



A DEEP SEARCH FOR ADDITIONAL SATELLITES AROUND THE DWARF PLANET HAUMEA

LUKE D. BURKHART^{1,2}, DARIN RAGOZZINE^{1,3,4}, AND MICHAEL E. BROWN⁵¹Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA; darin.ragozzine@gmail.com²Yale University, Department of Physics, 217 Prospect Street, New Haven, CT 06511, USA³University of Florida, 211 Bryant Space Science Center, Gainesville, FL 32611, USA⁴Florida Institute of Technology, Department of Physics and Space Sciences, 150 West University Boulevard, Melbourne, FL 32901, USA⁵California Institute of Technology, Division of Geological and Planetary Sciences, MC 150-21, Pasadena, CA 91125, USA

Received 2015 July 10; accepted 2016 March 27; published 2016 May 27

ABSTRACT

Haumea is a dwarf planet with two known satellites, an unusually high spin rate, and a large collisional family, making it one of the most interesting objects in the outer solar system. A fully self-consistent formation scenario responsible for the satellite and family formation is still elusive, but some processes predict the initial formation of many small moons, similar to the small moons recently discovered around Pluto. Deep searches for regular satellites around Kuiper belt objects are difficult due to observational limitations, but Haumea is one of the few for which sufficient data exist. We analyze *Hubble Space Telescope* (*HST*) observations, focusing on a 10-consecutive-orbit sequence obtained in 2010 July, to search for new very small satellites. To maximize the search depth, we implement and validate a nonlinear shift-and-stack method. No additional satellites of Haumea are found, but by implanting and recovering artificial sources, we characterize our sensitivity. At distances between $\sim 10,000$ and $\sim 350,000$ km from Haumea, satellites with radii as small as ~ 10 km are ruled out, assuming an albedo ($p \simeq 0.7$) similar to Haumea. We also rule out satellites larger than $\gtrsim 40$ km in most of the Hill sphere using other *HST* data. This search method rules out objects similar in size to the small moons of Pluto. By developing clear criteria for determining the number of nonlinear rates to use, we find that far fewer shift rates are required (~ 35) than might be expected. The nonlinear shift-and-stack method to discover satellites (and other moving transients) is tractable, particularly in the regime where nonlinear motion begins to manifest itself.

Key words: Kuiper belt objects: individual (Haumea) – methods: data analysis – techniques: image processing

Supporting material: data behind figure

1. INTRODUCTION

The dwarf planet Haumea (Rabinowitz et al. 2006), its two moons (Brown et al. 2005, 2006), and its collisional family (Brown et al. 2007) provide important constraints on the formation of the Kuiper belt and the outer solar system. This well-studied object is the fastest-rotating large body in the solar system (Rabinowitz et al. 2006) with rotational variability in color (Lacerda et al. 2008; Lacerda 2009), an unexpectedly high density (Lockwood et al. 2014), and large albedo (Elliot et al. 2010). It has two moons on dynamically excited orbits (Ragozzine & Brown 2009, hereafter **RB09**) which have scaled mass ratios and distances similar to the Earth–Moon system. Dynamical, photometric, and spectroscopic observations of objects in the vicinity of Haumea clearly indicate a collisional family of icy fragments with similarly high albedos (Ragozzine & Brown 2007; Schaller & Brown 2008; Snodgrass et al. 2010). However, though the expected dispersion velocity of these fragments is of the order of several hundred meters per second, the observed dispersion is well constrained within $\sim 150 \text{ m s}^{-1}$. The apparent lack of high velocity ejecta is confirmed by observational surveys and dynamical studies (e.g., Fraser & Brown 2012; Lykawka et al. 2012; Volk & Malhotra 2012), though it is possible that some high velocity ejecta would be unrecognizable dynamically (Marcus et al. 2011) and/or compositionally (Brown et al. 2012; Carry et al. 2012).

There is no simple high-probability formation scenario that naturally explains all of these observational constraints: Haumea’s rapid near-breakup rotation rate, the two moons on distant and dynamically warm orbits, and a collisional family

that is an order of magnitude smaller in velocity dispersion than expected. Though multiple explanations and variations have been proposed (e.g., Levison et al. 2008; Schlichting & Sari 2009; Leinhardt et al. 2010; Marcus et al. 2011; Ortiz et al. 2012; Čuk et al. 2013), none have adequately and self-consistently explained all of the unique features of this interesting system and its family.

Attempting to place the formation of the Haumea system in context with other similar systems in the Kuiper belt quickly leads to comparisons with Kuiper belt objects (KBOs) of similar sizes, particularly Eris, Pluto, and Makemake. Of these, Pluto is the best understood due to a wealth of observational data and the recent flyby by the New Horizons Mission (Stern et al. 2015). Furthermore, there are similarities between some of the theories for the formation of Haumea’s satellites (e.g., Leinhardt et al. 2010) and for the formation of Pluto’s satellites (e.g., Canup 2011): both suggest a relatively large impactor with very low incoming velocities that undergo a grazing collision to form a satellite system. With the discovery of a retinue of small satellites exterior to Charon’s orbit—now dubbed Styx, Nix, Kerberos, and Hydra—there is a renewed interest in observational constraints on the formation of the Pluto system (Weaver et al. 2006; Showalter et al. 2011, 2012; Showalter & Hamilton 2015). Standard explanations for the formation of Nix and Hydra were already problematic (Ward & Canup 2006; Lithwick & Wu 2008), and the characteristics of Styx and Kerberos are even more puzzling (Pires dos Santos et al. 2012; Cheng et al. 2014b; Kenyon & Bromley 2014). For example, in the current orbital configuration, the dynamical stability of Styx requires that Charon’s eccentricity at its

present semimajor axis was never above ~ 0.035 , using the circumbinary stability criterion of Holman & Wiegert (1999, see Equation (3)). Thus, the discovery of Styx combined with dynamical stability immediately precludes some of the more extreme proposed orbital histories of Cheng et al. (2014a) if Styx formed concurrently with Charon (Cheng et al. 2014b). Long-term dynamical stability can also place some of the best constraints on the masses of these small moons (Youdin et al. 2012; Ćuk et al. 2013; Porter & Stern 2015; Showalter & Hamilton 2015).

The discovery of small moons around Pluto and their ability to add constraints to the understanding of this system suggests that all asteroid and KBO binaries and triples be searched for additional moons. We recommend the continuation of this standard practice, even when an initial companion is identified. For KBOs, satellite searches are observationally difficult for multiple reasons. First, acquiring data of sufficient depth and resolution to identify faint moons of faint KBOs usually requires a considerable amount of time at the best telescopes in the world, such as the *Hubble Space Telescope* (HST) or 8–10 m class telescopes with Laser Guide Star Adaptive Optics. The only KBOs with large amounts of continuous high-quality data are Pluto and Haumea.

Second, the discovery of small moons can be frustrated by their a priori unknown satellite orbital motion during long exposures. Faint, fast-moving moons can then evade detection even with the best data, using standard analysis methods. Therefore, an enhanced methodology to search for faint moving moons is required.

In an attempt to better understand the formation of the Haumea system, we use a large set of consecutive HST observations to perform a search for very small moons around Haumea similar to those discovered around Pluto (Section 2). To search for faint, fast-moving moons well below the single-exposure limit, we implemented the nonlinear shift-and-stack method proposed by Parker & Kavelaars (2010, hereafter PK10) for the discovery of KBOs (Section 3). Adapted to the problem of finding additional satellites, this method was both efficient and effective. Though no additional satellites of Haumea were detected (Section 4), with careful characterization of this null detection, we set strong limits on the size and location of possible undiscovered moons (Section 5) and discuss the implications for understanding Haumea’s satellite system (Section 6).

2. OBSERVATIONS

In determining the ideal set of observations for a deep satellite search, a balance must be struck between including the largest number of observations and considering the motion of putative satellites during the total observational baseline. The standard stacking method of adding images that have been co-registered to the position of the primary to enhance sensitivity to faint satellites is limited to observational arcs where the satellite’s relative position remains within a region not much larger than ~ 1 point-spread function (PSF) FWHM. Our use of nonlinear shift-and-stack can mitigate this problem significantly and allows us to perform a sensitive search on longer timescales. In particular, our HST Program 12243 observed with a wide filter for 10 consecutive orbits and is an excellent data set for a deep satellite search; with the technique discussed below, we can search these observations even for close-in satellites which traverse a significant fraction of an orbit during

the 15 hr baseline, corresponding to several PSF widths. These observations are our main focus as they are clearly the best for a deep satellite search (Section 2.1); however, we also inspected other observations for the additional satellites of Haumea (Section 2.2).

2.1. HST 12243: 10 Orbits in 2010 July

HST Program 12243 obtained 10 orbits of observations over the course of ~ 15 hr in 2010 July. This program used the Wide Field Camera 3 (WFC3) UVIS imager with the F350LP (long-pass) filter. The primary goal of these observations was the detection of a mutual event between the inner moon Namaka and Haumea (RB09). In order to produce high cadence time series photometry of the proposed mutual event (and to avoid saturation), exposures were limited to ~ 45 s. To prevent a costly memory buffer download, only a 512×512 subarray of the full WFC3 camera was used with a field of view of ~ 20.5 arcsec. HST tracked at Haumea’s rate of motion (except for controlled dithering) to maintain its position near the center of the field of view throughout the observations. The geocentric distance to Haumea at the time of observations was 50.85 au. At this distance, $1''$ corresponds to 36,900 km, 1 WFC3 pixel (0.04 arcsec) corresponds to 1475 km, and the entire subarray field of view corresponds to $\sim 750,000$ km.

Parts of the last two orbits were affected by the South Atlantic Anomaly. This caused a portion of these orbits to lose data entirely, and another portion was severely affected by cosmic rays and the loss of fine pointing precision. The most offensive frames were discarded for the purpose of the satellite search, leaving 260 individual exposures.

The center of Haumea was identified by eye in combination with a 2D Gaussian fitting routine. With these well-determined preliminary Haumea locations, all images were co-registered to Haumea’s position. In this Haumea-centric frame, cosmic rays and hot pixels were identified by significant changes in brightness at a particular position using robust median absolute deviation filters. A detailed and extensive image-by-image investigation of cosmic rays by eye confirmed that this method was very accurate at identifying cosmic rays and other anomalies. Furthermore, the automatic routine did not flag the known objects (Haumea and its two moons: Hi’iaka and Namaka) as cosmic rays, nor were any other specific localized regions identified for consistent masking (i.e., putative additional satellites were not removed).

TinyTim software⁶ was used to generate local PSF models. As in RB09, these PSF models were then fit to the three known objects using standard χ^2 minimization techniques (Markwardt 2009). This identified the best-fit locations and heights of the scaled PSFs, with bad pixels masked and thus not included in the χ^2 calculation. The astrometric positions of the known satellites relative to Haumea seen by this method were in clear accord with their projected orbital motion from RB09. Despite Namaka’s nearness to Haumea (which was purposely chosen, as the goal of the observations was to observe a mutual event), it is easily distinguishable in the first few orbits.

The best-fit PSFs are visually inspected and found to be good fits for all images. These PSFs are then subtracted, removing a large portion of the three signals, but leaving non-negligible residuals; these residuals are caused by imperfect PSFs (note that Haumea may be marginally resolved in some images) and

⁶ <http://www.stsci.edu/hst/observatory/focus/TinyTim>

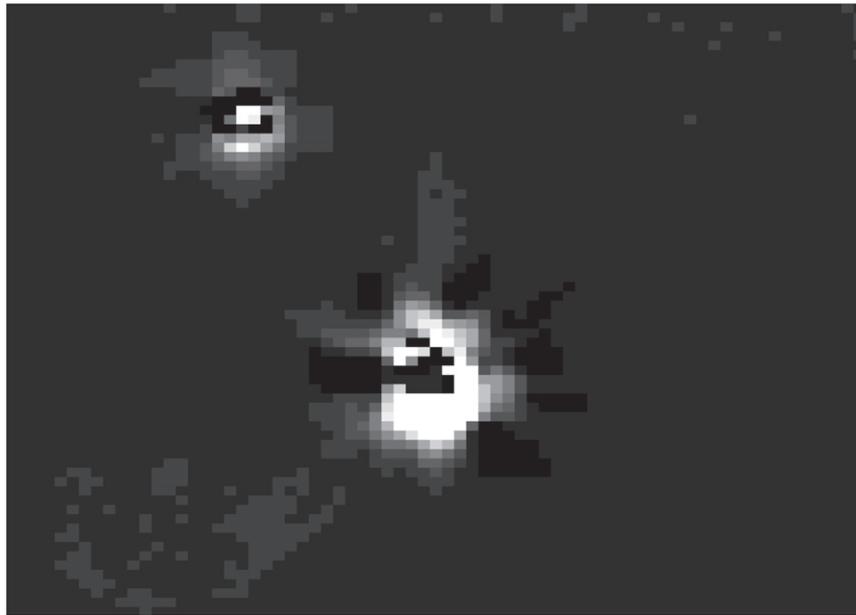


Figure 1. Portion of the median stack of 260 images from 10 orbits of HST WFC3 data (Program 12243). Individual images are co-registered to be stationary in the Haumea-centric frame, with best-fit TinyTim PSFs for Haumea, Hi'iaka, and Namaka subtracted. The brightness has been stretched significantly to highlight the residuals. These residuals will limit sensitivity near Haumea, but the diffraction spikes and the majority of the PSFs have been removed. Above and to the left of Haumea lie the residuals from Hi'iaka which is $1''.23$ away (45,600 km projected distance) in this stack. Vertical darker columns are due to minor uncorrected pixel sensitivity. The image is aligned so that Astronomical North is up. The data used to create this figure are available.

standard Poisson noise. Using the updated best-fit centers of Haumea, these PSF-subtracted images are re-coregistered to Haumea's best-fit location (though the additional shifts from the preliminary 2D Gaussian centers are very small). As described below, these same PSF-subtracted images are used to perform the nonlinear shift-and-stack.

Moons with negligible motion relative to Haumea can be identified in this image by stacking all the observations in the Haumea-centered reference frame (including fractional pixel shifts implemented by IDL's `fshift`). Throughout this investigation, we create stacks by performing a pixel-by-pixel median of the images, which is less sensitive to cosmic rays, bad pixels, and errors in the PSF subtraction of the known bodies. This results in a small decrease in sensitivity, as, assuming white noise, the noise level grows a factor of $\sqrt{\frac{\pi}{2}}$ faster when using the median over the mean. This will correspond to only an $\sim 20\%$ difference in brightness sensitivity, or an $\sim 10\%$ difference in radius, and we find this acceptable so that the other effects mentioned above can be mitigated. A portion of this stacked image around Haumea is shown in Figure 1. Detailed investigation of this deep stack by eye by each co-author yielded no clear satellite candidates. The median stacks were also automatically searched using the IDL routine `find`, which uses a convolution filter with an FWHM of 1.6 pixels to identify positive brightness anomalies. Detections with signal-to-noise ratios (S/Ns) of five or greater were examined; none were found that were consistent with an additional satellite (e.g., having a PSF-like shape). Scaling from the S/N of Haumea and using a more conservative detection limit S/N of 10, this places a limit on non-moving satellites as fainter than about $V \simeq 27.1$, corresponding to Haumea satellites with radii less than about 8 km (see Section 5).

2.2. Other Observations

HST has observed Haumea during many programs for multiple reasons. Program 11518 was proposed to obtain astrometry of both moons and is five independent orbits of observations spread over two weeks (RB09). Although it would be interesting to investigate the possibility of combining these data in a long-baseline nonlinear shift-and-stack, given the existence of other more sensitive data sets, we investigated only the single-orbit median-stacked images. Motion during a single 45-minute *HST* orbit is small compared to the PSF width, even for the shortest satellite orbital periods.

HST Program 11971 was five consecutive orbits and *HST* Program 12004 was seven consecutive orbits, both attempts to observe the last satellite-satellite mutual events. The latter program was within a few weeks of the *HST* 4th Servicing Mission but was still executed. Unfortunately, for 6.5 of the 7 orbits, the STIS shutter was closed and no on-sky data were taken.

For the 5-orbit Wide Field Planetary Camera 2 observation of Program 11971, we median-stacked images centered on Haumea and searched for additional sources by eye and using IDL's `find` as described above. We investigated stacks of individual orbits and of the entire five-orbit sequence and found no sources consistent with additional satellites. Though the nonlinear shift-and-stack method below could fruitfully be applied to these observations, the WFC3 observations are considerably deeper and we opted to focus on our best data set.

Finally, we obtained some long-duration (~ 5 hr) observations of Haumea using the Laser Guide Star Adaptive Optics system at the Keck Observatory. Co-registered stacks of this data also showed no clear additional satellites, though the known satellites were very easily detected.

3. METHODS

For the detection of faint bodies, with S/N per image of $\lesssim 5$, a useful approach is the co-addition (“stacking”) of multiple images. With the 260 images in our data set, this method can increase the S/N by $\sim \sqrt{\frac{2}{\pi}} \sqrt{260} \approx 13$, thereby searching for satellites with radii $\sim \sqrt{13} \approx 3.6$ times as small as could be detected in a single image.

If the object does not remain apparently stationary (within $\lesssim 1$ FWHM) over the course of the observation, the simple co-addition will result in insufficient overlap between images to yield the expected increase in S/N. If the motion of the object is known, images can be first shifted to compensate for this motion, and the images added with the object localized regaining nearly the full sensitivity: this is the meaning of “shift-and-stack.” Linear searches with shift-and-stack have been used to discover satellites in the past (Holman et al. 2004; Kavelaars et al. 2004; Showalter et al. 2013) though these searches did not need to use the nonlinear shift-and-stack method we employ below.

In a situation where the motion is unknown, such as a broad search for KBOs, a large set of possible paths on the sky can be considered, and each path independently used as the basis for a shift-and-stack, as described by PK10. A composite image results from each proposed orbital path (which we call a “sky track”), and each stack can be searched for faint satellites, which emerge from the noise due to shifting the image accurately enough to (mostly) compensate for its motion.

To minimize statistical false-positives and to increase computational tractability, it is important to identify a near-minimal number of sky tracks that will faithfully reproduce all the possible motions without performing redundant searches. PK10 suggested an algorithm for identifying the most important non-redundant set of sky tracks, which we fruitfully employ: we generate a large number of random sky tracks based on the full range of expected motion (within desired search parameters) and then remove tracks that are similar to one another. We have adapted this technique for our search.

There is a distinction between a general KBO search and a satellite search, which, it is important to note, is largely ignored in the method presented here. This distinction is that, in a broad KBO search, a sky track could be valid for any part of an image; that is, there is little correlation between position and motion. This is not the case for a satellite orbiting a given primary, in which a specific motion only applies to a small spatial region. The more highly curved tracks are, the more specific to a particular region they are—a curved orbital arc translated to the other side of Haumea would not make physical sense. The method described below involves shifting and stacking the entirety of each image, and searching the whole of the composite image, when in fact the track upon which the shift-and-stack is based applies to only a small subset of each image. In addition to the computational cost of shifting and searching larger images than is necessary, this overuse of the images could potentially result in an increase in statistical false-positives. However, neither of these effects manifest in a noticeable way—neither computation time nor an abundance of false-positives limit our search method. This suggests that we are near the optimal minimum number of sky tracks searched, or have at least reached an acceptably small number.

Discussed in greater detail below, an overview of our search algorithm is as follows.

1. Generate a large bank of physically reasonable putative sky tracks by randomly selecting from plausible Keplerian satellite orbital parameters.
2. Fit each sky track with nonlinear polynomials in time (shift rates). If the shift rates for two distinct sky tracks are similar enough (quantified below), discard one.
3. Continue searching for sky tracks until a nearly complete non-redundant set is identified.
4. For each track, create a composite image. This is done by overlaying the data set (in our case, 260 images) upon itself, with the images shifted by the appropriate shift rates such that an object on that track will appear in the same place in each image. Co-add the images into one composite.
5. Search each composite image for satellite candidate sources.

The use of nonlinear polynomial fits allows the shift rates to more accurately capture curved orbits than simple linear fits. For the motion of even the fastest detectable Haumean satellites over the timescale of our observations, we find that quadratic fits to the x and y positions are always sufficient. Note that the polynomial fits are included for convenience in describing the sky tracks; the actual positions of a putative satellite could be used, but the difference between the actual positions and the best-fit quadratic approximate was negligible. Including nonlinear rates is often expected to greatly expand the number of dissimilar shift rates to the point of computational impracticality, but we find that an appropriate criterion for similarity of shift rates easily permits the inclusion of quadratic rates.

3.1. Generation of Sky Tracks

In a typical shift-and-stack search for KBOs, the putative sky tracks are selected from a grid of the six degrees of freedom needed to describe an object in a Keplerian orbit (PK10, Bernstein et al. 2004). For the purposes of a KBO satellite, particularly that of a primary with other known satellites, it is convenient to instead sample the space of Keplerian orbital elements relative to the primary: semimajor axis (a), eccentricity (e), inclination (i), longitude of the ascending node (Ω), argument of periapse (ω), and mean anomaly at epoch (M). Sampling in this space allows for direct control over the types of orbits that are searched, making it straightforward to exclude unphysical motions. In our case, we also benefit from a well-known mass of the primary; if this is not known, a variety of plausible values could be sampled for the generation of sky tracks. For this search, a and e were randomly sampled from orbits with semimajor axes between 5310 and 368,000 km and eccentricities less than 0.5, while the orbital angles i , Ω , ω , and M were allowed to assume any value. All parameters were chosen from the sample space uniformly, with the exception of a , which was sampled on a log scale to increase the likelihood of sampling an orbit in the regime of fast-moving satellites.

The lower bound on a is a constraint imposed by our sensitivity of detection. This limit corresponds to 3.75 pixels (15 mas) on the WFC3, at which distance from the center of Haumea the subtraction noise is considerable enough to make reliable detections difficult (see Figure 1). The upper bound on a is set much larger than the semimajor axes needed to shift images (as opposed to investigation of the unshifted stack). At a distance where the satellite’s maximum velocity would cause it to travel less than one PSF FWHM over the course of the

15 hr observational period (here $\sim 27 \text{ m s}^{-1}$), shift-and-stack is unnecessary, giving the upper limit of $a \simeq 150,000$ in our search. This semimajor axis is ~ 3 times the semimajor axis of Hi'iaka, whose motion in these frames is detectable, but $\lesssim 0.5$ pixels. For the upper limit on a , we doubled this number to be conservative.

Much of this orbital parameter space can be excluded on physical grounds, reducing the number of shift rates necessary to well-sample the space. Any putative orbit that crossed paths with the known satellites was rejected, as was any orbit with periapsis less than 3000 km. These weak restrictions on orbital elements did not appreciably affect the selection of shift-and-stack parameters and additional tests (described below) show that we are sensitive to objects on practically any orbit with a semimajor axis $\gtrsim 10,000$ km.

3.2. Nonlinear Fitting and Shift Rate Similarity

Having created a bank of physically plausible orbits, we then generate a set of shift rates with which to create composite images to search for satellites. Orbital parameters were converted into sky coordinates relative to Haumea, right ascension ($\Delta\text{R.A.}$) and declination ($\Delta\text{decl.}$), for each image, as described in RB09. We assumed an instantaneous Keplerian orbit for the position of the satellite as this is an excellent approximation over the course of our observations. In our case, orbital acceleration was quite important, as we desired to search orbital periods down to ~ 40 hr, of which the 15 hr observational arc is a sizable fraction. Therefore, the sky positions $\Delta\text{R.A.}$ and $\Delta\text{decl.}$ were fit with quadratic polynomials in time, which we found were sufficient to accurately describe the nonlinear motion in every case.

In order to minimize the number of sky tracks, we eliminated tracks which were similar to one another, as suggested by PK10. To determine if two tracks were similar, we focused on the final requirement that the shift rates localize the flux of satellite so that it can be identified in the stacked image. If the flux of a satellite traveling along the second orbit would be well-localized by the shift rates of the first, then there would be considerable overlap of the flux between images when shifted according to the rates of the first. This criterion can be quantified by calculating the overlap fraction between two shift-and-stack rates using the reasonable assumption that the WFC3 PSF is nearly Gaussian, with FWHM of 0.067 arcsec (≈ 1.7 pixels). For a pair of shift rates, the overlap was defined for each image in the data set as the integral of the product of two such Gaussians separated by the difference in the two rates ($\Delta\text{R.A.}$ and $\Delta\text{decl.}$) at the time of that image. We call this the overlap between two orbits as it is calculated from the product of two overlapping PSFs, but it is distinct from the concept of the overlap in co-added images. If the median overlap (normalized to 1 for perfect coincidence) was greater than a pre-specified threshold, it was considered that a sufficient fraction of the flux of the proposed satellite would have been collected by the stack of an existing sky track, and the new track was rejected as unnecessary.

The goal is to build up a bank of sky tracks known to be mutually distinct. After accepting the first track into the bank, each subsequent track was compared to the previously selected tracks in the bank using the above overlap criterion. We experimented with different overlap threshold criteria and found the overall results mostly insensitive to the specific value chosen. In general, we required a median overlap of less than

0.7 with each previously accepted shift-and-stack tracks to accept the proposed track as distinct enough to add to the bank.

By drawing from a large set of orbits covering the desired search space, this method efficiently builds a bank of mutually distinct shift rates that are also the most relevant (PK10). However, unlike a grid search, random orbital draws can continue indefinitely. Thus, we also require a “stopping criterion” to decide when the bank is large enough for practical use. To determine the upper limit on the number of necessary shift rates, we noted that the sample space saturates quickly; that is, the rate of acceptance drops off drastically after 10–15 shift rates are chosen. Consequently, the number of orbits rejected between successive accepted rates grows very quickly. Our criterion for a dense sampling was that the number of rejected rates between successive accepted rates was at least equal to the total number of rates rejected so far. Put another way, the selection was stopped when the acceptance of each new shift rate required doubling the number of sampled orbits, which typically occurred after testing hundreds of thousands of random orbits. This is an exponentially slow process, which suggests that once past this threshold, we have reached the limits of our rate selection method. Practically speaking, we found that this stopping criterion still generated enough shift rates to recover injected sources with a complete variety of orbits.

Together with the above ranges for satellite orbital parameters, this method yielded only ~ 35 sufficiently distinct orbits over which to search. Considering the size of the parameter space (linear and quadratic terms for shifts in both x and y directions), it might seem surprising that so small a set of potential orbits spans the space. However, a large portion of the space consists of short, almost entirely linear tracks, where the quadratic corrections are of limited importance. The only strongly quadratic orbits are very near to Haumea, which also have large linear rates. In other words, there are strong correlations between the allowed linear and quadratic coefficients of physically plausible tracks. The result is a relatively small number of nonlinear shift rates that efficiently cover the desired search space (PK10); these tracks are illustrated in Figure 2. Contrasted with the orbital element motivated sampling presented here, the number of shift rates in the case of a quadratic sky motion grid search would have been much larger.

3.3. Creation of Composite Image

With our bank of non-degenerate sky tracks, we can now perform the nonlinear shift-and-stack procedure. Each track corresponds to a specific set of $\Delta\text{R.A.}$ and $\Delta\text{decl.}$ values of a putative satellite relative to Haumea. We used `adxy`, a routine from the IDL Astro Library which uses astrometric data from the image headers, to convert these sky coordinates into on-image pixel positions, thus yielding the desired pixel shifts.

The prepared images were shifted (including fractional pixel shifts implemented by IDL's `fshift`) and stacked using the pixel-by-pixel median of the images as described above. In preparation for the automated search, many images were investigated in detail by eye.

3.4. Sensitivity

To test the sensitivity of this method, artificial sources were implanted into the images with a range of random brightnesses.

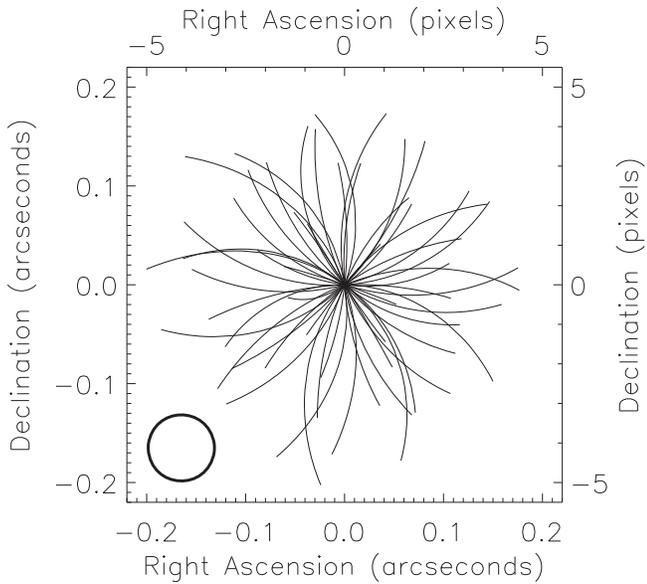


Figure 2. Nonlinear shift-and-stack rates. Arcs show the displacement of images (relative to position at the middle image) over the course of the 15 hr observation. Each arc represents a different shift rate or “sky track.” Horizontal and vertical axes show differential right ascension (Δ R.A.) and declination (Δ decl.) in arcseconds and pixels (1 pixel = 0.04 arcsec = 1475 km). The circle at the bottom left has a diameter of 0.067 arcsec, the FWHM of WFC3’s PSF. Following the method suggested by PK10, we generate random nonlinear shift rates out of Keplerian orbital elements. We reject as duplicates rates for which the overlapping PSFs would catch at least 70% of the flux if moving at the same rate as an existing orbit (see Section 3.2). This method requires only ~ 35 nonlinear rates to cover the vast majority of parameter space. The “sky tracks” associated with these rates are mostly symmetric about the origin as seen above, with slight asymmetries arising from the projection of eccentric orbits into the skyplane, and the variation in orbital speed throughout an orbit. Almost all rates are substantially quadratic, which shows the importance of the nonlinear approach. As can be seen, the use of quadratic shift rates allowed us to probe the region near Haumea where satellites would execute sizable fractions of an orbit during the 15 hr observation. Implantation of artificial sources on orbits randomly drawn from the same Keplerian elements showed an excellent recovery rate (see Figures 3 and 4).

Their positions and rates of motion were determined by orbits randomly drawn from the same space mentioned above (but without restriction of non-crossing orbits with the known moons). The implants were generated by scaling from the actual PSF of Haumea (when brightest). This source was implanted into the images at the pixel positions corresponding to the randomly chosen orbit. A subimage of 200×200 pixels was used for the search: the outer reaches of this subimage have objects that are practically not moving (see Section 3.1) and any object beyond this region would have the same detectability threshold as stationary objects.

This was done for a large number of orbits, with a new set of images created for each. Stacks were generated in the exact same manner as the real images, with the same ~ 35 shift rates making new median stacks for each new set of images. These stacks were inspected using the same automated search routine (IDL’s find). To distinguish detections of implanted sources from the detection of the three known bodies, we examined the output of the search for the sets of stacks with no sources implanted. All detections here were due to known bodies, and the positions were used to establish a mask with three regions, one for each known body, to reject detections that were not due to implanted sources. In this way, detections could be automatically classified as a recovery of an implant or as a

false positive due to the known bodies. These automated classifications were extensively verified with an investigation that included searching by eye and were found to be very robust.

Due to the application of a threshold S/N by the find routine, objects in the vicinity of Haumea, while still far enough and bright enough to be seen by eye, may be rejected by the routine itself (not our masks). The presence of the nearby primary leads to an artificially high computed background noise level, which reduces the computed S/N significantly, causing the object to appear below threshold. Any stacks with sources at risk of being left undetected due to this effect were searched by eye by multiple coauthors, and any that were detected in this processes were considered to be recovered for the purposes of our results, shown below.

The success rates of finding implanted objects places constraints on additional satellites of Haumea: any recovered implanted source represents a satellite that we can say with reasonable certainty is not present in the Haumea system.

4. RESULTS

The implantation and successful recovery of faint moving sources clearly indicated the effectiveness of our nonlinear shift-and-stack method. Nevertheless, we did not detect any additional satellites around Haumea and no candidate satellites were found that were worthy of additional investigation.

A careful characterization of this null result is able to place strong limits on the brightness and separation of undiscovered Haumean satellites. These limits are summarized by Figures 3 and 4, which show the results of our search for each implanted source. The source is either recovered, rejected (for being too close to one of the three known objects, usually Haumea), or “missed” because it was too faint to be detected, or because it fell off the 200×200 subimage that was searched. Note that the “rejected” category is primarily composed of objects that were not clearly detected by the automated routine, but were detected in a blind search by eye by multiple coauthors; these consist entirely of objects that are $\lesssim 0''.2$ (5 pixels) from Haumea. Figure 3 plots semimajor axis against brightness of the sources as a fraction of that of Haumea, while Figure 4 shows the projected distance (in arcseconds) of the moving sources versus the brightness. Assuming the same albedo ($p \simeq 0.7$) as Haumea, the relative brightness corresponds to the radius of a spherical satellite, which is also shown.

5. DISCUSSION

The constraints on undiscovered Haumean satellites can be divided into three categories based on orbital semimajor axis: close-in satellites ($a \lesssim 10,000$ km), intermediate satellites ($10,000 \text{ km} \lesssim a \lesssim 350,000$ km), and distant satellites ($a \gtrsim 350,000$ km).

5.1. Limits on Close-in Satellites

At a semimajor axis of $\sim 10,000$ km, the maximum separation of a satellite from Haumea would be 6.9 WFC3 pixels. Within 7 pixels ($\lesssim 4$ PSF FWHM) of Haumea, it is very difficult to recover objects due to imperfect subtraction of Haumea’s PSF. It is possible that an empirical PSF subtraction would perform better for recovering very close-in satellites, but we do not consider such an approach here. As can be seen in Figure 4, there is the expected anti-correlation between the

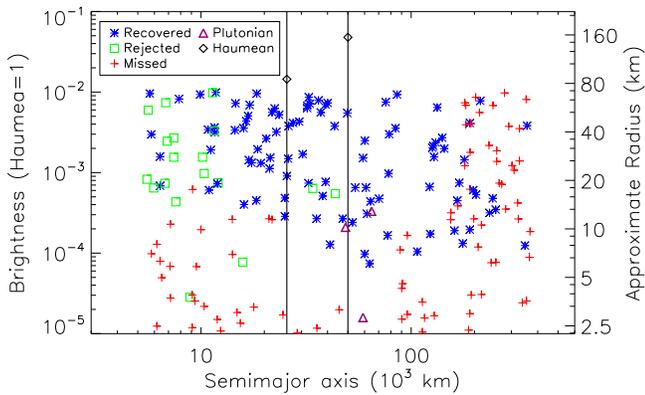


Figure 3. Results from the sensitivity survey. The figure shows implanted sources that were either recovered (blue stars), not recovered (red crosses), or recovered but rejected due to confusion with existing sources (green squares). The horizontal axis is the semimajor axis of implanted objects in thousands of kilometers. The left-hand vertical axis is brightness relative to Haumea (when brightest). The right-hand vertical axis is radius of a spherical satellite assuming an albedo (~ 0.7) similar to Haumea. Diamonds represent known satellites Namaka and Hi'iaka; the vertical lines are guides to the eye at their respective semimajor axes. Purple triangles represent the moons of Pluto—Nix, Hydra, and Kerberos—according to brightness relative to the primary. (The smallest moon Styx, with brightness approximately 6×10^{-6} that of Pluto, is below the range of brightness represented in this figure.) Because of differences in geocentric distance and albedo, the approximate radius does not directly apply to these three points. Figure 4 is similar, but shows distance in projected separation instead of semimajor axis. Bounds on brightness and semimajor axis were chosen as described in Section 3.1. The unrecovered implantations at a semimajor axis $\gtrsim 200 \times 10^3$ km are not found because their distance from Haumea often places them outside the subimages searched. This figure shows that satellites with radii as low as ~ 8 km would be detectable in much of the space searched, and that our lower detection limit on semimajor axis is limited by the properties of the data set, not by the sensitivity of the nonlinear shift-and-stack technique. Nix- and Hydra-like objects would be detected around Haumea, while Styx- and Kerberos-like objects would still be too faint, mostly due to Haumea's further distance (50 au compared to Pluto's 30 au).

brightness of an object that can be recovered and the separation from Haumea: close in, only brighter objects can be found.

However, there are dynamical reasons to expect that this region is nearly devoid of satellites. Due to Haumea's highly triaxial shape, the orbital region near Haumea is strongly perturbed and long-term stable orbits are difficult to maintain. According to Scheeres (1994), periods less than about 10 times the spin period are unlikely to be stable due to primary-spin-satellite-orbit resonances. In Haumea's case, this is exacerbated by the additional effects of tidal evolution and other dynamically excited satellites (Canup et al. 1999; Ćuk et al. 2013; Cheng et al. 2014b). An orbital period that is 10 times the spin period corresponds to a semimajor axis of about 5000 km (about five times the long-axis radius of Haumea). While about twice as distant as the Roche radius, for long-term dynamical stability, we consider this the inner limit.

Even if satellites were originally found in such short orbits, it is possible that long-term tidal evolution would have moved them to a more detectable distance. A detailed analysis by Ćuk et al. (2013) calls into question the idea first proposed that the satellites tidally evolved outward from orbits near the Roche lobe. While extensive tidal evolution might not have taken place, it is worth noting that scaling the tidal evolution from the properties of the other satellites (Brown et al. 2005), indicates that even for the smallest satellites we could have detected (which evolve the shortest distance due to tides), tidal evolution

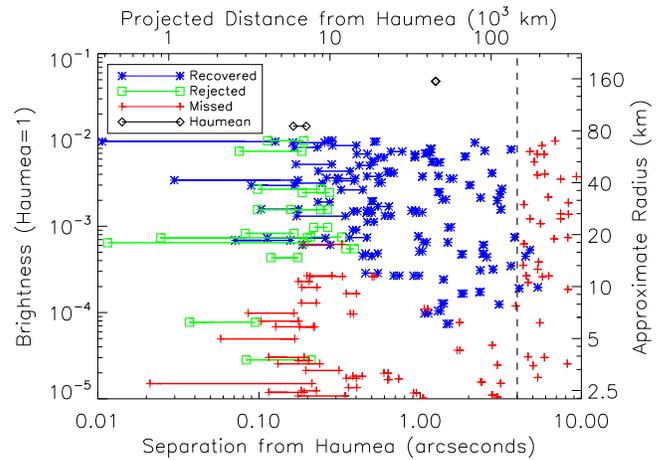


Figure 4. Results from the sensitivity survey. The lower horizontal axis is the sky-planet projected separation from Haumea in arcseconds, while the upper axis gives approximate projected distance from Haumea in thousands of kilometers. The vertical axes are brightness and radius of implanted sources, as described in the caption of Figure 3. Symbols connected by horizontal lines show the maximum and minimum apparent distance from Haumea of the implanted object during the 15 hr “observation.” As in Figure 3, implanted sources were either recovered (blue stars), not recovered (red crosses), or recovered but rejected due to confusion with existing sources (green squares). Note that implantations at separations greater than 4 arcsec are unrecovered because they fall outside the region of the $\sim 4''$ subimages that were searched (see Section 5.3). The vertical dashed line is a guide to the eye for this rough cutoff. No sources were implanted at separations larger than 10 arcsec, corresponding to the upper limit on semimajor axis shown in Figure 3. Diamonds represent Hi'iaka and Namaka as they appear in the observation; Namaka's separation is only given for the first four orbits where its presence is measured reliably enough for precise astrometry; as these observations were designed to catch Namaka in a mutual event, its projected separation would approach very low values if all 10 orbits were included.

would have placed them near or beyond the ~ 5000 km detection threshold.

There remains a range of semimajor axes from 5000–10,000 km that could potentially harbor very small undetected satellites, which would be somewhat protected from dynamical and tidal instability. By lying well within Haumea's PSF, these satellites also generally evade detection. Furthermore, some satellites would not have been detected if they had an orbital phase placing them at undetectably small distances (although this is mitigated somewhat by observations at a variety of times). Overall, it is difficult to hide stable inner ($a \lesssim 10,000$ km) Haumean satellites with radii $\gtrsim 30$ km.

5.2. Intermediate Satellites

At semimajor axes between about 10,000 and 350,000 km lies the region of satellites near where the other two moons are detected (at semimajor axes of 25,600 km for Namaka and 49,900 km for Hi'iaka, RB09). At this distance, contamination from Haumea is negligible and the main limitation to detecting satellites is insufficient S/N or falling beyond the edge of the image. By using the nonlinear shift-and-stack, we maximize the search depth, particularly closer to Haumea.

The search depth can be reported as relative brightness (in magnitudes and flux) and as the radius of a spherical satellite assuming the same albedo as Haumea. As is usual for such deep searches, the recovery rate is a function of magnitude (Figures 3 and 4). We reliably detect satellites at -9.2 mag (0.0002 relative brightness, radius of 10 km), our recovery rate

Table 1
Summary of Estimated Properties of Dwarf Planet Satellites

Object	Satellite	Relative Brightness (magnitudes)	H_{sat}^a	V_{sat}^a	Radius ^b (km)	a^c (10^3 km)	Ref
Haumea	Hi'iaka	−3.3	3.4	20.5	200	50	1
Haumea	Namaka	−4.6	4.7	21.8	150	26	1
Haumea	“close” upper limit	−6.7	6.8	23.9	30	$\lesssim 10$	2
Haumea	“intermediate” upper limit	−10.0	10.1	27.6	8	10–350	2
Haumea	“distant” upper limit	−6.2	6.3	23.4	40	$\gtrsim 350$	2
Pluto	Charon	−2.6	1.9	16.6	350	20	3
Pluto	Hydra	−8.7	8.0	22.7	41	64	3
Pluto	Nix	−9.2	8.5	23.2	35	49	3
Pluto	Kerberos	−12	11	26	12	59	4
Pluto	Styx	−13	12	27	11	42	5
Eris	Dysnomia	−6.7	5.5	25.4	60	37	6
Eris	“close” upper-limit	−5.8	4.6	24.5	80	$\gtrsim 18$	6
Eris	“distant” upper-limit	−8.2	7.0	26.9	30	$\gtrsim 37$	6
Makemake	S/2015 (136472) 1	−7.5	7.1	25.1	25	~ 100	7
Makemake	upper-limit	−10	9.6	26.9	8	$\gtrsim 30$	7, 8

Notes. Magnitudes and semimajor axes of bodies in KBO systems. The relative magnitude of the faintest detectable bodies in our search is -10 , comparable to that of Hydra and Nix. For Eris and Makemake, values are more approximate and/or interpolated from published estimates. We do not list the large number of KBO binaries (e.g., Noll et al. 2008) or KBO triple 1999TC36 (Benecchi et al. 2010) since the formation of these systems appears to be distinct from processes associated with dwarf planets. In particular, these binaries tend to be nearly equal in brightness without known small additional companions.

^a Approximate absolute magnitude (H) or approximate apparent magnitude in a typical optical filter (V) of the satellite. These are calculated combining the relative magnitude with the absolute and typical apparent magnitudes of the KBOs from JPL Horizons. These are meant mostly for illustration purposes and generally have significant uncertainties of $\lesssim 1$ mag.

^b Radius estimate in kilometers, listed for illustration purposes only. Quoted radii for the highly ellipsoidal small satellites of Pluto are volumetric means (Weaver et al. 2016). Note that these have albedoes of 0.5, somewhat less than assumed for Haumea’s moons. For simplicity and ease of inter-comparison, observed moons of Eris and Makemake are given an estimated albedo of 0.7 like the Haumea moons. The actual albedo and size of these moons is not well constrained.

^c Approximate semimajor axis in units of thousands of kilometers. For upper limits, this is the approximate range of semimajor axes where this limit applies.

References. (1) RB09 (Ragozzine & Brown 2009); (2) Section 4, this paper; (3) Weaver et al. (2006); (4) Showalter et al. (2011); (5) Showalter et al. (2012); (6) Brown & Schaller (2007); (7) Parker et al. (2016); (8) Brown (2008). The discovery of S/2015 (136472) 1 by Parker et al. (2016) within the magnitude and distance “upper-limit” quoted by Brown (2008) is easily attributed to the difficulty of detecting moons with small semi-major axes and/or edge-on orbits in single-epoch observations when the actual on-the-sky separation is often small enough to render the moon indistinguishable from the primary Parker et al. (2016). The upper-limits reported here should be understood with that caveat.

is roughly 50% at -10 mag (0.0001 relative brightness, radius of 8 km), and our best case recovery is at -10.4 mag (0.00007 relative brightness, radius of 6 km). Following typical practice, we summarize the recovery depth using the 50% recovery rate. Note that it is possible that the albedo of the satellites is even higher; using Haumea family member 2002 TX300’s measured albedo of 0.9 (Elliot et al. 2010) instead of Haumea’s presumed 0.7 albedo (Lockwood et al. 2014) would imply a radius detection threshold of only ~ 7 km (or ~ 5 km in the best case).

While close approaches to Hi’iaka and Namaka as projected on the sky would result in a missed detection for faint objects, this is generally unlikely (even for orbits coplanar with the known satellites which are near edge-on, RB09). Close approaches to Hi’iaka and Namaka are negligibly unlikely to happen at more than one epoch,⁷ thus any missed detection would be mitigated for moderately bright objects by the non-detection of satellites in other data sets. Thus, we expect that this region of the Haumean system does not contain undiscovered satellites larger than ~ 8 km in radius.

Our results compare favorably with the current state of knowledge regarding the small satellites of Pluto. From the New Horizons flyby, we now have detailed knowledge of the albedoes (about 0.5) and sizes of the small satellites: ~ 10 km

for Styx and Kerberos and ~ 40 km for Nix and Hydra (Stern et al. 2015). As Figure 3 shows, we predict that a satellite of apparent magnitude relative to Haumea, similar to that of Hydra or Nix around Pluto (-8.7 and -9.2 mag, respectively), would fall above our detection limit. With the higher expected albedo (0.7) of Haumean satellites, we would have detected objects as large as Styx and Kerberos. We conclude that Haumea very likely does not contain small satellites similar to Pluto’s.

5.3. Distant Satellites

Satellites with semimajor axes beyond 350,000 km may not have been detected in the Program 12243 WFC3 data due to the small field of view employed for the subarray observations. Other HST and Keck observations that were not as deep covered a larger area and were also searched for satellites. We estimate that satellites larger than about 40 km in radius (again assuming an albedo similar to Haumea’s) would have been detected even several tens of arcseconds away by, e.g., the WFC2 observations (with a field of view of $162''$). Because the motion of satellites in this region is negligible over the relevant timescales, the shift-and-stack method is not necessary.

Using half the size of Haumea’s Hill sphere at perihelion as an estimate of the full region of stable satellites (Shen & Tremaine 2008), the semimajor axis of the most distant stable

⁷ Unlike irregular satellites of the giant planets, long-term tidal stability precludes Hi’iaka or Namaka from being binaries themselves.

satellites would be about 4.6×10^6 km or $124''$. About half of this volume has been covered down to 40 km in radius.

For comparison, the Program 12243 deep observations covered separations up to about $10''$ around Haumea or 350,000 km. Thus, this limit on very small intermediate-range satellites corresponds to about 0.5% of the stable region radius.

6. CONCLUSIONS

By efficient application of the **PK10** method for nonlinear shift-and-stack and recovery of known implanted sources, we have strongly limited the possibility of undetected satellites in orbit around Haumea. As Figure 4 shows, we detect no satellites larger than ~ 8 km in radius with separations between 10,000 and 350,000 km. This same region around Pluto contains Charon and four small satellites which, by size, would all have been detected in this search.

Nearer to Haumea, diffraction limits make distinguishing small satellites difficult, but there are dynamical reasons to expect that this region is mostly unpopulated. Further from Haumea, other observations would have detected satellites larger than ~ 40 km in radius within much of the entire region of possible stable satellites. Significant improvement in the detection limits on smaller satellites would require extensive observations that are unlikely in the foreseeable future until, perhaps, deep observations with the *James Webb Space Telescope*.

Though Pluto contains multiple small moons and some formation theories (e.g., Lithwick & Wu 2008) predict that they are in the Haumea system, we find no additional Haumean moons. Considering upper limits from other studies (summarized in Table 1), Nix/Hydra analogues would have been discovered if present around Makemake and they would be near the detection threshold around Eris.

As the properties of the dwarf planet satellite systems differ significantly, it was not anticipated that Pluto's small satellites would necessarily find counterparts around Haumea, though it seems that Makemake may have a satellite of similar size (Parker et al. 2016). Our null result affirms that, for the time being, Pluto is the only known KBO with a retinue of small satellites, though such could have been detected or nearly detected around all four dwarf planets. This implies that the satellite systems may result from somewhat different formation pathways, although all the dwarf planet satellites are probably connected with a collisional formation. Pluto's small satellite system may be connected with Charon since, from a dynamical perspective, the other dwarf planet satellites are more like small moons compared to the near-equal-sized Pluto–Charon binary.

We demonstrate that the nonlinear shift-and-stack is a valuable tool for satellite searches. Utilizing the application techniques developed herein, this method can sufficiently capture the nonlinearity of the orbits of fast-moving satellites close to the primary. We have applied this technique to the regime of searching for sub-threshold satellites around Haumea, but it could also be used for other long-observation data sets (**PK10**). Besides the discovery of new moons, this technique has promise for improving astrometric parameters for known faint moving satellites (e.g., precovery observations of Styx and Kerberos). The tractability of the nonlinear shift-and-stack also promotes the possibility of applying this to the general search for KBOs, as originally proposed by **PK10**. Other applications for improving sensitivity are also possible, e.g., searching for moving exoplanets in direct imaging campaigns (Males et al. 2013). To facilitate further analyses,

all data and source codes used in this project are available upon request.

The sensitivity and tractability of the method presented in this work suggests that, when appropriate, it should be applied to other satellite searches in the solar system. The non-detection of small satellites around Haumea increases our understanding of this intriguing object and contributes to our understanding of the formation and evolution of multiple KBO systems.

We thank Alex Parker, Danielle Hastings, and the anonymous referee for discussions and suggestions that improved the manuscript. D.R. acknowledges the support of a Harvard Institute for Theory and Computation Fellowship. This work is based on NASA/ESA *Hubble Space Telescope* Program 12243. Support was provided by NASA through grants HST-GO-12243 from the Space Telescope Science Institute (STScI), which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-26555.

REFERENCES

- Benecchi, S. D., Noll, K. S., Grundy, W. M., & Levison, H. F. 2010, *Icar*, **207**, 978
- Bernstein, G. M., Trilling, D. E., Allen, R. L., et al. 2004, *AJ*, **128**, 1364
- Brown, M. E. 2008, *The Solar System Beyond Neptune*, ed. M. A. Barucci et al. (Tucson, AZ: Univ. of Arizona Press), 335
- Brown, M. E., Barkume, K. M., Ragozzine, D., & Schaller, E. L. 2007, *Natur*, **446**, 294
- Brown, M. E., Bouchez, A. H., Rabinowitz, D., et al. 2005, *ApJL*, **632**, L45
- Brown, M. E., & Schaller, E. L. 2007, *Sci*, **316**, 1585
- Brown, M. E., Schaller, E. L., & Fraser, W. C. 2012, *AJ*, **143**, 146
- Brown, M. E., van Dam, M. A., Bouchez, A. H., et al. 2006, *ApJL*, **639**, L43
- Canup, R. M. 2011, *AJ*, **141**, 35
- Canup, R. M., Levison, H. F., & Stewart, G. R. 1999, *AJ*, **117**, 603
- Carry, B., Snodgrass, C., Lacerda, P., Hainaut, O., & Dumas, C. 2012, *A&A*, **544**, A137
- Cheng, W. H., Lee, M. H., & Peale, S. J. 2014a, *Icar*, **233**, 242
- Cheng, W. H., Peale, S. J., & Lee, M. H. 2014b, *Icar*, **241**, 180
- Čuk, M., Ragozzine, D., & Nesvorný, D. 2013, *AJ*, **146**, 89
- Elliot, J. L., Person, M. J., Zuluaga, C. A., et al. 2010, *Natur*, **465**, 897
- Fraser, W. C., & Brown, M. E. 2012, *ApJ*, **749**, 33
- Holman, M. J., Kavelaars, J. J., Grav, T., et al. 2004, *Natur*, **430**, 865
- Holman, M. J., & Wiegert, P. A. 1999, *AJ*, **117**, 621
- Kavelaars, J. J., Holman, M. J., Grav, T., et al. 2004, *Icar*, **169**, 474
- Kenyon, S. J., & Bromley, B. C. 2014, *AJ*, **147**, 8
- Lacerda, P. 2009, *AJ*, **137**, 3404
- Lacerda, P., Jewitt, D., & Peixinho, N. 2008, *AJ*, **135**, 1749
- Leinhardt, Z. M., Marcus, R. A., & Stewart, S. T. 2010, *ApJ*, **714**, 1789
- Levison, H. F., Morbidelli, A., Vokrouhlický, D., & Bottke, W. F. 2008, *AJ*, **136**, 1079
- Lithwick, Y., & Wu, Y. 2008, arXiv:0802.2951L
- Lockwood, A. C., Brown, M. E., & Stansberry, J. 2014, *EM&P*, **111**, 127
- Lykawka, P. S., Horner, J., Mukai, T., & Nakamura, A. M. 2012, *MNRAS*, **421**, 1331
- Males, J. R., Skemer, A. J., & Close, L. M. 2013, *ApJ*, **771**, 10
- Marcus, R. A., Ragozzine, D., Murray-Clay, R. A., & Holman, M. J. 2011, *ApJ*, **733**, 40
- Markwardt, C. B. 2009, in ASP Conf. Ser. 411, *Astronomical Data Analysis Software and Systems XVIII*, ed. D. A. Bohlender, D. Durand, & P. Dowler (San Francisco, CA: ASP), 251
- Noll, K. S., Grundy, W. M., Chiang, E. I., Margot, J.-L., & Kern, S. D. 2008, *The Solar System Beyond Neptune*
- Ortiz, J. L., Thirouin, A., Campo Bagatin, A., et al. 2012, *MNRAS*, **419**, 2315
- Parker, A. H., Buie, M. W., Grundy, W. M., & Noll, K. S. 2016, *ApJL*, submitted (arXiv:1604.07461)
- Parker, A. H., & Kavelaars, J. J. 2010, *PASP*, **122**, 549
- Pires dos Santos, P. M., Morbidelli, A., & Nesvorný, D. 2012, *CeMDA*, **114**, 341
- Porter, S. B., & Stern, S. A. 2015, arXiv:1505.05933
- Rabinowitz, D. L., Barkume, K., Brown, M. E., et al. 2006, *ApJ*, **639**, 1238

- Ragozzine, D., & Brown, M. E. 2007, *AJ*, **134**, 2160
Ragozzine, D., & Brown, M. E. 2009, *AJ*, **137**, 4766
Schaller, E. L., & Brown, M. E. 2008, *ApJL*, **684**, L107
Scheeres, D. J. 1994, *Icar*, **110**, 225
Schlichting, H. E., & Sari, R. 2009, *ApJ*, **700**, 1242
Shen, Y., & Tremaine, S. 2008, *AJ*, **136**, 2453
Showalter, M. R., de Pater, I., French, R. S., & Lissauer, J. J. 2013, *BAAS*, **45**, 206.01
Showalter, M. R., & Hamilton, D. P. 2015, *Natur*, **522**, 45
Showalter, M. R., Hamilton, D. P., Stern, S. A., et al. 2011, *IAUC*, **9221**, 1
Showalter, M. R., Weaver, H. A., Stern, S. A., et al. 2012, *IAUC*, **9253**, 1
Snodgrass, C., Carry, B., Dumas, C., & Hainaut, O. 2010, *A&A*, **511**, A72
Stern, S. A., Bagenal, F., Ennico, K., et al. 2015, *Sci*, **350**, aad1815
Volk, K., & Malhotra, R. 2012, *Icar*, **221**, 106
Ward, W. R., & Canup, R. M. 2006, *Sci*, **313**, 1107
Weaver, H. A., Buie, M. W., Buratti, B. J., et al. 2016, *Sci*, **351**, aae0030
Weaver, H. A., Stern, S. A., Mutchler, M. J., et al. 2006, *Natur*, **439**, 943
Youdin, A. N., Kratter, K. M., & Kenyon, S. J. 2012, *ApJ*, **755**, 17