



# DELIGHT: Deep Learning Identification of Galaxy Hosts of Transients using Multiresolution Images

Francisco Förster<sup>1,2,3,4</sup> , Alejandra M. Muñoz Arancibia<sup>2,3</sup> , Ignacio Reyes-Jainaga<sup>2,3</sup> , Alexander Gagliano<sup>5,6,24</sup> , Dylan Britt<sup>7,8</sup> , Sara Cuellar-Carrillo<sup>9</sup> , Felipe Figueroa-Tapia<sup>10</sup> , Ava Polzin<sup>11</sup> , Yara Yousef<sup>12,13</sup> , Javier Arredondo<sup>2,3</sup> , Diego Rodríguez-Mancini<sup>14</sup> , Javier Correa-Orellana<sup>2,15</sup> , Amelia Bayo<sup>16</sup> , Franz E. Bauer<sup>2,15,17</sup> , Márcio Catelan<sup>2,15,17</sup> , Guillermo Cabrera-Vives<sup>2,18</sup> , Raya Dastidar<sup>2,19</sup> , Pablo A. Estévez<sup>2,20</sup> , Giuliano Pignata<sup>2,19</sup> , Lorena Hernández-García<sup>2,10</sup> , Pablo Huijse<sup>2,21</sup> , Esteban Reyes<sup>22</sup> , Paula Sánchez-Sáez<sup>2,16</sup> , Mauricio Ramírez<sup>19</sup> , Daniela Grandón<sup>23</sup> , Jonathan Pineda-García<sup>19</sup> , Francisca Chabour-Barra<sup>4,23</sup> , and Javier Silva-Farfán<sup>4</sup>

<sup>1</sup> Data and Artificial Intelligence Initiative (IDIA), Faculty of Physical and Mathematical Sciences, Universidad de Chile, Chile

<sup>2</sup> Millennium Institute of Astrophysics (MAS), Nuncio Monseñor Sótero Sanz 100, Providencia, Santiago, Chile

<sup>3</sup> Center for Mathematical Modeling (CMM), Universidad de Chile, Beauchef 851, Santiago 8320000, Chile

<sup>4</sup> Department of Astronomy, Universidad de Chile, Las Condes, Santiago, Chile

<sup>5</sup> Department of Astronomy, University of Illinois at Urbana-Champaign, 1002 W. Green Street, IL 61801, USA

<sup>6</sup> Center for Astrophysical Surveys, National Center for Supercomputing Applications, Urbana, IL, 61801, USA

<sup>7</sup> Department of Physics, Stanford University, 382 Via Pueblo Mall, Stanford, CA 94305, USA

<sup>8</sup> Kavli Institute for Particle Astrophysics & Cosmology, P.O. Box 2450, Stanford University, Stanford, CA 94305, USA

<sup>9</sup> Escuela de Ingeniería Eléctrica, Pontificia Universidad Católica de Valparaíso, Valparaíso, Chile

<sup>10</sup> Instituto de Física y Astronomía, Facultad de Ciencias, Universidad de Valparaíso, Gran Bretaña 1111, Valparaíso, Chile

<sup>11</sup> Department of Astronomy and Astrophysics, The University of Chicago, Chicago, IL 60637, USA

<sup>12</sup> Barcelona School of Informatics, Polytechnic University of Catalonia, E-08034 Barcelona, Spain

<sup>13</sup> University of Barcelona, E-08028 Barcelona, Spain

<sup>14</sup> Data Observatory Foundation, Santiago, Chile

<sup>15</sup> Instituto de Astrofísica, Facultad de Física, Pontificia Universidad Católica de Chile, Santiago, Chile

<sup>16</sup> European Southern Observatory, Karl-Schwarzschild-Strasse 2, D-85748 Garching bei München, Germany

<sup>17</sup> Centro de Astroingeniería, Facultad de Física, Pontificia Universidad Católica de Chile, Santiago, Chile

<sup>18</sup> Department of Computer Science, Universidad de Concepción, Concepción, Chile

<sup>19</sup> Instituto de Astrofísica, Facultad de Ciencias Exactas, Universidad Andres Bello, Las Condes, Santiago, Chile

<sup>20</sup> Department of Electrical Engineering, Universidad de Chile, Santiago, Chile

<sup>21</sup> Instituto de Informática, Universidad Austral, Valdivia, Chile

<sup>22</sup> Fintual Administradora General de Fondos S.A., Santiago, Chile

<sup>23</sup> Departamento de Física, Facultad de Ciencias Físicas y Matemáticas, Universidad de Chile, Santiago, Chile

Received 2022 August 5; revised 2022 September 5; accepted 2022 September 8; published 2022 October 17

## Abstract

We present DELIGHT, or Deep Learning Identification of Galaxy Hosts of Transients, a new algorithm designed to automatically and in real time identify the host galaxies of extragalactic transients. The proposed algorithm receives as input compact, multiresolution images centered at the position of a transient candidate and outputs two-dimensional offset vectors that connect the transient with the center of its predicted host. The multiresolution input consists of a set of images with the same number of pixels, but with progressively larger pixel sizes and fields of view. A sample of 16,791 galaxies visually identified by the Automatic Learning for the Rapid Classification of Events broker team was used to train a convolutional neural network regression model. We show that this method is able to correctly identify both relatively large ( $10'' < r < 60''$ ) and small ( $r \leq 10''$ ) apparent size host galaxies using much less information (32 kB) than with a large, single-resolution image (920 kB). The proposed method has fewer catastrophic errors in recovering the position and is more complete and has less contamination ( $< 0.86\%$ ) recovering the crossmatched redshift than other state-of-the-art methods. The more efficient representation provided by multiresolution input images could allow for the identification of transient host galaxies in real time, if adopted in alert streams from new generation of large -etendue telescopes such as the Vera C. Rubin Observatory.

*Unified Astronomy Thesaurus concepts:* [Supernovae \(1668\)](#); [Galaxies \(573\)](#); [Astroinformatics \(78\)](#); [Astronomical object identification \(87\)](#); [Classification \(1907\)](#)

## 1. Introduction

A new generation of astronomical surveys is revolutionizing the study of the dynamic sky thanks to their large etendues, i.e., the product between the field of view and collecting area, and their

massive data processing capabilities. Surveys such as the Zwicky Transient Facility (ZTF; Bellm et al. 2019; Graham et al. 2019), Gaia (Gaia Collaboration et al. 2016; Breidt 2019), the Asteroid Terrestrial-Impact Last Alert System, (ATLAS; Tonry et al. 2018), and the Legacy Survey of Space and Time (LSST; LSST Science Collaboration et al. 2009; Ivezić et al. 2019) are detecting or will detect objects that change their position and/or brightness at rates going from hundreds of thousands to tens of millions of objects per night in almost real time. These changes are reported in massive alert streams, where alerts contain image cutouts and other relevant information about objects of solar system, Galactic,

<sup>24</sup> National Science Foundation Graduate Research Fellow.



or extragalactic origin, but whose nature is not necessarily known in advance. Consequently, a new generation of astronomical alert brokers: the Automatic Learning for the Rapid Classification of Events (ALeRCE; Förster et al. 2021); the Alert Management, Photometry and Evaluation of Lightcurves (AMPEL; Nordin et al. 2019); the Arizona-NOAO Temporal Analysis and Response to Events System (ANTARES; Narayan et al. 2018); Babamul;<sup>25</sup> Fink (Möller et al. 2021); Lasair (Smith et al. 2019); and Pitt-Google<sup>26</sup> are being developed to provide a fast and reliable filtering and classification to enable the real-time characterization of objects in the stream using the available follow-up resources, including massively multiplexed spectroscopic instruments such as the 4-meter Multi-Object Spectroscopic Telescope (4MOST; de Jong et al. 2012), the Maunakea Spectroscopic Explorer (MSE; Simons et al. 2014), the Multi-Object Spectrograph for the Very Large Telescope (MOONS; Cirasuolo et al. 2014), the Subaru Prime Focus Spectrograph (PFS; Takada et al. 2014), and the Dark Energy Spectroscopic Instrument (DESI; DESI Collaboration et al. 2016).

An important ingredient for the correct classification and characterization of the objects contained in the alert stream is contextual information. In particular, the possible association with a host galaxy in the vicinity of an alert is key for the classification of extragalactic transients. A correct host association can allow the inference of transient properties based on the properties of its host, notably the distance, but also other global and local host properties such as age, stellar mass, star formation activity, and metallicity, that may provide clues about the nature of the transient’s progenitor system.

This host-galaxy-alert association can be done via cross-matches with a comparison source catalog, selecting for example the nearest normalized-distance object (Sullivan et al. 2006; Gupta et al. 2016), or by directly studying deep reference images in the same location of the sky (Gagliano et al. 2021). Both methods have advantages and disadvantages: while catalog association can be very fast, it can lead to catastrophic errors due to the incorrect representation of galaxies with complex geometries; and while direct image association is generally considered to be more reliable, it can require large input images that are slower to retrieve and process than a catalog entry.

In this work, we propose a method that is both fast and reliable to directly infer the host galaxies of transient candidates from images. It uses a combination of catalog information (source positions and characteristic ellipses); a compact, multiresolution image at the location of the transient; and a convolutional neural network trained with a sample of 16,791 visually selected transient candidate hosts. This method is able to identify the position of the most likely host galaxy using a small amount of input data (32 kB) in a timescale of about 60 ms using a CPU, fast enough to be applied in future alert streams such as that from LSST. Moreover, this method is able to indicate the degree of confidence of the predicted association by using rotation and flip data augmentation in real time.

This paper is organized as follows: in Section 2 we review some of the existing host association methods and give a brief introduction to the proposed method. In Section 3 we discuss the training set preparation and show some examples. In

Section 4 we give a detailed description of the proposed method, its architecture, and the hyperparameter optimization. In Section 5 we show the main results of the method, including a comparison with other methods, as well as an analysis of both the errors in angular separation and redshift association. Finally, in Section 6 we present the conclusions of this work and discuss future steps.

## 2. Host-galaxy Identification Methods

Historically, the most common method to do host-galaxy association has been the visual inspection of images (e.g., the Asiago catalog; Barbon et al. 1999). This method is generally reliable and up to now has been seen to outperform other methods in host galaxies with complex geometries, but it does not scale easily to large numbers of transients and neither is it failproof. First, the identification via visual inspection involves some degree of subjectivity that poses challenges for reproducibility. Second, in some instances an image may not be deep enough to detect the host (in the so-called *hostless* supernovae), or it may not provide enough information to break association degeneracies between multiple host candidates. In the former case, deeper images or catalogs may be required to get the correct association; and in the latter, spectroscopic data may be necessary to derive the redshift of the transient directly from narrow spectral lines. In this work we focus on those cases where a visual inspection is not possible, because of the large volume of events where the association is needed and where no additional information is available, because a fast association is required to decide whether to obtain spectroscopic observations of a transient candidate.

This type of automatic host association can be done in two ways: (1) selecting a source from a reference catalog based on the position of the transient, or (2) using an image centered in the position of the transient to predict a host center and then select a source from a reference catalog based on the predicted position. In both cases, it is possible to train a machine-learning algorithm to predict the position given the input, e.g., using a boosted decision tree to select the best source among nearby catalog matches (Smith et al. 2020). We will now discuss some of the methods that are publicly available.

### 2.1. Catalog-based Association Methods

*Nearest neighbor (1-NN)*—The simplest association method is to find the closest source in projection to the transient in a reference catalog. This has the advantage that no information other than the position of the source is needed, but it can lead to catastrophic errors when the nearest source (in projection) is a star or a smaller, distant galaxy unrelated to the transient.

*Nearest normalized elliptical SExtractor distance (SEx)*—An improved version of the previous method is to use normalized angular separations to associate the transient to the source, which has the disadvantage that information about the typical size and orientation of every source is required. The normalized angular separation can be computed as the dimensionless elliptical distance  $R$  defined in Sullivan et al. (2006),  $R^2 = CXX(x_{tr} - x_{gal})^2 + CXY(x_{tr} - x_{gal})(y_{tr} - y_{gal}) + CYY(y_{tr} - y_{gal})^2$ , where  $x_{tr}$ ,  $y_{tr}$  and  $x_{gal}$ ,  $y_{gal}$  are the transient and galaxy positions in pixels, respectively, and where  $CXX$ ,  $CXY$ , and  $CYY$  are the SExtractor (Bertin & Arnouts 1996) ellipse parameters.<sup>27</sup> The main

<sup>25</sup> <https://docs.babamul.dev/>

<sup>26</sup> <https://pitt-broker.readthedocs.io/>

<sup>27</sup> See <https://sextractor.readthedocs.io/en/latest/Position.html>.

limitation of this method is that galaxies with complex geometries can be broken down into several different ellipses.

*Nearest normalized Directional Light Radius (DLR)*—As in SEx, this requires information about the typical size and orientation of every source. In this case, a normalized angular separation is inferred as the ratio between the transient to host angular separation and a directional light radius (DLR; Gupta et al. 2016; Sako et al. 2018) computed using the source Stoke parameters as explained in Equations (4)–(9) of Gagliano et al. (2021).

## 2.2. Image-based Association Methods

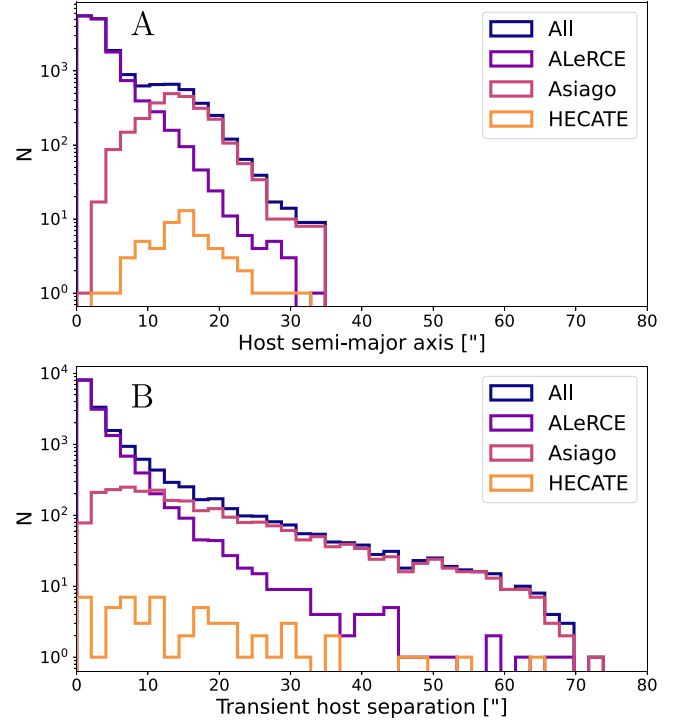
*Gradient Ascent method (Grad)*—Alternatively, one can infer the position of the host directly from the input images. This is done in the gradient ascent method described in Gagliano et al. (2021), in which after some preprocessing of Panoramic Survey Telescope and Rapid Response System (PanSTARRS) images, including the removal of stellar sources, can infer the position of the host iteratively starting at the position of the transient and gradually moving in the direction of the local surface brightness gradient until a local maximum is reached.

*Deep Learning Identification of Galaxy Hosts of Transients (DELIGHT)*—This is the method proposed in this work, that uses as input multiresolution images of the transient’s location, including small and large angular scales, and predicts the position of the host using a convolutional neural network invariant to rotations and flips. The input images are light-weight, taking about 3.5% of the space that single-resolution images of the same region would require.

## 3. Data Preparation

For the proposed method, DELIGHT, we use a total of 16,791 visually selected host galaxies, restricting the sample to galaxies whose candidates are within 1 arcmin from the host center. This is equivalent to the angular radius of the Milky Way seen at a distance of 57 Mpc. The sample can be divided into two families:

1. *ALeRCE sample*: a collection of 14,111 transient candidate host galaxies visually identified by the ALeRCE broker team between 2019 August and 2022 May. The transient candidates were first identified in the ZTF public alert stream using the ALeRCE stamp classifier (Carrasco-Davis et al. 2021) and then filtered and vetoed with the help of the ALeRCE SN Hunter tool.<sup>28</sup> Resulting candidates were used to visually select the most likely host-galaxy position among the positions of NASA Extragalactic Database (NED; Helou et al. 1995); Set of Identifications, Measurements, and Bibliography for Astronomical Data (SIMBAD; Wenger et al. 2000); and Sloan Digital Sky Survey (SDSS; Abdurro’uf et al. 2022) sources superimposed on PanSTARRS data release (DR) 1 color images using the ipyaladin tool.<sup>29</sup> This information was reported to the Transient Name Server (TNS<sup>30</sup>) daily for those candidates that had not been reported based on ZTF data before.



**Figure 1.** Angular semimajor axis of transient host galaxies (A) and transient to host angular separation (B) distributions of the input samples. Note that the isophotal semimajor axis is approximately 3 times the reported semimajor axis (Bertin & Arnouts 1996). The ALeRCE sample is significantly larger than the nearby sample (14,111 vs. 2680 galaxies, respectively), but it is dominated by smaller and more distant galaxies. Combining both samples helps to reduce the imbalance in host-galaxy size (see discussion in Section 2).

2. *Nearby sample*: a collection of 2563 visually identified host galaxies from the Asiago catalog (Barbon et al. 1999) and 117 host galaxies that have TNS reported SNe in the Heraklion Extragalactic CAtalogUE (HECATE) catalog (Kovlakas et al. 2021) of nearby galaxies. Thus, our nearby sample includes 2680 visually selected host galaxies that are in general closer and with a larger angular size (median SExtractor derived semimajor axis of 14.''1 versus 2.''4) than those obtained in the ALeRCE sample. Note that there is no overlap between the ALeRCE and nearby samples.

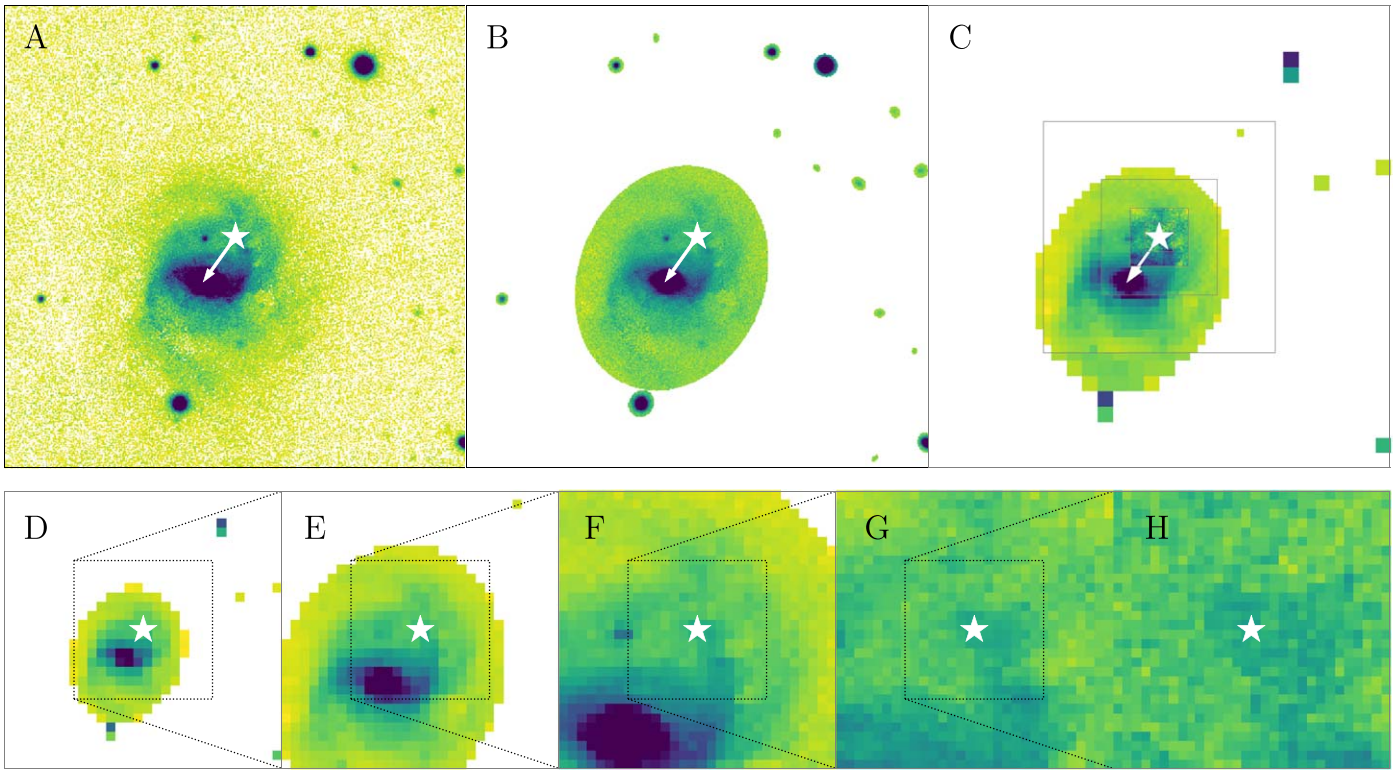
The distribution of host-galaxy angular semimajor axes and angular separations between the transient and the host-galaxy center is shown in Figure 1. We measured angular semimajor axes using the method explained in Appendix A. Note that the isophotal semimajor axis is approximately 3 times the SExtractor reported semimajor axis (Bertin & Arnouts 1996). One can see that the ALeRCE sample is much larger, but that it is dominated by smaller separations as its host galaxies tend to be more distant and smaller in angular size. The Asiago catalog includes transients detected by shallower surveys and in a much larger time span, and the HECATE catalog only contains galaxies within 200 Mpc. By combining both samples we attempt to reduce the imbalance between infrequent transients that have large-angular-size hosts and more frequent transients that have smaller-angular-size hosts.

We further address the imbalance between the sizes of the ALeRCE and nearby samples by selecting a loss function that naturally gives more weight to larger apparent size galaxies and by generating balanced batches when training the model (see Section 4.1).

<sup>28</sup> <https://snhunter.alerce.online/>

<sup>29</sup> <https://github.com/cds-astro/ipyaladin>

<sup>30</sup> <https://www.wis-tns.org/>



**Figure 2.** Training set preparation process: (A) initial PanSTARRS  $r$ -band  $480 \times 480$  pixels ( $0''.25 \text{ pixel}^{-1}$ ) image centered around the position of the transient candidate (white star, in this example SN2004eh), where the output of the model should be the two-dimensional vector connecting the transient and the host-galaxy center (white arrow); (B) the images are sky-subtracted, and sources are detected using SExtractor, masking out all pixels that are not within four elliptical dimensionless radii of any source; (C) the image is resampled into five images of  $30 \times 30$  pixels, all centered around the transient position, but with different fields of view, going from larger to smaller scales as shown in frames (D)–(H). Note that the final multiresolution image requires 3.5% of the memory space used by the original image and about 60 ms per source to generate.

### 3.1. Input Images

For every transient candidate, we obtain a  $120'' \times 120''$  ( $480 \times 480$  pixels with  $0''.25 \text{ pixel}^{-1}$ ) archival  $r$ -band PanSTARRS image centered at its location using the `panstamps` tool.<sup>31</sup> We use only one band in order to ensure a faster image retrieval (it takes about 4 s per image, where each image file size is 920 kB) than would be obtained using multiband images and we use the  $r$  band because it is a better tracer of the bulge, and therefore of the center of galaxies (Kim et al. 2016). The input images are processed using SExtractor with the SEP tool (Barbary et al. 2016) in order to subtract the sky emission and detect significant sources. Pixels that are not within four normalized elliptical radii of any source were masked out (see Section 2.1).

It is important to note that the catalog positions are not perfect, as many times one can observe a deviation from the true position by several arcseconds and sometimes by tens of arcseconds. Thus, all resulting images were visually examined by a team of 12 people in order to remove incorrectly selected hosts and those where the angular separation between the labeled position and the visually estimated host center was thought to be too large. This step is included when defining the sample of 16,791 galaxies.

The  $120'' \times 120''$ ,  $480 \times 480$  pixels, sky-subtracted masked images were used to build sets of multiresolution images made of  $n_{\text{levels}}$  images. This forces the input to be a more compact representation of the data, enabling its direct inclusion in future

enriched alert streams. All  $n_{\text{levels}}$  images are centered around the position of the candidate and have the same size in pixels, they are squared images with side  $480/2^{n_{\text{levels}}-1}$  pixels, but they have different fields of view, with sides  $120''/2^{i-1}$  with  $i \in 1..n_{\text{levels}}$  going from larger to smaller scales. Pixels from level  $i$  can be built taking the average of  $2 \times 2$  pixels at the resolution of level  $i+1$ . All images are normalized between zero and one using the minimum and maximum values between all the levels of a given object. The typical time taken to generate this set is 60 ms per image. An example set of sky-subtracted, masked, normalized, five-level multiresolution images is shown in Figure 2. In this example, the final image memory size is only 32 kB, much less than the original 920 kB. The detailed algorithm used to construct these images can be found in our public GitHub repository at <https://github.com/fforster/delight>.

### 4. Host Identification Method

The host identification method is posed as a regression problem where the input data are the multiresolution images described in Section 3.1 and where the outputs are two-dimensional vectors that connect the transients to their hosts in projection.

The model is based on a convolutional neural network whose architecture is described in Section 4.1. For the loss function we use the mean squared error between the predicted and labeled offset vectors in units of pixels of the highest-resolution image. Using the units of the highest-resolution image gives more weight to larger apparent size galaxies, as a small error

<sup>31</sup> <https://panstamps.readthedocs.io/en/latest/index.html>

relative to the apparent size of the galaxy in those cases will result in a much larger relative contribution to the loss function, addressing the imbalance between the host-galaxy angular sizes of the ALERCE and nearby samples.

In order to make the model invariant to rotations and flips we let the neural network rotate and flip the input images to produce eight different output vectors, which can be averaged after derotating and deflipping to get a final predicted position. Both the dispersion of the previous predictions and the image values at the predicted positions can be used as indicators of when the network may be giving incorrect answers: if the dispersion is too large or if the images at the predicted positions are masked indicates that the predicted position may be incorrect. Both can occur when there are two galaxies that appear equally likely to be the host (see Appendix B).

#### 4.1. Neural Network Architecture

The neural network architecture is based on a convolutional neural network that receives  $n_{\text{levels}}$  input images of the same size in pixels that are centered around the position of the candidate, as described in Section 3.1. Each set of images is rotated by  $90^\circ$  four times, flipping the image with the fourth rotation, and then further rotating them by  $90^\circ$  three more times. This results in seven more images that are an augmented version of the original set.

We trained the network using Tensorflow 2.9.1 and Keras using a mean squared error loss function and Adam optimization (Kingma & Ba 2014). The format of the input images, the architecture of the network, and the training strategy were chosen using  $k$ -fold cross validation in the training set with  $k = 5$ , after performing a train-test split using 4784 samples as the test set.

For the input images, we tried varying the number of levels and how we masked the images, either using a SExtractor mask, a median absolute deviation threshold mask, or no mask. For the architecture, we tried using either shared or independent convolutional layers for all the levels; using rotations and flips or not using them; using a dropout layer with different dropout factors; varying the number of filters and size of the convolution kernels; varying the number and size of the fully connected layers; and changing the type of activation functions. For the training strategy, we tried different learning rates and batch sizes.

After an initial exploration we chose to share convolutional layers among levels, use rotations and flips,  $3 \times 3$  convolution kernels, and rectified linear unit (ReLU), and tanh activation functions for the convolutional and fully connected layers, respectively. Although our analysis was limited to galaxies within 1 arcmin from the transient, we found that including host galaxies further away from the transient (in the corners of a square of side 2 arcmin centered around the transient) during training helped to improve the validation loss at large angular separations by more than a factor of 2.

We then performed a detailed hyperparameter exploration. We assumed 4, 5, or 6 levels for the input images; the three masking strategies described before; a log-uniform distribution for the learning rate, between  $10^{-4}$  and  $10^{-1.5}$ ; a log-uniform distribution for the batch size, between 16 and 64; a log-uniform distribution for the number of convolution filters, between 16 and 64 for each convolutional layer; a log-uniform distribution for the number of neurons in the fully connected

layer, between 128 and 512; and a uniform distribution for the dropout regularization factor, between 0 and 0.4. We explored 200 parameter combinations in Ray Tune (Liaw et al. 2018) using the ASHA scheduler (Li et al. 2020), with up to 50 epochs during training and a grace period of 20 epochs, where a random fold is chosen at each iteration.

We then selected the three best combinations of parameters and used the one with the best average validation loss between the five folds. Upon examination of the errors of the largest-angular-diameter galaxies, we decided to balance each batch during training by oversampling the galaxies from the nearby sample and repeating the hyperparameter search. In this second hyperparameter search we fixed the number of levels to 5 (the best number of layers according to the previous optimization) and used a fixed SExtractor masking strategy. We used the same distributions of parameters previously described, but we changed the range of possible number of neurons in the dense layer to be between 256 and 2048. We again selected the three best combinations of parameters and used the one with the best average loss between the five folds.

Finally, we retrained the network with the previous hyperparameters, but using the entire training set. The predicted positions in the test set were then used for all the comparisons with other host-galaxy identification methods, and the dispersions in the metrics of the five folds were used as metric errors. The best model training batch size was 40, and its learning rate was 0.0014. The resulting best architecture is summarized in Table 1 and Figure 3. The typical prediction time per candidate for a set of multi-resolution images was 60 ms using a CPU.

## 5. Results

### 5.1. Host Association Error

Using the best model described in Section 4.1, we evaluate the results in the test set, using the average of the eight derotated and deflipped output displacement vectors as prediction. The resulting rms deviation between the labeled and predicted positions was  $1''.836 \pm 0''.051$ , the mean deviation was  $0''.783 \pm 0''.009$ , the median deviation was  $0''.468 \pm 0''.008$ , and the mode was  $0''.427 \pm 0''.051$ .

In Figure 4 we show a comparison between the distributions of errors using DELIGHT, SEx, Grad, and DLR methods, compared to the distribution of angular separations between the transient and their host. We use only those transients where a comparison could be made, as in a fraction of the cases some of the methods do not yield a predicted host-galaxy position (2.7% and 9.5% of the cases have no predicted position in DLR and Grad, respectively). In the left and right panels of this figure, we show errors in linear and logarithmic scales, respectively; and in the top and bottom panels, histograms and cumulative distributions of errors, respectively. It can be observed that DELIGHT leads to fewer catastrophic errors in general. DELIGHT is also the least precise method at subarcsecond scales.

In order to understand when the different methods lead to catastrophic errors, we study the distribution of errors as a function of the angular separation to the host. In Figure 5 we show scatter plots of the errors versus the angular separation to the host in the cases where the methods used in this analysis can be compared. It can be seen that most other methods show

**Table 1**  
Summary of the DELIGHT Neural Network Architecture

	Type	Description	Activation	Output	# Parameters
Rotations + flips <sup>a</sup>	Input	Five images of $30 \times 30$ pixels		$5 \times 30 \times 30$	
	Rotation+flip	Rotate and flip input images		$8 \times 5 \times 30 \times 30$	
	Levels <sup>b</sup>	Convolution	ReLU	$52 \times 28 \times 28$	520
		Max-pooling		$52 \times 14 \times 14$	
		Convolution	ReLU	$57 \times 12 \times 12$	26,733
		Max-pooling		$57 \times 6 \times 6$	
		Convolution	ReLU	$41 \times 4 \times 4$	21,074
		Flatten		656	
	Concatenate	Merge five flattened tensors		3,280	
	Fully connected	685 neurons	tanh	685	2,250,766
	Dropout	Regularization of 0.06 at training			
	Fully connected	Two neurons		2	1,372
	Output	Collect & undo rotations/flips		$8 \times 2$	

**Notes.** The five input images are first rotated and flipped to produce eight image sets that are processed independently. Within each set, every image is processed independently by a combination of convolutional and max-pooling layers as shown above, that are then flattened, merged and fully connected into two additional layers of 685 and two neurons, respectively. Finally, all outputs are collected and the rotations and flips are undone to produce a set of eight displacement vectors, which can be averaged to produce a single predicted position.

<sup>a</sup> Independent processing for eight rotated and flipped sets of five images (levels), sharing weights and biases.

<sup>b</sup> Independent processing for five levels, sharing weights and biases.

a small fraction of cases near the identity line, which corresponds to cases where the predicted position is very close to the transient position, but the true host is much further away in the plane of the sky, i.e., the host is confused with a smaller angular size, probably more distant in the line-of-sight galaxy. We do not see this behavior in DELIGHT, suggesting that the proposed method is able to accurately predict the host position and leads to fewer catastrophic errors even in large nearby galaxies with complex geometries.

In Figure 6 we show the two-dimensional distribution of prediction errors versus host semimajor axes in logarithmic scale. We present all the test sample, without excluding examples where the comparison methods fail to predict a position as in Figures 4 and 5. We use hexagonal bins and show those bins containing more than one example, and the rest of the data as a cloud of points. We also include the marginalized distributions in the top and right axes. In this figure, we also indicate the lines where the errors equal the host semimajor axis reported by SExtractor, and where the errors equal the inferred isophotal semimajor axis, a value that is in better agreement with what a person would visually estimate to be the size of the host (see Appendix A). It can be concluded that in galaxies smaller than approximately  $10''$ , the errors are typically below  $1''$  in absolute value, and in larger galaxies the distribution of errors has a significant tail extending up to a few arcseconds in absolute value. This tail of large errors is consistent with the variable pixel size used in the input images, as only large host galaxies can have centers in regions where the resolution is lower and pixel sizes range from  $0''.25$  to  $4''$  from the highest-resolution to the lowest-resolution images, respectively.

Finally, in Appendix B we show 10 random examples in the ALERCE (Figure B1) and nearby samples (Figure B2) to visualize how the proposed method typically performs, as well as the 10 worst normalized error cases (Figure B3) and the 10 worst normalized standard deviation cases of the DELIGHT

predictions (Figure B4) in order to give a sense of when this method fails.

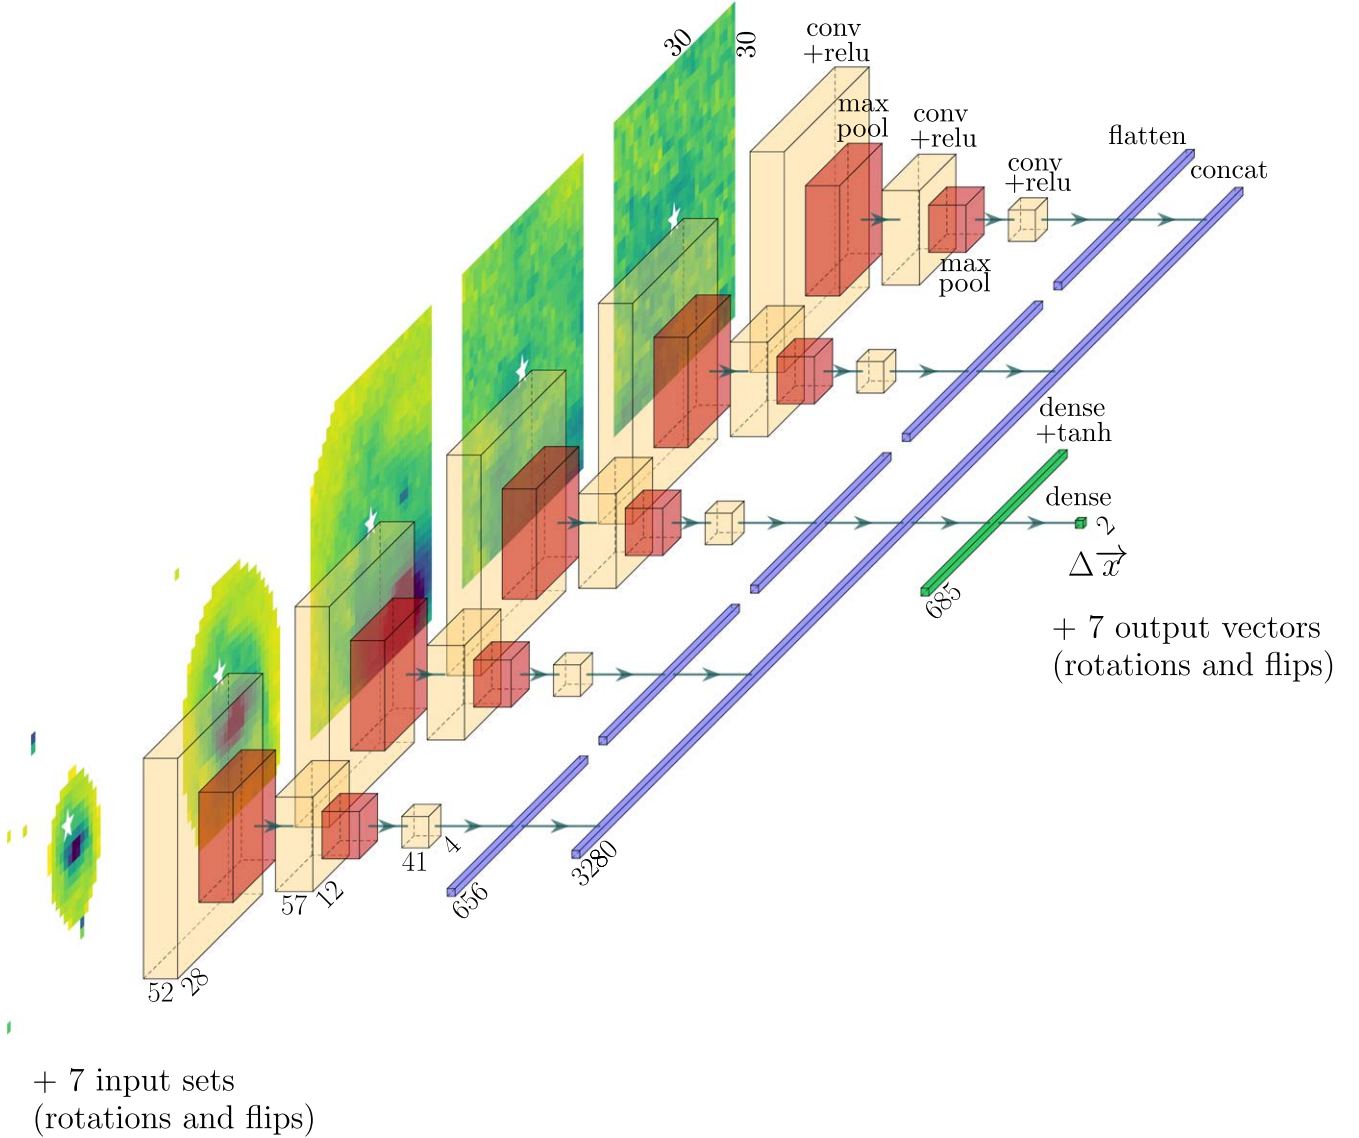
## 5.2. Redshift Association Error

An additional test that was conducted was to study the redshift association error that results when using the predicted host positions and crossmatching them with a catalog of host-galaxy redshifts. This is important because it is generally more important to recover the redshift of the host rather than its exact position. We used the SDSS DR17 Object CrossID service<sup>32</sup> to obtain the closest galaxy source with spectroscopic or photometric redshift within 0.5 from the predicted test-set positions. The number of matches using different methods for host determination is shown in Table 2. Note that DELIGHT is the method that yields more spectroscopic or photometric redshift matches other than using the true host position.

A comparison between photometric and spectroscopic errors suggests that the photometric redshift error is typically of the order of 0.051 (or 0.026 if we reject two outlier samples with  $|\Delta z| > 0.5$ ), as shown in Figure 7. The distribution of all the crossmatched spectroscopic and photometric redshifts in the test set is also shown in Figure 8. A bimodality is observed in the distribution that originates from the combination of both the Nearby and ALERCE samples.

Finally, in order to compare the accuracy in the redshift determination methods, in Figure 9 we show the distribution of redshift errors using different host identification methods and crossmatching with SDSS DR17. The top row shows the crossmatched photometric redshift inferred from the position predicted by the different methods using SDSS DR17 versus its difference with the redshift inferred from the labeled host position using SDSS DR17, indicating the crossmatch completeness at the top left corner. The bottom row shows the histogram of differences in logarithmic scale with bins of  $\Delta z = 0.05$ , indicating the fraction

<sup>32</sup> <https://skyserver.sdss.org/dr17/en/tools/crossid/crossid.aspx>



**Figure 3.** DELIGHT architecture diagram. The input of the network are five  $30 \times 30$  pixel images centered around the position of the candidate with different resolutions and fields of view (see Figure 2). This set of images is rotated and flipped to generate seven more sets used as parallel inputs to the network. Each image within each set is processed by a combination of convolutional layers with  $3 \times 3$  kernel size, no padding, and ReLU activation functions (conv+relu), shown as yellow rectangular cuboids; and  $2 \times 2$  max-pooling layers (max pool), shown as red rectangular cuboids. At this stage we use the same weights and biases for every image given that the problem is approximately scale-free. Each output is then flattened into 656 neuron vectors (purple rectangular cuboids), and the five flattened outputs are concatenated into a layer of 3280 neurons (purple rectangular cuboid). This layer is fully connected with a layer of 685 neurons using a tanh activation function (green cuboid), which is then fully connected to the two neuron output layer (small green cuboid) without activation function representing the predicted host center in pixels. Each of the eight sets of five images produces an output vector, which can be rotated/flipped to produce a cloud of eight predicted host centers per candidate.

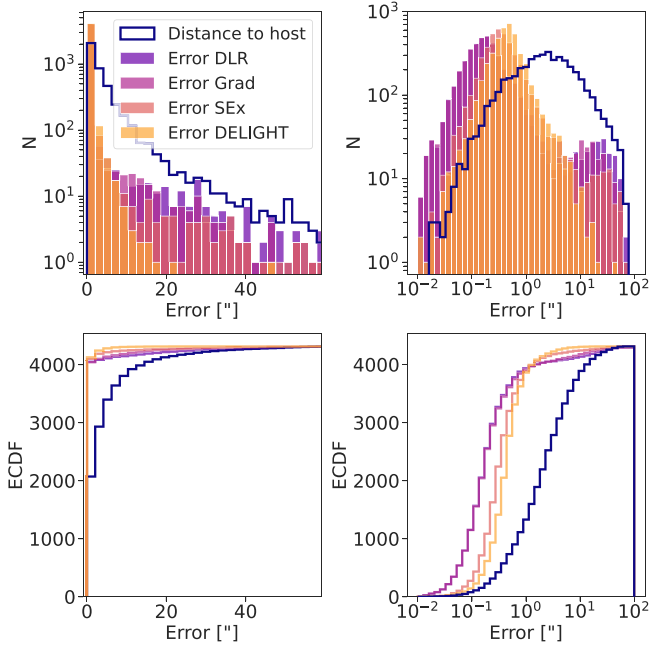
of redshift differences whose absolute value is greater than 0.05 (larger than the dispersion found in Figure 7), a proxy for contamination. Note that in the top row most incorrect redshifts lie close to the identity line and with positive redshift differences, which is what would be expected if large-angular-separation, nearby hosts are being confused with small-angular-separation, more distant, unrelated hosts.

It can be seen that DELIGHT outperforms all the methods in both completeness and contamination in the test set. A summary of the completeness and contamination of the different methods using different redshift types and samples is shown in Table 3. This shows that DELIGHT is the most complete method for all redshift types and samples, and the least contaminated method in almost all

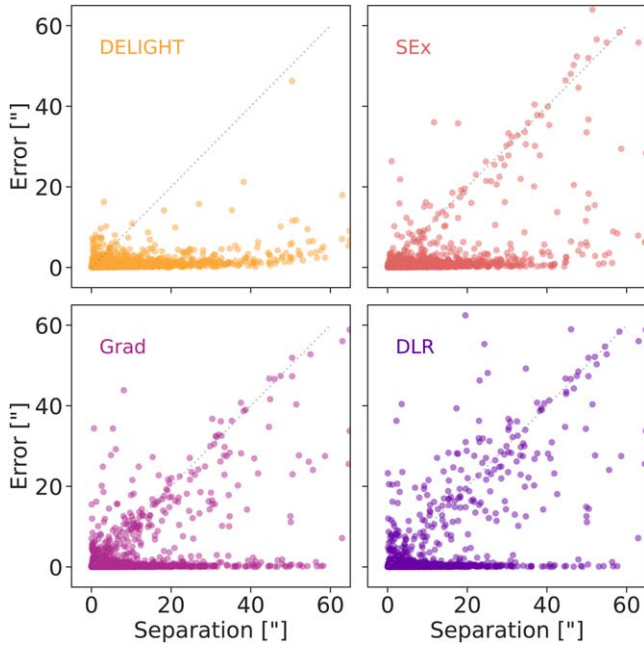
**Table 2**  
Number of Spectroscopic and Photometric Redshift Crossmatches in SDSS DR17 using Different Host Identification Methods as Input Positions

Method	# Spec-z Matches	# Photo-z Matches
True position	1079	3031
DELIGHT	1063	3010
SEx	1041	2999
DLR	988	2889
Grad	940	2752
1-NN	736	2655

**Note.** DELIGHT is the method that yields more spectroscopic and photometric redshift crossmatches other than using the true host position.



**Figure 4.** Distribution of errors between the predicted and labeled positions of the hosts using different methods (see Section 2) compared to the distribution of host angular separations from the position of the transient. We use both linear and logarithmic scales in the errors (left and right, respectively), and a regular and cumulative histogram (top and bottom, respectively).

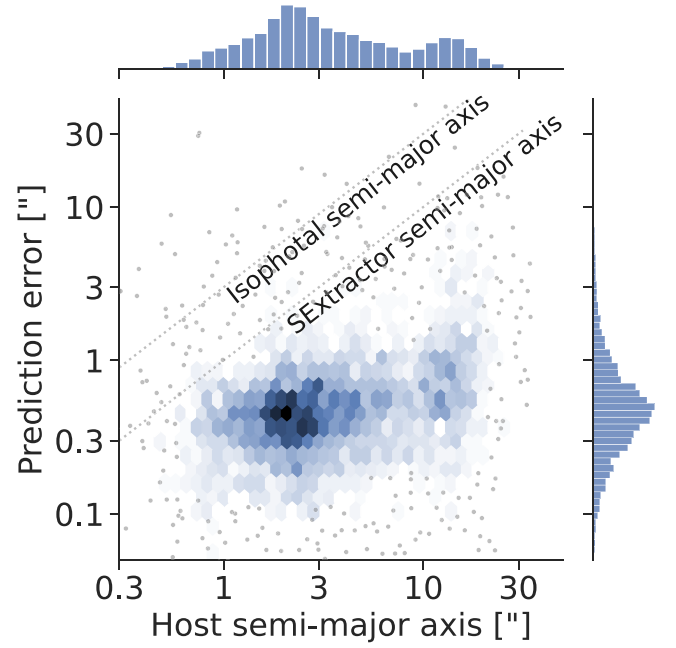


**Figure 5.** Prediction error vs. angular separation distance between the transient candidate and the host-galaxy center for the different methods tested in this work. Catastrophic errors are clustered near the identity line (gray, dotted lines). These errors correspond to cases where a large nearby host galaxy is confused with a smaller, distant galaxy, which is closer in projection.

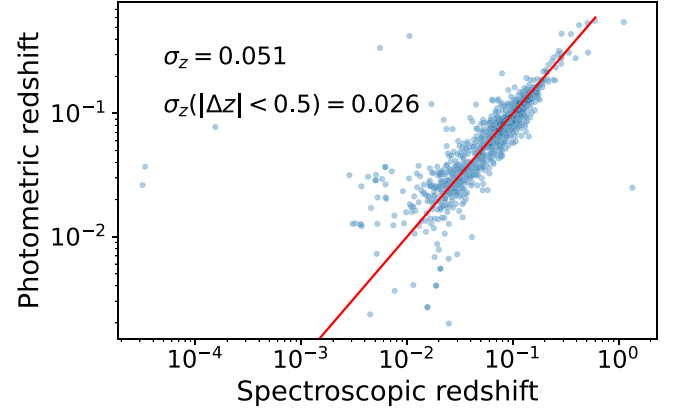
redshift types and samples. Our method is only outperformed by SEx in contamination for the ALerCE sample using photometric redshifts.

## 6. Conclusions

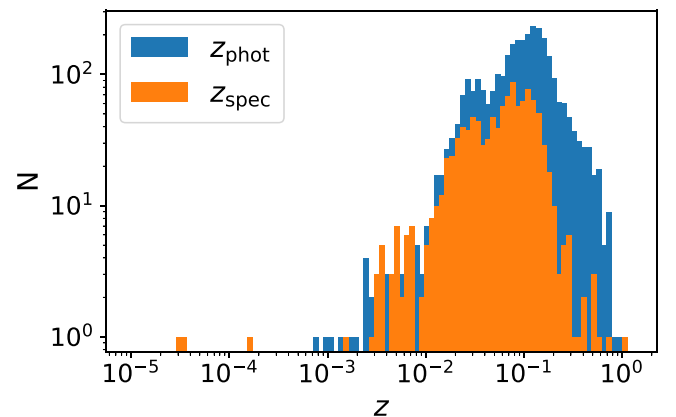
We have introduced DELIGHT, a novel host-galaxy identification method that uses multiresolution archival images



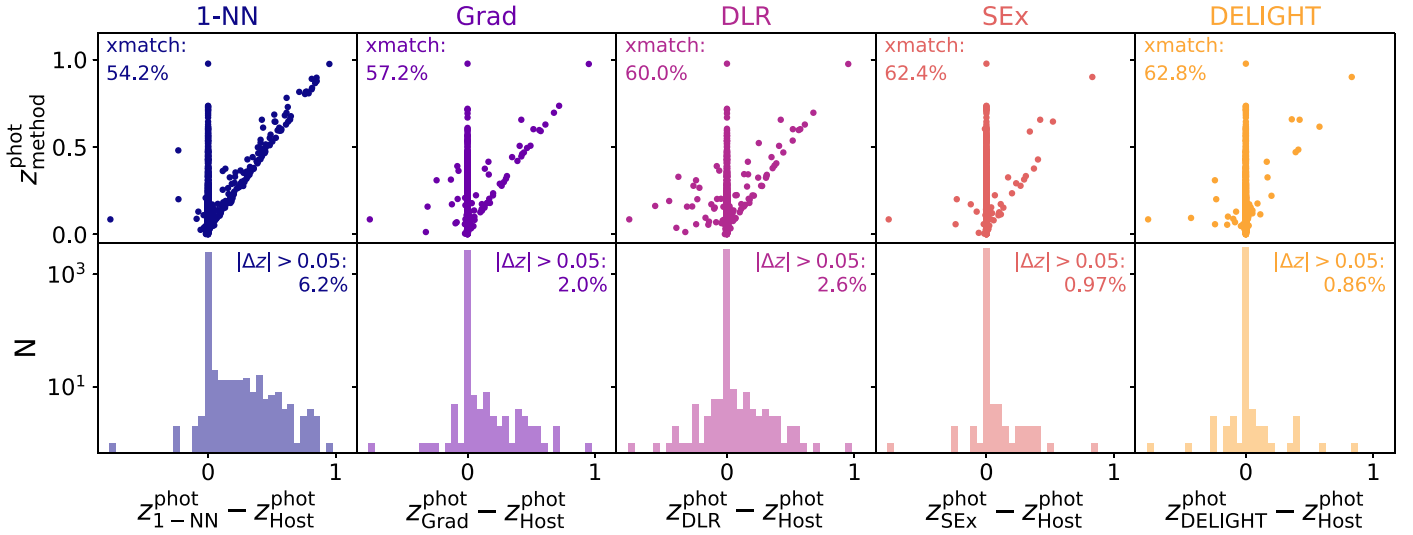
**Figure 6.** Distribution of the DELIGHT prediction errors and SExtractor-inferred host semimajor axes using logarithmic bins. Dotted lines indicate the region where errors would equal the SExtractor semimajor axis and the inferred isophotal semimajor axis (3 times the SExtractor semimajor axis).



**Figure 7.** Comparison between spectroscopic and photometric redshifts in the test-set crossmatched sample. The photometric error dispersion considering all the samples is 0.051, and the dispersion after removing redshift errors with absolute value larger than 0.5 is 0.026.



**Figure 8.** Histogram of test-set spectroscopic and photometric redshifts obtained after crossmatching the labeled host positions with SDSS DR17 using a search radius of 0.5.



**Figure 9.** Crossmatched photometric redshift differences. Top row: scatter plots of the redshifts determined using the method’s predicted position vs. its difference with the redshift determined using the true host position, indicating the fraction of crossmatches with the given method (a proxy for completeness). Bottom row: histograms of the redshift errors, indicating the fraction of errors whose absolute value is larger than 0.05, a proxy for contamination.

**Table 3**  
Completeness and Contamination when Recovering the Host Spectroscopic and Photometric Redshifts using Different Methods and Samples

Method	Spectroscopic Redshifts						Photometric Redshifts					
	Completeness			Contamination			Completeness			Contamination		
	All	ALeRCE	Nearby	All	ALeRCE	Nearby	All	ALeRCE	Nearby	All	ALeRCE	Nearby
True position	22.6	20.4	32.6				63.4	65.9	51.1			
DELIGHT	<b>22.2</b>	<b>20.2</b>	<b>31.7</b>	0	0	0	<b>62.8</b>	<b>65.6</b>	<b>49.9</b>	<b>0.86</b>	0.93	<b>0.48</b>
SEx	21.8	20.1	29.7	0	0	0	62.4	<b>65.6</b>	47.5	0.97	<b>0.66</b>	3.0
DLR	20.7	19.7	25.4	0	0	0	60.0	63.5	43.7	2.6	1.9	7.1
Grad	19.6	19.1	22.0	0	0	0	57.2	61.1	38.6	2.0	1.2	8.1
1-NN	15.4	16.8	8.4	0.14	0	1.4	54.2	60.9	22.9	6.2	3.4	40.8

**Note.** We assume that redshifts that differ by more than 0.05 from the one obtained using the true position are equivalent to a misclassification. We use bold letters for the best metric in each category other than using the true position. Note that DELIGHT outperformed all other methods in completeness and purity, except in purity for the ALeRCE photometric redshift sample using SEx.

centered at the position of extragalactic transient candidates using a convolutional neural network in order to predict the position of the host. We have compared this method to other state-of-the-art host-galaxy identification methods and found that it outperforms all of them in completeness and contamination when recovering host-galaxy redshifts using the predicted position and crossmatching with SDSS DR17, especially with large-angular-diameter, nearby galaxies. Although this method is the least precise in identifying the host-galaxy center at subarcsecond scales, it is the method with fewer catastrophic association errors, which explains the better inferred redshifts.

To the best of our knowledge, this simple but competitive model serves as a first example of the application of machine-learning techniques for the problem of direct and automatic transient host identification based solely on images, yielding the correct host position even if the host galaxy has not been cataloged. The model is currently being used by the ALeRCE broker team to automatically annotate the host galaxies of transient candidates from the ZTF public stream submitted to TNS using archival *r*-band PanSTARRS images and simple cuts to ensure that the predicted positions are matched to the

correct sources in external catalogs, and visually inspecting these annotations in order to ensure that no mismatches are sent to the community. We foresee developing more advanced versions of this model in the future, including returning multiple host candidates in the cases where the identification is degenerate.

Our method is very competitive in terms of speed if multiresolution images are available, taking about 60 ms per transient to determine the host position using a CPU after the image is loaded in memory (cf. 1.1 s in Grad, the other image-based method). This is fast enough to be applied in real time to the entire alert streams of ZTF and LSST with current broker computing capabilities. Additionally, the input images have a size of only 3.5% the size of the input images with the same scale and a constant resolution (32 kB versus 920 kB, respectively), which is smaller than the LSST alert packets (32 kB versus 80 kB), although larger in size to the postage-stamp images contained in these alerts (about 20 kB). This could be decreased even further by removing the central parts of the larger field-of-view images in the multiresolution set (they are redundant with the next, higher-resolution level), by

changing the total field of view and number of levels of the set, or by using different data compression techniques.

While the ZTF postage-stamp images are  $63'' \times 63''$  ( $63 \times 63$  pix), the LSST postage-stamp images are expected to be only  $6'' \times 6''$  ( $30 \times 30$  pix) in size, which will make the problem of automatically classifying postage-stamp images very difficult (see Carrasco-Davis et al. 2021). Thus, we highlight the real possibility of using some form of large field-of-view multi-resolution postage-stamp images for both the problems of rapid postage-stamp image classification and host association in the Vera C. Rubin Observatory alert stream. This would require that the multiresolution images are generated while the raw LSST images are in memory, before the alert packets are distributed to the different alert brokers. Alternatively, a special service for obtaining multiresolution images for LSST transient candidates could aid with the host-galaxy identification, especially in large-angular-separation hosts as a post-processing step close to real time.

A working version of DELIGHT is available via the Python Package Index (PyPI) in <https://pypi.org/project/astro-delight/> and can be installed using the command line `pip install astro-delight`. The source code and example notebooks are available in GitHub in <https://github.com/fforster/delight>. A static version of the code at publication time can be found in Förster et al. (2022).

The authors acknowledge support from the Chilean Ministry of Economy, Development, and Tourism’s Millennium Science Initiative through grant ICN12\_009, awarded to the Millennium Institute of Astrophysics (AMA, FEB, FF, GP, IR, JCO, LHG, MC, PE, PH, PSS, RD), and from the National Agency for Research and Development (ANID) grants: BASAL Center of Mathematical Modeling Grant PAI AFB-170001 (FF, AMA, IR), BASAL project FB210003 (MC, FEB), FONDECYT Regular 1200710 (FF), FONDECYT REGULAR 1200495 (FEB, JCO), FONDECYT REGULAR 1220829 (PAE), FONDECYT REGULAR 1211374 (PH), FONDECYT Postdoctorado 3220449 (RD), National Doctorate Grant 21221393 (SCC), ANID-PFCHA/Doctorado Nacional/2020-21202606 (MR), ANID-PFCHA/Doctorado Nacional/2019-2119188 (DG), and infrastructure funds QUIMAL140003 and QUIMAL190012. We acknowledge support from REUNA Chile. We are thankful to the La Serena School of Data Science (Bayo et al. 2021), where the first proof-of-concept version of this method was implemented. We are thankful for the support provided by Fintual. A.G. is supported by the National Science Foundation Graduate Research Fellowship Program under grant No. DGE-1746047. A.G. further acknowledges funding from the Center for Astrophysical Surveys Fellowship at UIUC/NCSA and the Illinois Distinguished Fellowship. Based on observations by the Zwicky Transient Facility (ZTF; see full acknowledgments in [https://irsa.ipac.caltech.edu/data/ZTF/docs/releases/d%07/ztf\\_release\\_notes\\_dr07.pdf](https://irsa.ipac.caltech.edu/data/ZTF/docs/releases/d%07/ztf_release_notes_dr07.pdf)), the Pan-STARRS1 Surveys (PS1; see full acknowledgments in <https://ipp.ifa.hawaii.edu/>), and the Sloan Digital Sky Survey (SDSS; see full acknowledgments in <https://www.sdss.org/collaboration/citing-sdss/>). This research has made use of the

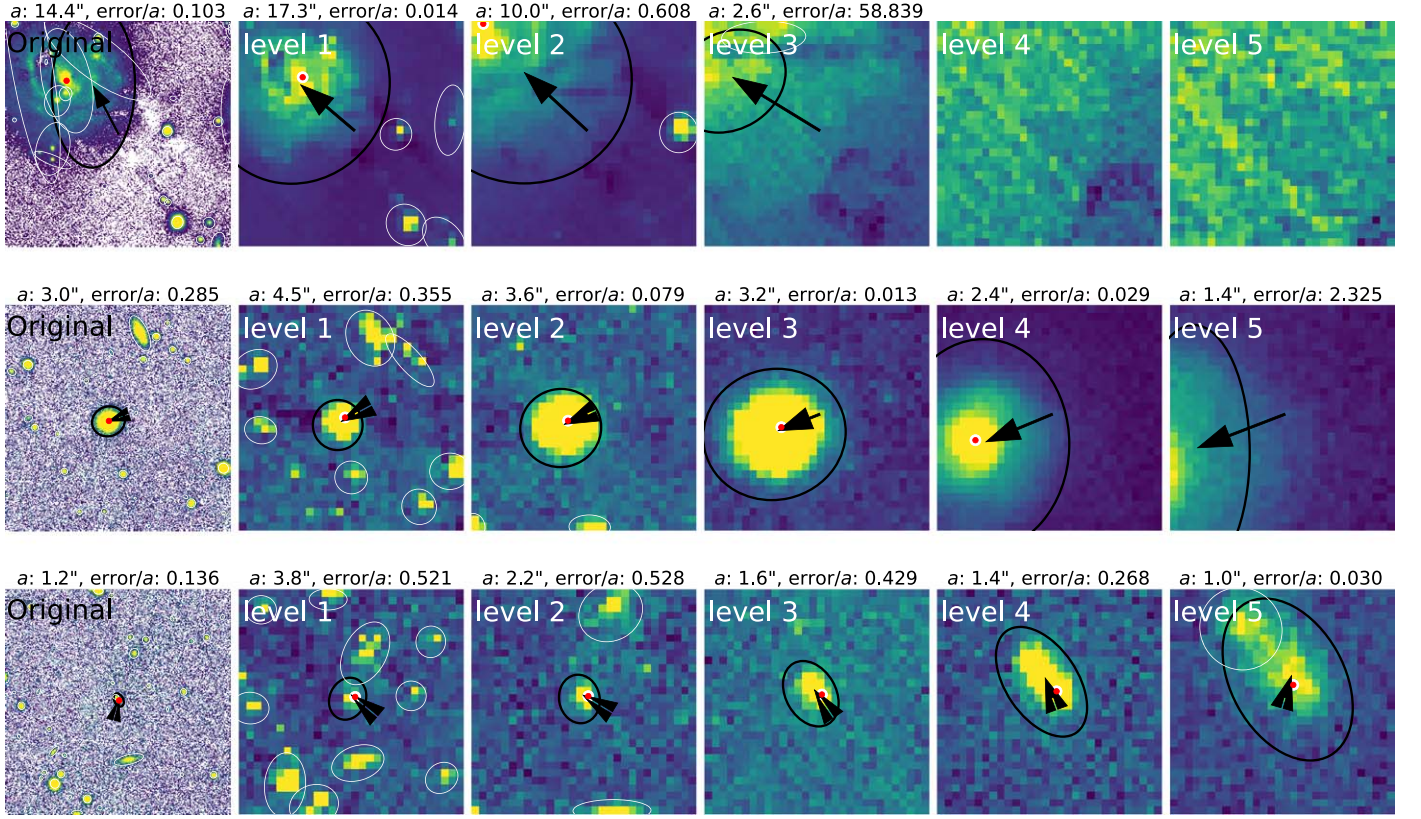
NASA/IPAC Extragalactic Database (NED) and the SIMBAD database.

*Software:* Astropy (Astropy Collaboration et al. 2013, 2018, 2022), ipyaladin (<https://github.com/cds-astro/ipyaladin>), NASA/IPAC Extragalactic Database (NED; 2019), Matplotlib (Hunter 2007), panstamps (<https://github.com/thespacedoctor/panstamps>), Python (Van Rossum & Drake 2009), Ray Tune (Liaw et al. 2018), SEP tool (Barbary et al. 2016), SExtractor (Bertin & Arnouts 1996), TensorFlow (Abadi et al. 2016), xarray (Hoyer & Hamman 2017), and DELIGHT (this work; Förster et al. 2022).

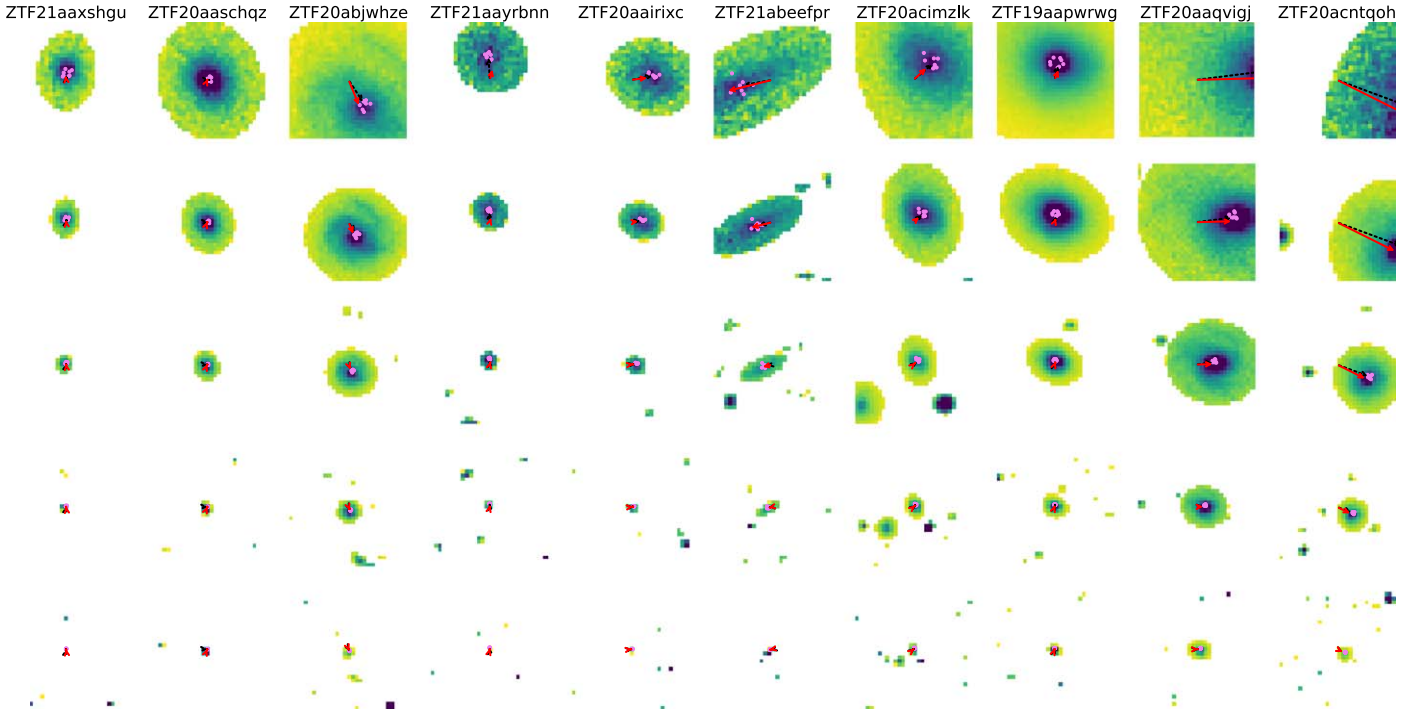
## Appendix A Measuring the Hosts Semimajor Axis using Multiresolution Images

In order to determine the host-galaxy semimajor axis we use SExtractor iteratively in a set of five multiresolution unmasked images centered around the transient. These are constructed in a similar fashion to the method described in Section 3.1, but without any masking. The motivation is to avoid SExtractor from seeing small-scale structure unless it is really needed. Starting from the largest-field-of-view image, we use SExtractor to identify all the sources visible in the image and select the semimajor axis of the one closest to the true galaxy position in normalized elliptical distance. As long as the distance from the inferred position to the labeled host center is larger than 5% the semimajor axis, we repeat the same measurement with a smaller field-of-view level. If the new normalized elliptical distance is smaller than in the previous level, we update the semimajor axis. The resulting semimajor axis is then used as a proxy for host size. The typical processing time per candidate is 150 ms. Three examples of this algorithm are shown in Figure A1, where we show ellipses using the isophotal semimajor and semiminor axes (3 times the SExtractor semimajor and semiminor axes, respectively).

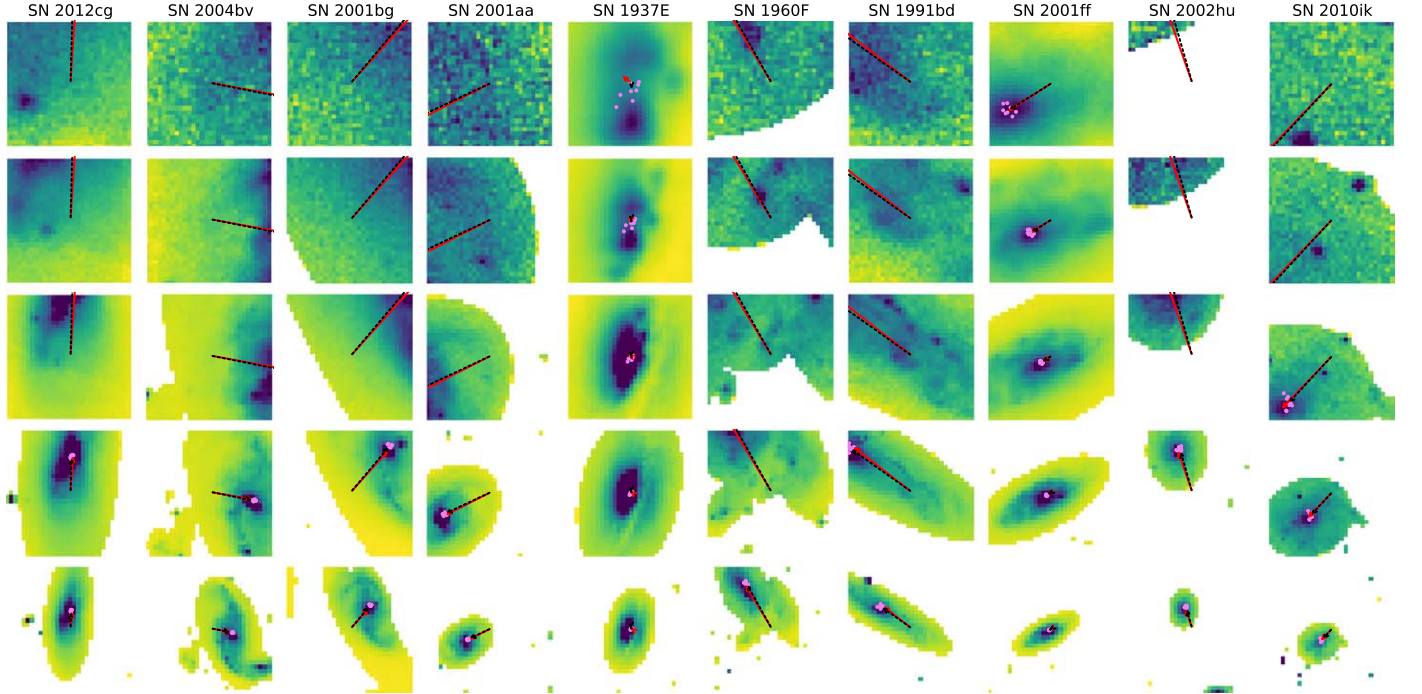
We tested this method via visual inspection of hundreds of cases, and it gives good estimations of the semimajor axis for both small- and large-apparent-size host galaxies, although with some underestimation of the semimajor axis for the largest galaxies (whose size is comparable or larger than the field of view of the image). We found that the SExtractor detected sources were coincident with the labeled host position within a radius of 5% of the semimajor axis in 98% of the cases. Only those cases were used to build Figure 1 to give a more robust representation of the distribution of semimajor axes. Upon visual inspection, we found that in the 2% of cases where this condition was not met, (1) either the host center was outside the field of view, (2) the labeled position was significantly off the visually identified host, or (3) the host was not detectable with SExtractor, suggesting that this method could be used to further clean the sample. In fact, up to 11% of the cases with no crossmatches and 31% of the cases with incorrect redshifts described in Section 5.2 are found within this 2% sample. We decide not to remove these examples to avoid biasing the results of this work to only those host galaxies that are well described with SExtractor ellipses.



**Figure A1.** Three examples of how the semimajor axis of the galaxies were determined. In the top row we show SN2009jy; in the middle row, ZTF19acbwqh; and in the bottom row, ZTF19acbwmd. We show the original 480 × 480 pixel PanSTARRS *r*-band image in the first column and the five central 30 × 30 pixel multiresolution images in the following columns. The SExtractor detected object ellipses are shown, with the ellipse in black having the closest normalized distance from its center to the labeled position. The black arrow points to the best ellipse center, and the white and red dot mark the labeled host position. In the title,  $a$  is the semimajor axis, and  $\text{error}/a$  is the distance between the object’s inferred center and the labeled position divided by the semimajor axis. In the top row, the nearest normalized-distance object detected in level 1 is coincident with the labeled position within 1.4% of the inferred semimajor axis, so the associated semimajor axis, 17''3, is accepted. In the middle row, the nearest normalized-distance object detected in level 3 is coincident with the labeled position within 1.3% of the inferred semimajor axis, so the associated semimajor axis, 3''2, is accepted. In the bottom row, the nearest normalized-distance object detected in level 5 is coincident with the labeled position within 3.0% of the inferred semimajor axis, so the associated semimajor axis, 1''0, is accepted. Note that the semimajor axis tends to be larger for lower-resolution images, that are more sensitive to faint surface brightness regions of the host galaxy. Also note that SExtractor cannot represent the largest galaxy correctly in the original resolution because of all the small-scale structures visible.



**Figure B1.** Comparison between the true host position and the DELIGHT-predicted position in 10 random examples from the ALerCE sample. Each column shows the set of five multiresolution images for a given galaxy, from smallest to largest field of view (highest to smallest resolution) from top to bottom. The red solid arrows represent the true host position (obtained via visual crossmatching with catalog positions in NED, SIMBAD, and SDSS DR16), the black dotted arrows represent the mean DELIGHT-predicted position, and the violet points represent the eight DELIGHT-predicted positions for the eight rotations and flips.



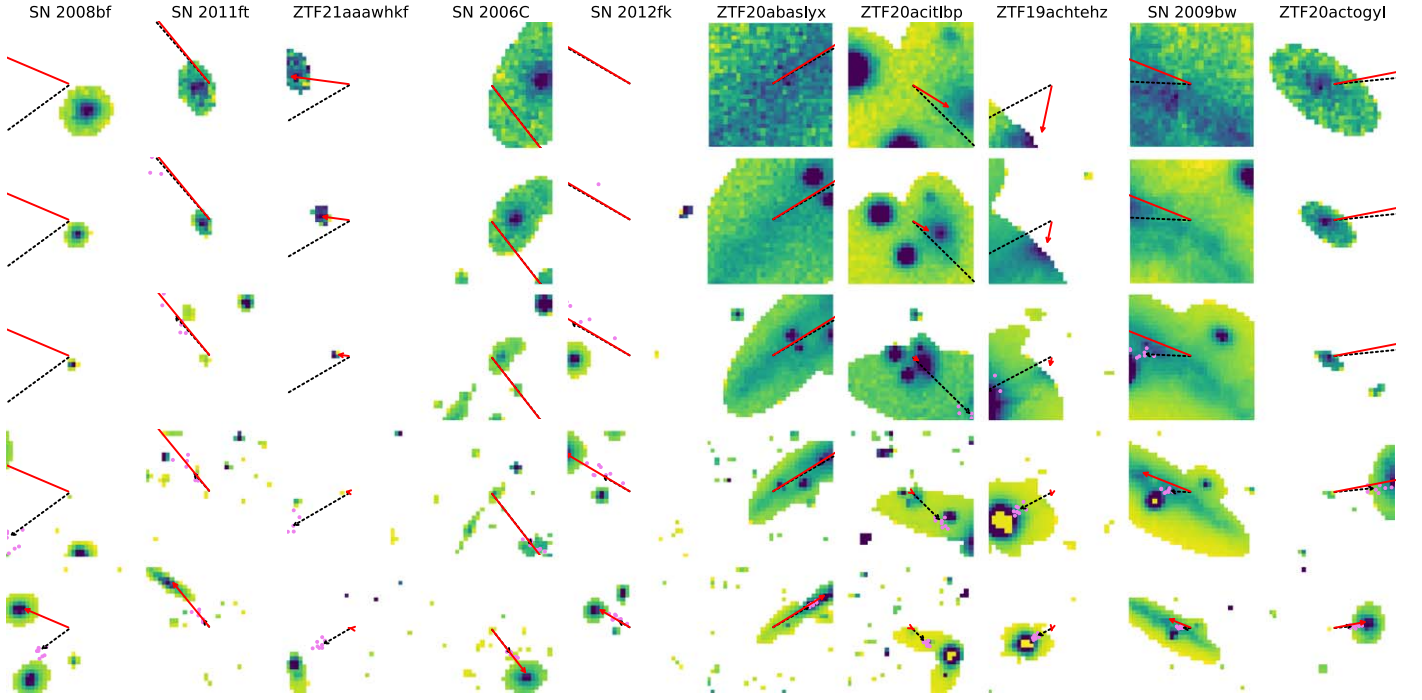
**Figure B2.** Same as Figure B1, but for 10 random examples from the nearby sample.

## Appendix B

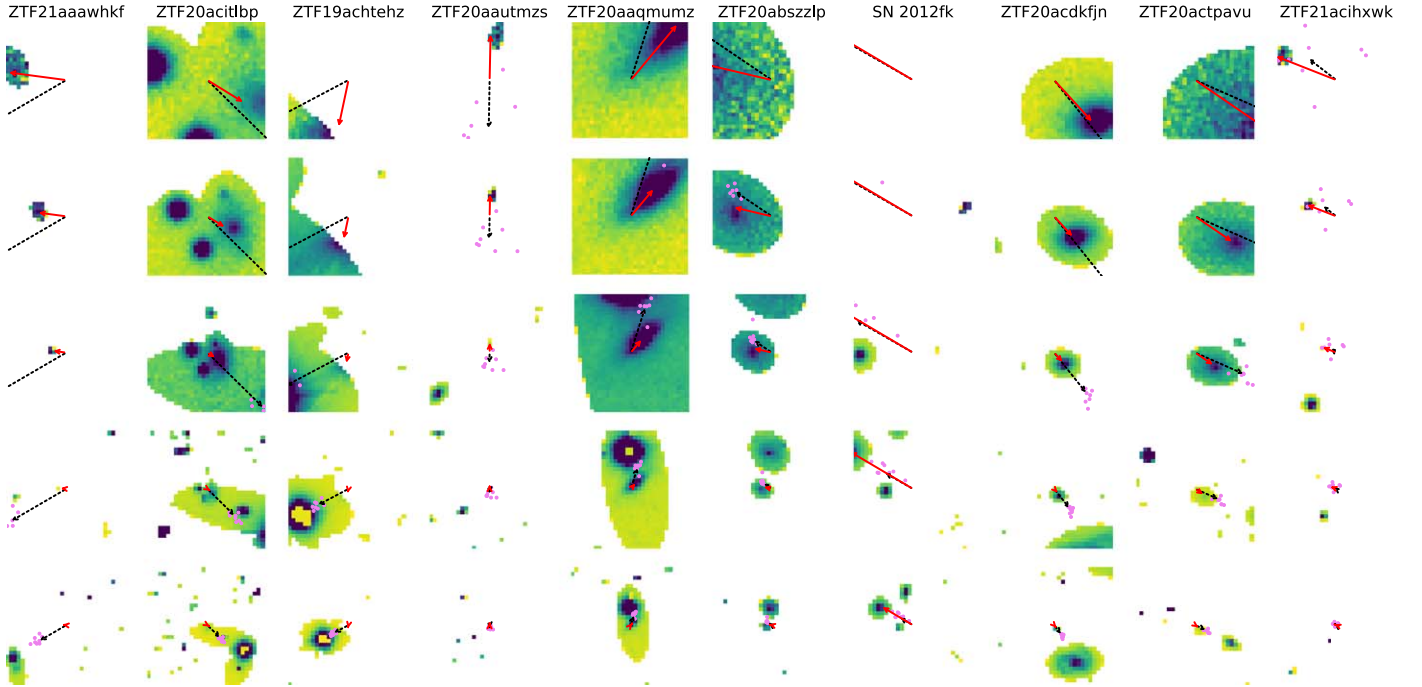
### Examples of Host-galaxy Predicted Coordinates

In this section we show individual predictions from DELIGHT compared to the true host position as visually selected from the NED, SIMBAD, and SDSS DR16 catalogs. First, in Figure B1 we show 10 random examples from the

ALerCE sample of galaxies, and in Figure B2 10 random examples from the nearby sample of galaxies. Then, in Figure B3 we show the 10 worst examples in terms of errors normalized by angular separation, and in Figure B4 the 10 worst examples in terms of standard deviation of the eight DELIGHT-predicted positions normalized by angular



**Figure B3.** Same as Figure B1, but for the 10 examples with the largest errors normalized by angular separation among galaxies with an angular separation larger than  $2''.5$ . Note that large normalized errors are seen in cases where there is a degeneracy between two or more galaxies as potential hosts, such as in SN2008bf, SN20011ft, ZTF21aaawhkf, and SN2012fk; when SExtractor masks H II regions in a given galaxy as different extended sources, such as in SN2006C and ZTF20actogyl; and when there are bright stars close to or in the line of sight of the real host, such as in ZTF20abaslyx, ZTF20acitlbp, ZTF19achtehz, and SN2009bw.



**Figure B4.** Same as Figure B1, but for the 10 examples with the largest standard deviation of the DELIGHT predictions normalized by angular separation, among galaxies with an angular separation larger than  $2''.5$ . They tend to be similar cases to those found in Figure B3, with four of these largest normalized standard deviations also being among the top 10 worst cases in terms of normalized errors.

separation. In the last two cases we restrict the sample to galaxies with a transient to host separation of at least  $2''.5$  to avoid division by small numbers. In almost all cases DELIGHT is able to correctly identify the host position, as

also shown in Figures 4 and 5, and for those cases in which it fails, the normalized standard deviation of the DELIGHT predictions can be used as a proxy for the normalized error.

## ORCID iDs

Francisco Förster  <https://orcid.org/0000-0003-3459-2270>  
 Alejandra M. Muñoz Arancibia  <https://orcid.org/0000-0002-8722-516X>  
 Ignacio Reyes-Jainaga  <https://orcid.org/0000-0003-3627-0216>  
 Alexander Gagliano  <https://orcid.org/0000-0003-4906-8447>  
 Dylan Britt  <https://orcid.org/0000-0001-9905-601X>  
 Sara Cuellar-Carrillo  <https://orcid.org/0000-0001-9532-9906>  
 Felipe Figueroa-Tapia  <https://orcid.org/0000-0001-9735-6051>  
 Ava Polzin  <https://orcid.org/0000-0002-5283-933X>  
 Javier Arredondo  <https://orcid.org/0000-0002-2045-7134>  
 Javier Correa-Orellana  <https://orcid.org/0000-0003-1643-0032>  
 Amelia Bayo  <https://orcid.org/0000-0001-7868-7031>  
 Franz E. Bauer  <https://orcid.org/0000-0002-8686-8737>  
 Márcio Catelan  <https://orcid.org/0000-0001-6003-8877>  
 Guillermo Cabrera-Vives  <https://orcid.org/0000-0002-2720-7218>  
 Raya Dastidar  <https://orcid.org/0000-0001-6191-7160>  
 Pablo A. Estévez  <https://orcid.org/0000-0001-9164-4722>  
 Giuliano Pignata  <https://orcid.org/0000-0003-0006-0188>  
 Lorena Hernández-García  <https://orcid.org/0000-0002-8606-6961>  
 Pablo Huijse  <https://orcid.org/0000-0003-3541-1697>  
 Esteban Reyes  <https://orcid.org/0000-0003-3455-9358>  
 Paula Sánchez-Sáez  <https://orcid.org/0000-0003-0820-4692>  
 Jonathan Pineda-García  <https://orcid.org/0000-0003-0737-8463>

## References

- Abadi, M., Agarwal, A., Barham, P., et al. 2016, arXiv:1603.04467  
 Abdurro'uf, Accetta, K., Aerts, C., et al. 2022, *ApJS*, **259**, 35  
 Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, **935**, 167  
 Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, **156**, 123  
 Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, **558**, A33  
 Barbary, K., Boone, K., McCully, C., et al. 2016, Kbarbary/Sep: V1.0.0, Zenodo, doi:10.5281/zenodo.159035  
 Barbon, R., Buondì, V., Cappellaro, E., & Turatto, M. 1999, *A&AS*, **139**, 531  
 Bayo, A., Graham, M. J., Norman, D., et al. 2021, in Proc. IAU 367, Education and Heritage in the Era of Big Data in Astronomy: The First Steps of the IAU 2020-2030 Strategic Plan, ed. R. M. Ros et al. (Cambridge: Cambridge Univ. Press), 458  
 Bellm, E. C., Kulkarni, S. R., Graham, M. J., et al. 2019, *PASP*, **131**, 018002  
 Bertin, E., & Arnouts, S. 1996, *A&AS*, **117**, 393  
 Breedt, E. 2019, in Southern Horizons in Time-Domain Astronomy, ed. R. E. Griffin, Vol. 339 (Cambridge: Cambridge Univ. Press), 12  
 Carrasco-Davis, R., Reyes, E., Valenzuela, C., et al. 2021, *AJ*, **162**, 231  
 Cirasuolo, M., Afonso, J., Carollo, M., et al. 2014, *Proc. SPIE*, **9147**, 91470N  
 de Jong, R. S., Bellido-Tirado, O., Chiappini, C., et al. 2012, *Proc. SPIE*, **8446**, 84460T  
 DESI Collaboration, Aghamousa, A., Aguilar, J., et al. 2016, arXiv:1611.00036  
 Förster, F., Arancibia, A. M. M., Reyes, I., et al. 2022, Deep Learning Identification of Galaxy Hosts in Transients (DELIGHT), Zenodo, doi:10.5281/ZENODO.7049083  
 Förster, F., Cabrera-Vives, G., Castillo-Navarrete, E., et al. 2021, *AJ*, **161**, 242  
 Gagliano, A., Narayan, G., Engel, A., Carrasco Kind, M. & LSST Dark Energy Science Collaboration 2021, *ApJ*, **908**, 170  
 Gaia Collaboration, Prusti, T., de Bruijne, J. H. J., et al. 2016, *A&A*, **595**, A1  
 Graham, M. J., Kulkarni, S. R., Bellm, E. C., et al. 2019, *PASP*, **131**, 078001  
 Gupta, R. R., Kuhlmann, S., Kovacs, E., et al. 2016, *AJ*, **152**, 154  
 Helou, G., Madore, B. F., Schmitz, M., et al. 1995, in Information & On-Line Data in Astronomy, ed. D. Egret & M. A. Albrecht, Vol. 203 (Dordrecht: Kluwer), 95  
 Hoyer, S., & Hamman, J. 2017, *JORS*, **5**, 10  
 Hunter, J. D. 2007, *CSE*, **9**, 90  
 Ivezić, Ž., Kahn, S. M., Tyson, J. A., et al. 2019, *ApJ*, **873**, 111  
 Kim, K., Oh, S., Jeong, H., et al. 2016, *ApJS*, **225**, 6  
 Kingma, D. P., & Ba, J. 2014, arXiv:1412.6980  
 Kovlakas, K., Zezas, A., Andrews, J. J., et al. 2021, *MNRAS*, **506**, 1896  
 Li, L., Jamieson, K., Rostamizadeh, A., et al. 2020, in MLSys 2020, Vol. 2, ed. I. Dhillon, D. Papailiopoulos, & V. Sze (Indio, CA: MLSys), 230, <https://proceedings.mlsys.org/paper/2020/file/f4b9ec30ad9f68f89b29639786cb62ef-Paper.pdf>  
 Liaw, R., Liang, E., Nishihara, R., et al. 2018, arXiv:1807.05118  
 LSST Science Collaboration, Abell, P. A., Allison, J., et al. 2009, arXiv:0912.0201  
 Möller, A., Peloton, J., Ishida, E. O., et al. 2021, *MNRAS*, **501**, 3272  
 Narayan, G., Zaidi, T., Soraisam, M. D., et al. 2018, *ApJS*, **236**, 9  
 NASA/IPAC Extragalactic Database (NED) 2019, NASA/IPAC Extragalactic Database (NED), IPAC, doi:10.26132/NED1  
 Nordin, J., Brinnet, V., van Santen, J., et al. 2019, *A&A*, **631**, A147  
 Sako, M., Bassett, B., Becker, A. C., et al. 2018, *PASP*, **130**, 064002  
 Simons, D. A., Crampton, D., Côté, P., et al. 2014, *Proc. SPIE*, **9145**, 914515  
 Smith, K. W., Williams, R. D., Young, D. R., et al. 2019, *RNAAS*, **3**, 26  
 Smith, K. W., Smartt, S. J., Young, D. R., et al. 2020, *PASP*, **132**, 085002  
 Sullivan, M., Le Borgne, D., Pritchett, C. J., et al. 2006, *ApJ*, **648**, 868  
 Takada, M., Ellis, R. S., Chiba, M., et al. 2014, *PASJ*, **66**, R1  
 Tonry, J. L., Denneau, L., Heinze, A. N., et al. 2018, *PASP*, **130**, 064505  
 Van Rossum, G., & Drake, F. L. 2009, Python 3 Reference Manual (Scotts Valley, CA: CreateSpace)  
 Wenger, M., Ochsenbein, F., Egret, D., et al. 2000, *A&AS*, **143**, 9