The Astrophysical Journal, 942:99 (19pp), 2023 January 10 https://doi.org/10.3847/1538-4357/aca480

A Light in the Dark: Searching for Electromagnetic Counterparts to Black Hole–Black Hole Mergers in LIGO/Virgo O3 with the Zwicky Transient Facility

Matthew J. Graham[1], Barry McKernan[2,3,4,5], K. E. Saavik Ford[2,3,4,5], Daniel Stern[6], S. G. Djorgovski[1], Michael Coughlin[7], Kevin B. Burdge[8,9], Eric C. Bellm[10], George Helou[11], Ashish A. Mahabal[12,13], Frank J. Masci[11], Josiah Purdum[14], Philippe Rosnet[15], and Ben Rusholme[11]
[1] California Institute of Technology, 1200 E. California Blvd, Pasadena, CA 91125, USA; mjg@caltech.edu
[2] Department of Science, CUNY Borough of Manhattan Community College, 199 Chambers Street, New York, NY 10007, USA
[3] Department of Astrophysics, American Museum of Natural History, Central Park West, New York, NY 10028, USA
[4] Physics Program, CUNY Graduate Center, 365 5th Avenue, New York, NY 10016, USA
[5] Center for Computational Astrophysics, Flatiron Institute, New York, NY 10010, USA
[6] Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA 91109, USA
[7] School of Physics and Astronomy, University of Minnesota, Minneapolis, MN 55455, USA
[8] Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[9] Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[10] DIRAC Institute, Department of Astronomy, University of Washington, 3910 15th Avenue NE, Seattle, WA 98195, USA
[11] IPAC, California Institute of Technology, 1200 E. California Blvd, Pasadena, CA 91125, USA
[12] Division of Physics, Mathematics and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA
[13] Center for Data Driven Discovery, California Institute of Technology, Pasadena, CA 91125, USA
[14] Caltech Optical Observatories, California Institute of Technology, Pasadena, CA 91125, USA
[15] Université Clermont Auvergne, CNRS/IN2P3, LPC, F-63000 Clermont-Ferrand, France

Received 2022 September 23; revised 2022 November 15; accepted 2022 November 19; published 2023 January 17

Abstract

The accretion disks of active galactic nuclei (AGNs) are promising locations for the merger of compact objects detected by gravitational wave (GW) observatories. Embedded within a baryon-rich, high-density environment, mergers within AGNs are the only GW channel where an electromagnetic (EM) counterpart must occur (whether detectable or not). Considering AGNs with unusual flaring activity observed by the Zwicky Transient Facility (ZTF), we describe a search for candidate EM counterparts to binary black hole (BBH) mergers detected by LIGO/Virgo in O3. After removing probable false positives, we find nine candidate counterparts to BBH mergers during O3 (seven in O3a, two in O3b) with a p-value of 0.0019. Based on ZTF sky coverage, AGN geometry, and merger geometry, we expect $\approx 3(N_{\rm BBH}/83)(f_{\rm AGN}/0.5)$ potentially detectable EM counterparts from O3, where $N_{\rm BBH}$ is the total number of observed BBH mergers and $f_{\rm AGN}$ is the fraction originating in AGNs. Further modeling of breakout and flaring phenomena in AGN disks is required to reduce our false-positive rate. Two of the events are also associated with mergers with total masses $>100\ M_\odot$, which is the expected rate for O3 if hierarchical (large-mass) mergers occur in the AGN channel. Candidate EM counterparts in future GW observing runs can be better constrained by coverage of the Southern sky as well as spectral monitoring of unusual AGN flaring events in LIGO/Virgo alert volumes. A future set of reliable AGN EM counterparts to BBH mergers will yield an independent means of measuring cosmic expansion (H_0) as a function of redshift.

Unified Astronomy Thesaurus concepts: Active galactic nuclei (16); Astrophysical black holes (98); Stellar mass black holes (1611); Supermassive black holes (1663); Gravitational waves (678); Galaxy accretion disks (562)

1. Introduction

The gravitational wave (GW) detectors Advanced LIGO (Aasi et al. 2015) and Advanced Virgo (Acernese et al. 2015), hereafter referred to as LIGO/Virgo, detected binary black hole (BBH) mergers in the local ($z < 1$) universe at a rate of about once per week during O3 (O3a: 2019 March 1–2019 September 30; O3b: 2019 November 1–2020 March 30; LIGO Scientific Collaboration & Virgo Collaboration 2019). BBH mergers can come from two broad classes of channels (for a recent review, see Mapelli 2021): a field binary origin (i.e., from the evolution of a field binary system consisting of two massive stars; e.g., Belczynski et al. 2010; de Mink & Mandel 2016) and a dynamical origin. Among dynamical BBH mergers, subchannels include mergers in globular clusters (e.g., Rodriguez et al. 2016a, 2016b), mergers in quiescent galactic nuclei (e.g., Antonini 2014; Antonini & Rasio 2016; Fragione et al. 2019), and mergers in the accretion disks of active galaxies (e.g., McKernan et al. 2019; Graham et al. 2020). Due to pair-instability supernovae (SNe), which leave no compact remnant, the explosive deaths of massive stars are not thought to produce black holes (BHs) in the "upper mass gap" range of ~50–130 $M_\odot$ (Woosley 2017). There is also a "lower-mass gap" range of ~3–5 $M_\odot$, corresponding to the observed absence of compact objects with masses between the most massive neutron stars (NSs) and the least massive BHs. Massive BBH merger progenitors detected by LIGO/Virgo in the upper end of the upper mass gap range strongly imply a hierarchical (i.e., dynamical) merger origin. Since BHs can receive a strong kick at merger (e.g., Varma et al. 2022), hierarchical mergers are more easily retained in deep gravitational potentials, such as in the nuclei of galaxies (e.g., Gerosa & Berti 2019; Gerosa & Fishbach 2021).

A promising location for hierarchical mergers are active galactic nuclei (AGNs; e.g., McKernan et al. 2012, 2014, 2018; Bellovary et al. 2016; Bartos et al. 2017b; Stone et al. 2017; Secunda et al. 2019; Tagawa et al. 2020, 2021; Yang et al. 2019; Secunda et al. 2020; Samsing et al. 2022). Graham et al. (2020) presented the first candidate counterpart for such an event. Merger kicks, even of large magnitude (Varma et al. 2022), are insufficient to escape an AGN environment where the Keplerian orbital velocity is $\mathcal{O}(10^4)\text{km s}^{-1}(R/10^3 r_g)$ at disk radius R, where $r_g = GM_{\rm SMBH}/c^2$ is the supermassive black hole (SMBH) gravitational radius and $M_{\rm SMBH}$ is the SMBH mass. This makes AGNs ideal for retaining and growing BHs via hierarchical mergers. AGNs are expected to dominate the rate of mergers in the deep potential wells of gas-free galactic nuclei (Ford & McKernan 2022). Besides BBH mergers at the upper end of the mass gap, other pointers to a significant contribution to observed BH mergers from the AGN channel include significantly asymmetric mass ratio BH mergers and the observed anticorrelation between BH mass ratio and BBH effective spin (Callister et al. 2021), which at present can only be explained in the context of the AGN channel (Wang et al. 2021b; McKernan et al. 2022a).

Unlike all other BH merger channels, significant detectable electromagnetic (EM) counterparts may develop due to compact object mergers in AGNs (e.g., McKernan et al. 2019; Graham et al. 2020; Wang et al. 2021a; Kimura et al. 2021; Perna et al. 2021). Furthermore, identified counterparts to BH mergers in AGN accretion disks provide a test of the dynamics of the merger, as well as a probe of fundamental AGN disk properties (Vajpeyi et al. 2022). In addition, if we can confidently associate particular GW mergers with specific AGNs at identified redshifts (e.g., Ashton et al. 2021; Calderón Bustillo et al. 2021; Palmese et al. 2021), the GW signal becomes a standard siren that provides a new, independent measurement of the Hubble constant, H_0, as a function of redshift (Mukherjee et al. 2020; Chen et al. 2022).

If we are optimistic about identifying EM counterparts to BH mergers in AGN disks, this approach promises to produce the merger locations of GW sources, an important new probe of AGN accretion disks and a powerful technique to measure the expansion history of our universe over a critical redshift range. However, one might also be pessimistic about identifying EM counterparts in AGN disks. Even if AGNs are responsible for most of the GW-detected BH mergers, we might not detect EM counterparts due to either the muffling of embedded EM signatures by optically thick disks, or the emerging EM flare might be too faint to detect against a bright, variable quasar disk. Even in the pessimistic case, however, a search for unusual AGN flares is valuable as a test of extreme-variability mechanisms in AGN disks. We also note that even if EM counterparts are never confidently detected from AGNs, a significant AGN contribution to the merger rate can still be estimated from a statistical approach (see, e.g., Bartos et al. 2017a; Veronesi et al. 2022).

The EM signature of a merged BBH in an AGN disk depends on the reaction of the surrounding disk gas to the merger. In general, a merged BBH in an AGN disk moves away from the merger site with recoil kick velocity v_k, which is a function of binary mass asymmetry and spin orientation. Gas that was gravitationally bound to the BBH will attempt to follow the merged product, but collides with surrounding AGN disk gas, heating it and producing a bright shock in the disk, possibly detectable in the optical/UV wave band in a thin or relatively dim disk (McKernan et al. 2019). An even brighter EM signature may result from the continued onward progression of the recoiling BH through the disk as surrounding gas produces a Bondi drag "tail" behind the kicked BH. This can generate a significant, detectable luminosity at super-Eddington accretion rates (Graham et al. 2020), as long as a jetted outflow allows radiation to escape and emerge (McKernan et al. 2019; Wang et al. 2021a). EM counterparts can emerge on the side of the AGN disk facing toward, or away from, the observer's sight line.

The search for EM counterparts to isolated NS-NS or BH-NS mergers typically requires rapid follow-up with coverage of as much of the LIGO/Virgo localization map as possible (Coughlin et al. 2019; Kasliwal et al. 2020; Anand et al. 2021). In contrast, the EM signal for BBH mergers in AGN disks only emerges days to weeks after the merger event (McKernan et al. 2019). The detection strategy is therefore different from the non-AGN case and requires regular monitoring of the AGN population within the LIGO/Virgo error volume rather than rapid scanning. Modern time-domain surveys which observe large areas of sky with day-to-week cadences therefore present an ideal data set for identifying possible counterpart events. However, this also illustrates the importance of correctly updating public localization maps: if the parameterization of a GW event changes over time, the two-dimensional and three-dimensional event localization can change significantly.

In this paper, we present a search with the Zwicky Transient Facility (ZTF; Bellm et al. 2019a; Graham et al. 2019) for EM counterparts in AGN disks to all BBH merger detections by LIGO/Virgo during the O3 run. Over this period, ZTF covered the visible sky above decl. = −30° from the Palomar Observatory every three nights in the g and r bands to ~20.5 mag (5σ detection limit). This provides a data set with the required large spatial coverage and sampling cadence to detect an association between a BBH merger and an AGN flare. The paper is structured as follows. In Section 2, we consider how a BBH merger in an AGN disk could generate an EM signal; in Section 3, we consider other events that could produce an equivalent signal; Section 4 describes our search procedure; and Section 5 presents our results. We discuss the implications of our results in Section 6 and detail our conclusions in Section 7. Throughout, magnitudes are reported in the AB system and we adopt the Planck 2015 cosmology, $H_0 = 67.7\ \text{km s}^{-1}\ \text{Mpc}^{-1}$, $\Lambda = 0.693$, and $\Omega_m = 0.307$ (Planck Collaboration et al. 2016).

2. Electromagnetic Counterparts to Binary Black Hole Mergers in Active Galactic Nuclei Disks

BBH mergers in AGN disks occur in the presence of gas and must always produce some EM radiation, though the detectability of the resulting EM signature depends on three basic factors:

1. *Can the EM counterpart escape from the midplane of a dense, optically thick disk on sufficient timescales?*

   If a BBH merger occurs in (and remains in) the midplane of a disk that is both optically and geometrically thick, the diffusion timescale for radiation produced at the midplane is many years and the signature is not detectable (McKernan et al. 2019). If the merger happens in the midplane of a razor-thin disk, an EM signature could emerge promptly. However, razor-thin disks are disfavoured since models of pressure-supported disks tend to generate modest disk aspect ratios (e.g., Sirko & Goodman 2003; Thompson et al. 2005). A BH remnant that is strongly kicked at merger could quickly emerge from an optically thick midplane into a diffuse, optically thin disk atmosphere. There is, however, also a geometrical problem: an EM counterpart to a merger in an edge-on AGN or a kick that pushes the remnant to the far side of the disk with respect to the observer will be obscured.

2. *Is the change in brightness relative to the already bright AGN disk detectable?*

   The brighter the AGN, the less likely it is that we can identify a flare associated with the BBH merger. The brightness of the flare also depends on how the EM emission emerges from the merger. A shock is usually too dim to show up against bright AGNs and will only be detectable against lower-luminosity AGNs (McKernan et al. 2019; Graham et al. 2020). However, a jet from a kicked, rapidly spinning and accreting merger product may be sufficiently bright (Wang et al. 2021a), or may be powerful enough to mechanically clear a low-optical-depth path out from the midplane, unless, for example, pre-merger outflows have excavated a Local Bubble (Kimura et al. 2021) or the accretion rate is insufficient to generate high luminosity (Pan & Yang 2021).

3. *Can we distinguish a resulting flare from false positives?*

   Even if the flare is bright enough to be observable, we still need to be able to distinguish the flare from other AGN variability events or known false positives. This requires models of intrinsic extreme AGN variability (Graham et al. 2017) as well as an understanding of light-curve and color evolution from embedded disk eruptions such as SNe or tidal disruption events (TDEs) in the presence of an AGN accretion disk (Chan et al. 2019; Yang et al. 2022), as well as microlensing events (e.g., Lawrence et al. 2016).

In this section, we address the first two of these points and review our model for a prompt EM counterpart to a BBH merger in an AGN disk (McKernan et al. 2019). This underpinned our reporting of the first plausible candidate EM counterpart to a GW BBH merger trigger (Graham et al. 2020). In particular, we consider what the model implies for the properties of the associated flare, including the characteristic timescales, luminosity, and flare shape. We discuss the third point, namely the statistical uniqueness of the EM signature, in Sections 3 and 4.

2.1. Parameters from the Initial Gravitational Wave Trigger

The initial notification of a candidate GW event from LIGO/Virgo reports the luminosity distance (d_L), the confidence interval for the luminosity distance (Δd_L), and the 90% confidence interval for the sky localization area (A_{90}) of the event (Veitch et al. 2015; Singer & Price 2016; Singer et al. 2016). The probability distribution for d_L is a convolution of source mass, detector orientation, and source angle to the observer line of sight. The localization area is proportional to the signal-to-noise ratio (S/N) of the event detection by LIGO/Virgo, $A_{90} \propto \mathrm{S/N}^{-2}$ (Berry et al. 2015), where $\mathrm{S/N} \propto M_c^{5/6}\, d_L^{-1}$ (Finn et al. 1993) and M_c is the chirp mass.[16] Thus, for any given GW event trigger at time $t = 0$, we can estimate the approximate source-frame BBH mass, M_{BBH}, and we have a search volume for AGN EM counterparts in the volume given by $A_{90} \times \Delta d_L$.

2.2. Parameters from the Bound Gas Shock

At merger, the new BH recoils with a kick velocity, v_k. In an AGN disk, gas at distance $R_{\mathrm{bound}} < GM_{\mathrm{BBH}}/v_k^2$ is bound to the merged BBH and attempts to follow the kicked merger product. In doing so, it collides with surrounding disk gas and a shock luminosity emerges on a timescale $t_{\mathrm{bound}} = R_{\mathrm{bound}}/v_k = GM_{\mathrm{BBH}}/v_{\mathrm{k}}^3$ (McKernan et al. 2019), which can be parameterized as

$$t_{\mathrm{bound}} \sim 20\ \mathrm{day}\left(\frac{M_{\mathrm{BBH}}}{100\ M_{\odot}}\right)\left(\frac{v_{\mathrm{k}}}{200\ \mathrm{km\ s^{-1}}}\right)^{-3}. \tag{1}$$

This is a low-luminosity effect ($\mathcal{O}(10^{42})$ erg s^{-1}) compared to other mechanisms discussed here and so we do not consider it any further. We note, however, that if v_k is very small ($<$50 km s^{-1}) then this low-luminosity prompt flare is likely the only EM counterpart.

2.3. Parameters from the Bondi Drag Accretion and Shock

Once the kicked BH leaves behind originally bound gas, the disk gas it passes through is accelerated around the BH, producing a shocked Bondi tail (e.g., Ostriker 1999; Antoni et al. 2019). This tail both acts as a drag on the BH and accretes onto it. We assume the associated Bondi–Hoyle–Lyttleton (BHL) luminosity is $L_{\mathrm{BHL}} = \eta \dot{M}_{\mathrm{BHL}} c^2$, where η is the radiative efficiency and the mass accretion rate is

$$\dot{M}_{\mathrm{BHL}} = \frac{4\pi G^2 M_{\mathrm{BBH}}^2 \rho}{v_{\mathrm{rel}}^3}, \tag{2}$$

where ρ is the local disk gas density, $v_{\mathrm{rel}} = v_k + c_s$, and c_s is the gas sound speed, assumed to be $c_s \sim 50\ \mathrm{km\ s^{-1}}$ (Graham et al. 2020). Then we can parameterize L_{BHL} as

$$L_{\mathrm{BHL}} = 2.5 \times 10^{45} \mathrm{erg\ s^{-1}} \left(\frac{\eta}{0.1}\right)\left(\frac{M_{\mathrm{BBH}}}{100\ M_{\odot}}\right)^2 \times \left(\frac{v_{\mathrm{rel}}}{200\ \mathrm{km\ s^{-1}}}\right)^{-3}\left(\frac{\rho}{10^{-10}\ \mathrm{g\ cm^{-3}}}\right). \tag{3}$$

Following Graham et al. (2020), the dynamical time in the source frame associated with the ram-pressure shock (or the time for the merger remnant to cross the sphere of bound gas) is $t_{\mathrm{ram}} = G\ M_{\mathrm{BBH}}/v_k^3 \sim 20\ \mathrm{days}(M_{\mathrm{BBH}}/100\, M_{\odot})\ (v_k/200\ \mathrm{km\ s^{-1}})^{-3}$. We assume that the luminosity of the flare rises linearly to L_{BHL} from $t = [t_{\mathrm{ram}}, 2t_{\mathrm{ram}}]$ as the disk gas rearranges itself around the kicked BH. At $t = 2t_{\mathrm{ram}}$, the flare luminosity is assumed constant at L_{BHL}. Note that this represents hyper-Eddington accretion as parameterized, and it is an open question as to whether enough radiation could escape from a hyper-Eddington accretion rate BH to justify the choice of $\eta \sim 0.1$. Simulations of super-Eddington accretion that reach up to 1500× Eddington imply that $\eta \to 0.01$ in this context (Jiang et al. 2019). In order for enough radiation to

---

[16] Defined as $M_c \equiv \mu^{3/5}\, M_{\mathrm{BBH}}^{2/5}$, where $M_{\mathrm{BBH}} = m_1 + m_2$ is the total binary mass and $\mu = m_1 m_2/M_{\mathrm{BBH}}$.

escape to produce a bright flare against a quasar, jetted or collimated outflows are required. In this study we have no constraints on the very-high-energy emission (X-rays) that we should expect from such outflows. We recommend that future work on simulations of hyper-Eddington accretion establish whether there is an upper limit to accretion which can choke off jets. This will help establish luminosity upper limits on any flares that emerge from kicked mergers in AGN disks.

Bondi drag slows down the kicked BH from an initial kinetic energy of $1/2\, M_{\mathrm{BBH}}\, v_k^2$. The drag force is $\dot{M}_{\mathrm{BHL}}\, v_k$ and is equal to $M_{\mathrm{BBH}} v_k/\tau_{\mathrm{dec}}$, where τ_{dec} is the source-frame deceleration timescale. τ_{dec} is $\mathcal{O}(10^2)$ yr for plausible BBH merger and disk parameters (Graham et al. 2020), and kicks are likely not exactly aligned with the (relatively thin) AGN disk. Therefore, it is most likely that a modest inclination kick (i.e., $\theta \neq 0°$, where $\theta = 0°$ is the disk midplane) lets the kicked BH exit the disk on a timescale $t_{\mathrm{exit}} \ll \tau_{\mathrm{dec}}$. We define t_{exit} as the time for the remnant to reach the $\tau = 1$ optical depth surface from the merger point and we assume a Gaussian atmosphere with scale height H, i.e., $\rho = \rho_0 \exp(-z^2/2H^2)$, where z is the height above the midplane, ρ is the disk gas density at height z, and ρ_0 is the midplane density. The most rapid exit will be for a vertical kick directly out of the midplane, i.e., $\theta = 90°$; if the inclination angle is (as is likely) smaller, then the relevant velocity is just the vertical component of v_k, i.e., $v_k \sin(\theta)$. In the absence of strong constraints on θ, we adopt the fastest exit time (see Appendix C for a full derivation):

$$t_{\mathrm{exit}} = \frac{H\sqrt{2\ln(\tau_{mp})}}{v_k}, \tag{4}$$

where τ_{mp} is the merger point optical depth (we have used the fact that $\rho \propto \tau$). The peak flare luminosity will occur on a timescale $\approx t_{\mathrm{exit}}$, and we can use this calculation to constrain the range of kick velocities we are sensitive to based on the temporal search window. Our ZTF search, described in Section 4, assumes a flare peak <200 days post-merger.

Assuming a gas-pressure-supported disk (which should be true for $a > 10^2\, r_g$, where a is the BH orbit semimajor axis in units of r_g), we expect $H = c_s/\Omega$, where $\Omega = v_{\mathrm{orb}}/a$, and so

$$v_k > \left(\frac{c_s a}{v_{\mathrm{orb}}\, t_{\mathrm{exit}}}\right)\sqrt{2\ln(\tau_{mp})}. \tag{5}$$

In general, for remnants kicked at $\theta = 90°$, the minimum kick velocity could be as small as a few kilometers per second; however, for more realistic parameters, we expect $v_k > \mathcal{O}(100\ \mathrm{km\ s^{-1}})$.

On the other hand, larger v_k will reduce L_{BHL}, so we can also establish an upper limit on v_k by approximating the total energy of the flare, E_{tot}, as the luminosity times the flare duration, $t_{\mathrm{flare}} = t_{\mathrm{exit}} - t_{\mathrm{st}}$. Here t_{st} is the start time of the EM flare and not the merger time. Rearranging, we find

$$v_k < t_{\mathrm{flare}}^{1/3}\left(\frac{\eta}{0.1}\right)^{1/3}\left(\frac{M_{\mathrm{BBH}}}{100\, M_\odot}\right)^{2/3} \times \left(\frac{E_{\mathrm{tot}}}{2\times 10^{45}\ \mathrm{erg\ s^{-1}}}\right)^{-1/3}\left(\frac{\rho}{10^{-10}\ \mathrm{g\ cm^{-3}}}\right)^{1/3}. \tag{6}$$

Thus we can reject candidate counterparts in cases where these two limits are mutually exclusive.

2.4. Flare Form as a Function of Observer Orientation

For kicked BBH mergers in AGN disks, which are assumed to be face-on (i.e., within ∼45° to the observer), we expect around half of all kicks to be directed away from the observer, out the far side of the disk, and thus not detectable. A merger that occurs off the disk midplane and directed away from the observer should appear to become redder and diluted on an increasing diffusion timescale, t_{diff}, as the source is carried through a deeper scattering screen with an increasing optical depth. Once a merger remnant crosses the midplane we consider it undetectable due the optical depth of the accretion disk.

3. False Positives from Active Galactic Nuclei Disks

AGNs are intrinsically variable (e.g., Matthews & Sandage 1963). Most of the optical signatures that ZTF will detect from AGNs will therefore be false positives. Here we summarize the properties of different kinds of known optical variability that may occur in AGNs.

3.1. Disk Variability

Optical/UV variability in AGNs is typically $\mathcal{O}(10\%)$ over a few months (e.g., Krolik 1999; Kasliwal et al. 2015) and is generally associated with the behavior of the accretion disk. Larger-scale optical/UV variability can also occur in quasars (e.g., Graham et al. 2017; Ross et al. 2018; Stern et al. 2018). There are a very large number of possible causes of accretion disk variability including (but not limited to) disk inhomogeneities, different instabilities, fronts, and magnetic reconnection. Stern et al. (2018) conveniently parameterizes timescales for orbital (T_{orb}), thermal (T_{thermal}), heating/cooling front propagation (T_{cool}), and viscous transport (T_ν) as a function of disk radius (R), aspect ratio (height over radius, $h \equiv H/R$), viscous parameter (α), and SMBH mass (M_{SMBH}). Following Stern et al. (2018), we write

$$T_{\mathrm{orb}} \sim 100\ \mathrm{day}\ M_8\left(\frac{R}{750 r_g}\right)^{3/2}, \tag{7}$$

$$T_{\mathrm{thermal}} \sim 100\ \mathrm{day}\ M_8\ \alpha_{0.03}^{-1}\left(\frac{R}{65 r_g}\right)^{3/2}, \tag{8}$$

$$T_{\mathrm{front}} \sim 100\ \mathrm{day}\ M_8\ h_{0.1}^{-1}\ \alpha_{0.03}^{-1}\left(\frac{R}{15 r_g}\right)^{3/2}, \tag{9}$$

$$T_\nu \sim 100\ \mathrm{day}\ M_8\ h_{0.2}^{-2}\ \alpha_{0.03}^{-1}\left(\frac{R}{6 r_g}\right)^{3/2}, \tag{10}$$

where $M_8 = M_{\mathrm{SMBH}}/10^8\ M_\odot$, $h_{0.1} = h/0.1$, and $\alpha_{0.03}$ corresponds to a viscous parameter $\alpha = 0.03$. From Equations (7)–(10), we can associate variability on timescales of <100 days in a given AGN disk to, for example, an embedded or interacting orbiter at $<10^3 r_g M_8$, thermal variability in the disk at $\leqslant 65 r_g M_8$, front propagation at $\leqslant 15 r_g M_8$ of the SMBH, or viscous changes near the innermost, stable circular orbit (ISCO; $\approx 6r_g$), assuming a moderately puffed-up and viscous inner disk.

3.2. Supernovae and Kilonovae

Supernovae (SNe) are expected to occur in AGN accretion disks, but the expected rate is small, $>2 \times 10^{-7}$ AGN^{-1} yr^{-1} in the Wide-field Infrared Survey Explorer (WISE) sample (e.g., Assef et al. 2018). SNe will also occur in the host galaxy, and appear consistent with an origin in the galactic nucleus if their separation is less than the survey angular resolution. For unobscured SNe we expect rise times of $\mathcal{O}(20\text{–}50)$ days and a decay time or plateau of ∼100–200 days (Kasen & Bildsten 2010) and we also expect an evolution in color over time (Foley et al. 2011). We were able to rule out a SN origin for the flare reported in Graham et al. (2020) based on its much shorter flare timescale and lack of color evolution.

3.3. Tidal Disruption Events

Tidal disruption events (TDEs) also occur in AGNs (e.g., Chan et al. 2019; Ricci et al. 2020; McKernan et al. 2022b). Main-sequence star disruptions can occur around the central SMBH in a galaxy, but only for $M_{\rm SMBH} \leqslant 10^8 M_\odot$ (for a nonspinning SMBH; Rees 1988; Ryu et al. 2020). These are typically characterized by a fast rise (i.e., several weeks), $\sim t^{-5/3}$ decay signatures, and will be false positives in our search for EM counterparts to any O3 GW triggers.

In addition, TDEs can also occur around small BHs in AGN disks, as NS or white dwarf (WD) disruptions by stellar-origin BHs (e.g., Yang et al. 2022). Thus, for BH-NS mergers, where the BH is $\leqslant$7–10 $M_\odot$ (depending on BH spin), the expected EM counterpart corresponds to a NS tidal disruption emerging from inside an AGN disk. We expect the rate of such events at $z < 0.5$ to span $\sim[4, 113]\,(f_{\rm AGN}/0.1){\rm yr}^{-1}$, where $f_{\rm AGN}$ is the fraction of BBH mergers expected from the AGN channel (McKernan et al. 2020). The expected integrated total energy of such events is $\mathcal{O}(10^{52}$ erg) (Cannizzaro et al. 2020). BH-WD disruptions lead to underluminous Type Ia SNe with integrated energy 10^{49-51} erg (Rosswog et al. 2009).

3.4. Microlensing

Microlensing is uniform in color at rest-frame UV/optical bands and is expected for AGNs with an expected rate of $\mathcal{O}(10^{-4})$ per AGN (Lawrence et al. 2016). However, the expected characteristic timescale for microlensing is $\mathcal{O}$(yr) (Lawrence et al. 2016), which generally is far longer than the flaring timescales considered here. The Einstein ring radius of a gravitational lens is

$$r_E = \sqrt{\frac{4GM_*}{c^2}\frac{D_L D_{LS}}{D_S}}, \tag{11}$$

where $GM_*/c^2 = 1.5\ (M_*/M_\odot)$ km, D_S is the distance to the source, D_L is the distance to the lens, and D_{LS} is the distance from the lens to the source. If the lens is in the AGN host $D_L \sim D_S$. The duration of a microlensing event is $t_\mu = r_E/v_{\rm rel}$, where $v_{\rm rel}$ is the relative velocity of the lens. The resulting magnification is

$$\mathcal{M} = \frac{u^2 + 2}{u\sqrt{(u^2 + 4)}}, \tag{12}$$

where $u = b/r_E$ is the impact parameter ($b \sim 1$) of the lensing event in units of r_E. Assuming the stellar orbits at distance D_{LS} in the bulge of the host galaxy are evenly distributed in the hemisphere facing the observer, they cover an area $A = \pi D_{LS}^2$ and each star sweeps out a microlensing area of $A_{\mu,*} = \pi D_{LS} r_E$ in half its Keplerian orbital time.

A fiducial $M_\odot$ lens in the source galaxy with $D_{LS} \sim 1$ kpc and $v_{\rm rel} \sim 200$ km s^{-1} gives us a timescale $\sim 2 \times 10^6$ s (i.e., ∼3 weeks) with magnification depending on the choice of $u(r_E)$. Assuming a population of $\mathcal{O}(10^{10})$ stars in random orbits, geometric considerations produce a rate of $\mathcal{O}(10^{-5})$ events yr^{-1} AGN^{-1}. We also note that the flare in Graham et al. (2020) has a noticeable asymmetry that is not expected for microlensing and therefore disfavors this explanation for that flare.

4. Method and Data Sets

4.1. Zwicky Transient Facility

The ZTF is a state-of-the-art time-domain survey employing a 47 deg^2 field-of-view camera on the Palomar 48 inch Samuel Oschin Schmidt telescope (Bellm et al. 2019a; Graham et al. 2019). Since 2018 March, it has operated a number of observing programs, including a public survey covering the visible Northern sky every 2–3 nights in the g and r bands to ∼20.5 mag, as well as boutique partnership programs such as higher-cadence coverage of specific regions and use of an i-band filter (Bellm et al. 2019b). Each ZTF observation is processed by an image-differencing pipeline (Masci et al. 2019), which generates real-time alerts for all 5σ detections of point-source transient events (Patterson et al. 2019). Each observation is also processed by a point-spread function photometry pipeline, which produces a single-epoch catalog covering *all* identified sources. These are archived and used to create light curves for all detected objects in ZTF data releases.

ZTF Data Release 5 (DR5),[17] released on 2021 June 6, provides data from public surveys through to 2021 January 31 (as well as partnership data through to 2019 December 31) and therefore covers the entire LIGO/Virgo O3 data release. A forced-photometry data set has been produced based on all difference images available for DR5 (P. Mroz et al. 2022, in preparation) and source positions from the Pan-STARRS1 Data Release 1 catalog (Chambers et al. 2016).

4.2. Active Galactic Nuclei Matching

We employed the set of spectroscopically confirmed AGNs and high-probability AGN candidates from the Million Quasar Catalog, v7.3 (MQC; Flesch et al. 2019) as the primary AGN catalog. Given the sensitivity of LIGO/Virgo in O3, we excluded sources at redshift $z > 1.2$ from this analysis, as well as known blazars. The catalog was crossmatched using a 3″ matching radius against the set of forced-photometry ZTF light curves, giving a data set of 524,666 sources.

For each LIGO/Virgo event, we identified all ZTF sources from the light-curve data set that were located within the 90% credible volume of the event using the `crossmatch` method from the `ligo.skymap` Python package.[18] Table 1 gives the waveform employed for each LIGO/Virgo event but, in summary, the NRSur7dq4 waveform was used for O3a alerts

---

[17] https://www.ztf.caltech.edu/page/dr5

[18] https://lscsoft.docs.ligo.org/ligo.skymap/

Table 1
A Summary of the 83 LIGO/Virgo BBH and Lower-mass-gap Merger Alerts with Associated ZTF AGN Flares

LIGO/Virgo Alert ID	S/N	50% Area (deg^2)	90% Area (deg^2)	Distance (Gpc)	$f_{\rm cover}$	n_s
GW190408_181802*	14.7	23	145	$1.58^{+0.40}_{-0.59}$	1.000	332
GW190412*	18.9	3	12	$0.74^{+0.14}_{-0.17}$	1.000	34
GW190413_052954	8.6	331	1488	$4.10^{+2.41}_{-1.89}$	0.522	33453
GW190413_134308	10.0	54	589	$5.15^{+2.44}_{-2.34}$	0.180	11961
GW190421_213856	10.6	296	1267	$3.15^{+1.37}_{-1.42}$	0.000	8185
GW190424_180648	10.0	10754	28618	$2.55^{+1.56}_{-1.33}$	0.563	207206
GW190425	13.0	2417	9958	$0.16^{+0.07}_{-0.07}$	0.377	⋯
GW190426_152155*	10.1	289	1395	$0.38^{+0.19}_{-0.16}$	0.496	620
GW190503_185404	12.1	26	94	$1.52^{+0.71}_{-0.66}$	0.000	503
GW190512_180714*	12.3	46	227	$1.49^{+0.53}_{-0.21}$	0.233	367
GW190513_205428	12.3	112	404	$2.16^{+0.94}_{-0.80}$	0.949	1151
GW190514_065416	8.3	444	2930	$4.93^{+2.76}_{-2.41}$	0.652	64784
GW190517_055101	10.6	53	425	$2.11^{+1.79}_{-1.00}$	0.322	1803
GW190519_153144	12.0	170	838	$2.85^{+2.02}_{-1.14}$	0.542	7880
GW190521	15.0	180	822	$4.53^{+2.30}_{-2.13}$	0.384	20225
GW190521_074359	24.4	129	546	$1.28^{+0.38}_{-0.57}$	0.836	797
GW190527_092055	8.9	1115	3628	$3.10^{+4.85}_{-1.64}$	0.396	44299
GW190602_175927	12.1	180	687	$2.99^{+2.02}_{-1.26}$	0.262	6173
GW190620_030421	10.9	764	7570	$3.16^{+1.67}_{-1.43}$	0.654	103774
GW190630_185205	15.6	204	1258	$0.93^{+0.56}_{-0.40}$	0.465	5198
GW190701_203306	11.6	14	46	$2.14^{+0.79}_{-0.73}$	0.891	652
GW190706_222641	12.3	111	770	$5.07^{+2.57}_{-2.11}$	0.690	29332
GW190707_093326*	13.0	248	1328	$0.80^{+0.37}_{-0.38}$	0.484	1390
GW190708_232457*	13.1	2325	14420	$0.90^{+0.33}_{-0.40}$	0.552	43181
GW190719_215514	8.0	219	874	$4.61^{+2.84}_{-2.17}$	0.856	31178
GW190720_000836*	11.7	63	574	$0.81^{+0.71}_{-0.33}$	0.345	2131
GW190727_060333	12.3	140	880	$3.60^{+1.56}_{-1.51}$	0.386	4941
GW190728_064510*	13.6	62	574	$0.89^{+0.25}_{-0.37}$	0.257	736
GW190731_140936	8.5	728	3203	$3.97^{+2.56}_{-2.07}$	0.269	41160
GW190803_022701	9.0	364	1497	$3.69^{+2.04}_{-1.69}$	0.910	16456
GW190814*	22.2	3	15	$0.24^{+0.04}_{-0.05}$	0.712	1
GW190828_063405	16.0	101	539	$2.22^{+0.63}_{-0.95}$	0.461	1609
GW190828_065509*	11.1	145	655	$1.66^{+0.63}_{-0.61}$	0.254	1265
GW190909_114149*	8.5	1355	4613	$4.77^{+3.70}_{-2.66}$	0.570	79634
GW190910_112807	13.4	2786	10120	$1.57^{+1.07}_{-0.64}$	0.516	35170
GW190915_235702*	13.1	76	370	$1.70^{+0.71}_{-0.64}$	0.863	3192
GW190924_021846*	13.2	105	348	$0.57^{+0.22}_{-0.22}$	0.895	1088
GW190929_012149*	9.9	524	2031	$3.68^{+2.98}_{-1.68}$	0.522	13618
GW190930_133541*	10.0	496	1723	$0.78^{+0.37}_{-0.33}$	0.789	2406
GW190403_051529	8.0	1111	4250	$8.28^{+6.72}_{-4.29}$	0.499	132911
GW190426_190642	9.6	2041	8031	$4.58^{+3.40}_{-2.28}$	0.496	145368
GW190725_174728	8.8	325	2436	$1.03^{+0.52}_{-0.43}$	0.684	5403
GW190805_211137	8.3	660	3089	$6.13^{+3.72}_{-3.08}$	0.543	77630
GW190916_200658	7.9	986	3573	$4.94^{+3.71}_{-2.38}$	0.539	116115
GW190917_114630	9.5	451	1801	$0.72^{+0.30}_{-0.31}$	0.419	2876
GW190925_232845	9.9	112	955	$0.93^{0.46}_{-0.35}$	0.195	2813
GW190926_050336	7.8	625	2212	$3.28^{3.40}_{-1.73}$	0.461	32379
GW191103_012549	$8.9^{+0.3}_{-0.5}$	667	2171	$0.99^{+0.50}_{-0.47}$	0.616	7377
GW191105_143521	$9.97^{+0.3}_{-0.5}$	100	641	$1.15^{+0.43}_{-0.48}$	0.375	2622
GW191109_010717	$17.3^{+0.5}_{-0.5}$	418	1649	$1.29^{+1.13}_{-0.65}$	0.364	5826
GW191113_071753	$7.9^{+0.5}_{-1.1}$	578	2483	$1.37^{+1.15}_{-0.62}$	0.444	13157
GW191126_115259	$8.3^{+0.2}_{-0.5}$	354	1378	$1.62^{+0.74}_{-0.74}$	0.473	7214
GW191127_050227	$9.2^{+0.7}_{-0.6}$	137	983	$3.4^{+3.1}_{-1.9}$	0.460	15259
GW191129_134029	$13.1^{+0.2}_{-0.3}$	208	856	$0.79^{+0.26}_{-0.33}$	0.409	2387
GW191204_110529	$8.8^{+0.4}_{-0.6}$	922	3675	$1.8^{+1.7}_{-1.1}$	0.495	39638
GW191204_171526	$17.5^{+0.2}_{-0.2}$	79	256	$0.65^{+0.19}_{-0.25}$	0.286	161

Table 1
(Continued)

LIGO/Virgo Alert ID	S/N	50% Area (deg^2)	90% Area (deg^2)	Distance (Gpc)	f_{cover}	n_s
GW191215_223052	$11.2^{+0.3}_{-0.4}$	139	586	$1.93^{+0.89}_{-0.86}$	0.297	2158
GW191216_213338	$18.6^{+0.2}_{-0.2}$	61	206	$0.34^{+0.12}_{-0.13}$	0.592	59
GW191219_163120	$9.1^{+0.5}_{-0.8}$			$0.55^{+0.25}_{-0.16}$	⋯	7120
GW191222_033537	$12.5^{+0.2}_{-0.3}$	503	2168	$3.0^{+1.7}_{-1.7}$	0.451	24418
GW191230_180458	$10.4^{+0.3}_{-0.4}$	302	1086	$4.3^{+2.1}_{-1.9}$	0.187	17057
GW200105_162426	$13.7^{+0.2}_{-0.4}$			$0.27^{+0.12}_{-0.11}$	⋯	596309
GW200112_155838	$19.8^{+0.1}_{-0.2}$	550	3200	$1.25^{+0.43}_{-0.46}$	0.486	8558
GW200115_042309	$11.3^{+0.3}_{-0.5}$	56	388	$0.29^{+0.15}_{-0.10}$	⋯	174
GW200128_022011	$10.6^{+0.3}_{-0.4}$	714	2415	$3.4^{+2.1}_{-1.8}$	0.435	34995
GW200129_065458	$26.8^{+0.2}_{-0.2}$	6	31	$0.90^{+0.29}_{-0.38}$	0.968	41
GW200202_154313	$10.8^{+0.2}_{-0.4}$	44	150	$0.41^{+0.15}_{-0.16}$	0.886	295
GW200208_130117	$10.8^{+0.3}_{-0.4}$	9	30	$2.23^{+1.00}_{-0.85}$	0.000	80
GW200208_222617	$7.4^{+1.4}_{-1.2}$	385	2040	$4.1^{+4.4}_{-1.9}$	0.505	36630
GW200209_085452	$9.6^{+0.4}_{-0.5}$	217	877	$3.4^{+1.9}_{-1.8}$	0.643	15681
GW200210_092255	$8.4^{+0.5}_{-0.7}$	290	1388	$0.94^{+0.43}_{-0.34}$	0.473	5193
GW200216_220804	$8.1^{+0.4}_{-0.5}$	727	2924	$3.8^{+3.0}_{-2.0}$	0.861	64883
GW200219_094415	$10.7^{+0.3}_{-0.5}$	88	781	$3.4^{+1.7}_{-1.5}$	0.392	14993
GW200220_061928	$7.2^{+0.4}_{-0.7}$	1065	4477	$6.0^{+4.8}_{-3.1}$	0.272	75639
GW200220_124850	$8.5^{+0.3}_{-0.2}$	855	3129	$4.0^{+2.8}_{-2.2}$	0.500	65013
GW200224_222234	$20.0^{+0.2}_{-0.2}$	11	42	$1.71^{+0.49}_{-0.64}$	0.000	178
GW200225_060421	$12.5^{+0.3}_{-0.4}$	150	498	$1.15^{+0.51}_{-0.53}$	0.511	1125
GW200302_015811	$10.8^{+0.3}_{-0.4}$	1604	6016	$1.48^{+1.02}_{-0.70}$	0.377	18761
GW200306_093714	$7.8^{+0.4}_{-0.6}$	965	3907	$2.1^{+1.7}_{-1.1}$	0.524	35435
GW200308_173609	$7.1^{+0.5}_{-0.5}$	3671	25292	$5.4^{+2.7}_{-2.6}$	0.400	596309
GW200311_115853	$17.8^{+0.2}_{-0.2}$	10	35	$1.17^{+0.28}_{-0.40}$	0.821	81
GW200316_215756	$10.3^{+0.4}_{-0.7}$	12	187	$1.12^{+0.47}_{-0.44}$	0.726	276
GW200322_091133	$6.0^{+1.7}_{-1.2}$	6250	28703	$3.6^{+7.0}_{-2.0}$	0.475	596309

Notes. f_{cov} is the fraction of the LIGO/Virgo 90% confidence area covered by ZTF with at least 15 observations in the 200 days following the LIGO alert. n_s is the number of Million Quasar Catalog (v7.3) sources within the LIGO/Virgo 90% confidence volume. The divisions indicate events from O3a (GWTC 2.0), O3a (GWTC 2.1), and O3b. The NRSur7dq4 waveform is used for O3a detections, except for those marked with an asterisk, where the SEOBNRv4PHM waveform is used. The IMRPhenomXPHM waveform is used for all O3b detections. S/N are the published values and are derived as follows: in O3a, from GstLAL (including Virgo) where possible and PYCBC BBH otherwise; in O3b, from IMRPhenomXPHM; and in GWTC 2.1, from PYCBC BBH where possible and GstLAL otherwise. LIGO/Virgo considers signals with S/N >8 to be confidently real.

Table 2
A Summary of the Parameter Ranges Used in the Flare Modeling

Parameter	Description	Prior Range	Search Range
t_g	Gaussian Rise Time (days)	$0 \leqslant t_g < 4000$	$5 \leqslant t_g \leqslant 100$
t_e	Exponential decay time (days)	$0 \leqslant t_e < 4000$	$10 \leqslant t_e \leqslant 200$
t_0	Time of flare peak (MJD)	$58{,}574 \leqslant t_0 < 59{,}000$	$0 \leqslant t_0 - \mathrm{MJD}_{\mathrm{LIGO}} \leqslant 200$
A	Flare amplitude (ADU)	$0 < A \leqslant 1000$	$\log_{10}(A/r_0 + 1) > 0.04$
r_0	Baseline flux (ADU)	$r_0 < 100$	⋯
f	Additional noise term	$-8 \leqslant \log_{10} f \leqslant 2$	⋯

when available and the SEOBNRv4PHM waveform as an alternative. The IMRPhenomXPHM waveform was used for all O3b events.

Each ZTF source may have a g- or r-band light curve, and we fitted each (in flux space) with a Bayesian block representation (Scargle et al. 2013). This provides an optimal segmentation of the data in terms of a set of discontinuous piecewise constant components and makes significant local variations more easily detectable. We identified candidate flares using the hill-climbing procedure proposed by Meyer et al. (2019): peaks are identified as blocks that are higher than both previous and subsequent blocks and then extended in both directions as long as succeeding blocks are lower. The data contained within such a set of blocks is then characterized by a Gaussian rise-exponential decay form:

$$
\begin{aligned}
y(t) &= r_0 + A \exp\left(-\frac{(t - t_0)^2}{2t_g^2}\right), \; t \leqslant t_0 \\
&= r_0 + A \exp\left(-\frac{(t - t_0)}{t_e}\right), \; t > t_0, \qquad (13)
\end{aligned}
$$

with rise and decay times t_g and t_e, respectively. The flare peaks at time t_0 with an amplitude A above a background flux of r_0. We added an additional random noise term, f, in quadrature to the model covariance to account for potential systematic errors in the model. We assumed uniform priors for each parameter (see Table 2) and maximized the posterior likelihood for the flare model.

Sources were rejected when the time between the LIGO/Virgo event and a flare peak was greater than 200 days or less than the Gaussian rise time from the model fit. Random noise in a ZTF light curve may be misidentified as a flare, particularly in low-S/N data, so a LIGO/Virgo event must have a corresponding flare in both filters (coincident in time) and all single-filter flares were rejected. Similiarly, the amplitude of the flare must be at least 10% greater than the background flux level (r_0).

4.3. Astrophysical Implications of Our Search Assumptions

The search criteria in Table 2 correspond to assumptions about the mergers. In particular, given $M_{\rm SMBH}$, a radial location in the disk of the merger a, and a local disk scale height (or, equivalently, local sound speed c_s), the search timescales correspond to a minimum v_k. For each source, we can measure $M_{\rm SMBH}$, but can only make plausible assumptions for a range of a and c_s. The total energy of the flare (which is related to the amplitude) sets a maximum v_k, given $M_{\rm BBH}$, η, and the local disk density ρ. As with the minimum v_k, we can use the GW measurements to find $M_{\rm BBH}$, but η and ρ are free parameters. Note that for some combinations of candidate flare and GW source, there may be no plausible sets of a, c_s, η, and ρ which yield consistent min/max v_k, given the observed $M_{\rm SMBH}$ and $M_{\rm BBH}$. For most pairs, we can construct a parameter set that yields an internally consistent scenario. However, some of those may be implausible astrophysically.

In particular, given our search window of only 200 days post-merger, our search tends to require small a and c_s, such that the remnant can escape the optically thick part of the disk sufficiently quickly to be observed yet still have a sufficiently small v_k that it can produce a detectable flare via BHL accretion. Similarly, to generate sufficiently bright flares that they will be detectable above the background AGN luminosity, we are pushed to large η and ρ; we are also more likely to detect larger $M_{\rm BBH}$ and find them more easily around smaller $M_{\rm SMBH}$. Large η is expected for thin-disk accretion onto a highly spinning BH (which would be expected for a BBH merger remnant); however, simulations imply that η decreases for highly super-Eddington accretion, as required if these flares are associated with the GW events. Large ρ may indeed be expected in the innermost regions of an AGN disk, where we are more sensitive to detecting mergers (due to the shorter length scales and shorter timescales); however, there are several cases where only a maximal density can produce a consistent pair of optical flare and GW signal.

We find the lowest minimum v_k assuming $a = 100\,r_g$ and $c_s = 10\,{\rm km\,s^{-1}}$. For $M_{\rm SMBH} = 10^8\,M_\odot$ and a 200 day search window, we are sensitive to $v_k > 1\,{\rm km\,s^{-1}}$ (with lower velocities probed around lower $M_{\rm SMBH}$). We probe the largest maximum v_k assuming $\eta = 0.3$ and $\rho = 10^{-9}\,{\rm g\,cm^{-3}}$, and larger flare energies imply smaller v_k at fixed $M_{\rm SMBH}$ and $M_{\rm BBH}$. We note that $a = 1000\,r_g$ and $c_s = 50\,{\rm km\,s^{-1}}$ are more consistent with our chosen density (Sirko & Goodman 2003), and our choice of η may be larger than warranted.

4.4. Discriminating False-positive Signals

In Section 3, we considered other potential sources of flaring activity from an AGN that might be misidentified as an EM counterpart to a GW event. Normal TDEs, SNe, and AGN flares can be distinguished on the basis of rise and fade timescales, $g - r$ color, and the rate of color evolution (van Velzen et al. 2021). The total observed energy of the event can also be added to this list. Since instantaneous colors are not available for ZTF light curves and data sometimes exist for only one filter for a few nights, we model the multiband light curve for each source as a two-dimensional surface with time and wavelength (passband) as the two axes. A suitable interpolation scheme over the irregularly sampled surface then allows fluxes to be predicted in each passband at each observed time.

Thin-plate splines have been used for similar purposes (e.g., fitting stellar spectral energy distributions as a function of temperature and extinction; Bailer-Jones 2011) but a more probabilistic approach is provided by a Gaussian process (GP), as demonstrated by Boone (2019) and (Villar et al. 2020, hereafter V20). Note that the thin-plate spline and GP interpolation schemes can be shown to be equivalent with a particular choice of kernel function. Following V20, we use a composite kernel function for the covariance between observations at times t_i and t_j in filters f_i and f_j, respectively:

$$K(t_i, t_j, f_i, f_j) = \sigma^2 l_t \exp\left[-\frac{|t_i - t_j|}{l_t}\right] \exp\left[-\frac{d(f_i, f_j)^2}{2l_f^2}\right],$$

where σ^2 is the variance, and l_t and l_f are the respective length scales along each axis. This describes crossband information via a squared exponential with a Wasserstein distance metric, $d(f_i, f_j)$, between each filter's normalized transmission curve (see Figure 1). V20 also describe the temporal variability with a squared exponential function; however, this gave inadequate fits to ZTF data, particularly during the seasonal observation gaps of a few months. Since AGN variability can be described as a damped random walk (DRW) process, an Ornstein–Uhlenbeck (Matérn 1/2) kernel was found to be a better choice. Note, however, that perfect correlation between bands is assumed in this model with no wavelength dependency in the variability amplitude or characteristic timescale.

Figure 2 shows projections of the discriminating space and the respective distributions of ZTF optically selected TDEs (Hammerstein et al. 2021; van Velzen et al. 2021), SNe from the ZTF Bright Transient Survey (Fremling et al. 2020; Perley et al. 2020), and a general sample of flares in AGNs identified by applying the flare model to AGN light curves without a LIGO/Virgo timing constraint. A balanced random forest classifier, BalRF04, was trained on this data, achieving ∼90% accuracy, and applied to candidate flares to determine whether they might be a TDE or a SNe. Note that these could be associated with an off-nucleus event rather than a TDE or SN embedded in the accretion disk of an AGN. If a source is not clearly classified then we do not reject it. Finally, we test whether the time series in both bands is better fit (ΔBIC > 10) with a microlensing profile than the flare (Equation (13)).

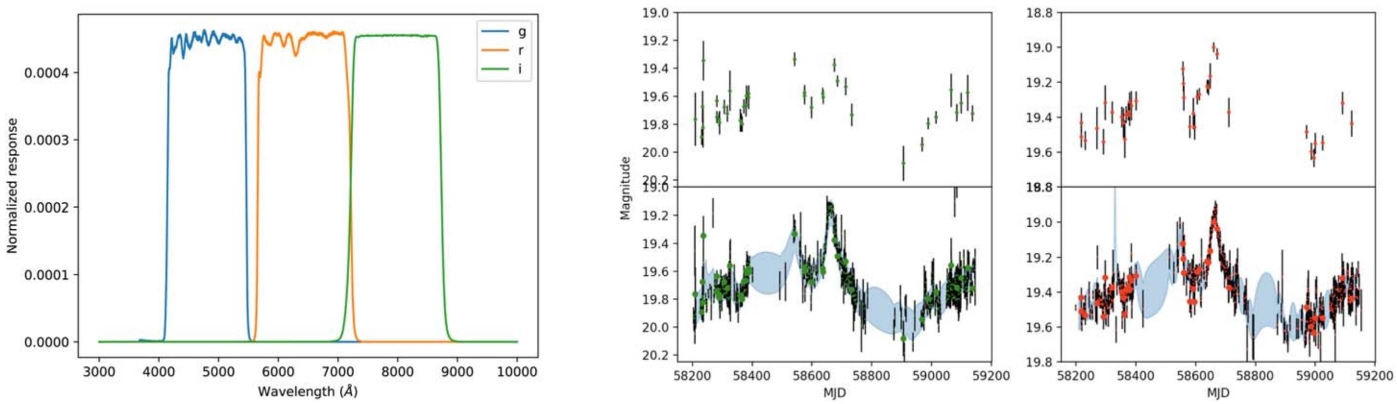

Figure 1. Left: the *g*-, *r*-, and *i*-band normalized filter response curves for ZTF (Bellm et al. 2019a). Right: the upper plots show *g*- and *r*-band subsampled light curves of an AGN; the lower plots show the interpolated light curves using a two-dimensional GP fit (blue shaded region). The larger points indicate the points used in the regression and the smaller points are from the full light curve.

4.5. Discriminating from Active Galactic Nuclei Activity

Variability is an inherent property of AGNs, often exhibited on short timescales, and the most likely origin of a false-positive signal, i.e., an identified flare is intrinsic behavior related to general gas physics and accretion disk activity rather than associated with a particular event in the accretion disk. Change-point detection is an aspect of statistical time series analysis that tries to identify when the probability distribution of a stochastic process changes, for example with the addition of a secondary temporary signal from an accompanying process.

We employed a GP-based algorithm (see Appendix A for details) to determine whether the flare seen in a given AGN is consistent with the type and level of intrinsic variability exhibited by the AGN as described with a DRW model (see Moreno et al. 2019, and references therein). We used a sliding window with a 50 day width to detect variations from the DRW and rejected all flares where the probability that the flare is just intrinsic AGN variability is greater than 0.5% in either filter. We note, however, that an incorrect choice of model for the intrinsic AGN behavior could be a source of false positives, i.e., intrinsic activity mistaken for significant flares. While a discussion of the correct model for AGN variability is outside the scope of this paper, the DRW model is seen as an adequate statistical description of AGN behavior and the fidelity of the data used here is such that higher-order autoregressive models are indistinguishable from it. A more promising avenue would be to use a multisurvey data set with extended temporal baselines to better constrain the models, but again there is work to be done on multivariate AGN models. The change-point algorithm, though, will work with any model provided.

5. Results

From the 83 LIGO/Virgo BBH and lower-mass-gap merger alerts, we find seven AGN flares that are statistically associated with one or more of nine LIGO/Virgo events (see Table 3 for details and Figure 3 for their light curves). The chance coincidence of this can be computed by considering that there are 20 AGN flares in the full ZTF data set that meet our selection criteria, i.e., they have the correct morphology, correct energetics, acceptable color evolution, and are inconsistent with being a TDE, SNe, or regular AGN activity. The total comoving volume probed by ZTF to $z = 1.2$ is 1.643×10^{11} Mpc^3. Assuming a mean flare lifetime of 100 days and 1000 days of ZTF data, the effective source density of flares in the survey volume at any given time is 1.321×10^{-11} Mpc^{-3}. The expectation value for the number of flares from O3 is the product of the effective density and the combined localization volume covered by ZTF (2.144×10^{11} Mpc^3), implying 2.83 expected random events. Using a Poisson distribution, the chance coincidence of finding nine matches is $p = 1.90 \times 10^{-3}$.

One of the sources, J053408.41+085450.6, has a quoted photometric redshift of $z = 0.5$ from MQC and a quoted spectroscopic redshift of $z = 1.62$ from Gaia Data Release 3 (Gaia Collaboration et al. 2022), though the latter has flags indicative of an ambiguous redshift identification (i.e., `FLAGS_QSOC` = 13). The higher redshift would place this source outside the LIGO/Virgo detection limit, though we still consider this candidate in our analysis given the low score of that redshift solution. We obtained a 600 s spectrum of the source from the Keck Observatory (see Appendix B), finding a relatively featureless spectrum and an inconclusive redshift determination. The source is also associated with a 32.1 mJy radio source at 21 cm from the NRAO Very Large Array Sky Survey (Condon et al. 1998), which suggests that this might actually be a blazar, although previous optical activity of the object is fairly quiet and it does not appear in any blazar catalogs.

As a comparison, we consider the original 45 BAYESTAR (Singer & Price 2016) or LALInference (Veitch et al. 2015) LIGO/Virgo alerts for BBH and lower-mass-gap events. This is a subset of the 83 such alerts eventually published from O3, and generally correspond to the higher S/N events. We use the original released skymaps rather than the final published skymaps, and find that three flares are associated with three LIGO/Virgo events. Of these, only J124942.3+344928.9 and its association with GW190521 is common to both sets of input. The other two alerts, GW190519_153544 and GW200112_155838, fall outside the 90% LIGO/Virgo confidence volumes of the final published skymaps, which reduced in volume by ∼15%–20% during further processing of the GW signal. This highlights the importance of providing accurate and revised volume localizations in O4 (preferably on the timescale of days after an event) to improve our ability to identify an EM counterpart to a BBH merger in an AGN disk.

The implied kick velocity constraint of each event is also broadly consistent with what LIGO/Virgo reported and so we cannot rule out any event on this basis here.

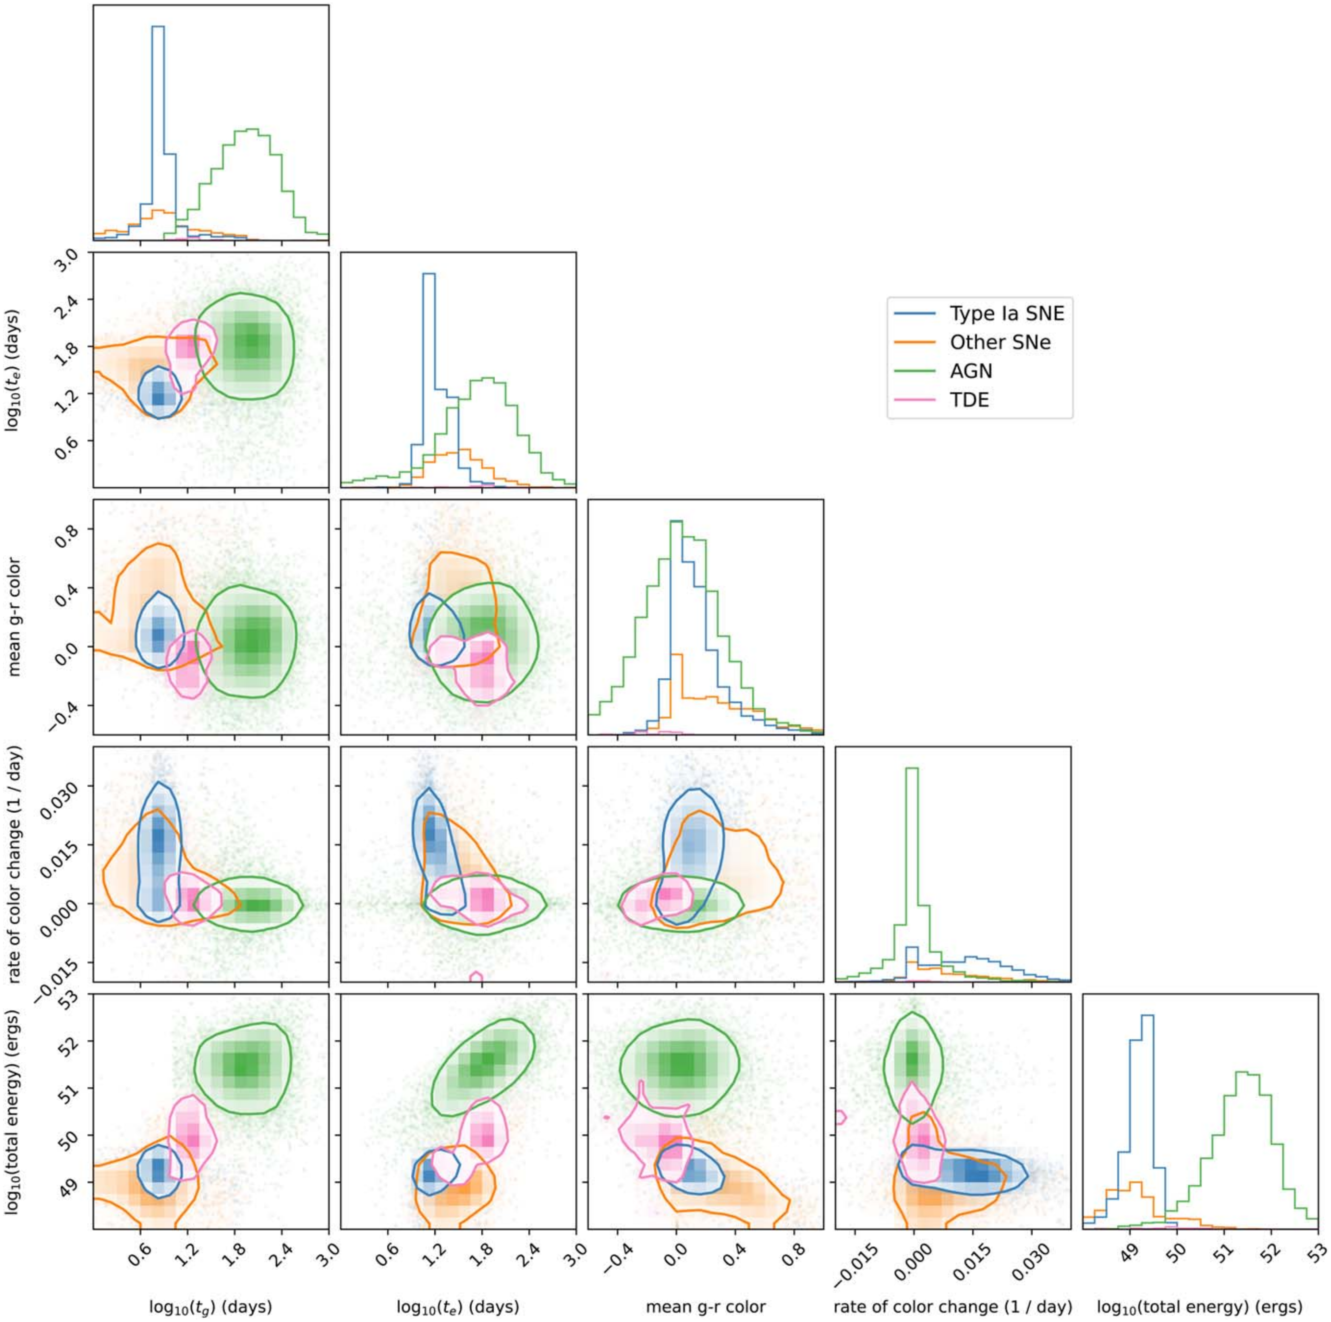

Figure 2. The distributions of rise time (t_g), decay time (t_e), mean $g - r$ color, rate of $g - r$ color change, and total flare energy for type Ia SNe (blue) and other SNe (orange) from the ZTF Bright Transient Survey (Fremling et al. 2020), optically selected TDEs in ZTF (pink; Hammerstein et al. 2021; van Velzen et al. 2021), and a general sample of AGN flares (green; see text for details). The timescales are measured in the rest frame of the source. The contours show 1σ levels.

To further assess the probability of finding a positive result, we have rerun our selection procedure with 1000 simulations of the full O3 LIGO/Virgo run. We apply a random rotation to the R.A. of each LIGO/Virgo error volume, which allows for random spatial localization but also maintains latitude-averaged dependencies. We also assign a random time to each LIGO/Virgo event drawn from the respective ranges: MJD = 58574–58756 for O3a events and MJD = 58788–58930 for O3b events. We note, however, that the localization volume of some events is large enough that a substantial range of rotations will still place a particular association found in the real data within the volume. Similarly, the post-event time window of 200 days used to identify associated flares is of the order of the durations of O3a and O3b and so a temporal match is also still possible with any time shift. This makes particular associations more likely in the simulations, and so we should consider these results as an upper bound on the results from the real data. Future studies could use the observing scenario skymaps simulated for the LIGO O4 and O5 runs (Petrov et al. 2022) to create a larger background set.

From the simulations, we can determine the expected rate of an association between an AGN flare and a LIGO/Virgo detection:

$$\langle n \rangle = n_s \cdot p_{\rm flare} \cdot f_{\rm cover} \cdot f_{\rm prof} \cdot f_{\rm FP} \cdot f_{\rm WISE} \cdot f_K \cdot f_{\rm CP},$$

where n_s is the number of sources in the LIGO/Virgo error volume, $p_{\rm flare}$ is the probability that an AGN has a flare within a

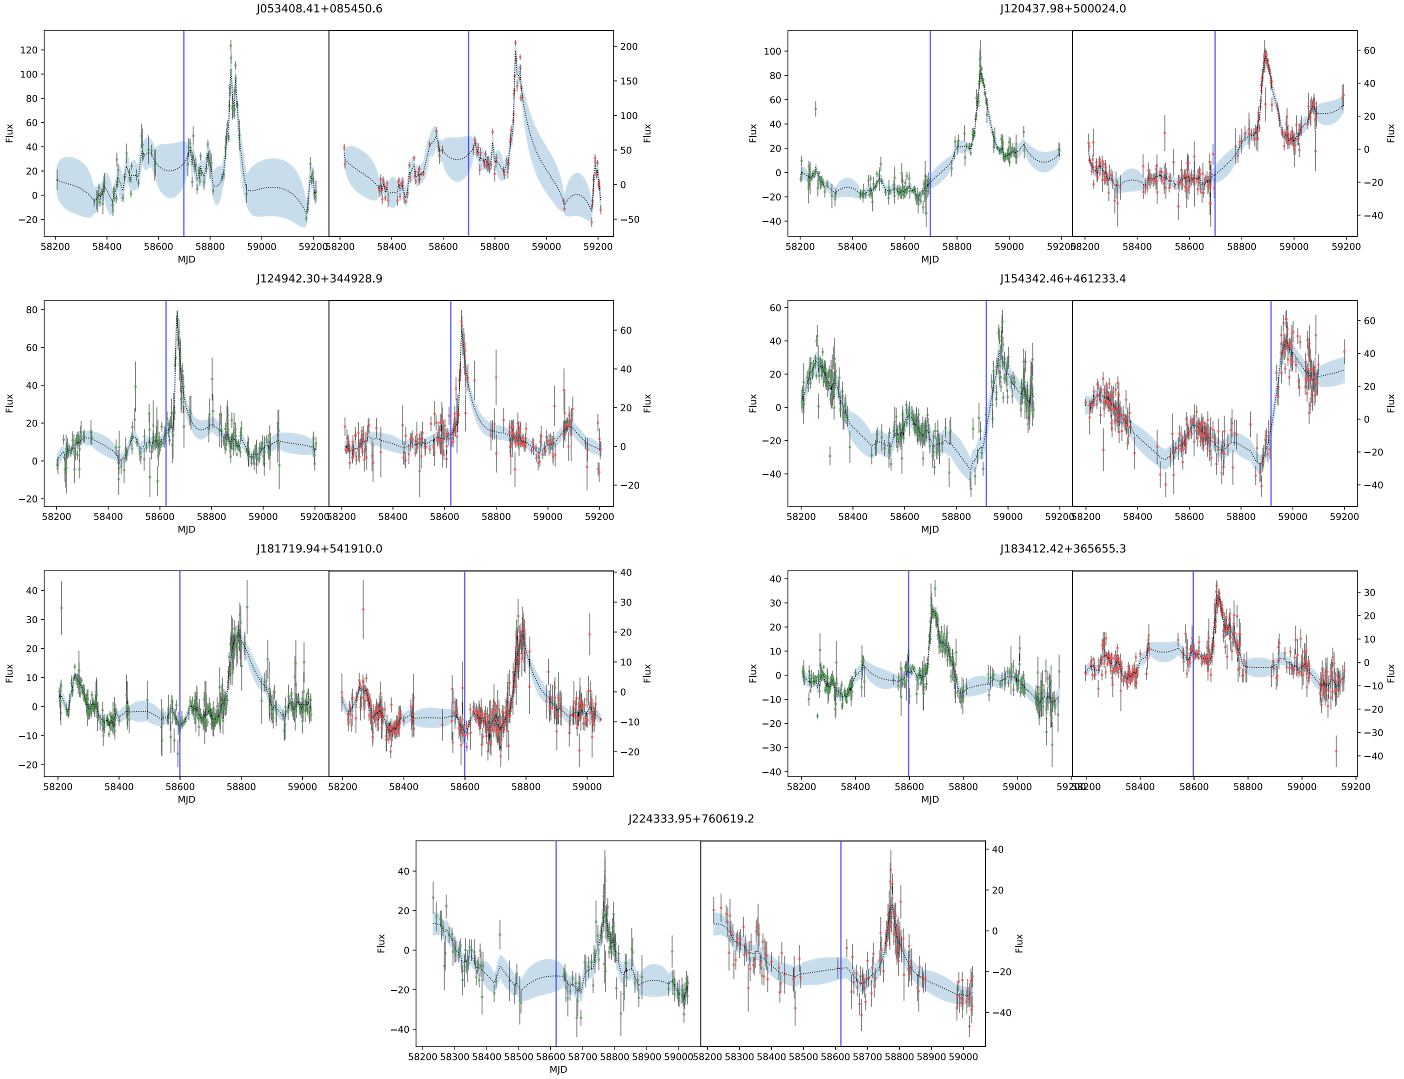

Figure 3. The ZTF g- and r-band light curves for the flares associated with LIGO/Virgo events. A GP fit to the data combining a mean flare function and a damped random walk kernel is shown with its predicted uncertainties (blue shaded region).

Table 3
A Summary of the Seven ZTF AGN Flares that Match LIGO/Virgo Events

LIGO/Virgo alert ID	Name	Redshift	$\log_{10}(M_{\rm BH})$ $(M_\odot)$	Conf. Limit	t_g (days)	t_e (days)	$\log_{10}(E)$ (erg s)	min v_k (km s^{-1})	max v_k (km s^{-1})
GW190403_051519	J124942.30+344928.9	0.438	8.6	0.606	11.7	45.3	51.6	5	990
GW190403_051519	J183412.42+365655.3	0.419	9.1	0.864	12.7	41.0	50.5	15	2300
GW190424_180648	J181719.94+541910.0	0.234	8.0	0.099	12.9	35.6	51.4	1	800
GW190514_065416	J124942.30+344928.9	0.438	8.6	0.754	11.7	45.3	51.6	5	740
GW190514_065416	J224333.95+760619.2	0.353	8.8	0.664	11.3	18.3	50.5	6	1400
GW190521	J124942.30+344928.9	0.438	8.6	0.596	11.7	45.3	51.6	5	1300
GW190731_140936	J053408.41+085450.6	0.5	(8.0)	0.754	7.6	27.4	51.0	1	990
GW190803_022701	J053408.41+085450.6	0.5	(8.0)	0.488	7.6	27.4	51.0	1	920
GW190803_022701	J120437.98+500024.0	0.389	8.0*	0.304	20.2	47.4	51.5	2	780
GW190909_114149	J120437.98+500024.0	0.389	8.0*	0.057	20.2	47.4	51.5	2	1100
GW200216_220804	J154342.46+461233.4	0.599	9.3	0.699	12.0	123.4	51.4	24	1300
GW200220_124850	J154342.46+461233.4	0.599	9.3	0.113	12.0	123.4	51.4	24	1100

Notes. The name used for each AGN is its position in sexagesimal format. BH masses are taken from PyQSOFit (Guo et al. 2018) fits to available spectra for the sources (see Appendix B) using the virial mass relationship of Ho & Kim (2015) for Hβ and Shen & Liu (2012) for Hα (asterisked values). The rise and decay timescales are measured in the rest frame of the AGN. min v_k and max v_k are the minimum and maximum kick velocities for the merged BH. "Conf. limit" is the percentile confidence contour within the 90% credible volume of the event at which the AGN is located. The redshift for J053408.41+085450.6 is a photometric redshift and is taken from the literature; its spectrum does not have any broad emission features to evaluate a virial mass so a fiducial mass of $\log_{10}(M_{BH}) = 8$ is used.

Table 4
A Summary of the Merger Parameters for the Nine LIGO/Virgo Events with Associated Flares

LIGO/Virgo Alert ID	Data Set	Total Mass ($M_\odot$)	M_1 ($M_\odot$)	M_2 ($M_\odot$)	Chirp Mass ($M_\odot$)	$\chi_{\rm eff}$
GW190403_051519	GWTC 2.1	111	85.0	20.0	34.0	0.68
GW190424_180648	GWTC 2	73	40.5	31.8	31.0	0.13
GW190514_065416	GWTC 2	68	39.0	28.4	28.5	-0.19
GW190521	GWTC 2	164	95.3	69.0	69.2	0.03
GW190731_140936	GWTC 2	70.1	41.5	28.8	29.5	0.06
GW190803_022701	GWTC 2	65	37.3	27.3	27.3	−0.03
GW190909_114149	GWTC 2	75	45.8	28.3	30.9	−0.06
GW200216_220804	GWTC 3	81	51	30	32.9	0.10
GW200220_124850	GWTC 3	67	28	39	28.2	0.10

200 day window, $f_{\rm cover}$ is the fraction that have at least 20 ZTF detections within 200 days of a LIGO/Virgo event, $f_{\rm prof}$ is the fraction that have the right flare profile, $f_{\rm FP}$ is the fraction where the flare is not consistent with a SNe or TDE, $f_{\rm WISE}$ is the fraction that have WISE colors consistent with an AGN, f_K is the fraction that are known flaring sources, and $f_{\rm CP}$ is the fraction where the flare is not consistent with general activity in the source. Calculated values for these parameters are given in Table 5. The appropriate values for n_s and $f_{\rm cover}$ for each LIGO/Virgo event are given in Table 1.

The expected number of associations from the simulated O3 LIGO/Virgo runs with 83 detections is 3.15. This is consistent with the rates-based estimate of 2.83.

Another way of approaching the reality of the associations is to consider the distribution of merger parameters, e.g., chirp mass, mass ratio, etc., for the flare events against those of nonflare events. BBH mergers occurring in an AGN disk are expected to favor certain regions of merger parameter space, for example higher masses or more extreme mass ratio (q) values, than mergers originating in other formation channels. We note that in this study we are dealing with small-number statistics and this line of argument will have more power when the overall population statistics of BBH mergers are better determined.

The relative distributions of total mass, mass ratio ($q = m_2/m_1$), and $\chi_{\rm eff}$ are shown in Figure 4 for both the flare events identified here and the total sample of merger events in O3. It can be seen that there is mass preference shown for the flare events, with no flare event having a total mass less than 65 $M_\odot$, and an Anderson–Darling test between the two distributions gives a significance level of 0.005, indicating that a low likelihood that the populations come from the same distribution. The values of q and $\chi_{\rm eff}$ are, however, not significantly different. If we consider 100,000 random subsamples of nine events from the LIGO/Virgo O3 full sample then only 282 show similar characteristics with a mass preference but no difference in q or $\chi_{\rm eff}$.

The association of the most likely merger event to originate in the AGN channel, namely GW190521, has already been discussed elsewhere (Graham et al. 2020) but we can also consider how many of the other events that would be considered as likely to originate in an AGN have an association or not. In particular, there are eight LIGO/Virgo events in O3 with total masses greater than 100 $M_\odot$. There is no inherent theoretical basis for this value other than as a fiducial value for hierarchical merger events favored in the AGN channel. Sufficient spatial coverage of the 90% confidence area of an event and sufficient sampling of the region in the 200 days following the event is required to detect any potential flare. Two events have slightly too low a coverage (∼25%) to expect a detection; the others all have about 50% coverage; so, if all events have a detectable flare then we should expect to get about half of the events. However, there is also an orientation effect to take into account since we expect half the events to be oriented away from us and so not observable. It is thus reasonable to expect two candidate flares associated with high-mass mergers in O3 if all high-mass mergers originate in the AGN channel. Indeed, we find that two of the events we have identified with associated AGN flares, GW190521 and GW190403_051519, have total masses greater than 100 $M_\odot$.

Table 5
Parameter Values for Determining the Mean Number of Associations for an Event

Parameter	Value	Uncertainty
$p_{\rm flare}$	0.051	…
$f_{\rm prof}$	0.006	0.017
$f_{\rm FP}$	0.288	0.277
$f_{\rm WISE}$	0.973	0.111
f_K	0.962	0.154
$f_{\rm CP}$	0.026	0.190
Total	2.12×10^{-6}	7.46×10^{-6}

Note. The value and uncertainty for each parameter are the mean fraction of events that pass the associated filter and its standard deviation from 1000 simulations of the full O3 LIGO/Virgo run.

6. Discussion

We can make simple inferences about the number of EM counterparts we should expect to observe in a study such as this. First, we should expect that EM counterparts are only in principle detectable for the fraction, $f_{\rm AGN}$, of BBH mergers that originate in the AGN channel. Second, only those BBH mergers that occur in the fraction of AGN that are Type I AGN could possibly have detectable optical signatures (i.e., unobscured AGNs showing broad emission lines, corresponding to those sources that are approximately face-on to the observer according to the standard AGN unification paradigm; $f_{\rm type-1} \sim 0.5$). Third, only the fraction of BBH mergers that are kicked out of the AGN on the side facing us ($f_{\rm side} \sim 0.5$) could yield a signature that would not be washed out by the optical depth of the accretion disk. Fourth, we require the AGN to lie in that fraction of the sky ($f_{\rm sky} \sim 0.5$) covered by ZTF. And fifth, we require the EM counterpart to be detectable against

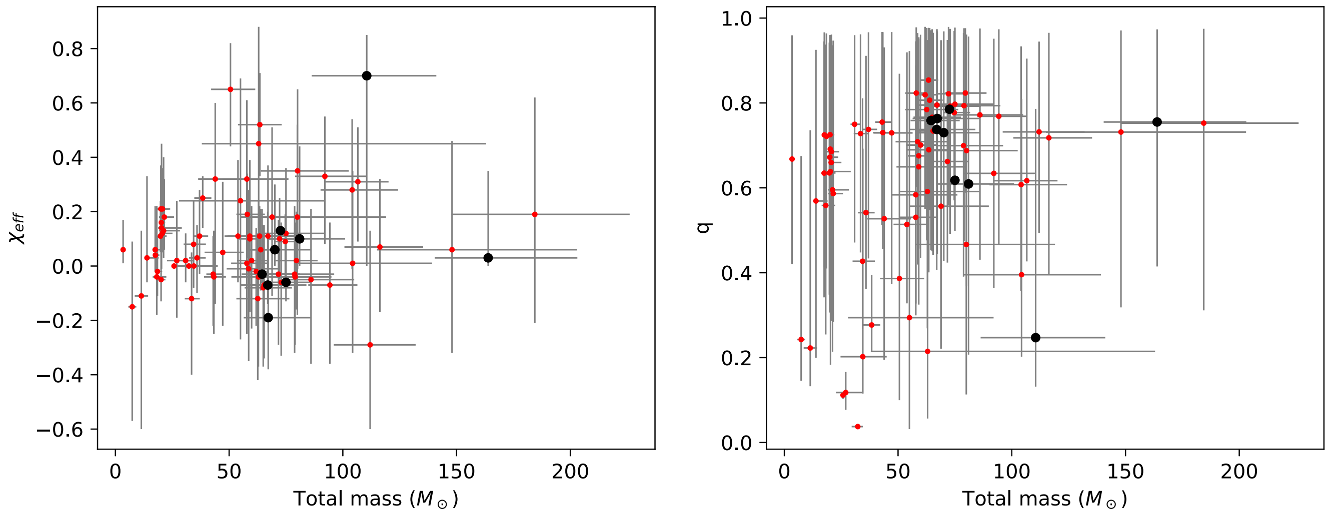

Figure 4. The distribution of merger parameters, (left) $\chi_{\rm eff}$ vs. total mass and (right) mass ratio (q) vs. total mass, for flare events (black points) and all O3 BBH mergers (red points).

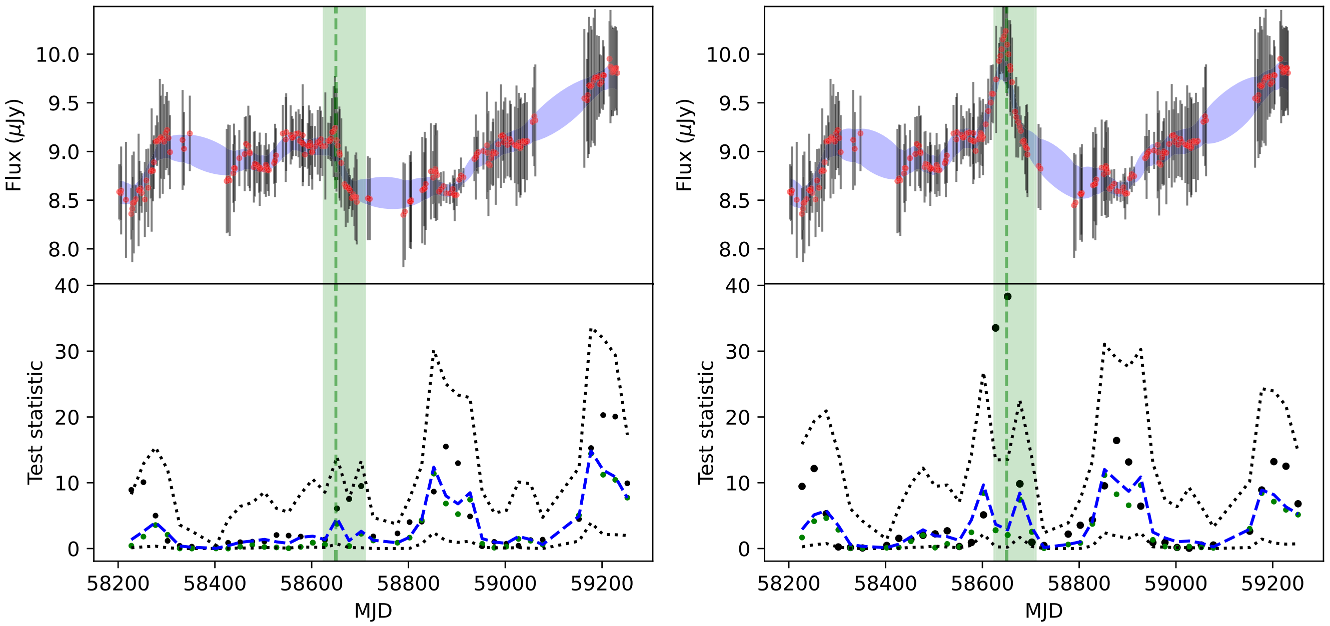

Figure 5. Left: a simulated ZTF AGN light curve from a DRW process and the change-point test statistic (black points in the lower plot) as a function of time. The full GP fit to the data and predicted uncertainties (blue shaded region) are shown. In the lower plot, the dashed blue line indicates the median value of the test statistic from 1000 posterior samples at each location and the dotted lines the 0.5th and 99.5th percentiles, respectively. Right: the same light curve with a flare peaking at $t = 58{,}650$ and with a rise time of 25 days, a decay time of 60 days, and an amplitude of 10% the median flux of the DRW model. The peak and duration of the flare are indicated in both plots by the dashed green line and shaded green region, respectively. The test statistic indicates that the flare is statistically inconsistent with the DRW model.

intrinsic variability and false-positive signatures, which includes signatures that may emerge on the far side of a face-on disk. Now suppose a fraction $f_{\rm AGN}$ of all $N_{\rm BBH}$ LIGO/Virgo-detected BBH mergers occur in AGN disks. Then, the number of EM counterparts potentially detectable by our survey, $N_{\rm EM,BBH}$, is

$$N_{\rm EM,BBH} \approx 3\left(\frac{N_{\rm BBH}}{83}\right)\left(\frac{f_{\rm AGN}}{0.5}\right)\left(\frac{f_{\rm type-1}}{0.5}\right)\left(\frac{f_{\rm side}}{0.5}\right)\left(\frac{f_{\rm sky}}{0.5}\right), \quad (14)$$

where the value of ∼3 is the expected number of associations from Section 5. In principle, for O4 and beyond we could increase $f_{\rm sky} \sim 1$ by including a large-sky survey in the Southern Hemisphere, thereby doubling the number of potentially detectable counterparts.

6.1. What if Our Sample Consists Entirely of False Positives?

We have isolated some extreme-variability events in AGNs that occurred in LIGO/Virgo O3 publicly announced event volumes. These events are not likely intrinsic disk variability according to models of stochastic AGN variability for these (or most) AGNs (see Section 4.5). These events are also not likely known false positives, such as microlensing events, SNe or TDEs (see Section 4.4). However, even if the AGN channel is responsible for most of the mergers LIGO/Virgo observes, it is possible that no EM counterparts will ever be detected from this channel due to muffling of EM counterparts by optically thick disks. In this case, only a statistical approach will reveal the AGN contribution to the observed rate (Bartos et al. 2017a; Veronesi et al. 2022). If we assume that all the candidate events in Table 3 are in fact drawn from a tail of rare disk variability events, then we can conclude that such events occur at a rate of $\mathcal{O}(10^{-6}\ \mathrm{yr}^{-1}\ \mathrm{AGN}^{-1})$. The rate of expected disk-crossing events expected is much higher than the rate here (Fabj et al. 2020). The short-timescale nature of the events implies that these events either represent short-lived embedded explosive breakout from within the disk that are not SNe, or explosive events very close to the ISCO. The possible candidates for such events can be observationally tested. For example, an off-center explosive event not associated with a kicked BBH merger must yield a temporary, asymmetric broad optical line profile on a timescale of days to weeks after the event, depending on the semimajor axis of the flaring event (McKernan et al. 2019). Unfortunately, we could not test this possibility for any of the flares in our sample as they were identified long after any such signature might have developed and decayed. Nevertheless, for future GW observing runs, spectral follow-up on interesting candidate AGN GW events with associated AGN flaring can be a powerful technique for identifying off-center flaring events.

By contrast, explosive events near the ISCO will cause the optical broad lines to reverberate symmetrically and so any change in broad optical lines that is not confidently asymmetric implies a flare origin close to the ISCO. Such flaring could be associated with disk instabilities or magnetic flux explosive release. If our sample of flaring events are due to such extreme effects near the central engine, this allows us to constrain models of magnetic field buildup around the SMBH due to accretion, as well as models of disk instabilities and the frequency of their occurrence.

6.2. Implications for Electromagnetic Follow-up in the Future: O4 and Beyond

EM follow-up of GW merger events is time and effort intensive. Therefore we suggest that the International Gravitational-Wave Network[19] (IGWN) update skymaps publicly once parameterization for an individual event has settled down. In particular, since we are searching for EM counterparts that might emerge on time frames of weeks after individual events, it would be very useful (and presumably low cost) for IGWN to automatically release public skymaps about one month after individual events.

As waveforms used in parameterizing individual events change, the resulting skymaps and error volumes can also change (even by small amounts). It would be very useful if IGWN were to publicly list the waveform models used to arrive at particular parameterizations (without necessarily revealing other information about mergers).

In order to optimize follow-up (including spectroscopy) during future IGWN operating runs, it would be useful for IGWN to list "AGN possible" flags in public data releases. We recommend that such a flag correspond to multiple merger parameters including (i) high-mass, i.e., $M_1 > 50\ M_\odot$, or intermediate-mass black hole (IMBH) formation events $M_{\mathrm{tot}} > 100\ M_\odot$, (ii) significantly asymmetric mass ratios ($q < 0.3$), and (iii) strongly misaligned spins such that a strong recoil kick (v_k) would be expected. By making an "AGN possible" flag multiparameter and one not tied to, for example, IMBH formation events, it allows us to cross-check EM flare parameters with GW measurements. For example, an "AGN possible" IGWN flag corresponding to a likely large recoil kick could help rule out a flare that is significantly delayed from the merger time as a false positive.

We also recommend that coordinating infrastructure for transient follow-up, such as community alert brokers and target and observation managers, may want to maintain watchlists of AGNs within the 90% confidence volumes of IGWN skymaps. Automatic follow-up, e.g., spectroscopy, could then be triggered for those sources which started to exhibit flaring activity.

7. Conclusions

Our picture of the accretion disk of an AGN is evolving from the simple Shakura–Sunyaev thin-disk model to a dynamic environment with encounters between gas, members of the nuclear star cluster, and clouds of stellar mass BHs in orbit around the central SMBH. In particular, AGN disks are a promising source of the stellar-origin compact object mergers being detected by GW observatories. They are also the only BH merger channel where an EM counterpart must occur (whether detectable or not).

We have conducted a systematic search for possible EM counterparts in AGNs to BBH object mergers detected by LIGO/Virgo in O3. We filtered out expected false positives, such as SNe and TDEs, and developed a change-point algorithm to test whether specific AGN flares are consistent with stochastic variability in their hosts or are more likely to be the result of some other mechanism. We found seven AGN flares associated with 12 merger events. This is statistically unlikely, with $p \sim 10^{-3}$. Simulation of random LIGO events and our selection procedure confirm the spatial coincidence rate.

However, our knowledge of the phenomenology of AGN flaring is as yet incomplete: for example, we expect TDEs and SNe embedded in the accretion disk but have no real idea of what these would look like in terms of a detectable signal. We are conducting work on the set of AGN flares detected in ZTF and other large optical time-domain surveys to identify categories of events and their respective rates as a way to resolve this unknown false-positive issue. Detailed numerical simulations of such events, involving full magnetohydrodynamics, general relativity, and radiative-transfer code, are also underway by other groups which will aid the search for EM counterparts in AGNs. If we consider that at least one of the associations we have identified is real, then this has significant implications for both GW and AGN physics.

The authors thank Will Farr and Colm Talbot for very useful discussions about LIGO/Virgo waveform choices and parameterization. M.J.G., B.M., and K.E.S.F. acknowledge the Center for Computational Astrophysics at the Flatiron Institute, New York for their hospitality and support.

This work was supported in part by the National Science Foundation grant Nos. AST-1815034, AST-1831412, and AST-2108402, the NASA grant No. 16-ADAP16-0232, and Simons Foundation grant No. 533845. The work of D.S. was carried out at Jet Propulsion Laboratory, California Institute of Technology, under a contract with NASA. M.W.C. is

---

[19] The O4 observing run will employ the LIGO, Virgo, and KAGRA gravitational wave detectors.

supported by the National Science Foundation with grant Nos. PHY-2010970 and OAC-2117997. P.R. acknowledges the support received from the Agence Nationale de la Recherche of the French government through the program "Investissements d'Avenir" (16-IDEX-0001 CAP 20-25).

This work made use of the Million Quasars Catalogue.

Based on observations obtained with the Samuel Oschin Telescope 48 inch and the 60 inch Telescope at the Palomar Observatory as part of the Zwicky Transient Facility (ZTF) project. ZTF is supported by the National Science Foundation under grant No. AST-1440341 and a collaboration including Caltech, IPAC, the Weizmann Institute for Science, the Oskar Klein Center at Stockholm University, the University of Maryland, the University of Washington, Deutsches Elektronen-Synchrotron and Humboldt University, Los Alamos National Laboratories, the TANGO Consortium of Taiwan, the University of Wisconsin at Milwaukee, and Lawrence Berkeley National Laboratories. Operations are conducted by COO, IPAC, and UW.

The ZTF forced-photometry service was funded under the Heising-Simons Foundation grant No. 12540303 (PI: Graham).

Facility: P O: 1.2 m (Zwicky Transient Facility), Hale (DBSP), Keck:I (LRIS).

Software: astropy (Astropy Collaboration et al. 2013, 2018), corner (Foreman-Mackey 2016), ligo.skymap (Singer & Price 2016), scikit-learn (Pedregosa et al. 2011).

Appendix A
Gaussian Process Change-point Detection

A Gaussian process (GP) is a random process produced by a collection of random variables such that any finite set of those variables follows a multivariate Gaussian distribution. A GP is completely specified by a mean function, $m(\cdot) = \mathbb{E}[f(\cdot)]$, and a kernel covariance function, $k = \mathrm{Cov}(f(\cdot),\ f(\cdot))$. These are parameterized with vectors θ_m and θ_k, respectively, with $\theta = (\theta_m,\ \theta_k)$ denoting a vector of hyperparameters for a given GP.

A time series, $\boldsymbol{y}_{1:N} = \{y_{t_i}\}_{i=1}^{N}$, consists of a set of N observations at times $\boldsymbol{t}_{1:N} = \{t_i\}_{i=1}^{N}$, which we model with a GP, f: $y_t = f(t) + \epsilon_t$, where $\epsilon_t \sim \mathcal{N}(\epsilon_t|0,\ \sigma^2)$ is white Gaussian noise with a zero mean and variance σ^2.

For a given set of N we can compute the posterior distributions of function values and observations, $\boldsymbol{f}_{1:N}|\boldsymbol{t}_{1:N},\ \theta \sim \mathcal{N}(\boldsymbol{f}_{1:N}|\mu,\ \boldsymbol{K})$, where $\boldsymbol{f}_{1:N} = f(\boldsymbol{t}_{1:N}) = \{f(t_i)\}_{i=1}^{N}$ are function f values at the given input times, $\mu = \{\mu_i\}_{i=1}^{N} = \{m(t_i)\}_{i=1}^{N}$ are realizations of the GP mean function at the input times, $\boldsymbol{K} = \{\boldsymbol{K}_{i,j}\}_{i,j=1}^{N} = \{k(t_i,\ t_j)\}_{i,j=1}^{N}$ are realizations of the GP covariance function at the input times, and

$$\boldsymbol{y}_{1:N}|t_{1:N},\ \theta \sim \mathcal{N}(\boldsymbol{y}_{1:N}|\mu,\ \boldsymbol{K} + \sigma^2\boldsymbol{I}),$$

where $\boldsymbol{I}$ is the identity matrix. The marginal log likelihood function of observed data is given by

$$\log p(\boldsymbol{y}_{1:N}|\boldsymbol{t}_{1:N},\ \theta) = -\frac{1}{2}(\boldsymbol{y}_{1:N} - \mu)^T(\boldsymbol{K} + \sigma^2\boldsymbol{I})^{-1}(\boldsymbol{y}_{1:N} - \mu) - \frac{1}{2}\mathrm{logdet}(\boldsymbol{K} + \sigma^2\boldsymbol{I}) - \frac{N}{2}\log 2\pi.$$

Estimates of the hyperparameters can be obtained by maximizing the marginal likelihood $\hat{\boldsymbol{\theta}}_a = \arg\max_{\theta_a} \log p(\boldsymbol{y}|\boldsymbol{t},\ \theta_a)$.

A *change point* represents a transition between different states in a process that generates the time series, i.e., a change in the latent probability distribution of observed data. For a time series described by a GP, this can mean a change at some $t = t_\star$ in hyperparameter values or even a change in the mean and/or covariance functions:

$$\begin{aligned} &y_t = f_0(t) + \epsilon_t^0,\ \theta = \theta_0,\ \mu = m_0(t),\\ &\boldsymbol{K} = k_0(t_i,\ t_j),\ t < t_\star,\\ &y_t = f_1(t) + \epsilon_t^1,\ \theta = \theta_1,\ \mu = m_1(t),\\ &\boldsymbol{K} = k_1(t_i,\ t_j),\ t \geqslant t_\star. \end{aligned} \tag{A1}$$

It can be shown that when $\boldsymbol{x}$ is a random vector distributed as a multivariate Gaussian and $\boldsymbol{A}$ is an arbitrary symmetric matrix, the quadratic form $\boldsymbol{x}^T\boldsymbol{A}^{-1}\boldsymbol{x}$ has a generalized χ^2 distribution. We can therefore consider a test statistic for change-point detection, λ, assuming a null hypothesis, $\mathcal{H}_0$, that the GP remains unchanged during the whole observation period and an alternative claim, $\mathcal{H}_1$, that there exists some window, $\boldsymbol{t}_L = \{t_\star \leqslant t < t_\star + L\}$, over which the GP has different hyperparameters or a different functional form:

$$\lambda = (\boldsymbol{y}_L - \mu_0)^T(\boldsymbol{K}_0 + \sigma^2\boldsymbol{I})^{-1}(\boldsymbol{y}_L - \mu_0), \tag{A2}$$

$$= -2\log p(\boldsymbol{y}_L|\boldsymbol{t}_L,\ \hat{\theta}_0) - \mathrm{logdet}(\boldsymbol{K}_0 + \sigma^2\boldsymbol{I}) - N\log 2\pi, \tag{A3}$$

evaluated for the observations $\boldsymbol{y}_L$ at times $\boldsymbol{t}_L$ within the window and where the hyperparameters $\hat{\theta}_0$ are evaluated for the time series excluding the window range, i.e., $\{t \leqslant t_\star,\ t > t_\star + L\}$. There is no closed-form expression for the distribution of λ but we can estimate a significance level for any measured value from the distribution of λ associated with sample observations in the window range drawn from the GP posterior distribution:

$$\hat{\boldsymbol{y}}_L|t_L,\ \hat{\theta}_0 \sim \mathcal{N}(\boldsymbol{y}_L|\mu_0',\ \Sigma_0'),$$

where

$$\mu_0' = \boldsymbol{\mu}_0(\boldsymbol{t}_L) + \boldsymbol{k}_L^T(\boldsymbol{K}_0 + \sigma^2\boldsymbol{I})^{-1}(\boldsymbol{y} - \boldsymbol{\mu}_0), \tag{A4}$$

$$\Sigma_0' = k_0(\boldsymbol{t}_L,\ \boldsymbol{t}_L) - \boldsymbol{k}_L^T(\boldsymbol{K}_0 + \sigma^2\boldsymbol{I})^{-1}\boldsymbol{k}_L, \tag{A5}$$

$$\boldsymbol{k}_L = k_0(\boldsymbol{t}_L,\ \boldsymbol{t}.) \tag{A6}$$

We simulate an AGN light curve via a DRW process characterized by a timescale τ and an amplitude σ^2. A (zero-centered) data point m_{i+1} at time t_{i+1} is given by

$$m_{i+1} = m_i e^{-\Delta t/\tau} + G[\sigma^2(1 - e^{-2\Delta t/\tau})],$$

where $G(s^2)$ is a Gaussian deviate with variance s^2 and $\Delta t = t_{i+1} - t_i$. We can add a flare of amplitude A peaking at time $t = t_0$ with rise and decay times, t_g and t_e, respectively, to this:

$$m(t) = m_{\mathrm{DRW}}(t) + A\exp\left(-\frac{(t - t_0)^2}{2t_g^2}\right),\ t \leqslant t_0, \tag{A7}$$

$$= m_{\mathrm{DRW}}(t) + A\exp\left(-\frac{(t - t_0)}{t_e}\right),\ t > t_0. \tag{A8}$$

Figure 5 shows the test statistic for both a plain DRW model generated with observation times taken from ZTF and the same model plus a flare with an amplitude equal to 10% of the median flux of the DRW process. A window with a width of 50 days was used and significance levels for the test statistic determined from 1000 samples drawn from the posterior at each window location. The test statistic in the vicinity of the flare peak ($t \sim t_0$) is a statistically significant deviate, indicating that this region of the light curve is not consistent with being generated by the same process as the rest of the light curve.

Note: we assume a particular kernel for the null hypothesis; however, it may be that a given time series is not well described by that model and so the observed values of the quadratic statistic do not generally match those sampled from the posterior. The hyperparameter estimates in the window generally match those for the full time series. It is also possible that the noise modeling is insufficient.

Appendix B
Spectra of Candidate EMGW-associated Active Galactic Nuclei

Spectroscopic observations for all candidate electromagnetic gravitational-wave-associated AGNs (except for J154342.46+461233.4 where there were two existing Sloan Digital Sky Survey spectra) were obtained with the Low Resolution Imaging Spectrometer (LRIS; Oke et al. 1995) on the 10 m Keck I telescope and the Double Spectrograph (DBSP) instrument on the 200 inch Palomar Hale telescope (see Table 6). The data were reduced with standard pipelines for both instruments. The reduced spectra are shown in Figure 6.

Table 6
Observing Log for Candidate AGNs Associated with LIGO/Virgo Events

Name	Date (UT)	Telescope	Instrument	Exposure (s)
J053408.41+085450.6	2022 February 25	Keck I	LRIS	600
J120437.98+500024.0	2022 February 25	Keck I	LRIS	600
J124942.30+344928.9	2020 January 25	Keck I	LRIS	600
	2022 May 27	P200	DBSP	900
J154342.46+461233.4	2003 April 2	SDSS	…	…
	2017 May 5	SDSS	…	…
J181719.94+541910.0	2022 February 25	Keck I	LRIS	600
J183412.42+365655.3	2021 September 10	Keck I	LRIS	600
J224333.95+760619.2	2022 April 28	P200	DBSP	900

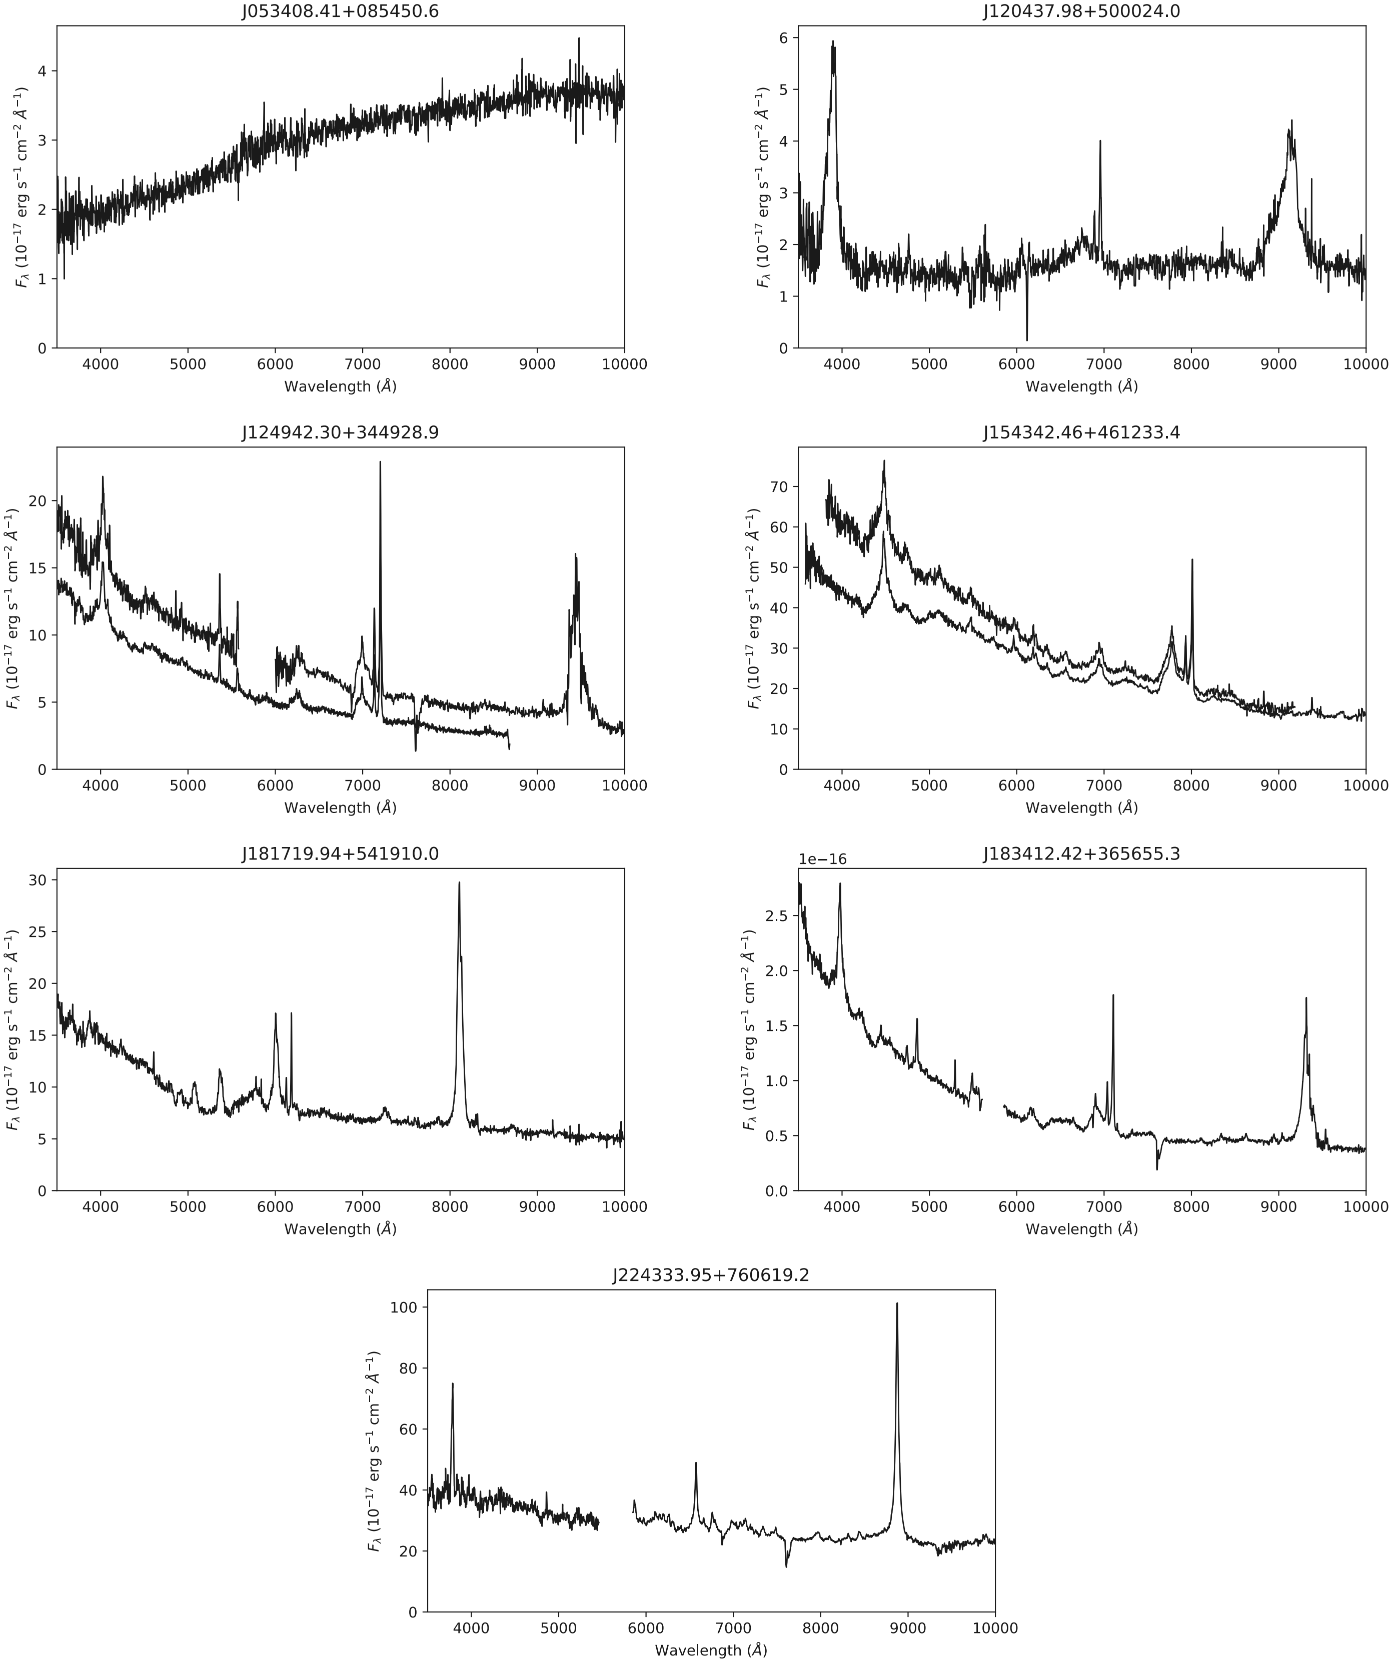

Figure 6. The spectra for the AGNs associated with LIGO/Virgo events. They have been smoothed with a 5 Å median filter.

Appendix C
Disk Exit Time Derivation

We assume the binary center of mass is initially orbiting in the midplane of the AGN disk with Keplerian velocity, $v_{\rm orb}$. At merger, the remnant experiences a kick velocity, v_k, in an arbitrary direction; the most rapid exit time will occur if the kick is directed perpendicular to the midplane of the disk. The time, t, to reach a height, z, above the disk midplane will be $t = z/v_k$. We assume a Gaussian atmosphere for the gas density away from the midplane, i.e., $\rho = \rho_0 \exp(-z^2/(2H^2))$, where ρ is the gas density, ρ_0 is the gas density at the midplane, z is the height above the midplane, and H is the scale height of the atmosphere. Rearranging, we find $z = H\sqrt{-2\ln(\rho/\rho_0)}$. If we say the height at which the remnant exits the disk is the height where the optical depth, τ, is unity, and use the relation that $\tau \propto \rho$, the vertical distance the remnant traverses to exit is $z_{\rm exit} = H\sqrt{-2\ln(1/\tau_{\rm mp})}$, where $\tau_{\rm mp}$ is the midplane optical depth. Thus, to find the time for the remnant to exit, we have

$$t_{\rm exit} = \frac{H\sqrt{2\ln(\tau_{mp})}}{v_k}, \tag{C1}$$

as in Equation (4).

ORCID iDs

Matthew J. Graham https://orcid.org/0000-0002-3168-0139
Barry McKernan https://orcid.org/0000-0002-9726-0508
K. E. Saavik Ford https://orcid.org/0000-0002-5956-851X
Daniel Stern https://orcid.org/0000-0003-2686-9241
S. G. Djorgovski https://orcid.org/0000-0002-0603-3087
Michael Coughlin https://orcid.org/0000-0002-8262-2924
Kevin B. Burdge https://orcid.org/0000-0002-7226-836X
Eric C. Bellm https://orcid.org/0000-0001-8018-5348
George Helou https://orcid.org/0000-0003-3367-3415
Ashish A. Mahabal https://orcid.org/0000-0003-2242-0244
Frank J. Masci https://orcid.org/0000-0002-8532-9395
Josiah Purdum https://orcid.org/0000-0003-1227-3738
Ben Rusholme https://orcid.org/0000-0001-7648-4142

References

Aasi, J., Abadie, J., Abbott, B. P., et al. 2015, CQGra, 32, 074001
Acernese, F., Agathos, M., Agatsuma, K., et al. 2015, CQGra, 32, 024001
Anand, S., Coughlin, M. W., Kasliwal, M. M., et al. 2021, NatAs, 5, 46
Antoni, A., MacLeod, M., & Ramirez-Ruiz, E. 2019, ApJ, 884, 22
Antonini, F. 2014, ApJ, 794, 106
Antonini, F., & Rasio, F. A. 2016, ApJ, 831, 187
Ashton, G., Ackley, K., Hernandez, I. M., & Piotrzkowski, B. 2021, CQGra, 38, 235004
Assef, R. J., Prieto, J. L., Stern, D., et al. 2018, ApJ, 866, 26
Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., Greenfield, P., et al. 2013, A&A, 558, A33
Astropy Collaboration, Price-Whelan, A. M., Sipocz, B. M., et al. 2018, AJ, 156, 123
Bailer-Jones, C. A. L. 2011, MNRAS, 411, 435
Bartos, I., Haiman, Z., Marka, Z., et al. 2017a, NatCo, 8, 831
Bartos, I., Kocsis, B., Haiman, Z., & Márka, S. 2017b, ApJ, 835, 165
Belczynski, K., Bulik, T., Fryer, C. L., et al. 2010, ApJ, 714, 1217
Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019a, PASP, 131, 018002
Bellm, E. C., Kulkarni, S. R., Barlow, T., et al. 2019b, PASP, 131, 068003
Bellovary, J. M., Low, M.-M.-M., McKernan, B., & Ford, K. E. S. 2016, ApJL, 819, L17
Berry, C. P. L., Mandel, I., Middleton, H., et al. 2015, ApJ, 804, 114
Boone, K. 2019, AJ, 158, 257
Calderón Bustillo, J., Leong, S. H. W., Chandra, K., McKernan, B., & Ford, K. E. S. 2021, arXiv:2112.12481
Callister, T. A., Haster, C.-J., Ng, K. K. Y., Vitale, S., & Farr, W. M. 2021, ApJL, 922, L5
Cannizzaro, G., Fraser, M., Jonker, P. G., et al. 2020, MNRAS, 493, 477
Chambers, K. C., Magnier, E. A., Metcalfe, N., et al. 2016, arXiv:1612.0556
Chan, C. H., Piran, T., Krolik, J.-H., & Saban, D. 2019, ApJ, 881, 113
Chen, H. Y., Haster, C. J., Vitale, S., Farr, W. M., & Isi, M. 2022, MNRAS, 513, 2152
Condon, J. J., Cotton, W. D., Greisen, E. W., et al. 1998, AJ, 115, 1693
Coughlin, M. W., Ahumada, T., Anand, S., et al. 2019, ApJL, 885, L19
de Mink, S. E., & Mandel, I. 2016, MNRAS, 460, 3545
Fabj, G., Nasim, S. S., Caban, F., et al. 2020, MNRAS, 499, 2608
Finn, L. S., & Chernoff, D. F. 1993, PhRvD, 47, 2198
Flesch, E. W. 2019, arXiv:1912.05614
Foley, R. J., Sanders, N. E., & Kirshner, R. P. 2011, ApJ, 742, 89
Ford, K. E. S., & McKernan, B. 2022, MNRAS, 517, 5827
Foreman-Mackey, D. 2016, JOSS, 1, 24
Fragione, G., Leigh, N. W. C., & Perna, R. 2019, MNRAS, 488, 2825
Fremling, C., Miller, A. A., Sharma, Y., et al. 2020, ApJ, 895, 32
Gaia Collaboration, Vallenari, A., & Brown, A. 2022, arXiv:2208.00211
Gerosa, D., & Berti, E. 2019, PhRvD, 100, 041301
Gerosa, D., & Fishbach, M. 2021, NatAs, 5, 749
Graham, M. J., Djorgovski, S. G., Drake, A. J., et al. 2017, MNRAS, 470, 4112
Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019, PASP, 131, 078001
Graham, M. J., Ford, K. E. S., McKernan, B., et al. 2020, PhRvL, 124, 251102
Guo, H., Shen, Y., & Wang, S. 2018, PyQSOFit: Python code to fit the spectrum of quasars, Astrophysics Source Code Library, ascl:1809.008
Hammerstein, E., Gezari, S., van Velzen, S., et al. 2021, ApJL, 908, L20
Ho, L. C., & Kim, M. 2015, ApJ, 809, 123
Jiang, Y. F., Stone, J. M., & Davis, S. W. 2019, ApJ, 880, 67
Kasen, D., & Bildsten, L. 2010, ApJ, 717, 245
Kasliwal, M. M., Anand, S., Ahumada, T., et al. 2020, ApJ, 905, 145
Kasliwal, V. P., Vogeley, M. S., & Richards, G. T. 2015, MNRAS, 451, 4328
Kimura, S. S., Murase, K., & Bartos, I. 2021, ApJ, 916, 111
Krolik, J. H. 1999, Active Galactic Nuclei : From the Central Black Hole to the Galactic Environment (Princeton, NJ: Princeton Univ. Press)
Lawrence, A., Bruce, A. G., MacLeod, C., et al. 2016, MNRAS, 463, 296
LIGO Scientific Collaboration & Virgo Collaboration 2019, ApJL, 882, L24
Mapelli, M. 2021, in Handbook of Gravitational Wave Astronomy, ed. C. Bambi, S. Katsanevas, & K. D. Kokkotas (Singapore: Springer), 16
Masci, F. J., Laher, R. R., Rusholme, B., et al. 2019, PASP, 131, 018003
Matthews, T. A., & Sandage, A. R. 1963, ApJ, 138, 30
McKernan, B., Ford, K. E. S., Bartos, I., et al. 2019, ApJL, 884, L50
McKernan, B., Saavik Ford, K. E., Bellovary, J., et al. 2018, ApJ, 866, 66
McKernan, B., Ford, K. E. S., Callister, T., et al. 2022a, MNRAS, 514, 3886
McKernan, B., Ford, K. E. S., Cantiello, M., et al. 2022b, MNRAS, 514, 4102
McKernan, B., Ford, K. E. S., Kocsis, B., Lyra, W., & Winter, L. M. 2014, MNRAS, 441, 900
McKernan, B., Ford, K. E. S., Lyra, W., & Perets, H. B. 2012, MNRAS, 425, 460
McKernan, B., Ford, K. E. S., & O?Shaughnessy, R. 2020, MNRAS, 498, 4088
Meyer, M., Scargle, J. D., & Blandford, R. D. 2019, ApJ, 877, 39
Moreno, J., Vogeley, M. S., Richards, G. T., & Yu, W. 2019, PASP, 131, 063001
Mukherjee, S., Ghosh, A., Graham, M. J., et al. 2020, arXiv:2009.14199
Oke, J. B., Cohen, J. G., Carr, M., et al. 1995, PASP, 107, 375
Ostriker, E. C. 1999, ApJ, 513, 252
Palmese, A., Fishbach, M., Burke, C. J., Annis, J., & Liu, X. 2021, ApJL, 914, L34
Pan, Z., & Yang, H. 2021, ApJ, 923, 173
Patterson, M. T., Bellm, E. C., Rusholme, B., et al. 2019, PASP, 131, 018001
Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, JMLR, 12, 2825
Perley, D. A., Fremling, C., Sollerman, J., et al. 2020, ApJ, 904, 35
Perna, R., Lazzati, D., & Cantiello, M. 2021, ApJL, 906, L7
Petrov, P., Singer, L. P., Coughlin, M. W., et al. 2022, ApJ, 924, 54
Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2016, A&A, 594, A13
Rees, M. J. 1988, Natur, 333, 523
Ricci, C., Kara, E., Loewenstein, M., et al. 2020, ApJL, 898, L1
Rodriguez, C. L., Chatterjee, S., & Rasio, F. A. 2016a, PhRvD, 93, 084029
Rodriguez, C. L., Haster, C. J., Chatterjee, S., Kalogera, V., & Rasio, F. A. 2016b, ApJL, 824, L8
Ross, N. P., Ford, K. E. S., Graham, M., et al. 2018, MNRAS, 480, 4468
Rosswog, S., Ramirez-Ruiz, E., & Hix, W. R. 2009, ApJ, 695, 404
Ryu, T., Krolik, J., Piran, T., & Noble, S. C. 2020, ApJ, 904, 98

Samsing, J., Bartos, I., D'Orazio, D. J., et al. 2022, Natur, 603, 237
Scargle, J. D., Norris, J. P., Jackson, B., & Chiang, J. 2013, ApJ, 764, 167
Secunda, A., Bellovary, J., Low, M. M. M., et al. 2019, ApJ, 878, 85
Secunda, A., Bellovary, J., Mac Low, M.-M., et al. 2020, ApJ, 903, 133
Shen, Y., & Liu, X. 2012, ApJ, 753, 125
Singer, L. P., & Price, L. R. 2016, PhRv, 93, 024013
Singer, L. P., & Price, L. R. 2016, PhRvD, 93, 024013
Singer, L. P., Chen, H. Y., Holz, D. E., et al. 2016, ApJL, 829, L15
Sirko, E., & Goodman, J. 2003, MNRAS, 341, 501
Stern, D., McKernan, B., Graham, M. J., et al. 2018, ApJ, 864, 27
Stone, N. C., Metzger, B. D., & Haiman, Z. 2017, MNRAS, 464, 946
Tagawa, H., Haiman, Z., Bartos, I., Kocsis, B., & Omukai, K. 2021, MNRAS, 507, 3362
Tagawa, H., Haiman, Z., & Kocsis, B. 2020, ApJ, 898, 25
Thompson, T. A., Quataert, E., & Murray, N. 2005, ApJ, 630, 167
Vajpeyi, A., Thrane, E., Smith, R., McKernan, B., & Saavik Ford, K. E. 2022, ApJ, 931, 82
van Velzen, S., Gezari, S., Hammerstein, E., et al. 2021, ApJ, 908, 4
Varma, V., Biscoveanu, S., Islam, T., et al. 2022, PhRvL, 128, 191102
Veitch, J., Raymond, V., Farr, B., et al. 2015, PhRvD, 91, 042003
Veronesi, N., Rossi, E. M., van Velzen, S., & Buscicchio, R. 2022, MNRAS, 514, 2092
Villar, V. A., Cranmer, M., Contardo, G., Ho, S., & Yao-Yu Lin, J. 2020, arXiv:2010.11194
Wang, J. M., Liu, J. R., Ho, L. C., Li, Y. R., & Du, P. 2021a, ApJL, 916, L17
Wang, Y.-H., McKernan, B., Ford, S., et al. 2021b, ApJL, 923, L23
Woosley, S. E. 2017, ApJ, 836, 244
Yang, Y., Bartos, I., Fragione, G., et al. 2022, ApJL, 933, L28
Yang, Y., Bartos, I., Gayathri, V., et al. 2019, PhRvL, 123, 181101