for our *k*-Nearest Neighbors (*k*-NN) classification algorithm. We vary the number of neighbors considered in voting (*x*-axis), whether distance is weighted (left and right plots), and how distance is determined using the Minkowski metric (the lines plotted). Each test is done using fivefold cross-validation

distance between objects is calculated (p). The distance between points is defined by the Minkowski distance of order p:

$$D(X, Y) = \left(\sum_{i=1}^{n} |x_i - y_i|^p\right)^{\frac{1}{p}},$$
(10)

where *X* and *Y* are two points in an *n*-dimensional parameter space. Euclidean distance is recovered for p = 2.

All algorithm optimization for this work is done by maximizing OH recall. In machine-learning classification, two metrics that are often considered when optimizing are precision and recall, both of which scale from 0 to 1 (1 being the best score). Precision is the fraction of positive identifications that are correct. By optimizing precision, the number of false positives (or Type I errors) is minimized. Recall, conversely, is the fraction of positives that were correctly identified. When recall is optimized, the number of false negatives (or Type II errors) is minimized (Sammut & Webb 2011). These terms correspond to the familiar astronomical concepts of sample purity and completeness. In our case, a positive identification is the classification of a galaxy as an OHM host. We choose to optimize OH recall due to the rarity of OHMs and the desire to not miss any potential candidates. Although this approach increases the number of false positives, this algorithm does add information, and any positive identification it makes can motivate follow-up observations for confirmation.

The algorithm parameter exploration is shown in Figure 8. The x-axis shows a wide range of choices for the number of neighbors used, the lines plotted show a few choices for Minkowski distance, and the two panels show the difference between weighting and not weighting neighbors by distance. Each unique combination of parameters is tested using a fivefold cross-validation test. This process involves randomly sorting our data into two sets: training and testing data. The training data build the algorithm and the testing data determine how successful the algorithm is. This split was done five times, randomly splitting data each time, for each combination of parameters, and the final OH recall was determined by averaging the five individual OH recall values. In total, 2000 k-NN algorithms were tested.

The results of these parameter tests give the highest OH recall for large numbers of neighbors $(k > 10^3)$. Although an OH recall of 1 would be ideal, this result comes at the cost of very low precision and defeats the purpose of the k-NN method by classifying based on the value of an algorithmic parameter instead of position relative to neighbors. Large numbers of neighbors also make for very computationally expensive algorithms. One common approach is to select $k = \sqrt{N}$, where N is the number of data points, but for our data ($k \approx 140$), this choice of k is near the lowest value of OH recall. Another common approach is to select k = 1 or another low number. For small data sets, always assuming the nearest object has the same classification can introduce noise. However, with a sufficiently large data set, this trend is less problematic. Figure 8 indicates that a small k achieves a recall of over 0.98 for distance-weighted learning. We stress that there exists no optimal k for all purposes, as each k-NN optimization varies based on the properties of the data (Altman 1992).

We choose our number of neighbors to be k=3, based on the above considerations. Increasing a small amount above k=1 also reduces noise while maximizing OH recall. For weighting and Minkowski metric, we choose to weight votes by distance and use standard Euclidean distance (p=2).

The final trained and tested k-NN algorithm results are shown in the top panel of Figure 9. The blue and red points show correctly identified OH and H I host galaxies, respectively, and the black stars are misidentified objects. The incorrect identifications concentrate where OH and H I sources overlap the most at the same observed frequency, and indicate where there will be the most confusion. The final OH recall is 0.985 and the OH precision is 0.974. In other words, for redshifts less than $z \sim 1.0$, we expect to identify 98.5% of OH lines in H I surveys, thereby mitigating the impact of contamination.

We repeat this process for another Suess et al. (2016) parameter space, W1-W2 versus W2-W3, as well as the observed line frequency. This alternative approach significantly limits the number of available detections because of the inclusion of the comparatively less sensitive W3 band. This approach, however, leads to a higher OH precision, because mid-IR data are relevant for distinguishing between these populations. The results of this test are shown in the bottom

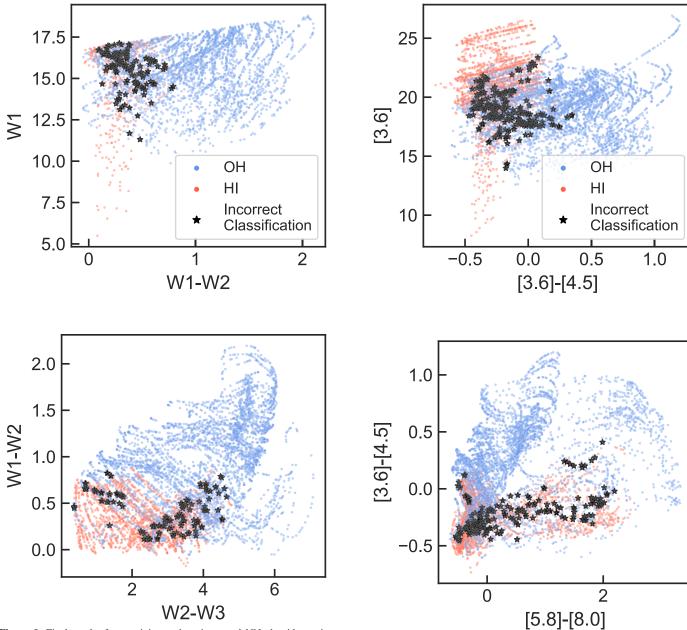


Figure 9. Final results from training and testing our k-NN algorithm using WISE W1 vs. W1 - W2 (above) and W1 - W2 vs. W2 - W3 (below). Blue points indicate OH host testing points that were correctly identified, and red points indicate the same for H I sources. Black stars show the misidentified objects, either OH misidentified as H I or vice versa (3.6% of objects in top panel, 1.9% of objects in bottom panel). Note that the OH and H I markers are partially transparent to show overlapping.

panel of Figure 9. Precision and recall from this test are compared to those for other tests in Table 3.

4.2.2. OH and H I Classification Using IRAC

As discussed previously, WISE is insensitive to galaxies at redshifts above $z\sim 1.0$. Although having significantly less sky coverage than WISE, Spitzer's IRAC bands 1 and 2 are very similar to WISE bands 1 and 2 but are much more sensitive and can detect OHM and H I host galaxies over the full redshift ranges probed by both LADUMA and SKA1. We therefore perform an exercise similar to that in Section 4.2.1 using IRAC

Figure 10. Final results from training and testing our k-NN algorithm using Spitzer IRAC [3.6] vs. [3.6] – [4.5] (above) and [3.6] – [4.5] vs. [5.8] – [8.0] (below). Blue points indicate OH host testing points that were correctly identified, and red points indicate the same for H I sources. Black stars show misidentified objects (3.4% of objects in top panel, 2.9% of objects in bottom panel). Note that the OH and H I markers are partially transparent to show overlapping.

data. Throughout this paper, IRAC bands are referred to by their wavelengths in microns (e.g., IRAC [3.6] denotes the $3.6 \mu m$ band magnitude).

We use a parameter space analogous to the first test for the IRAC k-NN algorithm ([3.6] versus [3.6] – [4.5]). (As a reminder, WISE uses Vega-based magnitudes, whereas IRAC uses AB magnitudes.) Because IRAC is sensitive to the entire redshift range of our H I and OHM hosts, we do not perform any detection cuts. Results from this exercise are presented in Table 3 and visualized in Figure 10. Comparing to the analogous WISE space, this test has the same OH precision, but OH recall suffers slightly. However, achieving an OH recall of

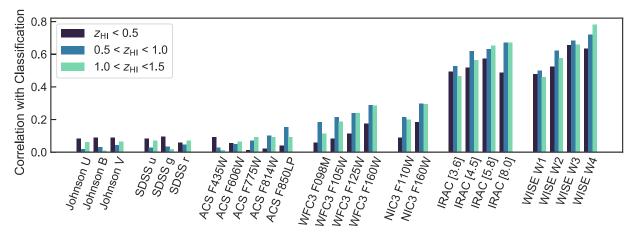


Figure 11. Correlation of band photometry with OH/H I classification using a Pearson correlation test. Bands on the *x*-axis are grouped by mission or type and then ordered by increasing wavelength. Each correlation is tested in three redshift ranges.

0.979 is still a powerful tool when it comes to sorting OHM hosts from H I hosts, and the ability to probe to higher redshifts has strong appeal.

We also consider IRAC [3.6] – [4.5] versus [5.8] – [8.0]. Stern et al. (2005) suggest that cuts in this space can separate active galaxies from normal galaxies. We perform algorithm optimizations similar to those mentioned previously before training and testing. We present the final results in Table 3 and Figure 10. Despite having information from the [8.0] band, this test performs slightly worse than the previous tests in both OH recall and precision, indicating that the overlap between OHM and H I hosts is greater in this parameter space than in the previous alternatives.

Being able to probe the full redshift range of LADUMA is beneficial but inevitably introduces more contamination, as indicated by the slightly reduced OH recall. However, it is worth noting that despite being lower, these recall values still exceed 95%. These IRAC and WISE tests create a new framework for the process of separating OH and HIhost populations.

5. Discussion

The methods presented in this paper will be crucial to mitigating OHM contamination of HI emission-line surveys. WISE provides the all-sky coverage needed for upcoming surveys that will be covering large portions of the sky, while IRAC has the deep-field coverage needed for surveys such as LADUMA, which will be the deepest HI emission-line survey to date.

One of the biggest shortcomings of these methods and calculations is that they are based on known OHMs, which currently extend to a highest redshift of $z_{\rm OH}=0.264$ (Darling & Giovanelli 2002a). This limitation has forced us to make some extrapolations to obtain predictions for higher-redshift surveys. This approach is unavoidable until we have higher-redshift data on both H I and OH populations. As more surveys are conducted, we will be able to update the OHLF and OH SED evolution as well as provide tighter constraints on these calculations and predictions.

As discussed in Section 1, H I and OH sources can be separated by spectroscopic redshift. It is therefore worth recognizing that some objects will be readily identifiable and that these objects will be crucial for helping classify those without redshifts. For LADUMA, the current largest source of

spectroscopic redshifts is the PRIsm MUlti-Object Survey (PRIMUS), with over 32,000 redshifts in the field and in the relevant redshift range (Coil et al. 2011). PRIMUS only detects galaxies out to $z \sim 1.2$, meaning that some of the most potentially contaminated (i.e., highest) redshift ranges will have few spectroscopic redshifts available. Another source of spectroscopic redshifts soon to come online is the Wide-Area VISTA Extragalactic Survey (WAVES), which will have two campaigns, WAVES-Wide (large-sky, low-redshift) and WAVES-Deep (small-area, high-redshift) (Driver et al. 2019). WAVES-Deep will have several small patches, including one on the LADUMA field. Slated to target 45,000 objects, WAVES-Deep will also be crucial in identifying objects; however, the current estimates only show detections out to redshift $z \sim 0.8$. For future all-sky untargeted H I surveys such as those on ASKAP or the SKA, WAVES-Wide aims to provide 880,000 spectroscopic redshifts out to redshift $z \sim 0.2$. Although we may have many redshifts for identifying objects as OH or HI sources, these redshifts are extremely limited where potential OH contamination is the greatest threat.

WISE and IRAC photometry were not the only data tested in Section 4 for the ability to separate OHM and HIhosts. We also tested other photometry for separability, focusing on data that have significant coverage in the LADUMA field. These include SDSS *ugr*, Johnson *UBV*, and *HST* ACS, WFC3, and NICMOS bands. Figure 11 shows how each of these bands correlates with OH/HI classification. For each band, a Pearson correlation test was done for three redshift ranges. Bands are grouped on the *x*-axis and then sorted by increasing wavelength. Figure 11 demonstrates the sorting value of bands in the near- to mid-IR. This distinction is due to the extreme star formation in OHM host galaxies, which is detected in the IR. Optical bands are poorer candidates for separation, because they are less sensitive to high star formation rates in dusty, gasrich systems.

6. Conclusions

We present predictions for the numbers of OH megamasers that will be detected in future untargeted H I surveys and explore how those numbers impact H I source confusion over a range of redshifts up to $z_{\rm H~I} = 1.5$. To assist in untangling these populations, we also present methods for estimating the likelihood that a line has been identified as H I or OH. Below,

we summarize our predictions and discuss the implications of this work for future H I surveys:

- 1. LADUMA will likely triple the number of known OHMs: we predict 83^{+21}_{-17} new detections. Larger surveys with telescopes such as the SKA1 will detect thousands more OHMs.
- 2. The contamination these OHM detections will impose on H I line surveys is highly dependent on redshift (and, secondarily, depth). In a line-flux-limited survey, OHMs are more abundant at higher redshift, while H I sources become sparser. For these high-redshift surveys, OH detections will outnumber H I detections near redshift $z_{\rm H~I} \sim 1.0$.
- 3. Near- and mid-IR observations can assist in separating H I from OHM emission lines, which we demonstrate using a k-NN machine-learning algorithm. We will be able to identify nearly 99% of OH lines for redshifts less than $z \sim 1.0\%$ and 96% of lines at higher redshifts.

Although OHM host galaxies represent a potential contamination for untargeted H I line surveys, these rare and interesting objects can be important scientific tools. As discussed in Section 2.4, OHM density can also provide an independent measurement on the major merger rate evolution parameter, γ , because OHMs serve as tracers of major galaxy mergers. These galaxies are signposts of the most extreme star formation in our universe, signaling where the most massive starbursts are happening (Briggs 1998), and can even offer a way to measure in situ magnetic fields using Zeeman splitting (Robishaw et al. 2008; McBride et al. 2014). The methods presented in this paper and follow-up observations will begin uncovering these galaxies and allowing us to characterize them at higher redshifts and potentially create better methods for mitigating contamination in H I surveys.

This work has been supported by the National Science Foundation through grants AST-1814648 (to H.R. and J.D.) and AST-1814421 (to A.B.). This research has made use of the NASA/IPAC Extragalactic Database (NED), which is funded by the National Aeronautics and Space Administration and operated by the California Institute of Technology. We thank Sarah Blyth, Natasha Maddox, and Aaron Stemo for helpful insights and conversations. We also thank the anonymous referee for their thorough and helpful comments which improved the presentation of this paper.

Software: Astropy (Robitaille et al. 2013), Matplotlib (Hunter 2007), Numpy (Van Der Walt et al. 2011), Scipy (Virtanen et al. 2020), Scikit-Learn (Pedregosa et al. 2011), MAGPHYS (Da Cunha et al. 2008), emcee (Foreman-Mackey et al. 2013), PYPHOT (https://github.com/mfouesneau/pyphot).

ORCID iDs

Hayley Roberts https://orcid.org/0000-0003-0046-9848
Jeremy Darling https://orcid.org/0000-0003-2511-2060
Andrew J. Baker https://orcid.org/0000-0002-7892-396X

References

Abdalla, F. B., Bull, P., Camera, S., et al. 2015, in Proc. of Advancing

```
Astrophysics with the Square Kilometre Array (AASKA14) (Trieste:
  PoS), 17
Altman, N. S. 1992, Amer. Statist., 46, 175
Beichman, C. A., Neugebauer, G., Habing, H. J., Clegg, P. E., & Chester, T. J.
   1988, IRAS, 1, 1
Blyth, S. L., Baker, A. J., Holwerda, B. W., et al. 2016, in Proceedings of
  MeerKAT Science: On the Pathway to the SKA (Cape Town: MeerKAT
  Science), 4
Braun, R., Bonaldi, A., Bourke, T., Keane, E., & Wagg, J. 2019, arXiv:1912.
Briggs, F. H. 1998, A&A, 336, 815
Brown, M. J., Moustakas, J., Smith, J. D., et al. 2014, ApJS, 212, 18
Coil, A. L., Blanton, M. R., Burles, S. M., et al. 2011, ApJ, 741, 8
Da Cunha, E., Charlot, S., & Elbaz, D. 2008, MNRAS, 388, 1595
Darling, J. 2005, in ASP Conf. Ser. 340, In Future Directions in High
  Resolution Astronomy: The 10th Anniversary of the VLBA, ed.
  J. D. Romney & M. J. Reid (San Francisco, CA: ASP), 216
Darling, J. 2007, ApJL, 669, L9
Darling, J., & Giovanelli, R. 2002a, AJ, 124, 100
Darling, J., & Giovanelli, R. 2002b, ApJ, 572, 810
Driver, S. P., Baldry, I. K., Brown, M. J. I., et al. 2019, Msngr, 175, 46
Duffy, A. R., Meyer, M. J., Staveley-Smith, L., et al. 2012, MNRAS,
  426, 3385
Duncan, K., Conselice, C. J., Mundy, C., et al. 2019, ApJ, 876, 110
Fernández, X., Gim, H. B., van Gorkom, J. H., et al. 2016, ApJL, 824, L1
Fernández, X., Van Gorkom, J. H., Hess, K. M., et al. 2013, ApJL, 770, 2
Ferreira, L., Conselice, C. J., Duncan, K., et al. 2020, ApJ, 895, 115
Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP,
Goldberger, J., Roweis, S., Hinton, G., & Salakhutdinov, R. 2005, in NIPS'04:
  Proc. of the 17th Int. Conf. on Neural Information Processing Systems, 17
  (Cambridge, MA: MIT Press), 513
Hess, K. M., Roberts, H., Dénes, H., et al. 2021, A&A, 647, A193
Hogg, D. W. 1999, arXiv:astro-ph/9905116
Hunter, J. D. 2007, CSE, 9, 99
Kessler, M. F., & Habing, H. J. 1996, AdSpR, 13, 485
Lockett, P., & Elitzur, M. 2008, ApJ, 677, 985
Lotz, J. M., Jonsson, P., Cox, T. J., et al. 2011, ApJ, 742, 103
Mantha, K. B., McIntosh, D. H., Brennan, R., et al. 2018, MNRAS, 475,
Martin, D. C., Fanson, J., Schiminovich, D., et al. 2005, ApJL, 619, L1
McBride, J., Quataert, E., Heiles, C., & Bauermeister, A. 2014, ApJ, 780, 182
Mundy, C. J., Conselice, C. J., Duncan, K. J., et al. 2017, MNRAS, 470, 3507
Obreschkow, D., Kloeckner, H.-R., Heywood, I., Levrier, F., & Rawlings, S.
O'Leary, J. A., Moster, B. P., Naab, T., & Somerville, R. S. 2021, MNRAS,
  501, 3215
Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, JMLR, 12, 2825
Robishaw, T., Quataert, E., & Heiles, C. 2008, ApJ, 680, 981
Robitaille, T. P., Tollerud, E. J., Greenfield, P., et al. 2013, A&A, 558, A33
Rodriguez-Gomez, V., Genel, S., Vogelsberger, M., et al. 2015, MNRAS,
  449, 49
Sammut, C., & Webb, G. I. 2011, Encyclopedia of Machine Learning (Berlin:
  Springer)
Skrutskie, M. F., Cutri, R. M., Stiening, R., et al. 2006, AJ, 131, 1163
Staveley-Smith, L., & Oosterloo, T. 2014, in Proc. of Advancing Astrophysics
   with the Square Kilometre Array (AASKA14) (Trieste: PoS), 167
Stern, D., Eisenhardt, P., Gorjian, V., et al. 2005, ApJ, 631, 163
Stoughton, C., Lupton, R. H., Bernardi, M., et al. 2002, AJ, 123, 485
Suess, K. A., Darling, J., Haynes, M. P., & Giovanelli, R. 2016, MNRAS,
  459, 220
Van Der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, CSE, 13, 22
Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, Nat. Method, 17, 261
Werner, M. W. 2005, AdSpR, 36, 1048
Wright, E. L., Eisenhardt, P. R., Mainzer, A. K., et al. 2010, ApJ, 140,
```