Wildlife@Home: Combining Crowd Sourcing and Volunteer Computing to Analyze Avian Nesting Video

Travis Desell*, Robert Bergman*, Kyle Goehner*, Ronald Marsh*, Rebecca VanderClute*, and Susan Ellis-Felege†
*Department of Computer Science
tdesell@cs.und.edu, robert.bergman@und.edu, kyle.goehner.2@und.edu, rmarsh@cs.und.edu, rebecca.vanderclute@und.edu
†Department of Biology
susan.felege@und.edu
University of North Dakota

***Abstract*—New camera technology is allowing avian ecologists to perform detailed studies of avian behavior, nesting strategies and predation in areas where it was previously impossible to gather data. Unfortunately, studies have shown mechanical triggers and a variety of sensors to be inadequate in capturing footage of small predators (*e.g.*, snakes, rodents) or events in dense vegetation. Because of this, continuous camera recording is currently the most robust solution for avian monitoring, especially in ground nesting species. However, continuous video footage results in a *data deluge*, as monitoring enough nests to make biologically significant inferences results in massive amounts of data which is unclassifiable by humans alone. In the summer of 2012, Dr. Ellis-Felege gathered video footage from 63 sharp-tailed grouse (*Tympanuchus phasianellus*) nests, as well as preliminary interior least tern (*Sternula antillarum*) and piping plover (*Charadrius melodus*) nests, resulting in over 20,000 hours of video footage. In order to effectively analyze this video, a project combining both crowd sourcing and volunteer computing was developed, where volunteers can stream nesting video and report their observations, as well as have their computers download video for analysis by computer vision techniques. This provides a robust way to analyze the video, as user observations are validated by multiple views as well as the results of the computer vision techniques. This work provides initial results analyzing the effectiveness of the crowd sourced observations and computer vision techniques.**

I. Introduction

Camera studies have become popular tools in the field of avian ecology as they can dramatically reduce researcher impacts on behavior and monitor animals in remote locations [1], [2]. However, many of these studies have been hampered by small sample sizes, where few have studied more than 100 nests [2], limiting the biological inferences that could be made due to lack of statistical significance. This limitation has been in part due to the lack of tools to swiftly analyze large amounts of video footage. In the summer of 2012, numerous cameras were set up across the western part of North Dakota, gathering over 20,000 hours of video footage of sharp-tailed grouse (*Tympanuchus phasianellus*), approximately 6 terabytes of data. In addition, 213 hours of test video has also been gathered for the interior least tern (*Sternula antillarum*), federally listed as an endangered species, and 682 for the piping plover (*Charadrius melodus*), federally listed as a threatened species. There are further plans to monitor these birds in future nesting seasons, which should result in a total of over 100,000 hours of video. The sharp-tailed grouse is considered an *indicator species*, meaning that the success of the species is closely tied to the health of the wildlife in the area. An analysis of this video will not only result in a wealth of biological knowledge about these species, but can also be used to examine the impacts of oil development in western North Dakota.

There are significant challenges in developing a purely computational analysis of this wildlife video, as shown in Figure 1). The species being studied, along with many of their predators, have evolved with *cryptic coloration*, or camouflage, making it difficult to distinguish them from their surroundings. Further, the video is taken from uncontrolled outdoor settings, with vegetation moving in the wind and changing weather conditions. Footage is recorded continuously with daytime video captured in color. Infrared light emitting diodes (LEDs) are used in low light and night conditions and recordings during this time are in black and white. This results in a wide variety of video quality and color.

This paper presents initial work developing a citizen science project called Wildlife@Home, which combines both *volunteer computing*, where people volunteer their computers to different computing projects, and *crowd sourcing*, where people volunteer their brain power, to aid in the analysis of this vast amount of video. To our knowledge, no citizen science project has combined the two, and crowd sourcing projects involving the analysis of video are limited. Wildlife@Home was used to compare the results of preliminary motion and feature detection algorithms to the validated observations made by users, and was able to detect a noticeable signal for the presence of piping plover using feature detection, and active events involving the sharp-tailed grouse video using motion detection; which is significant given the challenges of analyzing this uncontrolled 24-hour outdoor video.

Fig. 1. A sharp-tailed grouse in day, dusk and night conditions (top), and a piping plover in varying light conditions (bottom). Birds are circled in red. Given the cryptic coloration of the bird and lighting conditions, it can be very difficult to distinguish the bird from a rock, grass or other objects.

II. Related Work

A. Detecting Animals in Video

While more commonly used with humans and other objects where there is a large body of work (for surveys see [3], [4]), computer vision has also been successfully used to detect animals and animal events. Significant work with animals has been done in controlled laboratory settings, simplifying the task of gathering video and animal detection. A common approach is to subtract a uniform background from the animals, which has been used to track white mice on black backgrounds [5], [6], [7] or in water [8], [9], [10]. Tracking and detecting behavior of fruit flies (*Drosophilia*) [11], [12], [13] has been done in similar settings. Detection of particular actions or events has also been studied, such as vomiting of musk shrews [14], [15] using non-rigid body contour matching and various actions of a grasshopper [16] using spectral clustering [17].

Research has also been done in uncontrolled lab settings, without background subtraction or environmental controls. Sequential Monte Carlo methods, or particle filters [18], [19], [20], [21], have been used to to provide tracking with resiliency to unpredictable motion and non-linear measurement models, and have been mostly used with insects such as ants [22] and bees [23], [24]. Tracking outlines of animals in their stalls using active contours has been used for larger animals like cows [25] and pigs [26]. Using multiple features (image abstractions such as anatomical and cage characteristics) to track rats in reflective and potentially scratched cages [27] and determine mice behaviors [28], [29] has been successfully used as well. Also of note, Jhaung *et al.* developed a manually annotated video database for training and testing a computer vision system for detecting behavior in mice in cages [30], [31], [32].

Considerably less research has been done using video taken in uncontrolled natural settings. Particle filters have been used to track multiple birds in the sky [33]; and data association methods have tracked and counted extremely large numbers of bats in noisy infrared video, taken as the bats leave their caves at night [34]. Face detection has also been used to classify species of African great apes using footage taken from video traps [35]. In settings most similar to this work, BearCam has been used to detect bears in the arctic circle during four hour daytime periods [36] by taking low level features such as image gradients and background differences and combining them into a mid-level *motion shapelet* [37] using AdaBoost [38].

B. Citizen Science Projects

Volunteer computing has emerged as a viable and significant source of computing power. It is being successfully used to perform research in scientific applications ranging from astronomy [39], [40], [41], biology [42], [43], [44], [45], chemistry [46], and physics [47], [48], to climate modeling [49] as well as many other fields of enquiry. Berkeley's Open Infrastructure for Network Computing (BOINC) [50], [51] is the most widely deployed volunteer computing framework, in part due to its open source code and easy extension.

On the other hand, crowd sourcing has been successfully used by citizen science projects to tackle problems requiring human feedback. GalaxyZoo [52], [53] has had great success in using volunteers to classify galaxies in images from the Sloan Digital Sky Survey [54]; and PlanetHunters [55] has been used to identify planet candidates in the NASA Kepler public release data. More recently, Snapshot Serengeti [56] has been created to classify images from camera traps in the

Serengeti National Park. However, these projects focus on volunteers doing identification and classification of images, not video.

Cornell's NestCams project [57] has provided an outstanding resource for environmental education and gained popularity through the use of nest cameras to attract the public's interest in environmental science. NestCams primarily focuses on public outreach where video is collected opportunistically from cameras installed in bird houses, capturing a variety of cavity-nesting species. The CamClickr project has sparked applications of nest video archives for education in collegiate-level animal behavior courses [58].

III. Combining Crowd Sourcing and Volunteer Computing

Volunteer computing provides the required computing power to successfully analyze the video gathered for this work, while crowd sourcing provides access to enough people for the generation of training data and verification of the computer vision techniques. This combination of *volunteer computing* and *crowd sourcing* is a natural fit for tackling the problem of analyzing large amounts of avian nesting video, given its potential to provide massive amounts of human observation and computing power.

Figure 3 presents a detailed work flow for the databases, server side daemons and web pages utilized to manage the video and information generated by the crowd sourced users and volunteered computers. The following sections describe the details of the crowd sourcing and volunteer computing implementations.

A. Crowd Sourcing

The gathered avian nesting video gathered is converted to three minute segments for streaming to the crowd sourced users. These users can select the species and location they wish to view video from, after which videos are selected by the web services to be streamed to the user. This is done by first selecting other videos of that species from the same location which do have observations, but do not have a validated observation from other users. This is done in order to validate videos requiring another viewing and provide users with credit as fast as possible. If there are no other validated videos for the selected species and location, a random new video will be streamed to the user.

Users can then specify yes, no or unsure for a set of events that provide information for evaluating biological hypotheses about the birds (see Figure 2), and can also leave comments. When a user submits their observations, they are shown the observations made by other users, which aids in the learning process for observing the video, and are awarded credit if their observations are successfully validated. The project uses this credit to maintain a set of leaderboards for the users who have watched the most video. Observations from users are validated when a quorum of similar observations has been reached. Two observations match if there are no conflicting events, *e.g.*, a yes for presence by one user and a no for presence by another user. Unsure classifications count as a match to either yes or no; however a video will continue to be shown to users until either all events have validated non-unsure responses or observations have been reported by 5 different users.

B. Volunteer Computing

BOINC [50], [51] has been used for the volunteer computing framework, as it is easily extensible and already has a large user base which actively discover and participate in new volunteer computing projects. A motion detection application and a feature detection application which uses the SURF algorithm [60] have been implemented. The avian nesting video is recorded by the nest cameras into files which are typically 1.5 to 2 hours log. These are used as the basis for the *workunits* that the volunteered computing hosts process. When a host requests work, it downloads a set of videos to run the motion or feature detection applications on. In the case of feature detection, the host will also download a feature file to match the frames of the video against. In both cases, the applications report a likelihood of motion or the presence of the selected features for each three minute segment of the video, allowing for easy comparison to the observations reported by the crowd sourced users. The BOINC daemons handle validation of these results by sending the same video to multiple hosts, until a quorum of similar results is reached.

1) Motion Detection: The motion detection application utilizes an approach called average window differencing. In this approach every frame is compared to an average of a window of surrounding frames for each pixel component (*e.g.*, RGB or YUV). Pixel information was decoded from the video segments using *libavcodec*. Various window sizes were tested, and for the purposes of this work a window size of 10 seconds of frames (five before and five after) was used, which proved sufficient to minimize the effect of weather changes like clouds moving across the sky, as well as vegetation moving in the wind. For the first five seconds of video and for the last five seconds of video, the average window all the existing frames five seconds before and after (*e.g.*, the last frame would be compared only to the five previous) .

The likelihood of motion is calculated using a placeholder array (which effectively mimics a frame) which stores the sum of the red, green, and blue (RGB) components for each pixel in a frame. Another array is populated with the values from the summation array divided by the number of frames in the window (frames in 10 seconds + 1), representing the average values for the window. The RGB components of the middle frame of the window are subtracted from the calculated average, summed and then divided by the maximum possible response value of a frame (width x height x 3 x 255). This response value (the difference between the middle frame and the average of its window) is then summed for each frame in the three minute segment. By storing the frames in the window and subtracting the oldest frame and adding the newest frame to the placeholder array as a video is processed, this response can be calculated efficiently.

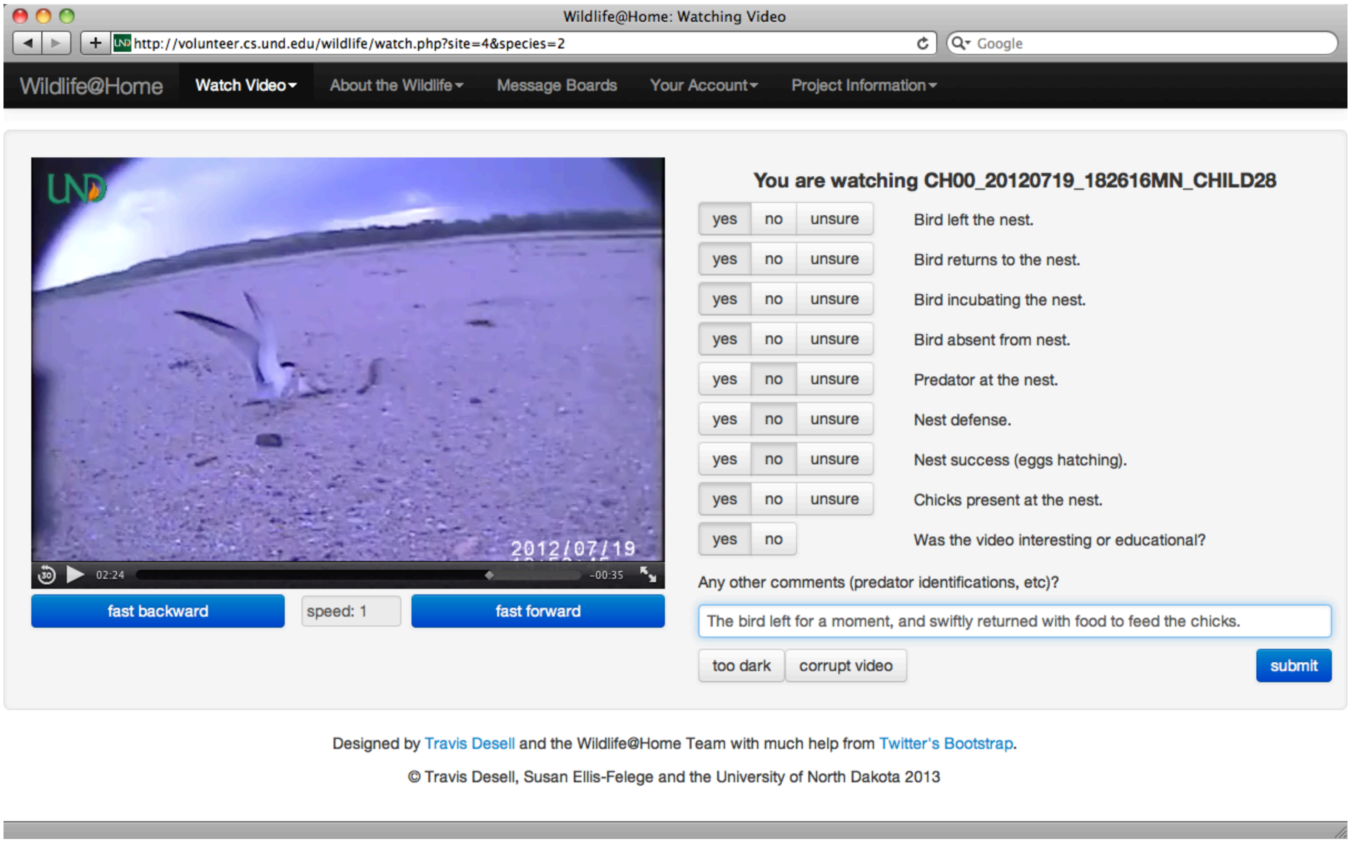

Fig. 2. An example of the interface used for crowd sourcing. The page has been developed using HTML5, jQuery and Bootstrap, allowing for easy use across different devices and web browsers.

2) Feature Detection with SURF: Feature detection consists of two parts. The generation of feature files is done offline by combining common features obtained using the SURF [60] algorithm on selected images of the event to be detected (*e.g.*., multiple angles of a bird at the nest, a bird in flight, an empty nest, etc). These feature files are then used by the feature detection application on videos downloaded by the volunteered hosts, who report a likelihood of the presence those features for each 3 minute segment of video.

a) Generating Feature Files: The feature collection program uses OpenCV [61] to load a video file and OpenCV's SURF implementation to extract the features. When the program loads it displays the first frame of the video and waits for the user to select the region of interest by dragging a box around the section. This generates a series of images containing just the bird to be used for generating the features. As each frame is processed, the newly collected feature descriptors are compared against the previous frames descriptors. Any new features that are smaller than twice a minimum threshold are discarded. This prevents features that are too similar from being added to the feature set. This greatly reduces the feature set size while maintaining a robust set of features. Once all of the features are collected from the video, they are stored in a feature file which can be used on many different videos of the target species and then combined with a feature consolidation program.

The consolidation program takes a series of feature files and combines them into a single feature file, removing any similar features using the same minimum threshold. This creates a single large feature set for matching a specific species of bird, or matching a specific bird position or action, such as flying, sitting, or standing. Once this robust feature set is collected, the main program is run against different videos by the volunteered hosts with a given feature set generate a likelihood of the event represented by that feature set occurring in each three minute segment of that video.

b) Detecting Feature Presence: The feature matching program works by extracting SURF features from each frame and comparing them against the given consolidated feature file. This is the basis for determining whether or not there is a match in the video. The success of this is primarily based on the quality of the feature set and how the matching data points are interpreted. Matched points are returned with a distance value from the given features, which is used to interpret the quality of any given match.

Matched features are not always well correlated. Some have very high difference values and when viewed they are not part of the bird or even the same type of texture as a bird. Others have fairly small difference values, but when viewed on the video do not match up with the bird but instead nearby patches

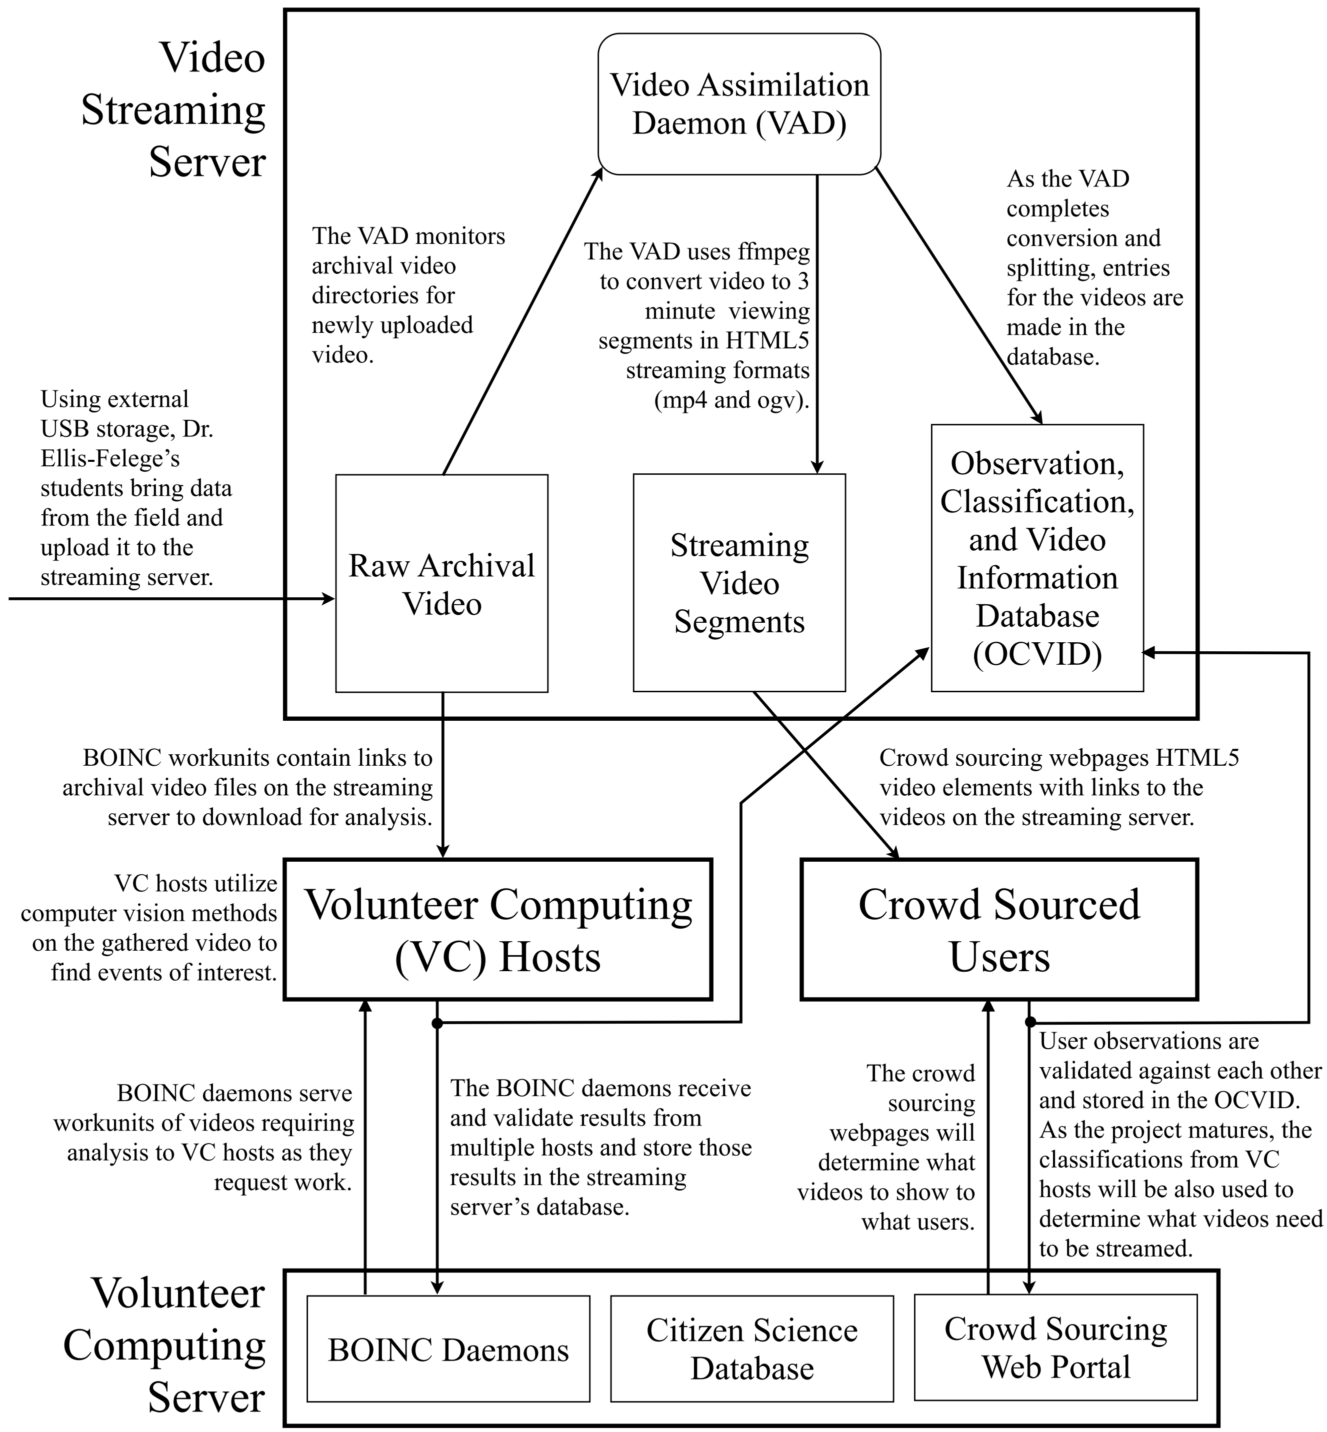

Fig. 3. Wildlife@Home utilizes two servers, two databases and a suite of web pages and server side daemons to manage the 20,000 hours of avian nesting video, as well as the crowd sourcing and volunteer computing results and services. Information related to the videos is stored in the OCVID on the streaming server, while information related to the citizen science services, *e.g.* user and volunteered host information, credit, etc., are stored in the citizen science database on the volunteer computing server.

of dirt or other dark places on the image. In order to handle the first problem of extreme outliers, the minimum distance value for all features from a frame is used and any features that are greater than three times this value are discarded.

With this a approach, if there is a high minimum value many features are matched (but they are of a high distance, so they are generally bad), and if there is a very small minimum value then very few features are matched (with a small distance so they are generally very good). By fitting a rectangle around the matched features, it is possible to get a likelihood of how well the target event was matched. If there are many matched features with high distance values, the rectangle tends to be quite large and does not match the event; on the other hand, if there are a few features clustered together they present a strong match to the event.

The likelihood, l, for each three minute segment of video is calculated as:

$$l = 1 - \frac{R_a}{R_f} \tag{1}$$

Where R_a is the average of the size of each feature bounding rectangle for the three minute period, and R_f is the frame size.

IV. Results

As of May 2013, approximately 70 volunteers have generated over 120 hours of validated observations from the collected video, with over 8400 three minute video segments having been watched with observations reported. Over 100 computers have participated in the project and were used to calculate the preliminary results for motion and feature detection. It should be noted the number of volunteered computers is limited as the client application is currently compiled only for Linux and OSX, however a Windows application is in development.

Motion detection was run across the entire video set (over 20,000 hours). With the motion detection application processing video at 120 frames per second on a typical host, Wildlife@Home and its volunteers were able to get motion likelihood values for the entire data (at 10 frames per second, approximately 1700 compute hours) set in a reasonable amount of time, 4-5 days. The SURF feature detection algorithm runs significantly slower due to its massive computational requirements, running on a similar machine at 1.7 frames per second. Even so, when generating feature detection results for the piping plover video, the 682 hours (at 10 frames per second, approximately 4000 compute hours) were also able to be analyzed in under a week.

A. Motion Detection Analysis

Initial motion detection analysis was done using video containing sharp-tailed grouse. This species was chosen due to the fact that it nests in dense vegetation and has cryptic coloration, making it highly difficult for feature detection to distinguish the nesting grouse from its environment. Over 60 hours of video has been observed and validated for the species, and this was used to examine the results of the motion detection application.

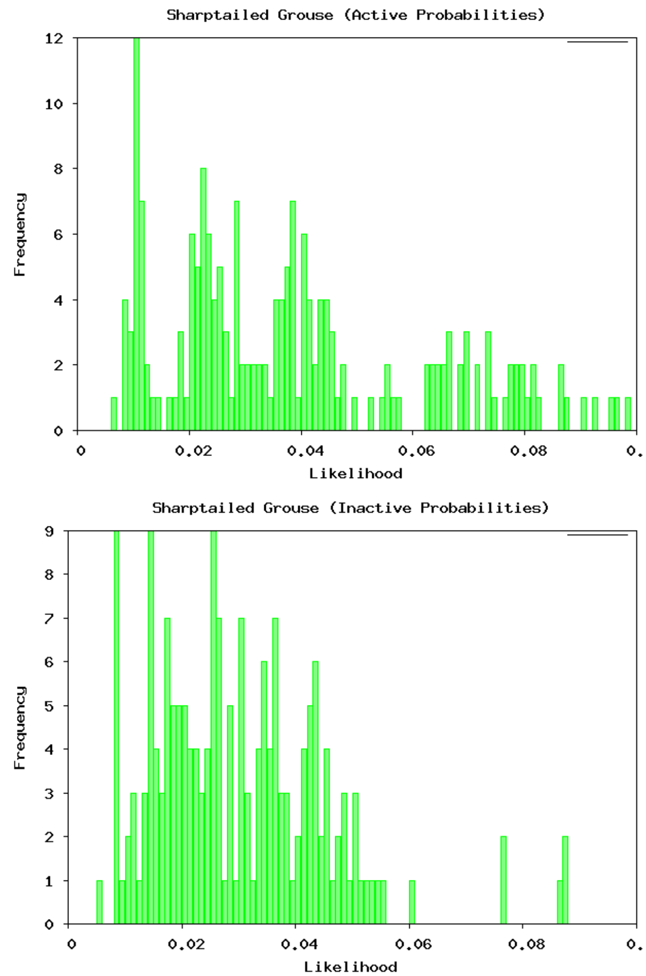

Fig. 4. Histograms for the motion detection likelihoods of video segments validated to contain active and inactive events.

Figure 4 presents the results for the motion detection compared to the users observations. Each three minute video segment with validated user observations was separated into two sets, one containing active events (a bird returning, leaving, presence of a predator, nest defense or if the user marked the video interesting) and the other containing inactive events (a bird simply incubating the nest, not being present, *i.e.*, any video containing no event that would flag it as active).

There were 188 video segments containing active events, and 179 video segments containing no active events. The average and median of the active event videos (0.039 and 0.035) were noticeably higher than the inactive event videos (0.030 and 0.028). Further, there are significantly more video segments with high (greater than 0.05) motion likelihoods containing active events than not. It is also important to note that the video segments occurred across night time, dawn, dusk and daylight, and that many videos contain significant amounts of vegetation moving due to the near constant winds of the plains of North Dakota. While the initial results do not

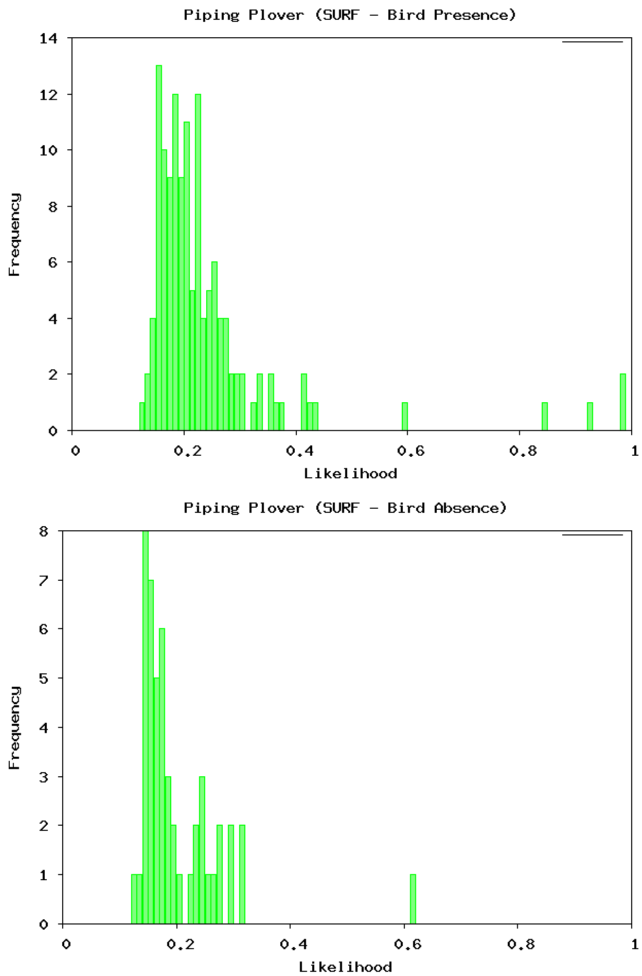

Fig. 5. Histograms for the likelihoods of video segments validated to contain or not contain a piping plover.

provide a strong enough signal to determine active events by computational analysis alone, the fact that there is positive correlation for this motion detection provides motivation to continue to improve upon this technique.

B. Feature Detection with SURF

The initial feature detection analysis was done using video containing piping plover. As this is a shore nesting species, there were significantly less obstructing objects in the landscape, unlike the sharp-tailed grouse video. As such, it provided a good data set for testing of the feature detection application. Over 20 hours of video has been observed and validated for this species, and this was used to examine the results of the feature detection application.

Figure 5 presents the results for detecting the presence of piping plover compared to the users observations. The piping plover shows bi-parental incubation, where one parent will leave the nest and the other will take its place after a short period of time, and therefore many segments contained observations with both the presence and absence of a bird marked. Because of this, the video segments used were restricted to those strictly containing only the presence or absence of a bird (and not both).

There were 133 videos containing bird presence, and 50 with bird absence (as there is bi-parental investment, there are not as many videos without a bird incubating the nest). The average and median of the presence videos (0.24 and 0.21) was also noticeably higher than average and median of the video segments with bird absence (0.20 and 0.17). Similar to motion detection, these videos were from any time of the day, and also contained varying weather conditions; so while they do not provide a strong enough signal to determine bird presence as of yet, the results provide motivation that further refinement of the feature sets and detection algorithm should be able to determine the presence of the birds and other events of interest.

V. Conclusions

This paper provides initial results describing the development of Wildlife@Home, a unique citizen science project which combines both volunteer computing and crowd sourcing. As of May 2013, over 70 volunteers have watched and provided observations on over 8400 three minute video segments, with over 120 hours of observed video having been validated by multiple users. Over 100 volunteered computers were used to calculate motion and feature detection likelihoods on thousands of hours of video to gather an initial analysis of these methods. Preliminary results show noticeable detection of bird presence using feature detection for the piping plover, and for active events within the sharp-tailed grouse video; providing a strong motivation that a citizen science project of this type will be able to successfully analyze the vast amount of avian nesting video. Further, this project can be easily extended and is open to use for video gathered by other wildlife biologists.

We expect to be able to improve the motion detection algorithms using filters and other methods of removing the impact of changing weather and lighting conditions as well as moving vegetation. With these negated it should be possible gain a strong signal for events of interest.

We also wish to investigate better localization of the objects detected by the SURF application. There is still significant noise with the feature matching, and running tests against the quality of the feature collection and the quality of the feature matching should aid in improving the feature detection methods utilized. Another improvement will be to analyze each frame's features based on their standard deviation to remove outliers, while also to giving features with small difference values priority in determining the presence of a bird. As a frame with very good feature matches may almost guarantee the presence of a bird, it is a much more valuable frame than one with poor feature matches. Further, this work lays the framework for examining other computer vision methods, such as SIFT [62] and its variants for detecting events within the video.

As the public is actively engaged in the analysis of this video, the project also provides a strong avenue for public education. Using these and further results, we will be able to perform outreach to future wildlife biologists as well as computer scientists by engaging them in the projects open source code development and in providing ways to better educate the volunteers on how to properly observe the video. As the motion detection and feature detection algorithms improve in accuracy, we also intend to utilize them to determine which videos to show to users; which should greatly increase participation and interest as only videos with interesting events occurring which require human observation will be streamed to them. This work on Wildlife@Home provides a strong example of the scientific and educational opportunities that citizen science can foster.

VI. Acknowledgements

Wildlife@Home has been supported by a collaborative research award from University of North Dakota's Office of Research Development and Compliance. The project's video streaming server is hosted and supported by UND's Computational Research Center and the volunteer computing server is hosted and supported by UND's Scientific Computing Center. We would also like to thank the Wildlife@Home volunteers for their feedback and participation.

References

[1] W. A. Cox, M. S. Pruett, T. J. Benson, J. C. Scott, and F. R. Thompson, "Development of camera technology for monitoring nests," *Video surveillance of nesting birds. Studies in Avian Biology*, vol. 43, pp. 185–210, 2012.

[2] S. N. Ellis-Felege and J. P. Carroll, "Gamebirds and nest cameras: present and future," *Video surveillance of nesting birds. Studies in Avian Biology*, vol. 43, pp. 35–44, 2012.

[3] D. Weinland, R. Ronfard, and E. Boyer, "A survey of vision-based methods for action representation, segmentation and recognition," *Computer Vision and Image Understanding*, vol. 115, no. 2, pp. 224 – 241, 2011. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1077314210002171

[4] T. B. Moeslund, A. Hilton, and V. Kruger, "A survey of advances in vision-based human motion capture and analysis," *Computer Vision and Image Understanding*, vol. 104, no. 23, pp. 90 – 126, 2006. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1077314206001263

[5] H. Ishii, M. Ogura, S. Kurisu, A. Komura, A. Takanishi, N. Iida, and H. Kimura, "Development of autonomous experimental setup for behavior analysis of rats," in *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*, 29 2007-nov. 2 2007, pp. 4152 –4157.

[6] W. Goncalves, J. Monteiro, J. de Andrade Silva, B. Machado, H. Pistori, and V. Odakura, "Multiple mice tracking using a combination of particle filter and k-means," in *Computer Graphics and Image Processing, 2007. SIBGRAPI 2007. XX Brazilian Symposium on*, oct. 2007, pp. 173 –178.

[7] H. Pistori, V. V. V. A. Odakura, J. B. O. Monteiro, W. N. Goncalves, A. R. Roel, J. de Andrade Silva, and B. B. Machado, "Mice and larvae tracking using a particle filter with an auto-adjustable observation model," *Pattern Recognition Letters*, vol. 31, no. 4, pp. 337 – 346, 2010. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0167865509001330

[8] Y. Nie, I. Ishii, K. Yamamoto, K. Orito, and H. Matsuda, "Real-time scratching behavior quantification system for laboratory mice using high-speed vision," *Journal of Real-Time Image Processing*, vol. 4, pp. 181–190, 2009, 10.1007/s11554-009-0111-7. [Online]. Available: http://dx.doi.org/10.1007/s11554-009-0111-7

[9] Y. Nie, I. Ishii, K. Yamamoto, T. Takaki, K. Orito, and H. Matsuda, "High-speed video analysis of laboratory rats behaviors in forced swim test," in *Automation Science and Engineering, 2008. CASE 2008. IEEE International Conference on*, aug. 2008, pp. 206 –211.

[10] T. Mukhina, S. Bachurin, N. Lermontova, and N. Zefirov, "Versatile computerized system for tracking and analysis of water maze tests," *Behavior Research Methods*, vol. 33, pp. 371–380, 2001, 10.3758/BF03195391. [Online]. Available: http://dx.doi.org/10.3758/BF03195391

[11] S. N. Fry, N. Rohrseitz, A. D. Straw, and M. H. Dickinson, "Trackfly: Virtual reality for a behavioral system analysis in free-flying fruit flies," *Journal of Neuroscience Methods*, vol. 171, no. 1, pp. 110 – 117, 2008. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0165027008001210

[12] H. Dankert, L. Wang, E. D. Hoopfer, D. J. Anderson, and P. Perona, "Automated monitoring and analysis of social behavior in drosophila," *Nature Methods*, 2009.

[13] K. Branson, A. Robie, J. Bender, P. Perona, and M. Dickinson, "High-throughput ethomics in large groups of drosophilia," *Nature Methods*, 2009.

[14] D. Huang, K. Meyers, S. Henry, F. De la Torre, and C. Horn, "Non-rigid tracking of musk shrews in video for detection of emetic episodes," in *Computer Vision and Pattern Recognition Workshops (CVPRW), 2011 IEEE Computer Society Conference on*, june 2011, pp. 17 –23.

[15] D. Huang, K. Meyers, S. Henry, F. D. la Torre, and C. C. Horn, "Computerized detection and analysis of cancer chemotherapy-induced emesis in a small animal model, musk shrew," *Journal of Neuroscience Methods*, vol. 197, no. 2, pp. 249 – 258, 2011. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0165027011001270

[16] M. M. Naeini, G. Dutton, K. Rothley, and G. Mori, "Action recognition of insects using spectral clustering," in *In IAPR Conference on Machine Vision Applications*, 2007.

[17] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, pp. 888–905, 2000.

[18] J. Carpenter, P. Clifford, and P. Fearnhead, "Improved particle filter for nonlinear problems," *Radar, Sonar and Navigation, IEE Proceedings -*, vol. 146, no. 1, pp. 2 –7, feb 1999.

[19] F. Dellaert, D. Fox, W. Burgard, and S. Thrun, "Monte carlo localization for mobile robots," in *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, vol. 2, 1999, pp. 1322 –1328 vol.2.

[20] N. Gordon, D. Salmond, and A. Smith, "Novel approach to nonlinear/non-gaussian bayesian state estimation," *Radar and Signal Processing, IEE Proceedings F*, vol. 140, no. 2, pp. 107 –113, apr 1993.

[21] M. Isard and A. Blake, "Contour tracking by stochastic propagation of conditional density," in *Proceedings of the 4th European Conference on Computer Vision-Volume I - Volume I*, ser. ECCV '96. London, UK, UK: Springer-Verlag, 1996, pp. 343–356. [Online]. Available: http://dl.acm.org/citation.cfm?id=645309.648900

[22] Z. Khan, T. Balch, and F. Dellaert, "Mcmc-based particle filtering for tracking a variable number of interacting targets," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 27, no. 11, pp. 1805 –1819, nov. 2005.

[23] A. Veeraraghavan, R. Chellappa, and M. Srinivasan, "Shape-and-behavior encoded tracking of bee dances," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 30, no. 3, pp. 463 –476, march 2008.

[24] Z. Khan, T. Balch, and F. Dellaert, "A rao-blackwellized particle filter for eigentracking," in *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, vol. 2, june-2 july 2004, pp. II–980 – II–986 Vol.2.

[25] A. Cangar, T. Leroy, M. Guarino, E. Vranken, R. Fallon, J. Lenehan, J. Mee, and D. Berckmans, "Automatic real-time monitoring of locomotion and posture behaviour of pregnant cows prior to calving using online image analysis," *Computers and Electronics in Agriculture*, vol. 64, no. 1, pp. 53 – 60, 2008. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0168169908001373

[26] R. Tillett, C. Onyango, and J. Marchant, "Using model-based image processing to track animal movements," *Computers and Electronics in Agriculture*, vol. 17, no. 2, pp. 249 – 261, 1997. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0168169996013087

[27] R. Farah, J. Langlois, and G. Bilodeau, "Rat: Robust animal tracking," in *Robotic and Sensors Environments (ROSE), 2011 IEEE International Symposium on*, sept. 2011, pp. 65 –70.

[28] P. Dollar, V. Rabaud, G. Cottrell, and S. Belongie, "Behavior recognition via sparse spatio-temporal features," in *Visual Surveillance and*

Performance Evaluation of Tracking and Surveillance, 2005. 2nd Joint IEEE International Workshop on, oct. 2005, pp. 65 – 72.

[29] S. Belongie, K. Branson, P. Dollár, and V. Rabaud, "Monitoring animal behavior in the smart vivarium," in *Measuring Behavior*, Wageningen, NL, 2005, pp. 70–72.

[30] H. Jhuang, E. Garrote, N. Edelman, T. Poggio, A. Steele, and T. Serre, "Trainable, vision-based automated home cage behavioral phenotyping," in *Proceedings of the 7th International Conference on Methods and Techniques in Behavioral Research*, ser. MB '10. New York, NY, USA: ACM, 2010, pp. 33:1–33:4. [Online]. Available: http://doi.acm.org/10.1145/1931344.1931377

[31] ——, "Vision-based automated recognition of mice home-cage behaviors." in *Workshop: Visual Observation and Analysis of Animal and Insect Behavior, in conjunction with International Conference on Pattern Recognition (ICPR)*, 2010.

[32] H. Jhuang, E. Garrote, X. Yu, V. Khilnani, T. Poggio, A. D. Steele, and T. Serre, "Automated home-cage behavioural phenotyping of mice." *Nature communications*, vol. 1, no. 6, p. 68, 2010. [Online]. Available: http://www.nature.com/doifinder/10.1038/ncomms1064

[33] D. Tweed and A. Calway, "Tracking many objects using subordinated condensation," in *Proceedings of the British Machine Vision Conference*, P. Rosin and D. Marshall, Eds. BMVA Press, October 2002, pp. 283–292. [Online]. Available: http://www.cs.bris.ac.uk/Publications/Papers/1000674.pdf

[34] M. Betke, D. Hirsh, A. Bagchi, N. Hristov, N. Makris, and T. Kunz, "Tracking large variable numbers of objects in clutter," in *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, june 2007, pp. 1 –8.

[35] A. Ernst and C. Kublbeck, "Fast face detection and species classification of african great apes," in *Advanced Video and Signal-Based Surveillance (AVSS), 2011 8th IEEE International Conference on*, 30 2011-sept. 2 2011, pp. 279 –284.

[36] J. Wawerla, S. Marshall, G. Mori, K. Rothley, and P. Sabzmeydani, "Bearcam: automated wildlife monitoring at the arctic circle," *Mach. Vision Appl.*, vol. 20, no. 5, pp. 303–317, Jun. 2009. [Online]. Available: http://dx.doi.org/10.1007/s00138-008-0128-0

[37] P. Sabzmeydani and G. Mori, "Detecting pedestrians by learning shapelet features," in *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, june 2007, pp. 1 –8.

[38] P. Viola and M. Jones, "Robust real-time object detection," in *International Journal of Computer Vision*, 2001.

[39] N. Cole, "Maximum likelihood fitting of tidal streams with application to the sagittarius dwarf tidal tails," Ph.D. dissertation, Rensselaer Polytechnic Institute, 2009.

[40] N. Cole, H. Newberg, M. Magdon-Ismail, T. Desell, K. Dawsey, W. Hayashi, J. Purnell, B. Szymanski, C. A. Varela, B. Willett, and J. Wisniewski, "Maximum likelihood fitting of tidal streams with application to the sagittarius dwarf tidal tails," *Astrophysical Journal*, vol. 683, pp. 750–766, 2008.

[41] T. Desell, B. A. Willet, M. Arsenault, H. Newberg, M. Magdon-Ismail, B. Szymanski, and C. Varela, "Evolving n-body simulations to determine the origin and structure of the milky way galaxy's halo using volunteer computing," in *The IPDPS'11 Fifth Workshop on Desktop Grids and Volunteer Computing Systems (PCGrid 2011)*, 2011.

[42] T. Desell, L. Newberg, M. Magdon-Ismail, B. K. Szymanski, and W. Thompson, "Finding protein binding sites using volunteer computing grids," in *2nd International Congress on Computer Applications and Computational Science (CACS 2011)*, November 2011.

[43] S. Pellicer, N. Ahmed, Y. Pan, and Y. Zheng, "Gene sequence alignment on a public computing platform," in *Parallel Processing, 2005. ICPP 2005 Workshops. International Conference Workshops on*, june 2005, pp. 95 – 102.

[44] C. Mony, M. Garbey, M. Smaoui, and M.-L. Benot, "Large scale parameter study of an individual-based model of clonal plant with volunteer computing," *Ecological Modelling*, vol. 222, no. 4, pp. 935 – 946, 2011. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0304380010005557

[45] A. L. Beberg, D. L. Ensign, G. Jayachandran, S. Khaliq, and V. S. Pande, "Folding@home: Lessons from eight years of volunteer distributed computing," *Parallel and Distributed Processing Symposium, International*, pp. 1–8, 2009.

[46] I. J. General, E. K. Asciutto, and J. D. Madura, "Structure of aqueous sodium perchlorate solutions," *The Journal of Physical Chemistry B*, vol. 112, no. 48, pp. 15 417–15 425, 2008. [Online]. Available: http://pubs.acs.org/doi/abs/10.1021/jp806269w

[47] C. B. Ries, "Comsolgrid - a framework for performing large-scale parameter studies using comsol multiphysics and the berkeley open infrastructure for network computing (BOINC)," *Applied Sciences*, p. 8, 2010.

[48] J. A. Lopez-Perez, J. Salt, and E. Ros, "Distributed computing and farm management with application to the search for heavy gauge bosons using the atlas experiment at the lhc (cern). oai:cds.cern.ch:1115355," Ph.D. dissertation, Valencia (Spain), University of Valencia, Valencia (Spain), 2008, presented on 21 Jan 2008.

[49] C. Christensen, T. Aina, and D. Stainforth, "The challenge of volunteer computing with lengthy climate model simulations," in *e-Science and Grid Computing, 2005. First International Conference on*, july 2005, pp. 8 pp. –15.

[50] D. P. Anderson, "Volunteer computing: the ultimate cloud," *Crossroads*, vol. 16, no. 3, pp. 7–10, 2010.

[51] D. P. Anderson, E. Korpela, and R. Walton, "High-performance task distribution for volunteer computing." in *The First International Conference on e-Science and Grid Techniques (e-Science 2005)*, 2005, pp. 196–203.

[52] C. J. Lintott, K. Schawinski, A. Slosar, K. Land, S. Bamford, D. Thomas, M. J. Raddick, R. C. Nichol, A. Szalay, D. Andreescu, P. Murray, and J. Vandenberg, "Galaxy zoo: morphologies derived from visual inspection of galaxies from the sloan digital sky survey," *Monthly Notices of the Royal Astronomical Society*, vol. 389, no. 3, pp. 1179–1189, 2008. [Online]. Available: http://dx.doi.org/10.1111/j.1365-2966.2008.13689.x

[53] C. Lintott, K. Schawinski, S. Bamford, A. Slosar, K. Land, D. Thomas, E. Edmondson, K. Masters, R. C. Nichol, M. J. Raddick, A. Szalay, D. Andreescu, P. Murray, and J. Vandenberg, "Galaxy zoo 1: data release of morphological classifications for nearly 900,000 galaxies," *Monthly Notices of the Royal Astronomical Society*, vol. 410, no. 1, pp. 166–178, 2011. [Online]. Available: http://dx.doi.org/10.1111/j.1365-2966.2010.17432.x

[54] J. e. a. Adelman-McCarthy, "The 6th Sloan Digital Sky Survey Data Release, http://www.sdss.org/dr6/," July 2007, apJS, in press, arXiv/0707.3413.

[55] D. A. Fischer, M. E. Schwamb, K. Schawinski, C. Lintott, J. Brewer, M. Giguere, S. Lynn, M. Parrish, T. Sartori, R. Simpson, A. Smith, J. Spronck, N. Batalha, J. Rowe, J. Jenkins, S. Bryson, A. Prsa, P. Tenenbaum, J. Crepp, T. Morton, A. Howard, M. Beleu, Z. Kaplan, N. vanNispen, C. Sharzer, J. DeFouw, A. Hajduk, J. P. Neal, A. Nemec, N. Schuepbach, and V. Zimmermann, "Planet hunters: the first two planet candidates identified by the public using the kepler public archive data," *Monthly Notices of the Royal Astronomical Society*, vol. 419, no. 4, pp. 2900–2911, 2012. [Online]. Available: http://dx.doi.org/10.1111/j.1365-2966.2011.19932.x

[56] Lion Research Center, University of Minnesota, [Accessed Online, 2012] http://www.snapshotserengeti.org/.

[57] R. Bonney, C. B. Cooper, J. Dickinson, S. Kelling, T. Phillips, K. V. Rosenberg, and J. Shirk, "A developing tool for expanding science knowledge and scientific literacy," *BioScience*, vol. 59, pp. 977–984, 2009.

[58] M. A. Voss and C. B. Cooper, "Using a free online citizen-science project to teach observation and quantification of animal behavior," *American Biology Teacher*, vol. 72, pp. 437–443, 2012.

[59] M. A. Voss, Personal communication, 2012.

[60] H. Bay, T. Tuytelaars, and L. Van Gool, "Surf: Speeded up robust features," *Computer Vision–ECCV 2006*, pp. 404–417, 2006.

[61] "OpenCV," [Accessed Online, 2012] http://opencv.org/.

[62] A. Abdel-Hakim and A. Farag, "Csift: A sift descriptor with color invariant characteristics," in *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, vol. 2. IEEE, 2006, pp. 1978–1983.