

**Stereo-Optic High Definition Imaging: A New Technology to Understand Bird and Bat
Avoidance of Wind Turbines**

Final Technical Report

July 21, 2017

Covering Period: December 15, 2016 to March14, 2017

Approved Project Period: December 15, 2015 to March14, 2017

Recipient: Biodiversity Research Institute
276 Canco Road
Portland, ME 04103

Website: www.briloon.org

Award Number: DE-EE0006803

Working Partners: University of Maine School of Computing and Information Science
New Media at the University of Maine

Cost-Sharing Partners: SunEdison
HiDef Aerial Surveying Limited

PI (s): Evan Adams, Ecosystem Modeler, evan.adams@briloon.org
Melissa Duron, Project Manager, melissa.duron@briloon.org
Wing Goodale, Deputy Director, wing_goodale@briloon.org
Phone: 207-839-7600

DOE Project Team: DOE Contracting Officer - Pamela Brodie
DOE Project Officer - Michael Hahn
Project Monitor - Stephanie Hodge

Department of Energy Office of Energy Efficiency and Renewable Energy Wind and Water Power
Technologies Office

Office of the Assistant Secretary
Energy Efficiency and Renewable Energy
Mail Stop EE-1
Department of Energy
Washington, DC 20585


Acknowledgments: This report is based upon work supported by the U. S. Department of Energy under Award No DE-EE0006803. Any findings, opinions, and conclusions or recommendations expressed in this report are those of the author(s) and do not necessarily reflect the views of the Department of Energy. We are deeply appreciative of the field support provided by SunEdison staff.

STEREO-OPTIC HIGH DEFINITION IMAGING: A NEW TECHNOLOGY TO UNDERSTAND BIRD AND BAT AVOIDANCE OF WIND TURBINES

Final Technical Report

Prepared by

Evan Adams, Biodiversity Research Institute
Steve Burns, HiDef Aerial Surveying Limited
Chris Dorr, University of Maine
Melissa Duron, Biodiversity Research Institute
Andrew Gilbert, Biodiversity Research Institute
Wing Goodale, Biodiversity Research Institute
Reinhard Moratz, University of Maine
Mark Robinson, HiDef Aerial Surveying Limited



Conducted in Partnership with
David Cowan, SunEdison



TABLE OF CONTENTS

List of Acronyms	6
List of Figures.....	7
List of Tables	11
Executive Summary	12
Introduction	13
<i>Potential Applications and Impacts.....</i>	<i>13</i>
<i>Brief Description of Technology.....</i>	<i>14</i>
<i>Commercialization Status and Plans</i>	<i>14</i>
Background	15
<i>Project Objectives.....</i>	<i>15</i>
<i>State of the Art</i>	<i>17</i>
<i>Approaches to Accomplish Objectives.....</i>	<i>18</i>
Results and Discussion	21
Phase 1: Camera System Manufacturing and Lab testing.....	21
<i>Objectives.....</i>	<i>21</i>
<i>Technical Approach</i>	<i>21</i>
<i>Results</i>	<i>21</i>
Phase 2: Preliminary Field-testing.....	23
<i>Objectives.....</i>	<i>23</i>
<i>Technical Approach</i>	<i>23</i>
<i>Results</i>	<i>26</i>
Phase 3: Data Collection and Maintenance at Target Sites.....	28
<i>Objectives.....</i>	<i>28</i>
<i>Technical Approach</i>	<i>28</i>
<i>Results</i>	<i>29</i>
Phase 4: Data Management and Video Processing Improvements	33
<i>Objectives.....</i>	<i>33</i>

<i>Data Management: Technical Approach and Results</i>	<i>33</i>
<i>Improving the System: Technical Approach and Results</i>	<i>34</i>
Phase 5: Video Review and Data Analysis	40
<i>Objectives</i>	<i>40</i>
<i>Testing to Detect Bats: Technical Approach and Results.....</i>	<i>40</i>
<i>Testing to Detect Birds: Technical Approach and Results.....</i>	<i>42</i>
<i>Improving the camera calibration: Technical Approach and Results</i>	<i>59</i>
<i>Improving the 3-D modeling: Technical Approach and Results</i>	<i>62</i>
Accomplishments	68
<i>Summary of Technical Accomplishments</i>	<i>68</i>
<i>Presentations.....</i>	<i>68</i>
<i>Project Milestones and Technical Accomplishments.....</i>	<i>69</i>
Conclusions.....	70
Testing the performance of the camera	70
<i>Receipt and initial testing.....</i>	<i>70</i>
<i>Testing for detection of bats</i>	<i>70</i>
<i>Testing for detection of eagles</i>	<i>70</i>
<i>Testing for overall bird detection probability.....</i>	<i>70</i>
Improving the system.....	71
<i>Hardware.....</i>	<i>71</i>
<i>Software</i>	<i>72</i>
Recommendations	74
<i>Testing the performance of the camera</i>	<i>74</i>
<i>Testing for overall bird detection probability</i>	<i>75</i>
<i>Improving the system</i>	<i>75</i>
Literature Cited	78

LIST OF ACRONYMS

ATOM - Acoustic-Thermographic Offshore Monitoring system

BGEPA – Bald and Golden Eagle Protection Act

BRI – Biodiversity Research Institute

DOE – Department of Energy

ESA – Endangered Species Act

FOV – Field of View

HiDef – HiDef Aerial Surveying, Limited

ITP – Incidental take permit

LCOE – Levelized cost of energy

MBTA – Migratory Bird Treaty Act

MS – Microsoft

nIR – Near-Infrared

RSR – Rotor Swept Zone

SOPo – Statement of Project Objectives

TADS - Thermal Animal Detection System

TIR – Thermal Infrared

TRL – Technology Readiness Level

T&E – Threatened and Endangered

UMaine SCIS – University of Maine Orono School of Computing and Information Science

WS – wing span

WETO – Wind Energy Technologies Office

LIST OF FIGURES

Figure 1. The marinized camera system deployed for testing.....	14
Figure 2. Testing at first site with bird-sized object (held at center of chest).	24
Figure 4. View from camera at second testing location with a Bald Eagle in the field of view.....	26
Figure 5. The change in detection probability with distance for the camera system during daytime of a bird-sized object.	26
Figure 6. The change in detection probability with distance for the camera system during nighttime of a bat-sized object.	27
Figure 7. Camera installation at terrestrial wind farm.....	28
Figure 8. Camera view at turbine test site.	29
Figure 9. Average size of video segments based on time of day showed that nighttime video segments are typically smaller in size.	30
Figure 10. Video segments with lower visibility resulted in smaller video sizes than video segments with clear visibility.....	31
Figure 11. Weather that resulted in lower visibility or lower bird activity resulted in lower average file size than clear days.....	31
Figure 12. Example images of masking differences between footage with little background motion (a) and considerable background motion (b). The second footage chunk was truncated at 273 frames, about half the amount usually saved in a chunk file, due to the excess amount of water coming through the masking.	35
Figure 13. Frame-pair view from the original HiDef viewer.	36
Figure 14. Frame-pair view of image calibration.	37
Figure 15. Image of the new GUI zoomed in on marked gulls. The full extent is visible in the box to the right of the image, with the zoomed portion illustrated with a red box. Marker numbers are adjacent to marker boxes drawn around gulls observed in the video footage.....	39
Figure 16. Nighttime image of a bat captured during field testing at the Reservoir/Eagle Test Site. 40	
Figure 17. An example image showing a Red Bat at 25 m from the camera system on September 9, 2015 in Parker River National Wildlife Refuge.....	41
Figure 18. Bald eagle adult observed on January 13, 2016. Identified as a definite Bald Eagle adult, older than Basic II plumage, at the Eagle Test Site. Pixel measure distance was 27–34 m at 1 st and ID frame. The stereo-optic position distance measurement at that point was 16.8 m.....	47
Figure 19. Bald Eagle adult observed on January 13, 2016. Identified as a probable Bald Eagle adult, plumage greater than basic II. No pixel distance possible. The stereo-optic position distance measurement in the identification frame was 113.9 m.	47

Figure 20. Bald Eagle adult observed on October 30, 2015. Identified as a possible Bald Eagle adult of greater than basic II plumage. Pixel distance measured at 225-271 m in the 1 st frame, 202-243 m in the ID frame. No stereo-optic position distance measurement was calculated.....	47
Figure 21. Bald eagle subadult observed on March 28, 2016. Identified as a definite Bald Eagle subadult, plumage less than Basic III. Pixel distance in first frame measured at 31-40 m, and identified at 26-33 m. The stereo-optic position distance measurement in the first frame was 25.8 m, and 22 m at the identification frame.....	47
Figure 22. Bald Eagle subadult observed on January 6, 2016. Identified as a definite Bald Eagle subadult, plumage less than Basic III. Pixel distance in first frame measured at 46-59 m, and was identified in the same frame. The stereo-optic position distance measurement was at 29.7 m.....	47
Figure 23. Bald Eagle observed on December 28, 2015. Identified as a definite Bald Eagle subadult of unknown plumage. Pixel distance when identified was 54-69 m. The stereo-optic position distance measurement when identified was 44.7 m.....	47
Figure 24. Bald Eagle observed on January 13, 2016. Identified as a probable Bald Eagle subadult of unknown plumage. Pixel distance when identified was 215-258 m. The stereo-optic position distance measurement when identified was 90 m.....	48
Figure 25. Bald Eagle observed on November 13, 2015. Identified as a possible Bald Eagle subadult with plumage less than III. Pixel distance when identified was 223-268 m. The stereo-optic position distance measurement was 123.8 m in the identification frame.....	48
Figure 26. Bald Eagle observed on December 25, 2015. Identified as a definite Bald Eagle of unknown age and plumage. Pixel distance for the first and identified frame was 43-51 m. The stereo-optic position distance measurement was 24.6 m in the identification frame.....	48
Figure 27. Bald Eagle observed on January 13, 2016. Identified as a probable Bald Eagle of unknown age and plumage. The pixel measurement for the identification frame was 47-60 m, and the eagle was measured at 49-63 m in another adjacent frame. The stereo-optic position distance measurement was 46.1 m in the identification frame and the measurement was 45.4 m in the adjacent frame.....	48
Figure 28. Bald Eagle observed on March 21, 2016. Identified as a possible Bald Eagle of unknown age and plumage. No pixel measurements were possible for the first or identification frames, but the eagle was measured at 473-568 m in another frame. The stereo-optic position distance measurement was 158.7 m in the same frame.	48
Figure 29. Distance from camera for 1st detection and identification detection of eagles. Measurements shown are for pixel distance ranges.	50
Figure 30. Outlier box plot of distances from the camera for birds first detected in the camera system and Distance identified to species (AMCR = American Crow, AMGO = American Goldfinch, CANG = Canada Goose). Distances are based on pixel count calculations. Single lines for AMGO and CANG indicate only one observation.....	51

Figure 31. Stereo-optic position estimates of a flock of crows flying through the turbine site. Each crow image represents an observation of a crow in a specific frame. The different colors are used to identify individual crows that were tracked.....	52
Figure 32. Distance from the camera as measured by two methods: the stereo-optic distance algorithm (blue) and using the known wing length of the species (green).....	53
Figure 33. A Box plot of distance from camera measured using pixel distance calculations and stereo-optic position distance estimations for birds at the Eagle Test and turbine sites. AMRO is american robin, AMGO is American Goldfinch, AMCR is American Crow, RBGU is Ring-billed Gull, HERG is Herring Gull, COLO is Common Loon, CANG is Canada Goose; in order from smallest to largest average animal size.....	54
Figure 34. The percentage of birds identified to the group or individual species level, broken out by weather condition during observations. Footage with no animals observed is also included (empty).....	55
Figure 35. Detection curve for Bald Eagles observed by the camera system (green) with 95% credible interval (grey). The theoretical identification limit for eagles is shown in blue.....	56
Figure 36. Detection curve for non-eagle birds observed by the camera system (green) with 95% credible interval (grey). The detection curve for the theoretical case of perfect detection within 400m from the camera system is shown in blue.....	57
Figure 37. Types of radial distortion. Source: http://www.mathworks.com/help/vision/ug/camera-calibration.html	60
Figure 38. Tangential distortion occurs when the image plane and camera lens are not parallel. Source: http://www.mathworks.com/help/vision/ug/camera-calibration.html	60
Figure 39. Matlab camera calibration. Shown is a list of calibration images (left), the pattern as identified by the calibration application (center) in one such image, and rough 3-D estimates (and errors) of the pattern's location relative to the camera (right).	61
Figure 40. Triangulation of a 3-D point x is performed by projecting intersecting lines from image points y_1 and y_2 with each camera's focal point O_1 and O_2	62
Figure 41. Normalized cross-correlation was used to successfully match the rear tire of a vehicle across a pair of images captured by the camera system. The template used is the area selected in blue on the left image, and the red region in the right image depicts the resulting match....	63
Figure 42. Example flight path of an eagle estimated using a stereo-optic position algorithm.	64
Figure 43. Simplified shapes with error of 3.01%, along with camera system images that the shapes were based upon.	66
Figure 44. Simplified shapes with error of 4.25%, along with camera system images that the shapes were based upon.	66
Figure 45. Simplified shapes with error of 5.3%, along with camera system images that the shapes were based upon.	66

Figure 46. Simplified shapes with error of 6.9%, along with camera system images that the shapes were based upon.	67
Figure 47. Simplified shapes with error of 8.5%, along with camera system images that the shapes were based upon.	67
Figure 48. Formula for equisolid angle mapping.....	72

LIST OF TABLES

Table 1. Milestone summary table.	16
Table 2. nIR and thermal imaging systems and object detection distances.....	18
Table 3. Animals identified in the video analysis at two observation locations recorded from September 15, 2015 to January 15, 2016.....	45
Table 4. Bald Eagles identified in the video footage from September 15, 2015 to January 15, 2016. Observations are broken down into age ranges (Adult, Subadult, and Unknown) and by observer confidence in the identification (Definite, Probable, Possible). See Figure 18-Figure 28 for example images for eagles in each category.	45
Table 5. Bird distances calculated using the pixel count and stereo-optic position estimate methods. Smaller birds include American Goldfinch and American Robin; Medium birds include American Crow, Common Loon, and Ring-billed Gull; and Large birds include Canada Goose and Herring Gull. WS = Wing span.	58

EXECUTIVE SUMMARY

There is a critical need to develop monitoring tools to track aerofauna (birds and bats) in three dimensions around wind turbines. New monitoring systems will reduce permitting uncertainty by increasing the understanding of how birds and bats are interacting with wind turbines, which will improve the accuracy of impact predictions. Biodiversity Research Institute (BRI), The University of Maine Orono School of Computing and Information Science (UMaine SCIS), HiDef Aerial Surveying Limited (HiDef), and SunEdison, Inc. (formerly First Wind) responded to this need by using stereo-optic cameras with near-infrared (nIR) technology to investigate new methods for documenting aerofauna behavior around wind turbines. The stereo-optic camera system used two synchronized high-definition video cameras with fisheye lenses and processing software that detected moving objects, which could be identified in post-processing. The stereo-optic imaging system offered the ability to extract 3-D position information from pairs of images captured from different viewpoints. Fisheye lenses allowed for a greater field of view, but required more complex image rectification to contend with fisheye distortion. The ability to obtain 3-D positions provided crucial data on the trajectory (speed and direction) of a target, which, when the technology is fully developed, will provide data on how animals are responding to and interacting with wind turbines. This project was focused on testing the performance of the camera system, improving video review processing time, advancing the 3-D tracking technology, and moving the system from Technology Readiness Level 4 to 5. To achieve these objectives, we determined the size and distance at which aerofauna (particularly eagles) could be detected and identified, created efficient data management systems, improved the video post-processing viewer, and attempted refinement of 3-D modeling with respect to fisheye lenses.

The 29-megapixel camera system successfully captured 16,173 five-minute video segments in the field. During nighttime field trials using nIR, we found that bat-sized objects could not be detected more than 60 m from the camera system. This led to a decision to focus research efforts exclusively on daytime monitoring and to redirect resources towards improving the video post-processing viewer. We redesigned the bird event post-processing viewer, which substantially decreased the review time necessary to detect and identify flying objects. During daytime field trials, we determine that eagles could be detected up to 500 m away using the fisheye wide-angle lenses, and eagle-sized targets could be identified to species within 350 m of the camera system. We used distance sampling survey methods to describe the probability of detecting and identifying eagles and other aerofauna as a function of distance from the system. The previously developed 3-D algorithm for object isolation and tracking was tested, but the image rectification (flattening) required to obtain accurate distance measurements with fish-eye lenses was determined to be insufficient for distant eagles. We used MATLAB and OpenCV to improve fisheye lens rectification towards the center of the image, but accurate measurements towards the image corners could not be achieved. We believe that changing the fisheye lens to rectilinear lens would greatly improve position estimation, but doing so would result in a decrease in viewing angle and depth of field. Finally, we generated simplified shape profiles of birds to look for similarities between unknown animals and known species. With further development, this method could provide a mechanism for filtering large numbers of shapes to reduce data storage and processing.

These advancements further refined the camera system and brought this new technology closer to market. Once commercialized, the stereo-optic camera system technology could be used to: a) research how different species interact with wind turbines in order to refine collision risk models and inform mitigation solutions; and b) monitor aerofauna interactions with terrestrial and offshore wind farms replacing costly human observers and allowing for long-term monitoring in the offshore environment. The camera system will provide developers and regulators with data on the risk that wind turbines present to aerofauna, which will reduce uncertainty in the environmental permitting process.

INTRODUCTION

Biodiversity Research Institute (BRI) collaborated with the University of Maine School of Computing and Information Science (UMaine SCIS), HiDef Aerial Surveying, Limited (HiDef), and SunEdison, Inc. (formerly First Wind) to test a stereo-optic camera system for documenting aerofauna interactions with wind turbines. The team also further developed the monitoring system, which uses an image-processing algorithm to detect and track daytime and nighttime aerofauna and assist video reviewers in parsing large amounts of visual data. Each of these components had not been tested in the field, so we had an opportunity to test and improve these components based on realistic deployment conditions. These developments allowed us to bring this technology closer to market and provide additional data for understanding aerofauna exposure to wind turbines. This project supports the Department of Energy (DOE) Wind Energy Technologies Office (WETO) goals of lowering the unsubsidized, levelized cost of energy (LCOE), and alleviates a significant market barrier to wind power development by reducing the uncertainty surrounding avian avoidance behavior and collision risk.

Potential Applications and Impacts

Current methods for modeling the effects of offshore wind energy development on birds account for uncertainty by taking a precautionary approach and incorporating conservative assumptions of avoidance behavior into collision risk models. These assumptions lack empirical support, and because avoidance rates are a primary driver of these models, lead to increased estimates of mortality, or “take” (Masden and Cook, 2016). The wind industry needs data on species-specific avoidance responses for both terrestrial and offshore turbines to increase the accuracy of collision risk models (Ferrer et al., 2012; Fox et al., 2006). Without species-specific and condition-specific avoidance probabilities (as avian behaviors differ under a range of conditions, e.g., breeding vs. recently fledged birds), the accuracy of collision risk models will remain poor (Chamberlain et al., 2006). The stereo-optic camera system, further developed in this project, directly addresses this uncertainty by providing developers and regulators with much needed data on how aerofauna interact with wind turbines, and thus provides more accurate estimates of avoidance response.

Specifically, when fully developed, the stereo-optic camera system will bring a higher level of precision to avian and eagle risk analysis by being able to monitor bird movements around turbines in three dimensions. The camera system will provide data on how eagles interact with turbines in different environmental conditions, which could be used to design effective mitigation options. Increased precision can reduce the financial and temporal uncertainty associated with determining whether an incidental take permit (ITP) is necessary under the Bald and Golden Eagle Protection Act (BGEPA), and if a permit is deemed necessary, what level of take can be estimated. Both Bald and Golden Eagles have been injured and killed at wind energy facilities across the United States, particularly Golden Eagles in the western U.S.; in part, this may be because terrestrial regions with high wind resources in the western U.S. also tend to include habitats important to Golden Eagles (Pagel et al., 2013). The stereo-optic camera system also has the potential to address issues related to Migratory Bird Treaty Act (MBTA) compliance, state and local wildlife protection standards, and permitting under the Endangered Species Act (ESA). In the future, when the camera system is commercialized it will provide detailed information on how different species respond to individual turbines in various seasons and weather conditions, potentially allowing for operational adjustments during time periods when eagles or other sensitive species are known to be at risk of collision.

Brief Description of Technology

The stereo-optic camera system uses two 29-megapixel digital video cameras with fisheye lenses installed in heavy-duty stainless steel housing able to withstand marine environments (Figure 1). The control system contains an uninterruptible power supply, a computer, a relay box, and an emergency buffer storage disk. The control computer preprocesses the data to enable transmission back to a base station, manages the data until it is successfully transmitted, and operates the camera cleaning system. The initial data processing includes masking stationary parts of the image, so that only sections of the video containing moving objects are recorded. The resulting image file sizes are typically reduced by >90% without compressing the video, requiring less than 10% of the bandwidth for transmission that would be needed to send the raw video in its entirety. A pair of full-frame color images is stored every 5 minutes to provide contextual information (weather, visibility, turbine orientation). Post-processing includes calibrating the images for object distance estimates, marking avian object pairs between the two cameras, calculating 3-D positions for each stereo pair, and then manually identifying birds to species or species grouping.



Figure 1. The marinized camera system deployed for testing.

Commercialization Status and Plans

There is a need to better understand bird avoidance around wind turbines. We lack reliable estimates of avoidance rates of animals around turbines offshore, which can significantly alter our understanding of the effects of offshore wind power on aerofauna. BRI has successfully responded to this need and advanced the stereo-optic camera system from Technology Readiness Level (TRL) 4 to 5. The stereo-optic camera system is now ready to be advanced to TRL 6-8 and to be validated at terrestrial and offshore wind energy projects. With further field-testing and development, the stereo-optic camera system could be delivered to the commercial market. Once the system is commercially ready, it could be used to reduce permitting time and inform or alter wind farm operational practices to reduce adverse effects of wind farms on aerofauna.

We have identified three key areas for further system development: (1) modify the system to accept different lens configurations, which would improve distance estimates and allow for site-specific camera system configurations; (2) further develop algorithms using shape recognition and machine learning, to reduce time spent on video review and advance the system towards automatic target detection and identification; and (3) incorporate use of a thermal imaging system instead of near-Infrared (nIR) to improve nighttime detection of aerofauna. Further testing in an offshore environment, while employing recommended improvements, would also be needed to move this technology from a TRL 5 to a TRL 8.

BACKGROUND

Project Objectives

The goal of this project was to further develop, and eventually bring to market, a stereo-optic camera system designed to document aerofauna interactions with wind turbines. Our objectives included: (1) increasing the resolution of a HiDef-manufactured prototype stereo-optic camera system to 29-megapixels; (2) pairing the cameras with a nIR light to detect aerofauna at night; (3) determining feasible methods for exporting data from remote wind farm locations; (4) refining the 3-D algorithm for object isolation and tracking; (5) determining the percentage of eagles that can be detected at specific distances from the camera; and (6) using a distance-based hierarchical model to estimate detectability of aerofauna by the camera system. To achieve these objectives, we conducted field trials with the camera to determine the size and distance at which aerofauna (particularly eagles) could be detected and identified, created efficient data management systems, validated three-dimensional position estimates, and explored new shape-processing algorithms to filter large amounts of data. These advancements allowed us to bring this technology closer to market. Milestones are described in Table 1 and detailed objectives are as follows:

- Objective 1: Build and test the physical components of the camera system, including the 29-megapixel high definition camera and the nIR lighting system, that detects all objects encountered in a laboratory setting.
- Objective 2: Validate detection rates of the camera system to assure that the total detection rate of targets of comparable size to birds and bats is at least 95% within 10 m of the camera system and 25% at 250 m during the day and at least 25% detection rates within 100 m at night.
- Objective 3: Evaluate three different methods of exporting data from the field remotely: remotely download via on-site internet connection, remotely download via point to point wireless radio transmission, and remotely download via cellular network.
- Objective 4: Video reviewers identifying Bald Eagles at 10-50 m result in 90% species-level identification, 50-100m=70% species level identification, 100-200m=50% species identification and 200-500m=30% species identification.
- Objective 5: Generate baseline information on species detection and identification of bats using the nIR camera system
- Objective 6: Improve the 3D algorithm for object isolation and tracking by reducing bandwidth needs for data transfer and including real-time distance estimates.
- Objective 7: Build an interactive bird event viewer that reduces the manual inspection and analysis time by 70% compared to the current software configuration.
- Objective 8: Use statistical modeling techniques to describe the effective strip width of the camera system and distance detection thresholds for target species.

Table 1. Milestone summary table.

Milestone Summary Table							
Recipient Name:		Biodiversity Research Institute					
Project Title:		Stereo-Optic High Definition Imaging: A new technology to understand bird and bat avoidance of Wind Turbines					
Task Number	Task or Subtask (if applicable) Title	Milestone Type (Milestone or Go/No-Go Decision Point)	Milestone Number* (Go/No-Go Decision Point Number)	Milestone Description (Go/No-Go Decision Criteria)	Milestone Verification Process (What, How, Who, Where)	Anticipated Date (Months from Start of the Project)	Anticipated Quarter (Quarters from Start of the Project)
1	Manufacture Camera & nIR	Milestone	1	BRI takes delivery of functional camera and nIR system.	BRI will test that the camera system functions for 24 hours at the Portland, Maine office prior to test deployment.	3	1
2	Preliminary data collection and systems refinement	Milestone & Go/No Go	2.1	Testing has an outcome where total detection rates of 20 bird and bat-sized targets are at least 95% within 10 m during the day, 25% within 250 m during the day, and at least a 25% detection rate within 100m at night.	BRI will place the camera system at a private test site in order to test the detection rates of the camera equipment.	4-6	2
3	Equipment deployment, maintenance, evaluation of data transfer methods and nIR limited trial	Milestone	3	Obtain sample size of 20 eagle images.	BRI will place the camera system in locations with concentrations of eagles to meet analysis goals as well as place the camera at a turbine site for extended video capture.	6-21	2015-Q2-4 2016-Q1-3
4	Data Management at BRI Video Lab	Milestone	4	The final method for data export results in no less than 95% of wildlife video images being transferred from the camera to a video catalog and associated database.	BRI will perform audits to ensure that data delivery is no less than 95% of wildlife video images.	8-15	2015-Q3-4 2016-Q1
5	Video Lab Analysis	Milestone	5	Video analysis of daytime eagle-sized objects at 10-50 m result in 90% species-level identification, 50-100m=70% species level identification, 100-200m=50% species identification and 200-500m=30% species identification.	BRI will work within the confines and successes of the camera system such that identification goals are reached.	9-18	2015 Q3-4 2016 Q1-2
6	Data Analysis	Milestone	6	Quantify the error in ID rates and estimating 3D position.	BRI will quantify the error in ID rates and estimating 3D position through data analysis.	12-18	2015-Q4 2016-Q1-2
6.3	Refine 3D model & build an interactive bird event viewer	Milestone	6.3	The 3D module reduces the information stream down to 1% of the original data and distinguishes bigger birds (>1m wingspan) from smaller birds with 90% reliability. Build an interactive bird event viewer that reduces the manual inspection and analysis time by 70% compared to the current software configuration.	BRI/UMaine will collect data at a private test site during Task 2 activities in order to work towards 3D module development goals.	12-18	2015-Q4 2016-Q1-2
7	Reporting	Milestone	7	Report to DOE	BRI will provide a final report to DOE.	15-21	2016 Q1-3

State of the Art

A number of methods exist to measure risk to, and effects of, wind energy development on birds and bats including: collecting carcasses (Pagel et al., 2013), posting human observers (Ferrer et al., 2012), conducting aerial and boat-based surveys (Williams et al., 2015), collision risk modeling (Masden and Cook, 2016), field experiments to assess turbine avoidance (Hull and Muir, 2013), and observation methods using video, radar, or thermal cameras (Cullinan et al., 2015; Groom et al., 2013; Maurer, 2016; Sinclair and Degeorge, 2016). Automated monitoring methods that are currently commercially available or in development are as follows:

- **HiDef Stereo-optic camera system** (this project): Prior to this DOE funded project, the camera system was at Technology Readiness Level (TRL) 4, with all the key physical components assembled and successfully interfacing with control software. We have now moved the system to TRL 5. TRL 5 is defined by DOE as “the basic technological components are integrated so that the system configuration is similar to the final application in almost all respects. Examples include testing a high-fidelity system in a simulated environment.”
- **DeTect**: The DeTect Merlin Avian Radar System¹ uses automated radar technology to detect birds and bats passing through the radar swept region to help improve collision risk assessment and provide warning when flocks are approaching.
- **DTBird**: DTBird² records video of birds in the daytime using an automated system that can detect flying birds and take programmed actions to reduce avian fatalities.
- **WT-Bird**: WT-Bird is a system used to detect bird collisions with wind turbines, using sound to detect the impact and video images to identify species (Wiggelinkhuizen et al., 2007).
- **ATOM**: Normandeau's Acoustic-Thermographic Offshore Monitoring system (ATOM™) is a bird and bat detection system that provides day and night monitoring in the offshore environment with infrared video and bird/bat acoustic monitors.³
- **TADs**: Thermal Animal Detection System (TADS) uses thermal cameras to detect aero fauna interactions with wind turbines.⁴
- **Laufer Wind Group**: Integration of radar and pan-tilt-zoom cameras.⁵
- **Oregon State University**: 360-degree cameras to detect birds and a blade-mounted collision detection system
- **IdentiFlight**: Wide-angle camera array integrated with high magnification stereo-optic cameras for detecting and classifying wildlife (specifically designed for Golden Eagles at the moment) and issuing deterrence signals.⁶
- **Thermal cameras**: Thermal and nIR cameras have been tested to record flight tracks autonomously and classify aero fauna as birds or bats (Table 2).

¹ <http://www.detect-inc.com/avian.html>

² <http://www.dtbird.com/>

³ <http://www.normandeau.com/pages/technology/atom.asp>

⁴ http://www.dmu.dk/1_viden/2_publicationer/3_fagrappporter/rapporter/FR440.pdf

⁵ http://www.designworldonline.com/three-ways-doe-research-will-help-eagles-coexist-wind-energy-deployment/#_

⁶ <http://www.res-group.com/en/services-products/identiflight/>

Table 2. nIR and thermal imaging systems and object detection distances.

PI	Illumination/ Camera Type	Distance to Object of Interest	Field of View (FOV)	Resolution (pixels)	Reference
Trevor Harrison, U. of Washington	Thermal Infrared (TIR)	14.6 m (mean)	15° x 12°	640 x 480	NWCC 2014 Abstract
Normandeau's ATOM	TIR	200m (max)		256 x 324	NWCC 2014 Abstract
Horn et al. (2008)	TIR	76 m	24°	320 x 240	http://www.bioone.org/doi/abs/10.2193/2006-465
Cryan et al. (2014)	TIR	80 m			http://www.pnas.org/content/111/42/15126.short
Cryan et al. (2014)	nIR with illumination	12 - 60 m (range)			http://www.pnas.org/content/111/42/15126.short
Elliott et al. (2005)	nIR with illumination	10 m (max)			http://www.nckms.org/2005/pdf/Papers/Elliott-mdc bats.pdf
Desholm et al. (2006)	TIR, 7-15 um Thermovision IRMV 320V	150 m waterbird 30 m passerine (max)	7-24°		http://onlinelibrary.wiley.com/doi/10.1111/j.1474-919X.2006.00509.x/full

Approaches to Accomplish Objectives

We executed the project in five phases: (1) HiDef built and laboratory tested the camera in the U.K.; (2) BRI conducted preliminary field tests after the camera was shipped and reactivated in the United States; (3) BRI collected video data at a terrestrial wind farm and at sites with large amounts of eagle activity; (4) BRI and UMaine developed data management and viewer processes; and (5) BRI and UMaine reviewed and analyzed the data, improved the 3-D modeling algorithm, and reported on the findings. Below is an overview of each phase. Detailed methods are presented in the results and discussion section of the report.

Phase 1: Camera system manufacturing and lab testing

HiDef built and tested two stereo-optic camera systems that consisted of two ultra high-definition, 29-megapixel cameras with fisheye lenses and near infrared (nIR) capabilities. Algorithms for removing non-static parts of video (e.g., motion segmentation) and software for reviewing video produced by the camera rig were provided by HiDef. In their U.K.-based laboratory, HiDef conducted the following tests of the system: proper operation of the electronics and mechanics of the camera system; proper operation of the software in a lab setting; radio frequency (RF) link patch antenna connectivity to the local area network; lens focus and calibration; day and night operation of cameras; successful collection of imagery and post-processing of test imagery; and 3-D track estimation. The nIR illumination was tested in HiDef's workshops using heated inanimate targets (bean bags). At the end of this phase, UMaine engineer Chris Dorr traveled to the U.K. to assist in initial field-testing and was trained in the operation of the camera system.

Phase 2: Preliminary field-tests

BRI researchers deployed the camera system at several test sites in Maine to understand the capabilities of the system in an environment more closely approximating future deployment conditions. The camera was first deployed at a local test site in order to determine the system's ability during daytime to detect 20 bird and bat-sized targets within 10 m at least 95% of the time, and to detect targets within 250 m at least 25% of the time. We also tested the camera at night to determine whether we had at least a 25% detection rate within 100m. The camera exceeded the daytime detection criteria, but initial nighttime testing resulted in 100% detection at 0-30 m, and no detections from 30-100 m. BRI modified the nIR hardware configuration and was able to improve the detectability distance to 60 m, but still did not successfully meet the Go/No-go criteria for this objective. Because of the poor results of the nIR system tests, and in consultation with DOE, it was decided to discontinue testing and further development of the nIR system in order to reallocate project resources to redesign the video post-processing event viewer.

Phase 3: Data collection and maintenance at target sites

BRI and UMaine researchers deployed the camera system at a wind energy site to evaluate camera performance and assess remote data downloading options. Researchers installed the stereo-optic camera at a turbine so that it had views of the rotor-swept zone of that turbine as well as two additional turbines within 500 m of the camera system. BRI and UMaine conducted routine maintenance of the system throughout the deployment period. Data was transferred from the field site to BRI's video review lab by swapping hard drives. In anticipation of a low sample size of eagle images at the turbine site, and to meet project objectives of obtaining at least 20 eagle images, BRI placed the second camera system at a non-turbine test site with high eagle activity.

Phase 4: Data management and video processing improvements

BRI developed data management and video review protocols based on field test results. Protocols addressed raw data collection, preliminary database fields, identification criteria, and guidance on video/image analysis. Species identification criteria addressed environmental factors, video quality, and target behavior. During video review, each analyst used the same identification criteria to ensure accuracy, identification confidence, and repeatability of the data. BRI biologists developed a storage system where raw data were archived and pertinent analyzable sections or subsets were stored on local drives or servers for data analysis. Data entry errors were reduced through data entry templates and predefined data entry lists (pick-lists).

In order to improve the speed of target detection and identification, UMaine developed a new post-processing video viewer with enhancements such as manual masking, black-and-white mode, video playback at varying speeds, object marking, and controllers for changing brightness and contrast levels. These features allowed quicker manual inspection of low-light video segments and reliable documentation of objects of interest.

Phase 5: Video review and data analysis

The focus of this phase was to conduct video review, analyze data, improve the 3-D modeling algorithm, and report on findings. BRI biologists identified birds in randomized video segments using established protocols and identification criteria. Video analysis results were evaluated for errors and audited. Stereo-optic algorithms were used in conjunction with pixel counting to determine object distances for identified objects. Using real bird detection data, we specified a model that allowed us to estimate the overall rate of identification of an animal, conditioned upon

detection. Using a distance detection framework, we used a hazard model to describe the decrease in identification rate for targets with increasing distance from the position of the camera. UMaine developed techniques for estimating stereo-optic position using extreme wide-angle lenses, and for identifying bird-like shapes that pass through the motion segmentation.

RESULTS AND DISCUSSION

PHASE 1: CAMERA SYSTEM MANUFACTURING AND LAB TESTING

Objectives

- Build and test the physical components of the camera system, including the 29- megapixel high definition camera and the nIR lighting system that detects all objects encountered in a laboratory setting.

Technical Approach

HiDef fabricated the stereo-optic camera systems from a combination of off-the-shelf and custom components during the winter 2014-2015 in the U.K. . HiDef completed initial camera system testing by mid-March 2015, with a focus on gross function in anticipation of shipment to UMaine and BRI for advanced lab and field-testing of the system. HiDef trained a collaborator, Ph.D. Candidate Chris Dorr from UMaine SCIS, on the hardware and software of the stereo-optic high definition cameras prior to shipment of the system to the U.S. in March 2015.

Results

Results of system testing in the U.K. showed that most components of the camera system, including the critical motion segmentation, worked as planned. However, near-infrared (nIR) lights illuminated a narrower and shallower field of view than anticipated, and so further testing was scheduled after arrival in the U.S. The stereo-optic cameras captured images as designed, but it was determined that one of the cameras in the stereo-pair intermittently captured images at half the frame rate of the other camera. This issue was identified as one of the most critical problems for successful implantation of the system in Phase 1, and further investigations were made after the camera shipped to the U.S. (see below).

Resolution of camera trigger issue

The project team identified and fixed the camera trigger issues first encountered in the U.K. The master camera was able to run at predicted speeds of ~1.8 frames per second (fps), but the second camera ('slave' camera) was often and unpredictably dropping frames and running at a much lower frame rate. Stable synchronized frame rates worked consistently when the triggering system was reset and when running cameras at night. The project team determined that there was a data bandwidth limitation caused by the very high-resolution cameras, which resulted in incomplete image storage prior to another trigger command being sent from the master to the second camera.

In the triggering phase, the master camera would both initiate camera exposure and simultaneously send a short pulse across to the second camera via a direct connection from the output data port from the master through the input/output (i/o) connector. The second camera was parameterized to use the incoming pulse as a trigger for the camera. We verified by oscilloscope that the trigger impulse was sent at regular intervals from the master camera and the second camera was correctly receiving triggers. However, because the master camera was set to continuous run, it captured images as rapidly as it could. The normal course of image capture was: trigger pulse -> exposure -> next trigger. Both cameras were set at automatic exposure and auto gain to be determined by each camera independently, such that the time from image to image could be different between cameras

depending on differences in camera settings. In consulting the IMPERX™ camera manual we discovered a cautionary note about triggering: “the time interval between trigger pulses must be greater than the corresponding frame time.” Therefore, if the trigger pulse occurred before the image was exposed fully, the cycle would be interrupted and part of the frame lost, presumably leading to image drop. The independent exposure settings and other parameters, and thus independent frame capture times, could lead to slightly different image capture cycles between cameras. If the second camera was exposed for less time than the master, however, it would be able to finish image capture ahead of the next trigger impulse from the master camera, allowing it to keep up the same frame rate.

The nighttime capture supported this theory; at night, the exposure was at the same maximum value for both cameras, allowing for both the master and second cameras to be in synchrony. Also, restarting the cameras led to camera synchrony, likely because a restart required the cameras to re-evaluate settings that could have led to the second camera processing faster than the master. Thus, it was possible that better control of triggering the system could alleviate the problem.

The goal was to slow down the triggering just enough to allow for a full exposure and processing of images on both cameras, and for each camera to revert to a wait-for-trigger mode. In this way, both cameras would be able to capture full images without frame drop. Together with engineers in the U.K., three possible approaches to repairing the trigger problem were proposed:

- 1) Use an external trigger impulse from a third party device connected to each camera. The trigger would be sent to both cameras simultaneously, removing the master and secondary relationship between cameras.
- 2) Use a computer-based trigger from the attached data collection computer to trigger each camera. The trigger would be sent to both cameras simultaneously, removing the master and secondary relationship between cameras.
- 3) Use the on-camera pulse-generation utility built into the IMPERX™ Bobcat cameras to trigger the master camera, which would also be passed to the second camera for triggering. This is similar to the current setup, except that the master is not in free-running mode, and instead the trigger is controlled by the pulse-generator and is simultaneously passed to the second camera.

The team experimented with the third solution first because this was the easiest to implement and test, requiring the fewest system modifications and no additional equipment. The other two solutions would have required input wire modification, which carried risk of damaging the connections and required more time to implement. The solution was initially tested by using the video software (Bobcat GEV player) provided by IMPERX™ to manually configure and view camera output. Two instances of the software were initiated to monitor both cameras simultaneously and to confirm that by using a fixed trigger period, both cameras were stable from 1.65-1.72 fps. This solution was the most parsimonious, and these parameters for signal generation and trigger input were easily passed to the cameras during startup from the proprietary acquisition software that controlled the cameras and processed imagery. Additionally, to test the response of the cameras to this change, the master camera was disconnected and reconnected to the second camera while the master camera was still generating pulses, and the second camera came back online quickly. After some additional testing, the fix was integrated into the acquisition code and installed on each camera system. Camera synchronization and trigger timing were no longer issues for the duration of the project.

PHASE 2: PRELIMINARY FIELD-TESTING

Objectives

- Validate detection rates of the camera system to assure that the total detection rate of targets of comparable size to birds and bats is at least 95% within 10 m of the camera system and 25% at 250 m during the day and at least 25% detection rates within 100 m at night.

Technical Approach

The project team installed the camera system at two sites that were approved by DOE through the NEPA process. The first site, located on a coastal estuary, allowed researchers to test daytime and nighttime detection rates of bird- and bat-sized objects and collect images of nontarget birds such as gulls, crows, and egrets. The second site, located near a freshwater reservoir, allowed the team to further test nIR detection rates, obtain a nighttime image of a bat, record bat calls with an acoustic recording unit, and obtain bald eagle images.

At the estuary test site, the camera system was deployed at a private location that had an unobstructed airspace at least 250 m from the camera location. The primary objective was to estimate the detection rates of bird and bat-sized objects with the camera system and determine if they were above the predefined thresholds outlined in the above objective.

To determine the detection rates of the camera system within 250 m (daytime) and 100 m (nighttime), standard methodology for point counts were used with detections categorized by distance away from the survey point. During the daytime, researchers tested whether the camera could detect a bird-sized object (a ball 10 cm in diameter; about the size of an American Robin) at intervals of 10, 25, 50, 100 and 200 m away from camera system center (Figure 2). Each distance was tested in the left, center, and right portions of the field of view. We used nIR lighting to test nighttime detections of the same object at five distances: 10, 20, 30, 40, and 50 m from the camera system center (Figure 3). The system was tested at least twice at each distance band, with the object stationary and the object being thrown. After the data were captured by the camera system as still images, a video analyst reviewed the images to determine if the object was present within the frame. The objectives of the project ask us to evaluate the accuracy of the system within 250 m during the daytime, so we will be conservative and assume no detections were possible from distances further away than we tested. The proportion of detections over attempts will be calculated for each distance band to estimate how detection probability changes over distance.

At a second private site overlooking a reservoir (Figure 4), we made improvements to the nighttime field of view by repositioning the nIR lighting to be more focused in the middle of the field. With the new configuration of nIR lights, testing was performed again to attempt 100 m detection of a bat-sized object. However, the best detection that could be achieved on a bat-sized object was 60 m.



Figure 2. Testing at first site with bird-sized object (held at center of chest).



Figure 3. Nighttime testing of bat-sized object (Softball in the person's hand) 20 m from the camera. Image is magnified and cropped




Figure 4. View from camera at second testing location with a Bald Eagle in the field of view.

Results

There was 100% detection of objects 50 m or less from the camera system during the daytime. Objects were detected out to 200 m away from the camera system, the furthest range tested, but detection probability in the 100-200 m distance bands dropped to 67% (Figure 5). It was assumed for modeling purposes that the detection probability for a robin-sized object would be 0% past 200 m due to lack of data, though this is likely an underestimate as such objects could be seen in some circumstances. This outcome was considered to meet our overall objective for detection probability.




Figure 5. The change in detection probability with distance for the camera system during daytime of a bird-sized object.

During nighttime testing, objects were detected out to 30 m. Unlike during daylight testing, detection within each distance band appeared to be binary: either the object was detected all the

time (100% detectability at a given distance) or it was not detected at all (Figure 6), so detectability was 100% within 30 m and 0% thereafter. This was not considered to meet our overall objective for detection for this system. Given the all or nothing aspects of detection, we were not able detect animals far enough out to be useful for our application as 30 m is often not far enough to reach even part of the turbine rotor swept zone from a ground-level mounted unit. Improvements to the nIR system would be necessary to allow detections far enough away to be useful in monitoring turbines blades in the field.




Figure 6. The change in detection probability with distance for the camera system during nighttime of a bat-sized object.

PHASE 3: DATA COLLECTION AND MAINTENANCE AT TARGET SITES

Objectives

- Collect data that video reviewers can use to assess camera performance in detecting and identifying birds (bats will be discussed in Phase 5).
- Obtain 20 eagle images.
- Evaluate three different methods of exporting data from the field to BRI's video review lab, so that no less than 95% of the wildlife video images are being transferred: remote download via on-site internet connection, remote download via point to point wireless radio transmission, and remote download via cellular network.

Technical Approach

In consultation with SunEdison staff, the project team chose a turbine area for testing that had unobstructed views and was in proximity to an eagle-use area. On September 16, 2015, the stereo-optic camera was installed underneath an operational turbine in this area. The stereo-optic camera system was mounted on a wooden frame to enhance stability and security at the base of the turbine. The camera system was positioned so that the angle of view would allow for filming the rotor swept zone, the area around the turbine, and other turbines within 500 meters of the camera system (Figure 7; Figure 8). Each camera on the camera system was adjusted for optimal daytime filming. The stereo-optic camera was left for 24 hours and checked the next day to ensure that all processes were working correctly. Another check was performed on September 29, 2015 to ensure that back-up power worked properly after a scheduled power down of the turbines. Data collection had not been interrupted and there were no issues noted on either visit. The camera system was removed from the site in December 2015.



Figure 7. Camera installation at terrestrial wind farm.

In November 2015, a second camera system was re-deployed at the Reservoir Site to capture images of eagles using the area. This site was named “Eagle Test Site” for the remainder of the project. We removed the system from this site in March 2016. Both sites were operating concurrently during November and December 2015.




Figure 8. Camera view at turbine test site.

Results

A total of 16,173 video segments were recorded at the Turbine Site over a period of 79 days. At the Eagle Test Site, 30,340 video segments were recorded over 114 days. A subset consisting of 10,354 segments were placed into the queue for video analysis (totaling >4 TB in size). Video analysis of 2,508 video segments, described in Phase 5 below, resulted in 1,140 objects of interest. These included a total of 51 Bald Eagle observations in 2,033 Bald Eagle images. Other species observed included American Crow, Common Raven, Common Loon, Herring Gull, and American Goldfinch.


Data Transfer

All 46,513 video segments, ranging in size from 66 MB to 2 GB, were transferred from the field to the BRI video lab through physical swapping of 1 TB hard drives. All video segments (n=26,527) slated for video analysis were saved to servers as well as archived on hard drives without incident. Extra video segments, recorded at the Eagle Test Site but not analyzed, were also archived. We had no loss of recorded data during transfer from the field to the lab.

BRI staff attempted to test a point-to-point wireless networking method to transfer data remotely, rather than using manual swapping of hard drives. This point-to-point method uses long-range Ligowave wireless transmitting equipment. In order to fully test this method, however, a clear line

of sight with some distance from Point A to Point B, a transmitting location, and a receiving location with internet are all required. Without a proper DOE-approved testing location that fulfilled these criteria, it was determined that long-range wireless testing could not be provided under the scope of this project. However, this method is commonly used for line-of-site wireless communications, and HiDef has successfully tested this method offshore in Europe. We are confident that this method of communication would be viable if distances were within equipment specifications and clear line-of-sight was available. More information is provided in the Recommendations section.

Overall, approximately 8 TB of video were recorded. A variety of factors influenced individual file size, such as time of day (Figure 9), visibility (Figure 10), and weather (Figure 11). Nighttime video segments were consistently lower in average size than daytime video segments (Figure 9). Clear conditions produced larger file sizes, whereas video segments with lower visibility (Figure 10) or inclement weather (Figure 11) resulted in lower file sizes on average.



*Three outliers were removed that were atypical for nighttime file size; nIR-lit swirling leaves in water activated motion segmentation throughout the entire video segments, leading to larger file sizes.

Figure 9. Average size of video segments based on time of day showed that nighttime video segments are typically smaller in size.




Figure 10. Video segments with lower visibility resulted in smaller video sizes than video segments with clear visibility.




Figure 11. Weather that resulted in lower visibility or lower bird activity resulted in lower average file size than clear days.

Camera system performance and maintenance

Overall, despite months in the field under highly variable weather conditions and temperature regimes, the camera systems performed well with few issues and near constant uptime. A number of issues were identified and addressed, and are discussed below.

Mechanical issues

Minor equipment and software maintenance has been ongoing since the start of the project. Aside from the aforementioned trigger issue during the first phase of the project, there were other issues encountered and overcome during deployment:

- Motherboard failure due to lightning strike: There was a system malfunction on September 8, 2015. Through troubleshooting with HiDef and UMaine, the project team concluded that the motherboard from the on-board computer failed, most likely due to storm activity that occurred at the Eagle Test Site. In order to stay on schedule with turbine deployment, the project team replaced the malfunctioning hardware with the on-board camera from the other camera rig. After re-establishing software relationships and Windows-based interactions, the switched motherboard/camera rig combination was ready for turbine deployment. The camera rig continued to perform as expected without data interruptions through the end of the quarter. HiDef was able to ship a new motherboard from the U.K. to the U.S. in seven business days. After changing a cable and installing the new motherboard, the camera rig was working properly and the project team was able to hold a demonstration in time for a DOE site visit on September 24, 2015. The second camera rig has performed as expected through the end of the project. While we do not anticipate another motherboard failure, we did order another motherboard in case this type of malfunction occurs again.
- Outages due to software crash: We had system outages from an unknown source during the turbine trial. The stereo-optic cameras were in operation, but data collection stopped with no information or errors in the camera or computer system logs. After consultation with UMaine and HiDef, it was decided that scheduling a nightly restart of the system would be the best course of action. This resolved the problem, but this issue might need further troubleshooting under a controlled environment to look for possible causes such as temperature anomalies, operating system disruptions, or onboard memory limitations.
- Wiper motor failure: Two of the wiper motors ceased to function properly. One shipped from the U.K. in a nonworking state, and another experienced mechanical failure in the micro-switch that resulted in near constant movement of the wiper arm without wiper fluid. We worried that the wiper would scratch the external lens surface once the rubber was degraded enough, so we disabled this wiper. We have not yet sourced a replacement wiper motor, but it is an easy replacement with the correct part.

PHASE 4: DATA MANAGEMENT AND VIDEO PROCESSING IMPROVEMENTS

Objectives

- Develop a Standard Operating Procedure (SOP) using a hierarchical system for: a) detecting an object; b) categorizing the object as bird or other; c) categorizing the bird into species group (e.g., waterfowl, raptor, passerine); and d) identifying to genus or species.
- Build an interactive bird event viewer that reduces the manual inspection and analysis time by 70% compared to the current software configuration.

Data Management: Technical Approach and Results

Technical Approach

Method for selecting images for analysis

BRI researchers developed a database to store detection and identification data. Data entry consisted of categorical values that characterized the video segment (visibility and weather) as well as object information such as identification to the lowest taxonomic level and behavior. On-screen position information was also recorded.

The on-board camera system software organized video segments in a folder hierarchy, based on system time. The file naming convention used epoch time (the number of seconds that have passed since January 1, 1970) as part of the file name, which created a unique filename for each video segment. This consistent file-naming system allowed for easy programming in Microsoft (MS) Access® to randomize, prioritize, and retrieve video segments. Dawn, daytime, and dusk video were prioritized for analysis. The remaining segments, such as nIR-activated segments, were archived on hard drives.

Audit

In order to determine the rate of agreement of identification between video analysts, a random audit was conducted by a second video analyst on 20% of objects. For instance, objects identified as Bald Eagle by the original reviewer were also identified as Bald Eagle by the auditor in order for that object to count as an agreement or match. With American Goldfinch, the auditor could choose American Goldfinch or Unidentified Passerine for that object to count as an agreement. A successful random audit was considered to result in $\geq 90\%$ agreement. All species identified by the first observer as a Threatened or Endangered (T&E) species were also audited, with success requiring 100% agreement for the T&E audit.

Results

A total of 26,527 video segments were added to the video analysis queue, of which 11,740 nighttime video segments were removed from further analysis. Dawn, dusk, and daytime video segments were then randomized and assessed by video reviewers. Due to time constraints, a combination of 12,273 video segments from both sites were not reviewed. Of the 2,514 that were reviewed, video reviewers further eliminated 951 video segments where nIR lighting was activated. Of the 1,563 video segments remaining, 337 video segments had objects detected (251 at the Eagle Test Site and 86 at the Turbine site).

A total of 228 objects were audited. The random audit resulted in 93% agreement, while the T&E audit resulted in 67% agreement. Disagreements primarily occurred for eagles at greater distances and in low light conditions. Identification criteria were amended and eagles in disagreement were arbitrated for final answers.

Improving the System: Technical Approach and Results

Technical Approach

Acquisition, housekeeping and viewing software

In order to collect and analyze imagery, the camera system employs two main software applications: an acquisition and storage application that communicates directly with the cameras, and a viewer application that allows users to view and interact with recorded data. There is also a smaller, third application that is responsible for various scheduled housekeeping tasks such as operating the washer-wiper systems, and transferring recorded data to external storage devices.

The acquisition software, written in a combination of C++ and C#, is responsible for opening and maintaining communication (via the GigE Vision framework) with both of the IMPERX Bobcat B6620 cameras. This software runs on the camera system's on-board computer, and is capable of recording motion-segmented frames (as a Bayer-encoded stream) from both cameras in real-time at a frame rate of approximately 1.7 frames per second. While there is still some room for improvement, the acquisition software operates at or beyond the speed at which the cameras can transmit data to the on-board computer, the most relevant bottleneck in the system.

Given the volume of data associated with pairs of 6600 x 4400 px (29-megapixel) images, the acquisition software is heavily focused on compression. To reduce the amount of storage required and the amount of time users spend analyzing imagery, the acquisition software uses a motion-segmentation algorithm. This algorithm captures a pair of full-frame images from the cameras once every five minutes. Subsequent frames are compared to the full-frame, and only changes that exceed a specified threshold are recorded. This process repeats for five minutes, until a new full-frame image is captured. As configured, the smallest measurable area of change is a 16x16 pixel "tile." At the end of five minutes, the acquisition software produces a "chunk file," which contains one pair of full-frame bitmap images, and 510 image "maps" (5 minutes * 60 seconds/minute * 1.7 frames/second = 510; note that different frame-rates or recording durations affect this figure). Chunk files contain the motion-segmented tiles from each frame along with indexing information used to reassemble tiles in the viewer.

The effective compression rate of this segmentation approach depends on the amount of movement or change in a given scene: anywhere from 100% (no differences from initial reference frame, completely masked) to 0% (large image differences or changes in illumination, no masking). In practice, scene features such as fast moving clouds or large reflective areas like waterbodies have a detrimental effect on compression (see Figure 12).




Figure 12. Example images of masking differences between footage with little background motion (a) and considerable background motion (b). The second footage chunk was truncated at 273 frames, about half the amount usually saved in a chunk file, due to the excess amount of water coming through the masking.




Figure 13. Frame-pair view from the original HiDef viewer.

The initial full-frame images are ready for display after conversion back to RGB format from Bayer encoding, but the motion-segmented frames require some reassembly. This process involves: (1) unpacking motion-segmented tiles as Bayer encoded byte data; (2) computing the pixel XY coordinates of the tile for display; and (3) converting the tile Bayer data to a 16x16 RGB tile. Once a pair of segmented frames is reassembled, they are presented to the user for review (Figure 13).

The viewer is also responsible for performing stereo-calibration between the cameras. The stereo-calibration process is critical in obtaining reliable 3-D location estimates from pairs of stereo images, and requires the user to select several pairs of "match points" between the initial full-frame images in a chunk (Figure 14). The goal of the viewer application is to find an optimal translation between each half of the stereo-pair. This translation, along with several assumptions about the camera, lens, and stereo basis, then permits the triangulation of objects in a scene.




Figure 14. Frame-pair view of image calibration.

Once the calibration is complete, users can obtain approximate distance estimates of objects from a stereo-pair by carefully marking the same object in each half of the pair. Similarly, size estimates can be computed by exporting and comparing several pair of points. Distance estimates are saved by the viewer as plain text in the following format:

[<left xy> <right xy> [estimated x, estimated y, estimated z]]

Users navigate a chunk file by either moving forwards/backwards one frame-pair at a time, or by using a slider to move more quickly.

Improving the viewer

While the original viewer application permitted user-review of recorded data, early discussion identified a number of changes that could improve both the efficiency and utility of the viewer application. We identified the following as needed improvements:

- (1) Provide the ability to bookmark frames of interest in a chunk.
- (2) Provide the ability to (a) ignore empty frames, and (b) mask out areas of a scene.
- (3) Provide the ability to mark and save locations of objects of interest.
- (4) Provide the ability to adjust image properties such as contrast, saturation, and brightness.
- (5) Provide a separate reference frame displaying the current zoom window relative to the full frame.

UMaine SCIS engineers, in cooperation with University of Maine ASAP Media Services, determined the best way to provide users with the desired functionality would be to create a new viewer application, with the improvements listed above. Beyond providing users with the desired

functionality, development of a new viewer permitted the construction of an application that could run on Windows, MacOS, or Linux systems.

To fulfill the objectives, a new viewer was developed in C++, with a Qt GUI to allow cross-platform portability. UMaine SCIS engineers initially focused on creating a standalone library responsible for opening, parsing, and reconstructing image data from chunk files, while ASAP developers designed and implemented a Qt GUI. A library API ([chunkreaderlib](#)) allowed the viewer GUI application to request ready-to-display frames, both simplifying the viewer code itself and enabling easier future modification.

Development of “chunkreaderlib”

The library, which was responsible for parsing chunk files, fetching full frames, and reconstructing image tile data, relied on the open source computer vision programming API, OpenCV (<http://OpenCV.org>) library of programming functions for much of its basic functionality, including:

- (1) OpenCV data types were widely used, from storing image data to representing 2-D point or area information.
- (2) OpenCV `cvtColor` method was used to convert Bayer encoded image data to RGB.
- (3) OpenCV `pointInPolygon` test was used to check for masked tiles.

As development of the viewer GUI progressed, the library also became responsible for handling chunk metafiles, which stored user frame annotation and mask information. These metafiles effectively addressed improvements 1, 2b, and 3, allowing users to easily share and store areas of interest across frames in a chunk.

With the inclusion of mask data from the metafiles, the library was also responsible for improvement 2a. Beyond ignoring completely empty frames, the library could additionally determine if a frame was empty after considering user-applied mask(s).

Initially, the library workflow was as follows:

- (1) Open chunk file and parse header information
- (2) Fetch full-frame reference images
- (3) Fetch frame locations and frame headers
- (4) Reconstruct each frame by unpacking and debayering tiles, and
- (5) Return all rendered frames for viewing.

In an effort to improve performance of both the library and viewer, changes to (4) and (5) have been made such that the library only returns RGB image tiles and their locations, leaving the viewer GUI responsible for reassembly/placing tiles in the correct location on screen.

Development of “chunkreader ui”

The viewer GUI was written around the Qt framework, and provided a consistent UI across all three target platforms. Additionally, the application used OpenGL (Open Graphics Library), to render image data on screen, perform various image manipulation tasks as proposed in improvement 4, and provide zoom window reference panes as proposed in improvement 5.

To improve the user experience, the viewer GUI also allowed users to browse the frame data in several ways. As with the initial viewer application, users could navigate forwards or backwards one frame at a time, or quickly move through the frames with a slider. Additionally, users could playback/pause, jump to specific frames by number, or quickly jump to the start or end of a chunk.

The viewer GUI also provided users with the ability to mask out areas of a scene, and to mark areas of interest in a frame (Figure 15). This information was saved in metafiles, which were processed by the library as described in “Development of chunkreaderlib,” above.




Figure 15. Image of the new GUI zoomed in on marked gulls. The full extent is visible in the box to the right of the image, with the zoomed portion illustrated with a red box. Marker numbers are adjacent to marker boxes drawn around gulls observed in the video footage.

Results

The new viewer application successfully provided all of the proposed improvements, along with several other features that added to the overall utility. These included:

- (1) the ability to "link" or "unlink" halves of the stereo pair view, such that separate areas or zoom levels can be examined in each half at the same time,
- (2) a "black and white" render mode, which renders tiles in a two-color mask mode instead of as RGB image data,
- (3) a reference frame toggle, which allows users to quickly see what the current view area is in the initial reference frame, and
- (4) a status bar displaying human-readable date-time information.

PHASE 5: VIDEO REVIEW AND DATA ANALYSIS

Objectives

- Generate baseline information on species detection and identification of bats using the nIR camera system.
- Video reviewers identifying Bald Eagles at 10-50 m result in 90% species-level identification, 50-100 m=70% species level identification, 100-200 m=50% species identification and 200-500 m=30% species identification.
- Quantify the error in ID rates and estimating 3-D position
- Improve the 3-D algorithm for object isolation and tracking by reducing bandwidth needs for data transfer and including real-time distance estimates.
- Use statistical modeling techniques to describe the volume of space the camera system surveys effectively and how identification rates vary with distance to the camera.

Testing to Detect Bats: Technical Approach and Results

Technical approach

During the Go/No-Go review, DOE determined that given the limitations of the nIR technology on the camera system, further nighttime testing of the system should be scaled back. However, we were encouraged by the collection of a nighttime bat image (Figure 16), and therefore, a limited nIR trial was proposed in which three species of bats, captured during a separate federally funded project, were released at varying distances from the camera system in order to obtain baseline photographic images using nIR technology and determine nighttime detectability.




Figure 16. Nighttime image of a bat captured during field testing at the Reservoir/Eagle Test Site.

The study on bat detectability was conducted at Parker River National Wildlife Refuge in northeastern Massachusetts, USA, September 8-9, 2015, in collaboration with another research effort focused on attaching nanotags to bats. The study site was selected along a refuge road to catch a large number of tree bats breeding on or migrating through the island. Bats were captured using standard protocols, which included mist-netting with multiple stacked nets along bat movement corridors. Approved safety protocols were implemented, including handling of bats with fresh nitrile gloves every capture, holding of bats in sanitary bags, and sanitization of nets between capture periods to prevent the spread of disease among bats. All protocols were compliant with the US Fish and Wildlife Service National White-nose Syndrome Decontamination Protocol designed to prevent the further spread of the disease by researchers and spelunkers (<https://www.fws.gov/midwest/endangered/mammals/BatDisinfectionProtocol.html>).

After a bat was captured, the researchers collected a range of morphometric and health data and a nanotag was attached to the animal for the concurrent study. If the individual was determined to be in good condition for further testing (as determined by reproductive status and body mass), we used the animal in a test to determine the detection range of the camera system. Each bat was brought to a set distance from the camera system (0.5, 10, 25, 50, and 100 m). At each distance, both the front and the back of the bat were shown to the camera system for 15 seconds while video was taken continuously over that time period. After video was collected from each of the distance treatments, the bat was released at a randomly selected distance step (e.g., 0.5, 10, 25, 50, and 100m) from the camera system. Because the bat was being held by a person who was also reflective to the cameras, each individual bat was positioned above the researcher's head (Figure 17). The nIR illuminators were positioned in the middle of the system to expose the area in front of the camera with direct nIR luminance to focus maximum illumination in the area of greatest overlap of the stereo camera pair.



Figure 17. An example image showing a Red Bat at 25 m from the camera system on September 9, 2015 in Parker River National Wildlife Refuge.

Results: detection rates and distance measurements for known bat species

Over the course of the two study nights, we tested six individuals with the camera system across the range of distances (0.5-100 m). We captured only two different species: the big brown bat (*Eptesicus fuscus*, n=5) and the eastern red bat (*Lasiurus borealis*, n=1). The big brown bats were caught on September 8 and the eastern red bat was caught on September 9. Only two bats were not tested at all distance treatments—two big brown bats on September 8. These bats were released before the 100 m treatment to ensure a safe departure for those individuals.

None of the tested bats were clearly seen further than 25 m from the camera system. Past the 25 m treatment, the researchers holding the bat and monitoring the experiment could often be seen, but the bat could not. The alignment of the nIR illuminators appeared to extend the range of the nighttime vision of the system more than the previous configuration (with the illuminators placed on the outside of the camera bar to give a more diffuse nIR illumination to the study area) but not so much so that animals could be clearly seen 50 m away (Figure 17). Placing the illuminators in the middle of the system also had another effect, where large objects in the images were reflecting a large amount of nIR illumination (in this case, the researchers holding the bats and administering the experiment). In some cases, this high reflectance could make the bat more difficult to detect in the image.

While the bats were detectable in the video footage in many cases, it was never possible to identify species. Defining marks or silhouettes were not observable on the images, and the size differences between the two species were not large enough for wingspan to be differentiated with the stereo-optic estimation processes. In addition, when the bats were released by the BRI researcher, the camera system recorded a distorted image of the flying bat. This motion blur is likely caused by the high flight speed of the bat, the low frame rate of the system (1.7 fps), and the low light conditions. Combined with the poor resolution of the nIR images, the motion blur made it impossible to identify the target as a flying bat.

In summary, bats could only be reliably detected within 25 m of the camera system, and it was difficult to assess differences between species. These results match our earlier results when using non-bat test objects to test the range of the system. While nIR illumination has proved useful in other studies monitoring bat movements around turbines (Cryan et al., 2014), they will likely not be useful for this system unless the nIR illuminators are mounted on the turbine structure or blades, due to the distance between the camera system mounting surface and the nacelle in current wind power technology.

Testing to Detect Birds: Technical Approach and Results

Technical approach

Video Analysis

When conducting video analysis, reviewers proceeded through video footage and recorded and marked any animals observed within the segment. The reviewer marked every instance where the animal was present in both cameras (though marking methods depended on the version of the review software used). In cases where the review took more than one hour for one 5-minute video segment, the reviewer would stop searching for, marking, and recording non-eagle birds to facilitate review of a large number of segments.

When Bald Eagles were observed, the identifier recorded the age and plumage of the bird. In cases where the identifier wasn't able to distinguish the exact plumage class, the categories <Basic III and >Basic II were used. Where no age and/or plumage categorization could be made, the identifier noted the age and/or plumage as "Unknown".

In order to assess identification criteria and analyst identification variance, an audit was conducted on 20% of the non-eagle animals observed with the aim to have agreement rates $\geq 90\%$. All Bald Eagles were audited as part of a complete Threatened and Endangered Species audit. Disagreements were arbitrated to come to a final assessment of species identification.

While fisheye image rectification methods were being refined, the team used an alternative method of calculating object distance from the camera by using pixel measurements of wingspan or body length from the image combined with camera specifications (sensor size and focal length). If the actual measurement of the target object was known, and the proportion of pixel measurements in an image of a certain size was known, then distance could be calculated (Kendal, 2007). There are few published accounts of Bald Eagle wingspans, and most wingspan ranges include the smaller, southern US population of Bald Eagles. In order to obtain the best calculated distances for the pixel count method, wing chord to wingspan ratios were calculated from the Imler study (Imler and Kalmbach, 1955) and applied to wing chords recorded for northern populations of Bald Eagles (Buehler, 2000). Pixel measurements of target object wingspans or body lengths were obtained for the first frame of camera detection and the first frame where species identification was possible. If neither were possible, a measurement was taken of a frame with an eagle with flat extended wings or a clear body length visible (see Figure 18 for a full body length example, and Figure 24 for a fully extended wingspan example). Pixel counts used for the measurements were made using the New Media video viewer developed for this project.

In the final weeks of the project, while UMaine SCIS continued to work on fish-eye rectification methods, stereo-optic position estimates were calculated to track animal movement throughout the video footage. Effort was focused on tracking those animals that had recorded pixel count measurements as a means for comparison, though additional animals were tracked where possible.

Bird flight tracks were estimated using the HiDef viewer software. Calibration is necessary first to allow for track estimation. However, calibration of the video review files required specific “partly cloudy” conditions, which was not available during the recording of every chunk. Therefore additional software development work was performed on the HiDef viewer software to allow calibration data to persist between chunk files, allowing a single calibration routine to be performed for each placement of the camera system. Markers were placed on the object of interest in both the left and right cameras in each frame in which the animal was observed, and then exported to a .csv file, including X Y and Z coordinates for the animal of interest.

Daytime Detection Probability

Detectability in camera systems is often assessed using objects that are positioned at a known distance from the system. This is a useful approach, and one we used in early testing of the system, but it is difficult to know if these results are representative of how the camera system performs with animals where the shapes and behaviors deviate from the test objects.

In an attempt to reconcile this issue, we decided to evaluate animal detection probability using distance sampling, a standard method for independently testing the proficiency of human surveyors. This method has been used for over 15 years in ecology, and was pioneered for use with surveys that employ line transects (e.g., seabird boat surveys, Buckland et al., 2001). In this approach, the distance from the observed animals to the transect center line can be used to estimate the probability of detecting any animal within a certain distance to the center line. The core idea is that the probability of detecting an animal will decrease with increasing distance between the animal and the transect center line. In the application of this idea to camera observations, we would say that the estimated true density of objects within a circular area around the camera system is expressed as:

$$\hat{D} = \frac{n}{2\pi \int_0^\omega r \hat{g}(r) dr}$$

Where the numerator n is the number of objects detected in the survey area and the integral represents the summed unconditional probability of detecting an animal across all distance classes (ω , which vary by distance r from the camera system). The detection function, $g(r)$, is flexible in the distance detection framework, allowing multiple key functions to describe the change in detection over r . In this study, we used a hazard detection function where:

$$g(r) = 1 - \exp\left(-\left(\frac{r}{\sigma}\right)^\theta\right)$$

In this relationship, σ and θ govern the decrease in detection probability with r . We choose this distance detection function because detection probability because a visual inspection of the data indicated a broad shoulder, which hazard functions can fit well. These types of curves have patterns where detection probability remains high for some period before decreasing fairly rapidly.

To parameterize this model, we gathered distance of first detection data from 59 eagles and other birds using the pixel count method due to its increased accuracy at larger distances from the camera (see *Results - Video Analysis* below). Note that because the distance to the animal was estimated only after the animal was identified, what we are really measuring is the probability an animal is detected by the system and can subsequently be identified by a researcher. Thus, we are not estimating detection probability, and the true detection probability is higher than the marginal detectability and identifiability estimate we are obtaining here.

We divided the area observed by the camera into five distance categories (0-50 m, 50-100 m, 100-200 m, 200-300 m, 300-500 m) and assigned each observation to a distance category. Additionally, we were interested in what the detection curve of the camera would look like if all large birds larger than 14 pixels were identified, so we used the average body size of large birds to set a distance threshold to meet that criteria (400 m), and simulated a data set where an observation was equally likely from 0-400m, corrected for the area of the distance bins, then estimated the detection curve for these data. Essentially, this simulated a 400 m strip transect where all animals were detected that we could provide context for our results given the assumed theoretical maximum distance that a large bird could be identified.

We used a Bayesian modeling framework to set uninformative priors for parameters of interest (σ and θ), then calculated the posterior probability based on the likelihood of the data given the parameters and the prior information (in this case uninformative priors that assumed minimal knowledge) using Markov Chain Monte Carlo (MCMC) sampling methods with the programs JAGS (<http://mcmc-jags.sourceforge.net/>) and R (R Core Team, 2016). The MCMC sampler was adapted for 500 iterations then a burn-in of 5,000 iterations was used before we sampled 10,000 iterations with a thinning rate of 3. Stability of the parameter estimates was assessed visually and by using the Gelman-Rubin convergence diagnostic.

Results

Video Review

Video review took place from November 16, 2015 to June 23, 2016 on video footage recorded from September 15, 2015 to January 15, 2016. During that time, 1,141 objects of interest were found in

331 video segments while 1,225 reviewed video segments contained no birds. A total of 51 bald eagles were found in 30 video segments (Table 3). Birds were categorized by observer confidence in the identification (Definite, Probable, and Possible), and eagles were also categorized by age (Adult, Subadult, and Unknown; see Table 4 and Figure 18-Figure 28). There were 49 Bald Eagles observed at the Eagle Test Site, and two observed at the Turbine Site.

Table 3. Animals identified in the video analysis at two observation locations recorded from September 15, 2015 to January 15, 2016

Species	Total
Common Loon	6
Unidentified Loon	4
Herring Gull	1
Ring-billed Gull	4
Unidentified Gull	338
Mallard	2
Common Eider	1
Canada Goose	17
Unidentified Duck	43
Bald Eagle	51
American Crow	33
Common Raven	1
American Crow or Common Raven	6
American Goldfinch	6
Unidentified Swallow	3
Cedar Waxwing	3
American Robin	1
Unidentified Passerine	28
Unidentified Bird	484
Unidentified Invertebrate	34
Unidentified flying object (animal-origin)	75

Table 4. Bald Eagles identified in the video footage from September 15, 2015 to January 15, 2016. Observations are broken down into age ranges (Adult, Subadult, and Unknown) and by observer confidence in the identification (Definite, Probable, Possible). See Figure 18-Figure 28 for example images for eagles in each category.

Species	Total	Number Observed
Adult	Definite	10
	Probable	6
	Possible	1
Subadult	Definite	5
	Probable	6
	Possible	4
Unknown	Definite	2
	Probable	8
	Possible	9
Total	Total	51

Review Time

In order to assess video review time, we compared number of video segments, number of objects, and number of marked frames per hour. Processing of video includes searching, marking, and identifying targets. The mean processing time was 4.4 video segments (5 minutes each)/hour and the analyst marked 62 pairs of marked objects/hour. Within a year (365 days) for daytime video segments (12 hours per day), the camera system can generate 144 video segments per day and 52,560 daytime video segments per year. Therefore, processing an entire year of data, in the software's current configuration (no recognition modules or specialized algorithms), would take an analyst 5.97 years, where a year is defined by 250 business days at 8 hours a day. Four months of

video segments would take 2 years. One goal of the project was to reduce manual inspection and video analysis by 70%. Tools in the new video review software, namely the unique marking, masking, and the video playback features, improved the chunk processing per hour by 16% from 3.8 video segments per hour to 4.4 per hour.

There were several factors that affected the amount of time it took to review video footage. Reviewing footage that did not contain any birds was a quicker process than reviewing footage that did contain birds. Review times were impacted by weather, as windier footage resulted in more movement of water and trees. Tree and water movement often impacted viewer playback speeds due to increased image file size. The amount of cloud cover also impacted video review time, as days with broken cloud cover and wind again caused more of the background image to come through the mask. Tree, water, and cloud movement also made searching for objects of interest more difficult, as less of the background was masked and the reviewer had to differentiate between background images and animal movement. For footage with birds present, the amount of time the individual or groups of birds were on the screen and their flight behavior both impacted the amount of time it took to review. Birds flying directly across the screen took less time to mark than birds that flew in circular milling or soaring patterns further from the camera. The number of birds present in video footage increased the amount of time it took to review, especially with multiple soaring or milling birds. Footage recorded on overcast or rainy days impacted the ability to identify birds to species, and often slowed the review process as any bird images would be of lower quality or color rendition.



Figure 18. Bald eagle adult observed on January 13, 2016. Identified as a definite Bald Eagle adult, older than Basic II plumage, at the Eagle Test Site. Pixel measure distance was 27-34 m at 1st and ID frame. The stereo-optic position distance measurement at that point was 16.8 m.




Figure 19. Bald Eagle adult observed on January 13, 2016. Identified as a probable Bald Eagle adult, plumage greater than basic II. No pixel distance possible. The stereo-optic position distance measurement in the identification frame was 113.9 m.




Figure 20. Bald Eagle adult observed on October 30, 2015. Identified as a possible Bald Eagle adult of greater than basic II plumage. Pixel distance measured at 225-271 m in the 1st frame, 202-243 m in the ID frame. No stereo-optic position distance measurement was calculated.



Figure 21. Bald eagle subadult observed on March 28, 2016. Identified as a definite Bald Eagle subadult, plumage less than Basic III. Pixel distance in first frame measured at 31-40 m, and identified at 26-33 m. The stereo-optic position distance measurement in the first frame was 25.8 m, and 22 m at the identification frame.



Figure 22. Bald Eagle subadult observed on January 6, 2016. Identified as a definite Bald Eagle subadult, plumage less than Basic III. Pixel distance in first frame measured at 46-59 m, and was identified in the same frame. The stereo-optic position distance measurement was at 29.7 m.



Figure 23. Bald Eagle observed on December 28, 2015. Identified as a definite Bald Eagle subadult of unknown plumage. Pixel distance when identified was 54-69 m. The stereo-optic position distance measurement when identified was 44.7 m.




Figure 24. Bald Eagle observed on January 13, 2016. Identified as a probable Bald Eagle subadult of unknown plumage. Pixel distance when identified was 215–258 m. The stereo-optic position distance measurement when identified was 90 m.




Figure 25. Bald Eagle observed on November 13, 2015. Identified as a possible Bald Eagle subadult with plumage less than III. Pixel distance when identified was 223–268 m. The stereo-optic position distance measurement was 123.8 m in the identification frame.



Figure 26. Bald Eagle observed on December 25, 2015. Identified as a definite Bald Eagle of unknown age and plumage. Pixel distance for the first and identified frame was 43–51 m. The stereo-optic position distance measurement was 24.6 m in the identification frame.



Figure 27. Bald Eagle observed on January 13, 2016. Identified as a probable Bald Eagle of unknown age and plumage. The pixel measurement for the identification frame was 47–60 m, and the eagle was measured at 49–63 m in another adjacent frame. The stereo-optic position distance measurement was 46.1 m in the identification frame and the measurement was 45.4 m in the adjacent frame.




Figure 28. Bald Eagle observed on March 21, 2016. Identified as a possible Bald Eagle of unknown age and plumage. No pixel measurements were possible for the first or identification frames, but the eagle was measured at 473–568 m in another frame. The stereo-optic position distance measurement was 158.7 m in the same frame.

Audit results

The 20% random audit passed at a 93% agreement rate. The T&E audit, which consisted of all identified Bald Eagles, resulted in a 67% agreement rate. Disagreements primarily occurred for eagles that were further away and in low light conditions. Identification criteria were amended and eagles in disagreement were arbitrated for final answers.

Detection and identification distances

We made separate measurements of wingspans for birds based on pixel count methods and stereo-optic estimation methods for comparison. Pixel distance measurements were possible for 41 birds identified to the species level. Results showed that small birds (wingspan < 400 mm) were identified to species at distances of 20–44 m; medium birds (wingspan 850–1,110 mm) were identified to species at 12–88 m; and larger non-eagle birds (wingspan 1,270 –1,485 mm) were identified to species at 32–123 m. Eagles with wingspans ranging 2,036–2,446 mm (Buehler, 2000; Imler and Kalmbach, 1955) were identified to species at 26–352 m. The furthest camera detection with motion segmentation was a Bald Eagle estimated at 509 m.

Pixel count distances for both the first stereo-optic frame (detection distance) and first frame of positive species-level identification were possible for eleven Bald Eagles and six other birds. These distances are represented in Figure 29 and Figure 30. For eight of the eagles, species-level identification was possible when the bird was closer to the camera, while the other three were detected and identified in the same frame. Eagles first detected at farther distances from the camera tended not to be identified at first detection, but rather to be identified once the individual had flown closer to the camera system (Figure 29).




Figure 29. Distance from camera for 1st detection and identification detection of eagles. Measurements shown are for pixel distance ranges.

Stereo-optic position estimates were made for 601 non-eagle bird images of 43 individual birds. Stereo-optic position estimates were made for 2,033 eagle images of 51 individual eagles. Using the stereo-optic position estimation method, small birds were identified to species from 3–73 m; medium birds were identified to species at 2–58 m; and larger non-eagle birds were identified to species at 29–67 m. Using the stereo-optic position estimation method, eagles were identified to species at 16–146 m. The furthest detection with motion segmentation measured using the stereo-optic position estimation method was 218 m. See Figure 31 for examples of stereo-optic position estimation tracks of American Crows moving through the study area.




Figure 30. Outlier box plot of distances from the camera for birds first detected in the camera system and Distance identified to species (AMCR = American Crow, AMGO = American Goldfinch, CANG = Canada Goose). Distances are based on pixel count calculations. Single lines for AMGO and CANG indicate only one observation.




Figure 31. Stereo-optic position estimates of a flock of crows flying through the turbine site. Each crow image represents an observation of a crow in a specific frame. The different colors are used to identify individual crows that were tracked.

As shown in Figure 32, there were 44 instances where eagle distances were measured using both the stereo-optic position estimate and pixel count methods. In Figure 33, there were 51 instances where distances were measured using both the stereo-optic position estimate and pixel count methods. In both cases, measurements became less accurate the further the birds were from the camera. There was a notable divergence between the distances calculated using the two methods when the pixel measurements indicated that eagles were observed further from the camera system, beginning around 80–100 meters. These differences highlight some of the problems we have had with the stereo-optic position estimation, which stem in large part from the fisheye lenses and the difficulty in rectifying these images accurately. Stereo-optic depth calculations become less precise the further the animal is from the camera, as dictated by the lateral distance between the two cameras (Yu et al., 2013) and the nature of fisheye camera lens resolution.




Figure 32. Distance from the camera as measured by two methods: the stereo-optic distance algorithm (blue) and using the known wing length of the species (green).




Figure 33. A Box plot of distance from camera measured using pixel distance calculations and stereo-optic position distance estimations for birds at the Eagle Test and turbine sites. AMRO is american robin, AMGO is American Goldfinch, AMCR is American Crow, RBGU is Ring-billed Gull, HERG is Herring Gull, COLO is Common Loon, CANG is Canada Goose; in order from smallest to largest average animal size.

Influence of weather

Weather affected the probability that an animal would be identified once detected (Figure 34). There were similar percentages of birds identified to species and the group level in clear, cloudy, partly cloudy, and snowy weather, with a lower percentage identified in drizzle, and the lowest percentages identified in rain and fog. Notably there were also fewer observations of animals during the periods of weather where identification was difficult, suggesting these conditions were relatively rare, fewer animals were flying under those conditions, or that they reduced the chances that animals would be detected by the system.




Figure 34. The percentage of birds identified to the group or individual species level, broken out by weather condition during observations. Footage with no animals observed is also included (empty).

Daytime Detection Probability

The overall detection probability of eagles within 500 m of the camera system was estimated to be 6% (95% credible interval: 1.2-13.2%) based on estimates of 72.6 (95% Credible Interval: 19.6, 130.2) and 2.6 (95% CI: 1.8, 3.9) for σ and θ , respectively. Detection/identification probability stayed near 100% for the first 40-50m then descended quickly (Figure 35). The strip half width was 104 m (95% CI: 38-156), which indicated 50% of all detections were before that distance.

The overall detection/identification probability was estimated to be about 20% over the entire 500 m range. The instantaneous detection probability approached zero around 400m, suggesting that detections were rare past 400 m. To understand how the rates of identification might influence these detection curves, we simulated data under ideal conditions where all animals detected at less than 400 m were successfully identified. The detection probability of the idealized data was 67%, which suggested that we were both detecting and identifying about 9% of the total number of eagles that we could within the 500 m sampling area.




Figure 35. Detection curve for Bald Eagles observed by the camera system (green) with 95% credible interval (grey). The theoretical identification limit for eagles is shown in blue.

Non-eagle detectability was lower than what we saw with eagles. The overall probability of detecting a non-eagle in a 500 m area around the camera was 1.5% (95% CI: 0.004–2.2%), which was governed by an estimate for σ of 48.3 (95% CI: 9.5–62.4) and θ estimate of 4.2 (95% CI: 2.8–6.2). Detection probability held near 100% up to 30–40 m from the camera system then decreased rapidly and approached 0% detection near 150 m (Figure 36). The strip half-width was estimated to be 55.1 m (95% CI: 13.7–70.6). Correcting for the idealized detection curve over 500 m, as above, suggests that we both detected and identified about 2.2% of animals that we could within the sampling area.




Figure 36. Detection curve for non-eagle birds observed by the camera system (green) with 95% credible interval (grey). The detection curve for the theoretical case of perfect detection within 400m from the camera system is shown in blue.

Overall, these results were lower than expected. While our ability to detect and identify animals within 100 m of the system appears reasonable, our probability to detect and identify an animal decreased rapidly past that point. The primary reason for this result was that we were not measuring detectability and identifiability separately, as we initially planned in the project. The stereo-optic position estimation process proved to be less accurate past 100-150 m than we anticipated, and we were not able to use these position (and thus size) estimates in our assessment of detectability or identifiability over distance from the camera. As such, we were left with pixel estimation as our source of distance information, which relies on species identification to calculate an estimate. Thus, we know that probability of detection and identification is fairly low, but it is difficult to know what each of these probabilities are in reality. Testing with simulated data with 100% detection within 400 m suggested what we could achieve under ideal identification and detection conditions, but it is impossible for us to disentangle our probabilities of detection and identification in the current analysis. Anecdotal evidence suggested that the camera detects objects consistently out to 350-400 m, so we can speculate that our identification rates are lower than the detection rates, but we lack data to strongly support this hypothesis.

The choice of lens of the camera system is another possibility for why our detection/identification probabilities were low. Fisheye lenses have many advantages (excellent depth of field and field of view) but range is not one of them. Objects far away from the camera are perceived as smaller we compared to similarly far away objects in rectilinear lenses, and thus make up a proportionally

smaller number of pixels for such lenses with a similar camera sensor. This could influence both detectability and identifiability for our system, and would be an easy change to implement and test in future projects.

Table 5. Bird distances calculated using the pixel count and stereo-optic position estimate methods. Smaller birds include American Goldfinch and American Robin; Medium birds include American Crow, Common Loon, and Ring-billed Gull; and Large birds include Canada Goose and Herring Gull. WS = Wing span.

Bird size	Pixel Count		Stereo-optic	
	ID distance range (m)	Furthest camera detection (m)	ID distance range (m)	Furthest camera detection (m)
Eagles (2036-2446 mm WS)	26-352	509	16.0-145.5	218.3
Large (1270-1485 mm WS)	32-123	194.4	29-67	125.2
Medium (850-1175 mm WS)	12-88	132.5	2-58	99.5
Small (<400 mm WS)	20-44	70.3	3-73	129.5

Improving the camera calibration: Technical Approach and Results

Technical Approach

We determined that the calibration routine employed by the HiDef Viewer used a simple model for the fisheye lens characteristics with multiple sample points and local pin-hole camera assumptions. The assumptions made in this model resulted in a coarse calibration with extremely low accuracy in depth estimates. In order to improve on this for the new viewer and get better 3-D estimates at distance, we used methods provided from OpenCV and Matlab (for more detailed information on the camera calibration processes, see the Matlab documentation⁷). While the initial (HiDef) viewer performed a basic calibration, a more thorough approach through OpenCV or MATLAB was thought to be likely to yield an improved set of calibration parameters. However, until recently neither toolset provided sufficient support for calibration of cameras equipped with fisheye lenses. OpenCV later supported such calibration, but initial results were mixed, so we were challenged in using these systems to achieve our goals.

Camera calibration is typically performed by capturing several sets of images containing a test pattern of known size and shape from multiple views. Given these images, the calibration process attempts to estimate a camera's intrinsics, extrinsics, and distortion parameters. The calibration process searches these test pattern images and attempts to compute a series of camera parameters further described below.

We used an enlarged checkerboard for our trials to cover the wider space necessary for extreme wide-angle lenses. These images were used by a learning/training algorithm to approximate the camera parameters. Once each camera was calibrated separately, the resulting parameters could be used to perform a stereo calibration. Most importantly, this process produced the rotation and translation matrices between the two cameras later used to perform 3-D reconstruction. Once each camera is calibrated separately, the resulting parameters can be used to perform a stereo calibration. Most importantly, this process produced the rotation and translation matrices between the two cameras, later used to perform 3-D reconstruction.

Individual camera calibration involved determining three primary sets of parameters: camera intrinsics, camera extrinsics, and distortion coefficients. Camera intrinsics effectively captured the focal length of the lens, the optical center of the image (principal point), and any skew introduced by differences in the camera's CCD (image capture sensor array). The extrinsic parameters related the camera's coordinate system to the image plane, and consisted of a rotation and translation matrix. Distortion coefficients reflected the two types of image distortion often introduced by the lens: radial distortion and tangential distortion. Radial distortion occurs with most lenses, and refers to distortion near the edges of the image. Tangential distortion occurs less frequently, but can be present in cameras where the lens and image planes are not parallel.

⁷ <http://www.mathworks.com/help/vision/ug/camera-calibration.html>




Figure 37. Types of radial distortion. Source: <http://www.mathworks.com/help/vision/ug/camera-calibration.html>




Figure 38. Tangential distortion occurs when the image plane and camera lens are not parallel. Source: <http://www.mathworks.com/help/vision/ug/camera-calibration.html>

Distortion coefficients were particularly important, as most camera calibration methods operate around a pinhole camera model, which has no lens. Camera calibration, once completed, enables the measurement of distances and sizes of objects in a scene.




Figure 39. Matlab camera calibration. Shown is a list of calibration images (left), the pattern as identified by the calibration application (center) in one such image, and rough 3-D estimates (and errors) of the pattern’s location relative to the camera (right).

Results

Although the OpenCV camera calibration tools now support a fisheye camera lens model, lenses with a high FOV (>170 degrees) produce more distortion than the model can compensate for. As such, efforts to obtain a reliable calibration using the standard approach proved impossible during our 2016 trials. Future versions of OpenCV and MATLAB modules might have better support.

Improving the 3-D modeling: Technical Approach and Results

Technical approach

The ability to compute three-dimensional geometric information from two or more traditional, two-dimensional images is central to a successful stereo-optic imaging system. In the context of this project, 3-D information was obtained via a triangulation process: given a complete stereo-camera calibration, the distance to an object (or its depth in a scene) could be computed by examining the difference in its location (relative to the optical center) in each half of the stereo-pair. The success of this disparity-based approach relies heavily on the accuracy of the calibration processes as described in the previous section (Improving the camera calibration).

The triangulation depicted in Figure 40 is simplified, and assumes a pinhole camera model. In the real world (and in our application), lens distortion introduces a rather significant amount of error. Much of this error can be compensated for if the camera parameters as described in the previous section are well defined. Lacking adequate camera parameters, this process is unlikely to produce reliable distance estimates.




Figure 40. Triangulation of a 3-D point x is performed by projecting intersecting lines from image points y_1 and y_2 with each camera's focal point O_1 and O_2 .

As noted, the triangulation process requires a pair of “matching” points – one from each camera. These points can be manually derived, by having a user select a point from each image, or more typically, through a process known as cross-correlation. Computing the cross-correlation of two images f and t is a way to measure the similarity of t to a segment of f . This is also known as a sliding dot product, since it calculates the dot product of the image vectors at all positions within f and finds the closest match to t . For image processing, it is useful to use normalized cross-correlation, which accounts for some exposure and perspective differences between images. In the normalized cross correlation formula⁸

⁸ https://en.wikipedia.org/wiki/Cross-correlation#Normalized_cross-correlation

$$\frac{1}{n} \sum_{x,y} \frac{(f(x,y) - \bar{f})(t(x,y) - \bar{t})}{\sigma_f \sigma_t}$$

x and y are the horizontal and vertical positions in the frame, σ is the standard deviation of the values of either the f or t function, and n is the sample size.

The goal of this approach is to take a template image—say the object of interest from one image—and find the matching region in the other image. Effectively, this is done by sliding a window the size of the template across the other image and comparing the contents, until a best match is found.



Figure 41. Normalized cross-correlation was used to successfully match the rear tire of a vehicle across a pair of images captured by the camera system. The template used is the area selected in blue on the left image, and the red region in the right image depicts the resulting match.

While the template matching approach is viable for many applications, it proved somewhat problematic with our camera system. Part of this was due to the size of the source images relative to the size of the templates of birds: there may be many similar 16x16 pixel regions in a 4400x6600 image. Additionally, given the masking and motion segmentation process, there may be cases where the corresponding match may not be fully or equally captured by both cameras. As such, initial efforts to triangulate points in 3-D were based around user-selected pairs of points. Even with carefully selected points, camera calibration issues prevented reliable distance (and hence size) estimates using the system.

Given a more successful stereo-calibration, researchers at BRI and UMaine SCIS proposed a number of additional improvements to the viewer software, focused on further increasing user efficiency and software utility. Beyond the suggestions posed in the section above on “[Improving the viewer](#)”, an improved 3-D module would offer the following:

- Ability to produce reliable distance and size estimates for objects of interest.
- Ability to model the locations of objects in 3-D across multiple frames.
- Potential to filter data by distance criteria.

The initial viewing software does provide some support for estimating 3-D locations of objects, but calibration issues limit the reliability of this information (Figure 42). Accurate 3-D position calculations would greatly improve the viewer, and could enable object size estimates or filtering of data by size or distance. Size estimates can be obtained by computing the linear distance between a pair of 3-D points, such as wingtips. Masking or filtering a scene based on distance criteria is more computationally intensive, but fundamentally similar: stereo calibration enables the creation of a disparity (depth) map, which effectively captures the distance to all points in a stereo image (see <https://www.mathworks.com/help/vision/examples/depth-estimation-from-stereo-video.html> for more information and an example of this approach).

Such improvements could allow users to effectively “query” a scene, requesting only objects or regions that meet a set of predefined criteria (e.g., objects larger than 2 m in size at distances between 200 and 400 m). Improvements to calibration would also further users’ ability to examine an object’s trajectory over time, as captured in a series of images.




Figure 42. Example flight path of an eagle estimated using a stereo-optic position algorithm.

Objective (3) is not strictly tied to the 3-D module, but stems from ongoing work by UMaine researchers focused on a qualitative shape representation approach. Unlike many traditional object recognition methods such as template matching (described above), which focus on heavily quantitative visual information such as color or texture, this approach instead centers on a flexible, qualitative shape representation. A qualitative approach provides several unique advantages: (1) it is robust to the scale and the rotation of the shape; (2) you can adjust the level of detail you want to describe in the shape; and (3) the shapes that are observed do not have to be perfectly described for them to be assigned a best match. In order to evaluate the shape similarity measures developed by the UMaine team, researchers from BRI assembled a corpus of approximately 100 images of birds, as captured by the stereo-optic imaging system. Each of these bird images was preprocessed using a standard vectorizing module that takes a black and white image, traces the outline of the dark object (e.g., Latecki discrete curve evolution; [Latecki and Lakamper 1998](#)), and takes this outline polygon and reduces the number of vertices to the 12 most relevant vertices for describing the original shape. Twelve vertices is a useful number for describing shapes of flying birds, as during flight, birds have four points projecting out from the center of the body (two wings, a head and tail). Broadly, this shape looks something similar to a cross, which is an object with four subparts that is easily described using twelve vertices. These primary characteristic edges of a bird outline are mapped on an abstract qualitative spatial description, which has a straightforward similarity measure that captures the relative similarity of sets of shapes. Such a shape similarity is then used in a pairwise comparison to determine the following for each example shape: (1) which other example matches most closely and (2) what are the n closest other matches?

Results

To explain this process, we will describe some examples from the first objective. We use a sample of five bird pairs and each shape is described using: (a) the simplified shape of the unknown source image, (b) the simplified shape of the best match from the known shapes archive, (c) the unknown source image as captured by the camera system, and (d) the best match of the known shape as captured by the camera system (Figure 43-Figure 47). Error between shapes is expressed as a percentage, and reflects both the angular error between equivalent vertices, as well as the distance error between equivalent edges. After analyzing 100 sample images, the average 'best match' had a percent shape-distance error of less than 7%. Even with shapes that have an error greater than 7% (e.g., Figure 47, in which one wing is mostly missing due to occlusion), the matching system does a reasonable job of matching the shapes.




Figure 43. Simplified shapes with error of 3.01%, along with camera system images that the shapes were based upon.



Figure 44. Simplified shapes with error of 4.25%, along with camera system images that the shapes were based upon.



Figure 45. Simplified shapes with error of 5.3%, along with camera system images that the shapes were based upon.




Figure 46. Simplified shapes with error of 6.9%, along with camera system images that the shapes were based upon.



Figure 47. Simplified shapes with error of 8.5%, along with camera system images that the shapes were based upon.

Using these bird pose shape clusters to get a wingspan estimate or a body length estimate for birds which are oriented or angled towards the camera enables these estimates to be straightforward. Please note that the wingspan/body length could only be reported as pixel length until a distance estimate of reasonable quality would be supported by stereo triangulation (e.g., the problem with the calibration prevents a wingspan/body length estimate in inches).

In summary, we have used imagery from the camera system to build simplified shape profiles that have been compared to reference images to look for similarities between unknown animals and known species. These techniques are extensions of prior methods and provide a mechanism for filtering large numbers of shapes and finding objects that are most similar to a reference shape (e.g., a Bald or Golden Eagle).

ACCOMPLISHMENTS

Summary of Technical Accomplishments

- The 29-megapixel camera system successfully captured 16,173 five minute video segments in the field.
- We demonstrated a long deployment time of >90 days in the field under variable weather and temperatures.
- Using nIR, the best detection distance for bat-sized object was 60 m. This led to a decision during Go/No-Go to focus research efforts exclusively on daytime monitoring and to redirect resources towards improving the video post-processing viewer.
- Bird event post-processing viewer was redesigned.
- The system detected eagles up to 500 m away using the fisheye wide-angle lenses when animals were independently identified by an observer, and up 350 m away when identified through the camera imagery alone.
- Data was transferred from the field to the BRI Video Lab with no known data loss.
- Standard operating procedure (SOP) was developed using a hierarchical system for a) detecting an object; b) categorizing the object as bird or other; c) categorizing the bird into species group (e.g., waterfowl, raptor, passerine); and d) identifying to genus or species.
- Using distance-sampling methods, a probability curve for detecting and identifying birds based on size and distance from the camera was developed.
- The 3-D algorithm for object isolation and tracking was tested. The image rectification (flattening) required to obtain accurate distance measurements with fish-eye lenses was determined to be insufficient.
- MATLAB and OpenCV were used to improve fish-eye lens rectification towards the center of the image.
- We developed new methodologies for describing shapes and their similarity to other shapes that can be useful in filtering the large amount of data this system could produce.

Presentations

- Dorr, Christopher H., Longin Jan Latecki, and Reinhard Moratz. Shape Similarity Based on the Qualitative Spatial Reasoning Calculus eOPRAM. International Conference on Spatial Information Theory. Springer International Publishing, 2015.
- Conference on Spatial Information Theory XII, Presentation (UMaine)
- AWEA Offshore, Poster (BRI)
- Adams, Evan M., Steve Burns, Emily Connelly, Christopher Dorr, Melissa Duron, Andrew Gilbert, Wing Goodale, and Reinhard Moratz. Stereo-optic high definition imaging: a technology to understand bird avoidance of wind turbines. NWWC 2016, Presentation (BRI)

Project Milestones and Technical Accomplishments

	FY2014	FY2015	FY2016
Quarter One	NA	Milestone 1: Manufacture and lab testing of new camera and near-infrared (nIR)	Milestone 4:-6 Data management at BRI Lab, video lab analysis, and data analysis. The final method for data export results in no less than 95% of wildlife video images being transferred from the camera to a video catalog and associated database. Video analysis of daytime eagle-sized objects at 10-50 m result in 90% species-level identification, 50-100m=70% species level identification, 100-200m=50% species identification and 200-500m=30% species identification. Quantify the error in ID rates and estimating 3-D position.
Quarter Two	NA	Milestone 2: Preliminary data collection and system refinements. Testing has an outcome where total detection rates of 20 bird and bat-sized targets are at least 95% within 10 m during the day, 25% within 250 m during the day, and at least a 25% detection rate within 100m at night.	Milestone 5-6: Video lab analysis, data analysis. Video analysis of daytime eagle-sized objects at 10-50 m result in 90% species-level identification, 50-100m=70% species level identification, 100-200m=50% species identification and 200-500m=30% species identification. Quantify the error in ID rates and estimating 3-D position.
Quarter Three	NA	Milestone 3: Equipment deployment, equipment maintenance and evaluating methods of data export, and a limited trial of nIR Obtain sample size of 20 eagle images.	Milestone 5-6: Video lab analysis, data analysis, refine 3-D model. Video analysis of daytime eagle-sized objects at 10-50 m result in 90% species-level identification, 50-100m=70% species level identification, 100-200m=50% species identification and 200-500m=30% species identification. Quantify the error in ID rates and estimating 3-D position. The 3-D module reduces the information stream down to 1% of the original data and distinguishes bigger birds (>1m wingspan) from smaller birds with 90% reliability. Build an interactive bird event viewer that reduces the manual inspection and analysis time by 70% compared to the current software configuration.
Quarter Four	Contract signed	Milestone 3: Equipment deployment, equipment maintenance and evaluating methods of data export, and a limited trial of nIR Obtain sample size of 20 eagle images.	Milestone 6-7: Data analysis, refine 3-D model, reporting. Quantify the error in ID rates and estimating 3-D position. Report to DOE.

CONCLUSIONS

TESTING THE PERFORMANCE OF THE CAMERA

Receipt and initial testing

While the camera system proved to be quite robust, initial testing in the U.S. demonstrated some unanticipated difficulties associated with changing hardware components without full system testing. Upgrades to the camera system, while ultimately very positive for improving image quality and therefore detection rates, resulted in trigger issues that slowed down initial testing. Fortunately, the cameras are very sophisticated and rather simple changes to firmware settings on the camera allowed us to overcome the trigger problems completely. Otherwise, receipt of the system was relatively smooth and allowed us to transition to full testing of the capabilities of the upgraded system.

Testing for detection of bats

Detection range for bats was low with the nIR system, between 25-30 m, which limits the usefulness of this system for monitoring their avoidance behavior, as the system would be placed >30 m below the rotor swept zone of a typical turbine. This poor detection range was likely due to a combination of low frame rate, which limited the number of opportunities to view fast-moving bats, and the fact that nIR illuminators only exposed a limited portion of the area viewable to the camera.

Testing for detection of eagles

Eagles in many ways represent one of the easiest animals to detect with such a camera system. Bald Eagles, the only eagle species documented in our study, are large, often soaring animals that can be detected from far away and generally move slowly enough that they appear in many frames. The system detected Bald Eagles up to 500 m away from the camera, but a positive identification of an eagle is much more likely at 100 m or less from the camera system.

Testing for overall bird detection probability

The probabilities of detection and identification were confounded in our assessment of the system due to the methods of position estimation used. Generally, the camera system could detect and identify eagles better within 100 m of the camera, but the ability to do both decreased rapidly further out. Part of the reason for this result is the nature of the camera system: the wide-angle lenses are designed to monitor a large volume of space around each turbine. The orientation of turbine blades change with prevailing winds, so the system has to be robust to withstand these changes. Sampling a large volume of space reduces image fidelity for objects further away from the system, however. We also expect detections and identifications to be less common in foggy and rainy conditions.

IMPROVING THE SYSTEM

Hardware

Communications and data transfer

Communications with the system is limited by the ability to communicate over a standard local area network LAN Ethernet connection. If fast intranet/internet speeds are available, then in addition to being able to remotely access the system computer via the network, data transfer could also be achieved. However, over standard internet connection speeds (5-25 Mb/s), data transfer may be very limited, especially over the lower end of the bandwidth. Fast fiber connections to turbines and allowances for data transfer over these connections would allow data to be transferred remotely without physically swapping hard drives. While swapping drives was easily accomplished in the field settings utilized for this project, offshore or remote deployments onshore would mean reduced access to units, increased difficulty retrieving data, and long periods between data retrieval. Moreover, while swapping drives was an effective means of transferring data with zero loss, accessing the interior of the system cabinet may risk exposure to the elements during poor weather conditions and elevate risk to electrical components.

While wired connections will likely always be the preferred communication method for this system, due to the large amounts of data to be transferred, wireless connectivity can also transfer data at high bandwidth, and this could be a valid alternative method for remote system checks and data transfers. We were not able to fully test the long-range wireless connection to determine if this approach was sufficient to perform data transfer, and this is an important area for future testing. In particular, weather can negatively affect wireless transmission (Sim, 2002), and it may be important to check the system or look at recent data during periods of adverse weather, as these may be the times when communications are poorest.

Applicability and limitations of the camera system

The camera system hardware and its low-level software is able to drastically reduce the video stream by only storing parts of the image that differ from the full frames that are stored every five minutes. The algorithm and processing speed of the computer attached to the cameras suffices to detect moving objects that contain at least 20 pixels. A bald eagle at 400 m distance would typically correspond to 30 to 50 pixels.

Using a suitable mask, as provided with the new viewer, one can mask out the movement of non-target objects like waves, tree branches, and leaves. In such a setting, only cloud pieces and birds would be moving. Using the mask in addition to the base compression algorithm, the system should be able to reduce the data flow by 99% or more.

Currently, there are deficits in the hardware design of the current system that result in less than optimal distance estimates by stereovision. Stereo-based distance estimates would be more accurate if we could mechanically align the secondary camera's optical axis to be highly parallel to the master camera's optical axis. In such a configuration both cameras of the stereo pair point exactly in the same direction. Then the formula for the equisolid angle mapping function (Figure 48) of our fisheye lens ($r = 2 f \sin(\theta / 2)$ with $f = 16$ mm as the focal length and θ being the angular distance to the optical axis) would deliver adequate estimates because both angles have the same deviation. Only the differences between the angles for master and secondary cameras count, so the

joint systematic errors cancel each other out, resulting in greater accuracy in distance estimates. Finally, calculations of the distance error showed that the size of the stereo basis is sub-optimal at only 1.4 m.




Figure 48. Formula for equisolid angle mapping.⁹

The volume that the system can effectively monitor is determined by the maximum distance for object detection and angle of view. For the fisheye lens, the compromise between wider angle and smaller pixel size of the same object resulted in a reduced distance for possible object detection as compared to a rectilinear wide-angle lens. Additionally, object optical angles towards the lateral sides of the stereo basis reduced the effective distance between the cameras (e.g., an object 90 degrees from the optical axis at one of the sides would have the cameras behind each other with zero functional stereo basis). A rectilinear wide-angle lens configuration would address many of these considerations, and if properly designed could have nearly as large a monitoring volume as the fisheye lens configuration.

Software

Improving the Viewer

We made substantial improvements to the viewer that increased the speed of detection and identification of targets. The new viewer allows the biologist to use a black and white mode to assist in finding targets, a customized filter to exclude areas such as tree-lines from analysis, and a target marking system that allows targets to be followed through multiple frames.

Improving camera calibration

Using software tools available to us from OpenCV and Matlab, we were only partially successful in performing a better stereo-calibration, due to the non-rectilinear characteristics of the fish-eye lens. Ultimately, it may not be possible to fully calibrate this system with the software solutions available currently. Likely hardware solutions (i.e., increasing the focal length of the lenses) will be the most viable path to improving calibration and the accuracy of distance estimates.

⁹ <http://imagebank.osa.org/getImage.xqy?img=QC5sYXJnZSxvZS0yMC0yMC0yMjI1Mi1nMDAx>

Improving the 3-D modeling

We originally planned to create a 3-D space filter, which would have allowed us to remove distant moving objects like pieces of clouds and aircraft from analysis. Due to the low accuracy of the distance component, only two-dimensional filtering (in the form of a user-defined mask) was possible. The wingspan estimates based on automatic stereo-matching were also not reliable. The integration of distance-based filtering into the masking module of the new viewer software would require fairly basic code modification, if accurate distance calculations were available.

Shape recognition algorithm

The shape recognition algorithm successfully matched sample poses of bald eagles. The algorithm itself was described in Dorr et al. 2017. A framework for defining shapes of animals of interest was created and a mechanism for comparing these shapes was implemented. This technique shows promise for filtering the large amounts of data that is acquired when the camera is operating.

Improving the video review processing time

If video review processing time encompassed search, mark, and identification, and it is assumed that 1/2 the time goes to searching/marking and half the time goes to identifying (sometimes less) then introducing a shape recognition module could cut time searching for targets in half. In addition, if the 3-D modeling could also be implemented and measurements were known at the time of video review, then the amount of time spent identifying targets could also be significantly reduced. These improvements could at minimum improve review time to 8-10 video segments an hour for segments with birds and almost completely eliminate the need to review segments without birds.

RECOMMENDATIONS

Testing the performance of the camera

Receipt and initial testing

The initial receipt and testing process were challenging, but all issues were resolved. Broadly, we would recommend allotting a fair amount of time for coordinating with customs for importing large camera systems across international borders, and time for testing all components of the system upon arrival.

Nighttime operations

The nIR lighting we used was too bright for near objects and too dim for far objects, suggesting that we place nIR lighting closer (but not too close) to the area of primary interest. Since interaction with the turbine itself is of greatest interest, we recommend testing the placement of nIR lighting on the nacelle or just below it on the support structure (or both). If the lighting was placed on both the top and bottom of the nacelle, the RSZ could be illuminated regardless of the rotation of the turbine. Alternatively, nIR lighting could be placed around the support structure, shining upwards to illuminate from below. However, this type of lighting could result in over-saturation of the subject at the bottom of the RSZ and under-exposure at the top of the RSZ.

Potentially more effective steps to better detect bats could include both boosting the frame rate of the camera and increasing the portion of the survey area that is illuminated. We suggest considering thermal cameras instead of nIR illumination to achieve both these objectives. High frame rates are easily achievable with such systems, and nIR illumination is not needed to detect animals, so the effective range would only be limited by the optics of the thermal camera lens and sensor density. The downside of switching to a thermal camera system would be increased cost for high-resolution cameras and the difficulty of using a wide-angle lens for lower sensor resolution cameras. Additionally, species identification is probably not possible with thermal cameras, so species-specific avoidance rates would need to be achieved with additional identification methods like acoustics.

Future work using stereo-optic thermal imaging should be considered as the next step for nighttime monitoring as the increased detection range of those cameras should be more similar to the range seen for the camera system during the daytime, and we believe thermal imaging will likely be a better tool for achieving the monitoring objectives of the system.

Detecting and identifying eagles by the camera system

We suspect that our ability to identify eagles is what limits the detection/identification curve that we estimated using distance sampling. Thus, further improvements in our ability to estimate position and filter for eagle-like shapes will be important for development in this area, as those should improve our ability to identify eagles significantly. Improvements would include changes to allow for better stereo-optic position estimates, to more accurately measure both distance and the size of the animals observed. Many distant, soaring birds were marked as “Unidentified Birds”, but may have been Bald Eagles. Accurate distance and size measurements of these animals may have allowed reviewers to identify them as Bald Eagles.

Improvements to achieve this would include an added mechanical adjustment for the secondary camera's alignment and the use of rectilinear lenses like the Nikon 28 mm to improve 3-D tracking estimates. Also, adding a pan-tilt camera to track and focus on the object of interest based on measured trajectory by the stereo camera system would greatly enhance the resolution of the object, thus improving ID rates and size measurements. This camera should have a lens with at least a 200 mm focal length.

With such a system, one could add software that uses advanced object recognition algorithms, such as methods involving deep learning and convolution neural networks, an area of active research for automated image classification (Jia et al., 2014; Rodner et al., 2016). Using a training set of images, it may be possible for onboard computers to provide predictive capabilities of images supplied from initial chunk file processing. These algorithms have proven to be very fast (~2.5 ms per image on a fast graphics processing unit) and development is being extended to mainstream computers (Jia et al., 2014). Adding such a processing step to this monitoring system would allow for real-time detection of eagles or other birds.

Integrating deterrent systems

In a configuration with an added focus camera on a pan-tilt unit, the pan-tilt unit points to the object of interest. In the case of deterrence measures that are directional, it would be recommended to link the deterrent system to the pan-tilt unit, which would be mounted parallel to the camera. Then in case of verification of an eagle, this deterrent system (i.e. flashing light or acoustic signals) would be automatically directed towards the bird, reducing the potential for collateral damage to unintended animals and/or excessive environmental pollution in cases where light or sound are used. Since the continuation of the bird's flight trajectory can be estimated by the stereo system, a very flexible/dynamic deterrent system could be integrated that reacts with the behavior of the bird. With a reactive system, you could design the deterrent system such that the deterrence output is proportional to the needs of deterrence. This may be effective at minimizing habituation to the deterrence system. Such a system could be tested with other large and more common birds with similar response profiles or captive animals.

Testing for overall bird detection probability

Detection, identification, and position estimation were all most effective from 0-150 m away from the camera system. These results suggest that the system could be useful in describing avoidance behavior at the micro-avoidance scale (within immediate proximity of the turbine), but would not be effective at other scales without system modifications. With the system's current capabilities, even micro-avoidance could be difficult to describe for turbines with long turbine blades. Thus we recommend that the system be used to sample the air space around the turbine and not be used, as currently configured, for monitoring turbine interactions entirely.

Improving the system

Hardware Recommendations

The following modifications would lead to high-precision stereo matching and better distance estimates:

- Add a high-precision mechanical adjustment to the secondary camera mount for precise 3-D alignment.
- Use a rectilinear wide-angle lens like the Nikon 28mm f/2.8 AI-s).
- Use a wider stereo basis of 3 m, which would double the distance for which good quality distance estimates would be available.

As mentioned above, we also recommend the addition of a third camera, mounted on a pan-tilt unit that focuses on the object of interest based on measured trajectory by the current stereo camera system. This camera might have a lens with 200 mm focal length and would deliver the object of interest (e.g., potential eagle) at a much greater resolution, leading to better object identification.

SCADA system interfacing

In order for the system to provide deterrence or turbine shutdown capability, the system must either issue its own deterrence mechanism or communicate with the wind farm supervisory control and data acquisition (SCADA) system to send a signal for shutdown or activation of the external deterrence device. The system could act as a remote terminal unit, providing a call to the supervisory computer system for issuance of the command for turbine shutdown or deterrence. Currently, we know of no approved deterrence methods for eagles, so a call for turbine shutdown would be the most likely method of mitigation. We would need to test a method of interfacing with this system to allow proper routing of commands to the supervisory computers. This would require a hardware interface, most likely a networked (LAN-type) architecture. Additionally, software changes will need to be implemented to support the networking protocol of the SCADA system and proper issuance of commands to the supervisory computers. Once in place, these systems would need to be rigorously tested.

The National Renewable Energy Lab (NREL) has an onsite SCADA system and engineering support to allow development and testing of an interface. Working with NREL would allow us to create an efficient system for wind farms without affecting operation. NREL could provide guidance on hardware architecture as well as appropriate calls to the system which could be confirmed with working wind farm operators. We could then test this system at the NREL test facility on working test turbines, first using just the software system and hardcoded calls to the system, followed by full testing with the operational stereo-optic camera system.

Remote wireless network connectivity

We were unable to properly set up and test long-distance wireless connectivity between the camera system and remote network due to a lack of appropriate test venue. We have a line-of-sight microwave transmitter that should allow long distance (>10km) communication with a remote network, but this system will need to be tested at an appropriate test facility. We recommend testing at NREL during testing of any SCADA system. Alternatively, testing could be performed if trials of the system are conducted from an offshore meteorological tower. This would allow appropriate line-of-sight testing and engineering support in an environment that has a realistic level of electromagnetic noise.

Remote data connectivity needs

Connection to the camera system will require, at a minimum, bandwidth to support remote access to the control computer, and maximally, bandwidth to support data transfers from the camera system to onshore file servers for further processing and backup. Ideally, a network interface would be made through the wind farm SCADA system or auxiliary network, but may have to be made through the wireless point-to-point system. We have determined that the mean file size for a 5-minute chunk is 301MB, which requires a minimum bandwidth for data transfer of 8 Mbps. The remote desktop connection should be at least at the high speed broadband level (>2 Mbps) to show contents of the screen sufficiently. Therefore, we recommend a minimum bandwidth of 10 Mbps but optimally 50-100% above that, to allow for constant transfer of data and the remote desktop connection without large lags in data transfer. Connection and transfer of data will need to be evaluated under real-world operating conditions, particularly using the point-to-point system, as environmental conditions could affect transfer rates.

Software Recommendations

Given a reliable stereo calibration, the integration of distance-based filtering into the masking module of the new viewer software would be a straightforward modification (see sections on “Improving the camera calibration” and “Improving the 3-D modeling” about the related 3-D model discussion). Future improvements to the viewer could also be made to integrate the shape recognition software, introducing further data filtering options. For example, the viewer could potentially automatically filter out portions of an image that do not appear to contain a bird. Combining improved 3-D information with the shape recognition software would also present new options, such as querying a scene for birds by size and distance.

LITERATURE CITED

- Buckland, S.T., Anderson, D.R., Burnham, K.P., Laake, J.L., Borchers, D.L., Thomas, L., 2001. Introduction to distance sampling: estimating abundance of biological populations, New York New York USA. doi:10.1007/s00531-004-0408-5
- Buehler, D.A., 2000. Bald Eagle (*Haliaeetus leucocephalus*), in: The Birds of North America, No. 506 (A. Poole and F. Gill, Eds.). The Birds of North America Inc., Philadelphia, PA.
- Chamberlain, D.E., Rehfisch, M.R., Fox, A.D., Desholm, M., Anthony, S.J., 2006. The effect of avoidance rates on bird mortality predictions made by wind turbine collision risk models. *Ibis* (Lond. 1859). 148, 198–202. doi:10.1111/j.1474-919X.2006.00507.x
- Cryan, P.M., Gorresen, P.M., Hein, C.D., Schirmacher, M.R., Diehl, R.H., Huso, M.M., Hayman, D.T.S., Fricker, P.D., Bonaccorso, F.J., Johnson, D.H., Heist, K., Dalton, D.C., 2014. Behavior of bats at wind turbines. *Proc. Natl. Acad. Sci.* 111, 15126–15131. doi:10.1073/pnas.1406672111
- Cullinan, V.I., Matzner, S., Duberstein, C.A., 2015. Classification of birds and bats using flight tracks. *Ecol. Inform.* 27, 55–63. doi:10.1016/j.ecoinf.2015.03.004
- Desholm, M., Fox, a. D., Beasley, P.D.L., Kahlert, J., 2006. Remote techniques for counting and estimating the number of bird-wind turbine collisions at sea: a review. *Ibis* (Lond. 1859). 148, 76–89. doi:10.1111/j.1474-919X.2006.00509.x
- Elliott, W.R., Kaufmann, J.E., Samoray, S.T., Gardner, S.E., 2005. The MDC Method: Counting Bats with Infrared Video. *Natl. Cave Karst Manag. Symp.* 147–153.
- Ferrer, M., de Lucas, M., Janss, G.F.E., Casado, E., Muñoz, A.R., Bechard, M.J., Calabuig, C.P., 2012. Weak relationship between risk assessment studies and recorded mortality in wind farms. *J. Appl. Ecol.* 49, 38–46. doi:10.1111/j.1365-2664.2011.02054.x
- Fox, A.D., Desholm, M., Kahlert, J., Christensen, T.K., Krag Petersen, I., 2006. Information needs to support environmental impact assessment of the effects of European marine offshore wind farms on birds. *Ibis* (Lond. 1859). 148, 129–144. doi:10.1111/j.1474-919X.2006.00510.x
- Groom, G., Stjernholm, M., Nielsen, R.D., Fleetwood, A., Petersen, I.K., 2013. Remote sensing image data and automated analysis to describe marine bird distributions and abundances. *Ecol. Inform.* 14, 2–8. doi:10.1016/j.ecoinf.2012.12.001
- Horn, J.W., Arnett, E.B., Kunz, T.H., 2008. Behavioral responses of bats to operating wind turbines. *J. Wildl. Manage.* 72, 123–132. doi:10.2193/2006-465
- Hull, C.L., Muir, S.C., 2013. Behavior and turbine avoidance rates of eagles at two wind farms in Tasmania, Australia. *Wildl. Soc. Bull.* 37, 49–58. doi:10.1002/wsb.254
- Imler, R., Kalmbach, E.R., 1955. The Bald Eagle and its economic status. U.S. Government Printing Office.
- Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T., 2014. Caffe: Convolutional Architecture for Fast Feature Embedding. *arXiv Prepr. arXiv1408.5093*.
- Kendal, D., 2007. Measuring distances using digital cameras. *Aust. Sr. Math. J.* 21, 24–28.

- Latecki, L.J., Lakamper, R. 1998. Discrete approach to curve evolution. In: Levi P., Schanz M., Ahlers RJ., May F. (eds.) Mustererkennung 1998. Informatik aktuell. Springer, Berlin, Heidelberg.
- Masden, E.A., Cook, A.S.C.P., 2016. Avian collision risk models for wind energy impact assessments. *Environ. Impact Assess. Rev.* 56, 43–49. doi:10.1016/j.eiar.2015.09.001
- Maurer, W.G., 2016. Bird and bat interaction vision-based detection system for wind turbines. Oregon State University.
- Pagel, J.E., Kritz, K.J., Millsap, B. a, Murphy, R.K., Kershner, E.L., Covington, S., 2013. Bald Eagle and Golden Eagle Mortalities at Wind Energy Facilities in the Contiguous United States. *J. Raptor Res.* 47, 311–315. doi:10.3356/JRR-12-00019.1
- R Core Team, 2016. R Development Core Team. R A Lang. *Environ. Stat. Comput.*
- Rodner, E., Simon, M., Fisher, R.B., Denzler, J., 2016. Fine-grained Recognition in the Noisy Wild: Sensitivity Analysis of Convolutional Neural Networks Approaches. arXiv:1610.06756 [cs.CV].
- Sim, C.Y.D., 2002. The propagation of VHF and UHF radio waves over sea paths. University of Leicester.
- Sinclair, K., Degeorge, E., 2016. Wind Energy Industry Eagle Detection and Deterrents: Research Gaps and Solutions Workshop Summary Report. Golden, CO.
- Wiggelinkhuizen, E.J., Barhorst, S.A.M., Rademakers, L.W.M.M., den Boon, H.J., Dirksen, S., 2007. WT-BIRD ® : BIRD COLLISION MONITORING SYSTEM FOR MULTI-MEGAWATT WIND TURBINES. Petten, The Netherlands.
- Williams, K.A., Connelly, E.E., Johnson, S.M., Stenhouse, I.J., 2015. Wildlife Densities and Habitat Use Across Temporal and Spatial Scales on the Mid-Atlantic Outer Continental Shelf: Final Report to the Department of Energy EERE Wind & Water Power Technologies Office, Award Number: DE-EE0005362. Report BRI 2015-11. Biodiversity Research Institute, Portland, Maine.
- Yu, Z., Lincheng, S., Dianle, Z., Daibing, Z., Chengping, Y., 2013. Camera Calibration of Thermal-Infrared Stereo Vision System, in: 2013 Fourth International Conference on Intelligent Systems Design and Engineering Applications. IEEE, pp. 197–201. doi:10.1109/ISDEA.2013.449