Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by **Battelle** *Since 1965*

Triton: Igiugig Fish Video Analysis

Project Report

August 2017

S Matzner
C Trostle
G Staines
R Hull
A Avila
GEL Harker-Klimeš

U.S. DEPARTMENT OF
ENERGY

TRITON

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor Battelle Memorial Institute, nor any of their employees, makes **any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights**. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or Battelle Memorial Institute. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Triton: Igiugig Fish Video Analysis

Project Report

August 2017

S Matzner
C Trostle
G Staines
R Hull
A Avila
GEL Harker-Klimeš

Pacific Northwest National Laboratory
Richland, Washington 99352

Summary

Tidal and instream turbine technologies are currently being investigated for power generation in a variety of locations in the US. An environmental permitting and consenting requirement parallels this exploration generating the need to ensure little or no harm, in the form of strike or collision, befalls marine animals from device deployments. Monitoring methods (e.g., underwater cameras, active acoustics, passive acoustics) around turbine deployments provide empirical data allowing regulators and other stakeholders to assess risk. At present, there is a high level of concern and limited data precluding robust conclusions, which creates a challenge to regulators who must make decisions based on perceived risk versus actual risk. However, the data that are currently available to the scientific community for analysis indicate the issue to be of low risk to date, and strike or collision to be rare events. One such dataset that provides insight to the rarity of strike and collision risk to fish came from an instream turbine deployment in Alaska that used underwater video as the monitoring method.

This document describes the analysis of video data collected around the Ocean Renewable Power Company's RivGen® device deployed in the Kvichak River during July and August 2015 to gain an understanding of the implications of using underwater video cameras as a fish monitoring technique. The data were analyzed manually and used to develop automated algorithms for detecting fish in the video frames and describing their interaction behavior relative to the device. In addition, Pacific Northwest National Laboratory (PNNL) researchers developed a web application, EyeSea, to combine manual and automated processing, so that ultimately the automated algorithms could be used to identify where human analysis was needed (i.e., when fish are present in video frames).

The goal of the project was to develop software algorithms that could identify video frames with fish present to inform and accelerate manual analysis. To achieve this, independent manual analysis was completed for specific video clips (i.e., visual analysis and annotation by a human observer was the standard for assessing the algorithms). The analysis process indicated that some confounding aspects of the algorithm development could potentially be solved with recommended improvements in the initial camera data collection methods.

The manual analysis began to look at all data from the start of deployment of the RivGen® device, primarily using video from Camera 2 that looked directly at the upstream side of the turbine so any interaction could be identified; this was to ensure rare events were seen, and initially focused on Nighttime Data when more fish were present. This process highlighted the amount of time it takes to identify fish, and ultimately only 42.33 hours of video were reviewed because of the time-consuming analysis. The data were classified as "Fish" when the reviewer was confident it was a fish, and "Maybe" fish when it was difficult to distinguish. The two classes were distinguished based on the movement, shape, and color characteristics. Fish Events were further classified by "adult", "juvenile", or "unidentifiable" age. Behavioral attributes were noted and were broadly divided into Passive and Avoidance activities. In over 42 hours of the data reviewed, there were only 20 potential contact interactions, of which 3 were Maybe classifications, 12 were juveniles, and 5 were adults. While only 11.5% of the video data were analyzed from Camera 2, these results are from the time when most fish were present over the turbine deployment period (from Alaska Department of Fish and Game data) and provide preliminary evidence that fish strike or collision of fish in the Kvichak River with an instream turbine is rare.

On only one occasion was an actual contact confirmed, and this was an adult fish that contacted the camera, not the turbine itself. This experience highlights the difficulties associated with confirming a strike or collision event as either having occurred or having been a near-miss. More interactions were

detected at night; this was probably biased by nighttime use of artificial light, which may have attracted fish, but also could have increased detection probability because the light is reflected from the fish itself.

For the algorithm development, background subtraction, optical flow, and Deep Learning techniques were considered. The Deep Learning approach was determined to need too much training data for this application, so its use was not continued. The optical flow analysis was considered promising, but did not give immediate results, so it needs further investigation. Therefore, background subtraction was the main focus in algorithm development. Three methods of background subtraction were tried: Robust Principal Components Analysis (RPCA), Gaussian Mixture Model (GMM), and Video Background Extraction (ViBE). A classification technique was then applied to the foreground images to determine fish presence. Using this combination, fish could be accurately identified when occupying a higher number of pixels (>200 pixels, 98.2% correct; 100–200 pixels, 99.6% correct; 5–100 pixels, 85.4% correct; 2–5 pixels, 66.3% correct).

In parallel, EyeSea was developed to convert the video data to a usable form and to enable manual and automated analysis of the data that would have a standardized output.

Recommendations for further research, and optimizing methods for enhancing data collection and analysis include the following:

- Research

  – Conduct more studies of the effect of lights on fish behavior.
  – Investigate the use of low light video applications as an alternative to using lights.
  – Further investigate optical flow techniques and their applicability for automated analysis.
  – Further refine the parameters for background subtraction in automated analysis.

- Standardized techniques

  – Include markings on the turbines to determine relative range and size of fish within the field of view.
  – Use a standardized (non-proprietary) video format that has a consistent frame rate of at least 25 frames per second.
  – Use a scientific camera designed for underwater measurement in low light environments that has a field of view appropriate for the observations and a pixel resolution high enough to determine fish within the given range.
  – Carefully consider the use of lights and how they illuminate the areas of interest.
  – Standardized and detailed record keeping and metadata collection
  – Use other monitoring technologies (e.g., strain sensors on turbine blades) to determine actual collision or strike events.

Acknowledgments

Acronyms and Abbreviations

DOE	U.S. Department of Energy
FY	fiscal year
GMM	Gaussian Mixture Model
MPC-HC	Media Player Classic-Home Cinema
MHK	marine and hydrokinetic
LGL	LGL Alaska
ORPC	Ocean Renewable Power Company
PNNL	Pacific Northwest National Laboratory
RPCA	Robust Principal Components Analysis
TRL	Technology Readiness Level
ViBE	Video Background Extraction

Glossary

Term	Definition
asynchronous architecture	A system that does not depend on strict arrival times of signals for operation.
avoidance	Used in all instances to encompass behaviors that showed some form of active change; no attempt was made to distinguish between avoidance and evasion.
background subtraction	A computer vision technique used to separate an image (or video frame) into background and foreground, where foreground means objects or regions of interest and is application-dependent.
bilateral filter	A non-linear, edge-preserving, and noise-reducing smoothing filter for images.
"blobs"	Groups of connected pixels of similar intensity.
canonical analysis	A method of regression analysis to determine relationships between a predictor variable and a criterion variable.
collision	When a fish swims into a static object.
compare/ comparison	Qualitative, nonstatistical assessment of the project video data.
contrast stretch	An image enhancement technique that improves the contrast in an image by increasing the range of intensity values.
Deep Learning	Application of learning tasks to artificial neural networks.
directed motion	Motion that demonstrated intended movement; used in this report to describe fish-like movement
EyeSea	A database-driven website for accessing video data files and analysis data.
Event	A place in time during manual video processing marked by a reviewer as having a fish-like object or fish in two or more frames
false positive	Detection of a fish by the algorithm when there was not a fish
Fish	An object that is deemed to be a fish during a manual analysis event
Fish Event	An event deemed to contain a fish during manual analysis
forward-stepping linear discriminant analysis	A method for finding a combination of features that characterizes two or more classes of objects.
histogram equalization	A technique for adjusting image intensities to enhance contrast.
July 22 Data	The full 24-hour manual analysis data of July 22, 2015.
Maybe	An object that during manual analysis is deemed to possibly be a fish, but not a definite identification.
Maybe Event	An event that during manual analysis is deemed to contain an object that could possibly be a fish
MySQL	An open-source relational database management system.
near-field	Relative term referring to an object or fish being relatively close the turbine within the video camera field of view.

neural network	A computer system based on how networks within biological brains work, which "learns", i.e., improves its performance, by considering examples that have been labeled with key parameters.
Nighttime Data	Data from hours 00:00 – 06:00 and 23:00 – 00:00 from July 19, 23:00 to July 23, 03:00
OpenCV	Open Source Computer Vision Library
optical flow	An image processing technique used to compute motion of an object based on changes in the individual pixels in an image.
Rayleigh distribution	A continuous probability distribution characterized by a shape parameter used to model the magnitude of a multi-component vector.
strike	When a fish is hit by a moving part of the turbine
true negative	An object classified as non-fish by automated analysis that was deemed to be a non-fish by a human reviewer, or a frame classified by automated analysis as containing no fish that was not included in the frames containing fish noted by human analysts.

Contents

Figures

Tables

1.0   Introduction

The Triton initiative is a U.S. Department of Energy (DOE)-funded capability at the Pacific Northwest National Laboratory's (PNNL's) Marine Sciences Laboratory in Sequim, Washington. It aims to support DOE-funded projects that are developing technologies for measuring and monitoring the environment around marine energy devices through the mid- to high-level Technology Readiness Levels (TRLs). Ultimately, the initiative is intended to facilitate the permitting process and reduce the overall cost of marine renewable energy.

As part of the initiative, the Igiugig Fish Video Analysis project described herein used video data collected by LGL Alaska around an Ocean Renewable Power Company (ORPC) RivGen® device deployed in Alaska. The data on fish interactions around the operating device were made available to PNNL for further manual/human processing and use in developing automated processing software.

This final report summarizing the project tasks and results follows two previous reports: a data quality report (Trostle 2016) and a project progress report (Avila et al. 2016). The ensuing sections of this report briefly describe the project, the manual analysis of the video data, development of an automated algorithm for detecting fish presence in video frames, development of the software to enable data processing, and finally present conclusions and recommendations for future projects. Appendices A and B, respectively, contain manual annotations and the video data set used to develop the algorithm.

2.0   Description of Project

The ORPC RivGen® device (Figure 2.1) was deployed in the Kvichak River (Figure 2.2) from 19 to 25 July and 19 to 28 August in 2015. The device's two-turbine turbine generator unit (TGU) is supported by a chassis incorporating a pontoon support structure. The structure acts as a foundation when the device is deployed on the riverbed and gives it self-deployment and retrieval capabilities. The system is designed to generate reliable, renewable electricity in rivers near remote communities that have no access to large, centralized power grids.

Figure 2.1. Photograph of the ORPC RivGen® device.

Figure 2.2. Location of the RivGen® device near Igiugig, Alaska.

Five video cameras were attached to the ORPC RivGen® device to monitor fish upstream and downstream of the turbine foils. While the system was deployed, LGL Alaska (LGL) monitored the video for fish-turbine interactions, subsampling 10 minutes per hour (at the top of the hour) (LGL 2015). After the deployment was completed, the raw data, metadata, and a spreadsheet with processed events were released to PNNL for further analysis. Specifically, this was to develop automated algorithms that detected fish within the frames so that manual analysis could focus only on times when fish were present. To do this, manual analysis was required to annotate the video so that it could inform the algorithm development.

3.0 Manual Analysis

3.1 Methods

The development of tidal current and in-stream river current turbines as an industry is relatively new. It is in the early research and development stages that require testing to determine ideal technology and resource choices. The required technologies for monitoring fish interactions around turbine installations have the same early research stage limitations. This project used an underwater optical camera data set that captured numerous instances of fish and a turbine in the same field of view. Cameras and lights were manufactured by IAS systems. Cameras were customized SeeMate™ color to monochrome units with a F2.9 angle lens. Lights were SeeBrite™ omnidirectional model 24L-SS-LED-350. Power came from shore

and data were stored on digital video recorders. Manual processing of the data provided a baseline for software algorithm development as well as qualitative comparisons of fish behavior near the device.

3.1.1   Data Set

The data set comprised underwater video data from five cameras aligned on one side of the RivGen® device (Figure 3.1)—two upstream of the rotor and three downstream—recording 24 hours per day from 19–26 July and 19–28 August. Illumination from two artificial light sources was used between approximately 2300 and 0600 each night. PNNL received the raw video data (6,418 files; 368 hours), along with supplemental reports from LGL in December 2015. LGL had previously processed the first 10 minutes of certain hour blocks of the data, typically coincident with the turbine spinning; observed events were recorded in a spreadsheet that was provided to PNNL with the data set.

For PNNL, the first step was to determine whether the data quality was good enough for the proposed analysis. The research team needed to be able to visually observe fish presence, behavior, species, and any adverse impacts. In February 2016, the data quality was deemed satisfactory, but not suitable for species determination/identification. For additional information regarding the usability and overall quality of the video, see the Quality Check Summary Report (Trostle 2016).

During an Advisory Committee meeting (including participants from ORPC, PNNL, the University of Washington, and the University of Maine) held in March 2016, it was decided that the data set should be manually processed giving priority to nighttime segments (00:00 – 06:00 and 23:00 – 00:00) from July, for which previously subsampled data from LGL showed the highest frequency of fish interactions with the turbine. Additionally, camera 2 (Figure 3.1) was given priority because it showed the upstream view of the rotor. This meant that it:

> In this study, "collision" refers to when a fish swims into a static object, and "strike" refers to a moving part hitting a fish.

- could show potential fish collision or strike interactions with the turbine,
- could show near-field avoidance behavior,
- had a sufficient light source, and
- could be used to coarsely estimate the size and distance of fish relative to the turbine and supporting structures.

As data processing progressed, team members realized that full manual analysis of both July and August nighttime video data would be excessively time-consuming. For every 1 hour of raw video data, it took reviewers approximately 13 to 15 hours to manually review and annotate the video. Due to the amount of time it took for manual review, only part of the July Nighttime Data (July 19, 20, 21, and some of July 23), all of the July 22 Data, none of the August data, and the data required for the test bed development were reviewed (see Section 4.1). Of 18 days of video data recorded by LGL, PNNL was able to review 1 full day, 3 nighttime segments, 4 half-hour blocks on July 23, and 16 five-minute sections for the test bed—a total of 42.33 hours. The PNNL team decided to concentrate on particular subsets of data that would allow for meaningful comparisons. Statistical analysis was not performed on any of the manual review data because of the relatively short period analyzed. Therefore, all uses of the term compare/comparison regarding manual processing for this report refer to qualitative, non-statistical assessment of the data. These comparisons have been grouped into four categories:

- July 22 Data: A full day, July 22 (hereafter referred to as July 22 Data), was analyzed because it was important to include a full day inclusive of daytime data for preliminary comparisons of the first 10 minutes of each hour to the full 60 minutes of each hour. This also allowed for diel differences to be qualitatively compared.

4

Figure 3.1. Schematic showing approximate locations and views of cameras by number on the RivGen® device.

Lights are represented by blue rectangles. (Not to scale. The pontoon structure is 19.8 m long, 11.5 m wide and 1.7 m high. The turbine [TGU] is 10.4 m long, 1.5 m wide, and 1.5 m high.) (Figure courtesy of LGL and ORPC.)

- Nighttime Data: Nighttime only (00:00 – 06:00 and 23:00 – 00:00) data from July 19, 23:00 to July 23, 03:00 (hereafter referred to as Nighttime Data) were processed as mentioned in the bulleted paragraph. The data from July 23 Nighttime Data consist of only the first half-hour of each hour block (e.g., 01:00 to 01:30). These data represent the majority of the processing effort for this study. The first night (July 19) of data collection had technical issues with the lighting system and data were collected without a light source.
- Light Effects Data: The hour block from 23:00 – 00:00 for July 19–22 were processed with varying artificial light operations to gain a preliminary understanding of any effect lights may have on fish detection probability and fish behavior.
- Collision/Strike Data: Events with fish interactions where possible collision or strike occurred were separated into a small subset and further analyzed. Reviewers observed a total of 20 events that had possible collision or strike interactions. These events were separated for further comparisons.

3.1.2   Manual Processing

The data were provided in a proprietary format that was difficult to manipulate for the proposed processing and analysis procedures. Data files were changed to .mp4 format for ease of use with minimal

change in data quality. Two reviewers worked together to establish processing protocols and definitions for parameters annotated for each event. A subsample of data was processed by each reviewer and compared for similarity to ensure data processing would be consistent and accurate throughout the analysis.

The reviewers visually processed the data in half-hour blocks using Media Player Classic-Home Cinema (MPC-HC). Reviewer 1 processed the first half-hour and Reviewer 2 the second half-hour. Whenever a reviewer visually assessed a fish or an object that had characteristics different from the surrounding water column debris (i.e., shiny and/or non-passive movement) that was present in the field of view for more than one frame, it was deemed an event. For each event, the numerical annotation method explained in the *Igiugig Video Analysis – FY16 Progress Report* (Avila et al. 2016) was used to describe the event characterizations. Parameters for these characterizations are described in Appendix A. Manual review did not distinguish between the terms "avoid" and "evade" throughout this report. Because the reviewers were unaware of the exact distance of the objects from the turbine, and because they did not use the behavioral responses of the objects and fish to decide between the two terms, "avoid" was used in all instances to describe behaviors that showed some form of active change assumed to be related to the turbine. Important classification annotations referenced throughout this report are whether or not an event was a Fish Event or a Maybe Event. A Fish Event meant that the reviewer was positive the object was a fish, whereas a Maybe Event meant the reviewer was not sure (hereafter referred to as Fish or Maybe Events, respectively). The designations between these two annotation descriptions are important for comparison and analysis purposes, as well as for informing the algorithm development.

Objects that were not definitively defined as fish were still deemed events and recorded as Maybe Events. It was important to include these for two reasons. First, video quality could possibly affect the reviewer's determination of whether or not an event contained a fish, and erring on the side of capturing all events with some false positives was preferable to missing some Fish Events. Second, software and automated algorithm development was a major objective of this research. Objects often had characteristics similar to fish and could be identified as events by the automated algorithms. This again allows more confidence in capturing all Fish Events at the cost of some additional false positives.

To keep the reviewers calibrated during review, both started with a training period to go over the parameters and define characterizations. They separately reviewed the same video and compared results for 2–3 weeks, addressing any discrepancies in annotations. As the reviewers began processing the data individually, they kept in regular contact, went over interesting or questionable interactions, reviewed each other's annotations, and discussed methods to ensure calibration at bi-weekly meetings.

Even with these checks, during analysis, an inconsistency was discovered between reviewers. While both reviewers saw a similar number of events, meaning they were stopping for the same objects, the distinction between calling an object a Fish Event or Maybe Event did not always correspond. This meant that some of the objects one reviewer deemed as a Fish, the other reviewer deemed as a Maybe fish.

Comparisons were made using the July 22 Data to show any similarities that exist in event occurrence and fish count estimates between processing of the first 10 minutes per hour, and processing of the full hour. Additionally, figures for visualization were also made to display

- differences between definite Fish Events and objects with non-passive behavior,
- fish count differences between day and night, and
- fish count differences between when the device was spinning and static.

For the Nighttime Data, the behavior types that are associated with different categories of events were compared. This categorization of events was based on the Appendix A annotation, "Fish?". This annotation was simply a question to the reviewer about whether the object observed during a designated event was definitely a Fish or a Maybe. Initially, All Events that included Fish and Maybe Events were considered. Categories separated All Events into Fish and Maybe Events. Fish Events were further

categorized by annotations from Appendix A designating the fish as juvenile (likely a salmon smolt), adult (likely a salmon), or unidentifiable as determined by the reviewer. The category separation flow chart is shown in Figure 3.2. After behavior types were attributed to the categories of events by percentage (Figures 3.4–3.10), behavior types were coarsely grouped. Each of the behavior types the reviewers used for the annotation description was designated as either Avoidance, Passive, or Other (Table 2).
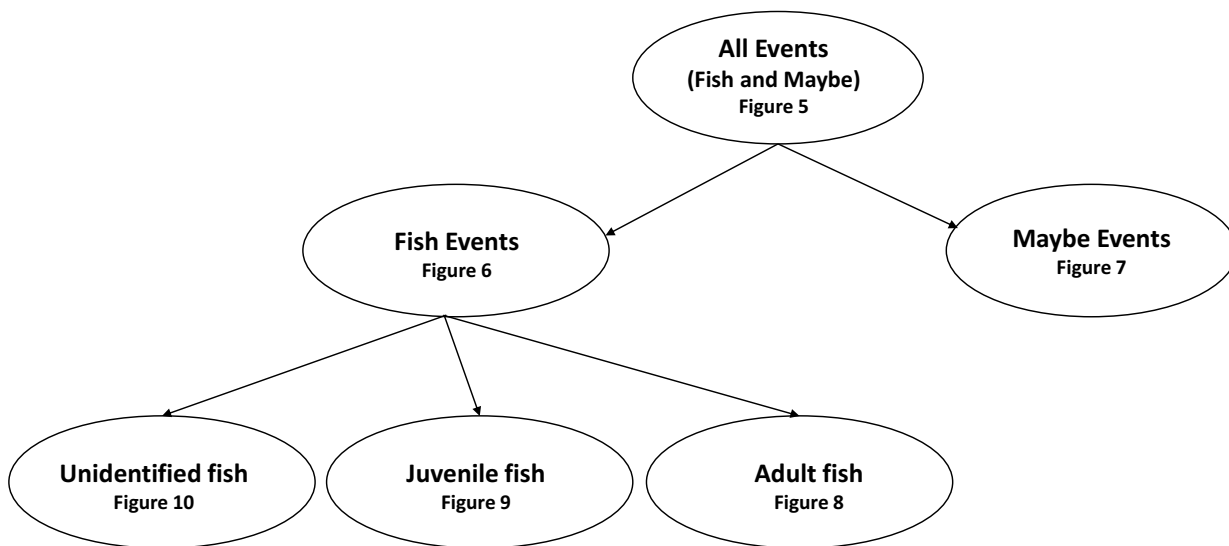
Figure 3.2. Flow chart showing the different categories of events used to visualize behavior types (Figures 3.4–3.9) attributed by data processing reviewers. All events that had potential fish collision or strike were placed in a separate subset (3.10).

Table 3.1. Grouping of annotated behavior types into Avoidance, Passive, and Other.

Avoidance	Passive	Other
Milling	Straight across (above or below)	Unable to tell
Pause	Through turbine	Other
Against current	Toward static parts	
Avoid reverse	Face first	
Avoid below		
Avoid above		
Avoid around		

A simple comparison of the Lights Effects Data was to determine whether more events were observed when the lights were on or off during varied light operations from July 19–22. This was possible because on July 19 the lights were off due to technical difficulties. The lights were turned on the next day and they were used for the remainder of Nighttime Data collection.

The last subset of data comparisons were the events when fish collision or strike may have occurred in the Nighttime Data. This data set includes only events that were positively determined by both data reviewers to be Fish and excludes the Maybe Events; hence, there is no disparity between reviewer determinations.

3.2　Results

Currently, no established video data analysis techniques exist for assessing fish interactions. Using the above methods, qualitative comparisons were made with the data set to highlight differences in 1) subsampling of the first 10 minutes of each hour and the entire hour; 2) nighttime behavior types; 3) possible collision and strike events; and 4) the effects of nighttime illumination. The data were further summarized between whether an object was a Fish or Maybe Event, and the categorical groupings associated with the observed behaviors.

3.2.1　Fish Presence/Behavior

3.2.1.1　July 22 Data Subsampling Comparisons

For the July 22 subset of data there were 2,538 events: 260 were Fish Events, 2,256 were Maybe Events, and 22 Events were a combination of Fish and Maybe occurrences. The majority (81%) of events occurred during nighttime hours. Only one Fish Event occurred during daytime—in hour 19 in the processed data. Fish abundance or frequency of events does not appear to be related to whether the turbine was spinning (hours 1–2) or static (hours 3–6), but does seem to coincide with low light levels (hours 1–6 and hour 24). To compare the 10-minute sampling regime with full analysis, the 10-minute counts were multiplied by six to produce an hour estimate. This assumes that the first 10 minutes is representative of the subsequent five 10-minute blocks. The comparison between these processing methods shows that when the first 10 minutes is subsampled and multiplied by 6 to approximate an hourly estimate the numbers are inflated for hours 3–6. Hour 1 is underestimated, for hour 2 the estimates are similar, and hours 3–6 are over-estimated for both the number of Fish and number of events (Figure 3.3).
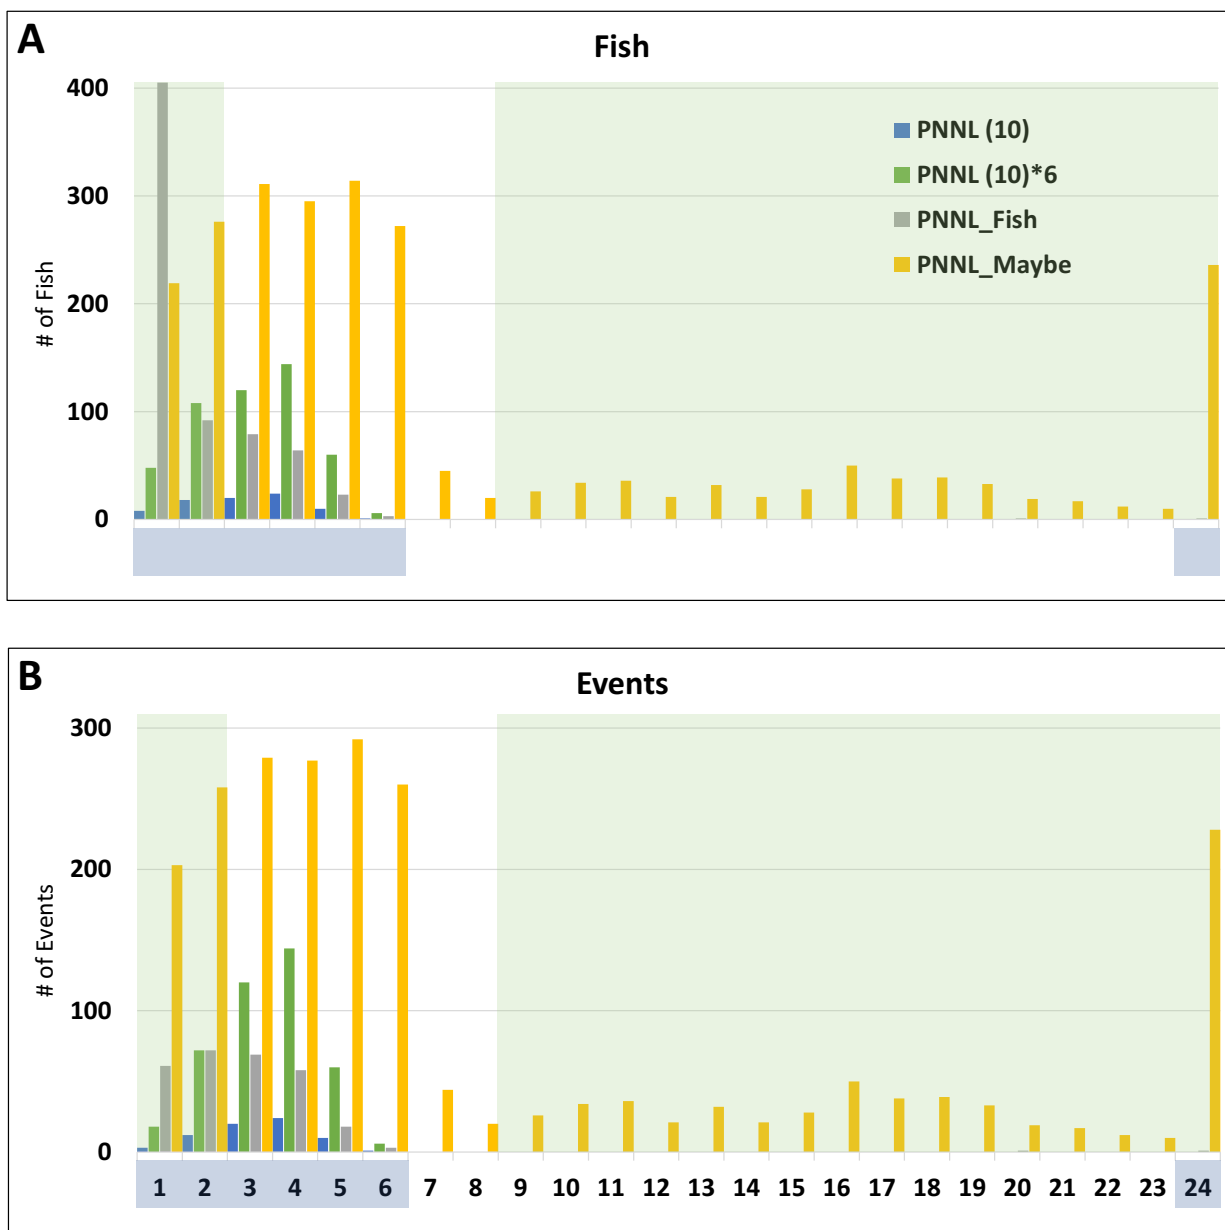
Figure 3.3. Bar graphs showing processed data for July 22, 2015.
The horizontal axis represents each hour block for both graphs and the dark shaded numbers represent hour blocks that are after sunset and before sunrise (Nighttime Data). The green shaded background of the plot areas show when the RivGen® was spinning. Graph A displays Fish and Maybe counts as observed during manual processing by hour blocks. Graph B displays the number of events. The blue bar, "PNNL (10)", represents the estimates from PNNL's processing of the first 10 minutes of each hour. The orange bar, "PNNL (10)*6", represents PNNL's estimate multiplied by 6 to be an approximation for the full hour. The gray bar, "PNNL_Fish", represents PNNL's count of "Fish" for Graph A or "events" for Graph B for the full hour. The yellow bar, "PNNL_Maybe", represents the count of objects (Graph A) or events (Graph B) with non-passive behavior but not determined to be fish.

3.2.1.2   Nighttime Data Behavior Types

Other than processing the July 22 Data, PNNL only processed Nighttime Data. For the category that includes both Fish and Maybe Events, there were 629 Fish Events, 4,149 Maybe Events, and 51

combination events that included both Fish and Maybe occurrences. Each event was broken down into the described behavior specific to the annotation list found in Appendix A. Grouping all of these described behaviors by annotation behavior type (e.g., avoid around, avoid above—see Appendix A) provides some evidence of what the dominant behaviors are within the camera field of view in front of the RivGen® during nighttime hours. The dominant behavior for Fish and Maybe Events was "through turbine", followed by "straight across", followed by "toward static parts" (Figure 3.4). Note that this comparison is not separated by Fish or Maybe Events or any other qualifier, and the Passive behavior group dominates with 80% of the behavior, compared to ~19% for the Avoidance group.
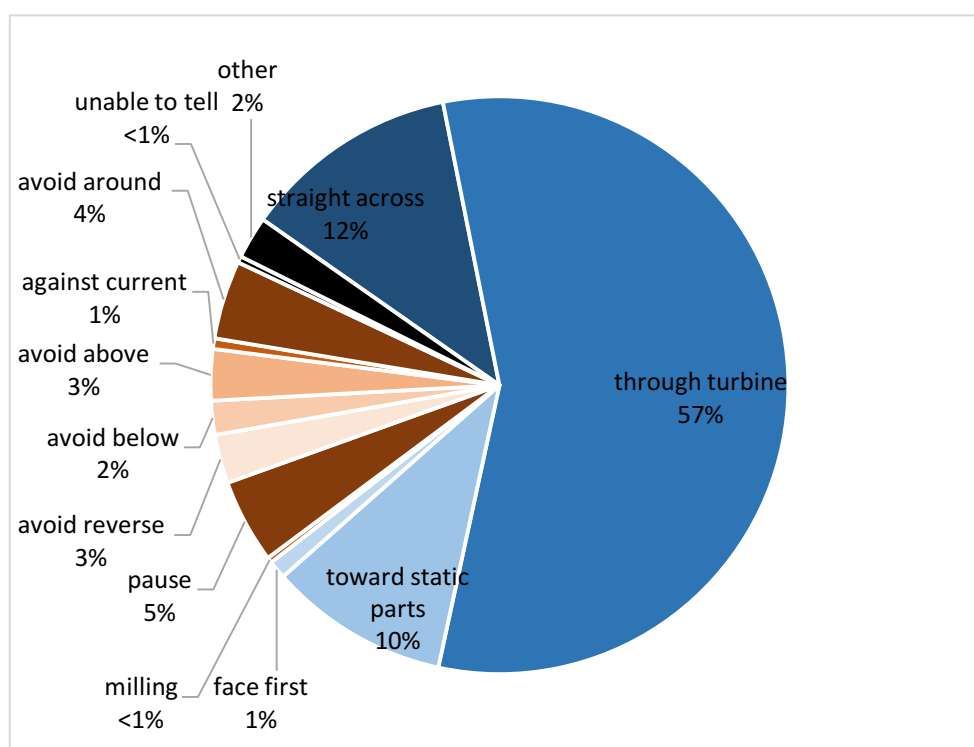
Figure 3.4. Behavior types recorded by reviewers for Nighttime Data. Data included are both Fish and Maybe Events for all sizes; n = 4,829.
The blue sections of the graph designate the Passive group of behaviors and the brown sections represent the Avoidance group of behaviors.

The Nighttime Data were separated into Fish and Maybe Events to visualize how reviewer-described behaviors may be different for each. Combination events (n = 51) that had both Fish and Maybe Events were removed because it was impossible to separate behavior annotations associated with a Fish object or a Maybe object during the event. For Fish Events the top three dominant behaviors were "through turbine", "avoid around", and "pause" (Figure 3.5), and for Maybe Events the top three behaviors were "through turbine", "straight across", and "into static parts" (Figure 3.6). There is a distinct qualitative difference in behavior types (Avoidance vs. Passive) between the Fish Events and Maybe Events. Figure 3.5 shows Fish Events dominated by the Avoidance group of behaviors (62%) and Figure 3.6 shows Maybe Events dominated by the Passive group of behaviors (80%). This abundance of passive behaviors is expected, because one of the qualifiers when distinguishing between Fish and Maybe Events was how the objects moved. It is important to note that there was some disparity between what the two reviewers designated as a Fish or Maybe Event as described in "Manual Processing."

Figure 3.5. Behavior types recorded by reviewers for Fish Events in Nighttime Data for all sizes (n = 618).
Maybe and Combination events were removed. The blue sections of the graph designate the Passive group of behaviors and the brown sections represent the Avoidance group of behaviors.

Figure 3.6. Behavior types recorded by reviewers for Maybe Events in Nighttime Data for all sizes (n = 4,149).
Fish and combination events were removed. The blue sections of the graph designate the Passive group of behaviors and the brown sections represent the Avoidance group of behaviors.

11

Within the Fish Events category, a separation was made to categorize juvenile, adult, and unidentifiable fish. Eleven Fish Events had a combination of adult, juvenile, or unidentifiable Fish Events that were removed from these comparisons. There were 174 adult Fish Events for which the dominant behaviors were "pause", "avoid around", and "avoid reverse" (Figure 3.7). There were 259 juvenile Fish Events, for which the dominant behaviors were "through turbine", "avoid around", and "pause" (Figure 3.8). Determining whether it was an adult or juvenile was sometimes impossible. This created the category of an unidentified Fish Event of which there were 185. The dominant behavior was "through turbine", "avoid around", and "pause" (Figure 3.9). Adult fish displayed Avoidance behavior 82% of the time compared to only 14% passive behavior. The behavior groups for juveniles were split, showing 50% Avoidance behaviors and 44% Passive behaviors. And the fish that were unidentifiable demonstrated dominant Avoidance behavior 57% of the time and Passive behavior 36% of the time.

Figure 3.7. Behavior types recorded by reviewers for adult Fish Events in Nighttime Data (n = 174). *Juvenile and unidentified Fish Events, Maybe Events, and combination events were removed. The blue sections of the graph designate the Passive group of behaviors and the brown sections represent the Avoidance group of behaviors.*

Figure 3.8. Behavior types recorded by reviewers for juvenile Fish Events in Nighttime Data. (n = 259).
Adult and unidentified Fish Events, Maybe Events, and combination events were removed. The blue sections of the graph designate the Passive group of behaviors and the brown sections represent the Avoidance group of behaviors.
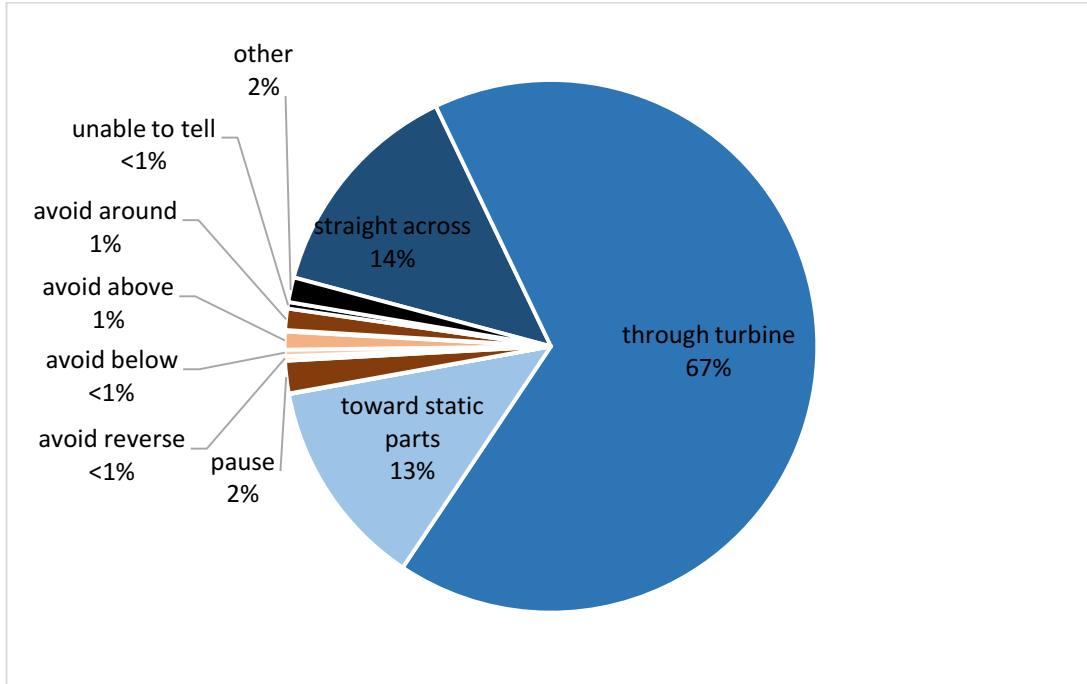

Figure 3.9. Behavior types recorded by reviewers for unidentified Fish Events in Nighttime Data (n = 185).

3.2.1.3   Fish Collision and Strike

Reviewers found a total of 20 events involving possible collision or strike (12 strike and 8 collision). All strike events in this data set refer to moving parts of the turbine hitting an object or fish, while collision refers to an object or fish coming into contact with a static part of the device (this could include the blade when it is not turning). Of these 20 potential events, 17 were Fish Events and 3 were Maybe Events. Of the 17 Fish Events, juveniles made up 12 of the events and adults made up 5 events. All but one of the juvenile Fish Events had multiple fish in the field of view, up to ~50. All of the adult events were single fish. All juvenile Fish Events occurred between 00:00 and 01:00, except two which occurred at 01:03 and 03:02. The turbine was spinning for all but one of the juvenile Fish Events and none of the adult Fish Events. Juveniles made up 11 of the strike events, while the remaining strike occurrence was a Maybe Event. No adults were involved in any of the strike events. Of the 8 collision events, adults made up 5, a single juvenile made up 1, and the rest were Maybe Events. The one juvenile collision and 4 of the 5 adult collision events occurred with a static blade. The last adult collision occurred with the camera and was the only confirmed collision. Behavior for these events was dominated by Avoidance group behaviors "through turbine", "avoid around", and "pause" (Figure 3.10).

Figure 3.10. Behavior types recorded by reviewers for potential collision or strike Fish Events in Nighttime Data (n = 17).
The blue sections of the graph designate the Passive group of behaviors and the brown sections represent the Avoidance group of behaviors.

3.2.1.4   Light Effects

On July 19, the lights remained off through the night, while on every other night the lights turned on as it became dark. A light operations record was not kept during deployment, so manual reviewers at PNNL watched the video and estimated the operational status of the lights (Figure 3.11). Events observed over four nights during hour block 23 (23:00–00:00) when light operations varied were compared to show fish presence while the lights were on and off (Table 3.2). Over the four-night comparison of hour block 23 (a total of 4 hours), the lights were off 45% and on 55% of the total time.

Only 5 events were recorded by the reviewers on July 19 when the lights remained off the entire hour. All 5 events on July 19 occurred in the first 9 minutes of the hour block, and they were all Maybe Events. On July 20, 1 Maybe Event occurred while the lights were off during the first 34 minutes of hour block 23, and 144 events (1 Fish Event, 143 Maybe Events) occurred while the lights were on in the last 26 minutes of hour block 23. On July 21, 2 Maybe Events occurred while the lights were off during the first 14 minutes of hour block 23, and 65 events (2 Fish Events, 63 Maybe Events) occurred while the lights were on during the last 46 minutes of hour block 23. On July 22, 135 Maybe Events were recorded by reviewers when the lights remained on during hour block 23. Over the four-night comparison, approximately 2% of the total events occurred while the lights were off, and 98% occurred while the lights were on.

Figure 3.11. Visual approximation of light operations determined by watching the video data and noting when the turbine and objects began to look brighter (lights on), or when the turbine and objects began to look darker (lights off).

Although a light operations record was not kept, PNNL staff watched the video data to estimate when the lights came on and turned off. The orange bars represent the duration of time the lights were estimated to be on, while the blue bars represent the duration of time manually reviewed by the two reviewers. The x-axis on Figure 3.11 shows a total range of 24 hours, from hour 22:00 on one day and up to hour 22:00 of the next day. This was done to show the light operations and review effort in a more continuous manner. The major y-axis labels list the dates reviewed with minor label divisions of when the lights were on and the hours reviewed on those dates. From the night of July 19 to the morning of July 20, the lights remained off, and the reviewers manually processed video data from 23:00–06:00, although only part of hour blocks 23 and 5 were visible. From the night of July 20 to the morning of July 21, the lights were estimated to be on from 23:35–06:17 and the reviewers manually processed video data from 23:00–06:00. From the night of July 21 to the morning of July 22, the lights were estimated to be on from 23:15–06:26 and the reviewers manually processed video data from 23:00–22:00. From the night of July 22 to the morning of July 23, the reviewers manually processed video data from 23:00–00:30, 01:00–01:30, 02:00–02:30, and 03:00–03:30. The lights were estimated to turn on at 23:00 and remained on during all of the manual review effort on July 23, but an estimation of when the lights turned off that day was not done.

Table 3.2. Comparison of the number of events recorded by reviewers during hour block 23 (23:00–00:00) over four nights when the lights were on and off, including the duration of light operation status and totals.

Hour Block 23	Lights Off		Lights On		Total
Date	Duration	Number of Events	Duration	Number of Events	Number of Events
7/19/2015	60 minutes	5	0 minutes	No data	5
7/20/2015	34 minutes	1	26 minutes	144	145
7/21/2015	14 minutes	2	46 minutes	65	67
7/22/2015	0 minutes	No data	60 minutes	135	135
Total	108 minutes	8	132 minutes	344	352

3.3   Discussion

Perceived risk and shortage of empirical data about fish interactions with tidal and in-stream turbines like ORPC's RivGen® means that monitoring is required during turbine deployment. For near-field interactions, optical cameras are the ideal choice because acoustic devices are limited at such close ranges because of transmitted sound scattering from the turbine blades. For this research, the use of cameras provided a useful data set that allowed the capture of hundreds of fish interactions with an operational commercial-scale device. These interactions included 17 possible collisions with static components and possible strike with dynamic components of the device. These 17 events accounted for 2.75% of all Nighttime Fish Events and 0.07 % of total hours processed. Only through intense manual processing effort was it possible to find the extremely rare events of collisions and possible strikes that were observed. These processed data also allowed comparison of a complete manually processed data set to the same subsampled processed data set. Of the 17 possible collision or strike events, only 1 was in the first 10 minutes of the hour. This means that 16 of the events would have been missed, pointing to the importance of full data set processing to ensure these rare events are observed. While strike and collision are of major concern, the behaviors used by fish as they approach these devices are important for continued research and to determine the need for monitoring around turbines. The types of behavior provide input parameters to models as well as identifying differences that may exist between different species or age classes of fish.

The previous analysis by LGL primarily processed data to coincide with times that the RivGen® device was spinning, which was typically during daylight hours. The PNNL research team concentrated on nighttime because 66% of all fish observed by LGL were observed during nighttime even though this time composed less than 10% of their total processed data.

3.3.1.1   July 22 Data Subsampling Comparisons

PNNL processed the first 10 minutes of each hour, illustrating the difference between subsampling and full analysis. Processing sub sets of data is common for researchers faced with the daunting task of large data sets, and it is considered a valid way to process large amounts of data. In this instance, when subsampled data [(10)*6] were compared to the fully processed (all 60 minutes of each hour) 24 hours of data on July 22, the number of Fish Events per hour was the same for hour 2, less for hour 1, and more for hours 3, 4, 5, and 6 (Figure 3.3B). While there was some discrepancy between the two, a larger sample size for comparison would be required for validation. The number of fish (counts) had similar results with the (10)*6 estimates being larger than the actual 60-minute counts for hours 2–6 (Figure 3.3A). Counts be skewed if large schools of fish are present. The first hour block in the July 22 fish count data (Figure 3.3A) is a good example of this; the full hour count data are more than four times the (10)*6 estimate. This was simply a case of several schooling Fish Event occurrences that were made up of tens of juvenile fish observed from minute 10 to 60. Subsampling may provide a valid estimate of Fish Events but fish counts may be biased low if events with large schools of fish are missed.

Having the entirety of the July 22 video processed provided evidence that the majority of Fish Events occur during nighttime. For this single day, it also indicated that the number of Fish Events did not dramatically increase or decrease based on whether the turbine was static or spinning. For the total data manually reviewed by PNNL, the turbine was spinning for 44% of the time, and not spinning for 56% of the time. After the RivGen® spinning ceased (typically around 01:00), the number of Fish Events decreased from then until 06:00. The occurrence of Fish Events is more likely to be related to light levels because Fish Events decrease temporally as sunrise approaches. If the driver for frequency of Fish Events is light levels, then use of artificial lighting to increase detection probability at night introduces a possible complication.

3.3.1.2   Nighttime Data Behavior Types

Describing the behavior of Fish and Maybe Events captured from a single camera is subjective for most of the descriptions (see Appendix A). While an observed movement upstream or downstream is definitive in nature, movement toward or away from the camera or attempting to use depth of field to describe an event is difficult and accuracy is impossible. Nevertheless, behavior was described for all Fish and Maybe Events during PNNL processing of the data. An extensive list of behavior types that described in detail the majority of observed fish behavior was used. Additionally, specific behaviors were qualified as being Avoidance or Passive behaviors (Table 3.1). For the manually processed data set, the extent to which behavior is addressed for each processed event is important to understand fish behavior in general as well as differences between behaviors of fish based on their size or age class.

The binary grouping of all specific behaviors into Avoidance and Passive behavior groups provided evidence of two important findings:

- First, the amounts of Avoidance and Passive behavior differ between Maybe and Fish Events. During the PNNL processing, both reviewers agreed that movement or behavior of the object during an event had a strong bearing on whether or not it was deemed to be a Maybe or Fish Event. More movement, especially those representing non-passive examples, typically led to classification as a Fish Event. The Avoidance group of behaviors is therefore important for separating Fish Events from Maybe Events. However, not all fish entering the field of view will necessarily change their behavior before exiting. Fish already in line to avoid the turbine may not change their trajectory and thus fall under one of the Passive group's behaviors.
- Second, the amount of Avoidance/Passive behavior differs between adult and juvenile Fish Events. Fish Events that consisted of adult fish had only 17% Passive behavior and of this amount only 4% were specifically "through turbine" (Figure 3.7). In contrast, juvenile Fish Events had a 50/50 split between Avoidance and Passive behavior and 34% of the Passive behavior was "through turbine". This comparison shows evidence that adult fish are better at avoiding the turbine than juvenile fish. Although juvenile fish behavior may consist of Avoidance behaviors, the juveniles tended to be less successful in actually avoiding the device and often went "through the turbine" even after attempting an Avoidance behavior.

Behavior types or groups may play an important role in algorithm development in the future. A variety of qualifiers are used in algorithm development, and behavior or movement is an important one for animal detection algorithms used with remote sensing devices and specifically optical cameras. Often, threshold metrics are used for initial investigations into automating an animal being in a frame. However, if this is successful and a variety of animal types have potential for being detected, then the next step is grouping them by some qualifying characteristic. Often size is the first characteristic for grouping followed by movement or behavior. Knowing the movements and behavior associated with the fish detected in these data has the potential to further general knowledge or inputs for modeling. Improved automated analysis to decrease the effort required to process and analyze these types of data and ultimately create cost-

effective methods. Use of these methods by developers and researchers can provide meaningful data accepted by regulatory bodies that require monitoring.

3.3.1.3   Fish Collision and Strike

As fish collision, strike, and near-miss events are generally accepted to be rare at marine and hydrokinetic (MHK) installations, it is important to process most, if not all, of the data collected to ensure these events are not missed. If an entire data set is not to be processed, then large-scale time blocks likely to coincide with the highest probability event occurrences, decided upon with expert opinion or existing empirical data and statistical analysis, should be processed. The sequence of the processing steps used for camera data set described herein is a good example of efficient gathering of useful information. The initial subset processing performed by LGL for the first 10 minutes of certain hour blocks made it clear that most Fish Events occurred during nighttime. This is a highly productive first step for a large data set for which no established processing methods exist, except for manually reviewing the data. As a logical first step, it saved time and provided the foundation for taking the next step to gather meaningful results. PNNL followed up and concentrated processing effort on nighttime hour blocks based on the LGL information that indicated more events occurred at night. This concentration on the Nighttime Data provided more meaningful comparisons of a variety of fish behaviors showing differences in adult and juvenile behaviors. The processed data also captured 17 events, out of a total of 618 Fish Events, with possible collision and strike between fish and a commercial-scale device, indicating how rare these events are and the difficulty associated with observing them. Even with capturing the events with possible collision or strike, actual contact is difficult to verify because uncertainty remains based mostly on the data quality specific to camera selection, lighting, placement, and field of view. Collision was only confirmed in one instance when an adult fish collided with the camera. Additionally, it is important to note that the outcome of a collision, strike, or near-miss event was not possible to determine because of data quality and the short duration that a fish was in the actual field of view.

Camera selection for underwater fish observations is not a trivial matter. The field of view, resolution, low light capacity, and frame rate are just a few of the parameters that are crucial to gathering high-quality, meaningful data. After data have been collected, the file type becomes important for effective processing that leads to useful analysis. The cameras used to collect the data presented in this report were customized SeeMate ™ color to monochrome units with a F2.9 wide angle lens, manufactured by IAS systems (North Vancouver, British Columbia). The images had a resolution of 352 × 240 pixels. Each camera had a variable frame rate (less than 10 per second), and a field of view of "approximately one-third of the area between the pontoons and the left (portside) one-third of the TGU"[1] (LGL 2015). Pixel resolution, field of view, and light capturing ability created limitations in the data set, and complications continued because the output files were of a proprietary format. Significant amounts of unplanned time and resources were required for file conversion to a non-proprietary format, followed by testing several video-file viewers to determine the one best suited for this analysis, which included requirements like moving forward and backward through each frame capture without skipping or freezing up. Based on this work, literature review, and discussions with other researchers, a brief set of guidelines for camera selection for future applications is given in the Recommendations section.

3.3.1.4   Light Effects

On every night except July 19 the lights turned on as natural light levels decreased to illuminate continued monitoring fish presence and interaction. A comparison was done to better understand the potential impact of artificial lights during this environmental monitoring effort during hour block 23 from July 19

---

[1] LGL Alaska Research Associates, Inc., 2015 *Fish and Wildlife Monitoring Plan for RivGen® Testing on the Kvichak River, Alaska in 2015*

to 22. As it became dark on July 19, the field of view began to fade into a grainy, grayscale image with portions of it becoming black over time. If fish were present during the last 15 minutes of hour 23:00 on July 19, it would have been very difficult for the reviewer to see or document their presence. When comparing the first half of hour 23:00 on July 19 to the same hour on July 20, the images of both nights seem similar, but when the lights appear to turn on at approximately 23:35 on July 20, the turbine is illuminated, potentially creating an opportunity for light to reflect off fish and be visible to the camera, as well as make the image sharper and clearer. In contrast, on July 19 the image degrades over time. Nighttime illumination probably affects the detection probability of fish by the reviewers, and may alter an avoidance/attraction response by the fish.

The number of events that occurred when the lights were on was considerably higher than when the lights were off (344 compared to 8, respectively), and with a similar operation duration (55% compared to 45%, respectively). The number of events when comparing lights on and off differs considerably, yet the reason for this in this application is not well understood. Artificial lights may have attracted fish, thereby causing more events, or more fish may have been present during the last half of hour 23 when the lights were generally on every night except July 19; the data from July 22 when the lights were on for the full hour show more fish during the second half hours (135 vs 94), but this does not account for the extreme difference observed overall. Alternatively, fish presence may be similar on all nights, but the artificial light provides the source needed to make them visible to the optical cameras and in turn, reviewers. Additionally, on July 22 when all 24 hours were reviewed, Nighttime events (when the lights were on from 00:00–06:26, and 23:00–00:00) made up 99.9% of the total events for that day, with only 1 daytime Fish Event overall (although 436 Maybe Events during the daytime). Due to this clear difference and the lack of baseline understanding of fish attraction or deterrence related to this variable, the role of artificial lights during environmental monitoring needs to be further investigated.

4.0   Automated Analysis

Automated analysis was investigated to develop algorithms for detecting fish presence in the video, so that an entire video data set could feasibly be analyzed automatically without the need for manual sampling. Reducing the volume of data to just those video segments during which fish were present would optimize human labor time, and the reduced-volume approach could also be used to perform a quick preliminary analysis of the data. Ultimately, the system could be fully automated, and the software optimized to run in real time as part of an underwater observation system for long-term monitoring of the effects of MHK devices on animal populations.

The vision for an automated video processing system consists of three main components: preprocessing, detection, and classification (Figure 4.1). The preprocessing component filters the raw video frame to improve its quality in terms of contrast, color balance, and smoothness. The detection component identifies objects that might be fish, and the classification component classifies the detections to filter out false positives such as kelp, shadows, or other objects that are not of interest. These three components interact, and each requires many design decisions in order to realize an effective system. Under this project, candidate algorithms were investigated for each of the components, and an infrastructure was developed to tie the components together into a web-based application developed by PNNL called EyeSea.
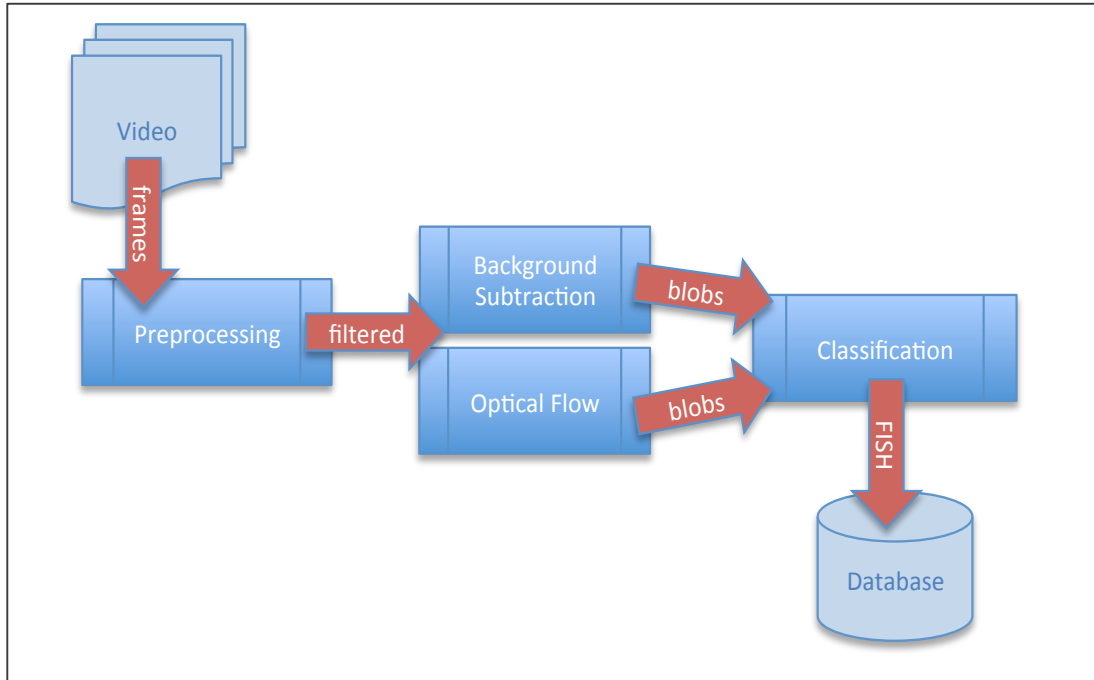
Figure 4.1. The automated processing chain.

4.1 Methods

A testbed was developed to evaluate the performance of different algorithms; it consisted of a development data set and a processing pipeline. This meant that algorithms could be evaluated in a consistent, reproducible manner. A development data set was assembled from a subset of the full Igiugig data set consisting of 16 five-minute video segments containing Fish Events (Appendix B). The video segments were selected to represent different lighting conditions, different camera views and different sizes of fish, individuals and schools (Figure 4.2). Each video consisted of a total of 7500 frames, and even though the segments were chosen to include fish, only 6% of the total frames did in fact contain fish (the presence of fish is not a common event). The data were annotated as described in the Manual Analysis section. The processing pipeline was adapted from the Fish4Knowledge (Boom et al. 2014) code[2] for fish detection with custom code. The pipeline was used to batch process all the development videos using a particular detection algorithm, and to calculate the resulting detection rate and false positive rate by comparing the detections to the manual analysis annotations (Figure 4.3).

---

[2] http://groups.inf.ed.ac.uk/f4k/

a) Camera 1, daylight

b) Camera 1, night

c) Camera 3, night illuminated

d) Camera 4, night illuminated

Figure 4.2. Example images of fish in the different cameras with different illumination.

For the detection algorithms, background subtraction and optical flow were investigated. Three different background subtraction techniques were evaluated: Robust Principal Components Analysis (RPCA) (Candès et al. 2011), Gaussian Mixture Model (GMM) (Lee 2005) and Video Background Extraction (ViBE) (Barnich and Van Droogenbroeck 2009). The optical flow analysis consisted of a dense optical flow calculation using the Farnebäck algorithm (Farnebäck 2003) and a sparse feature-based flow calculation using the Lucas-Kanade method (Lucas and Kanade 1981), both as implemented in OpenCV.[3] For classification, models were developed using forward-stepping linear discriminant analysis (Lotlikar and Kothari 2000) on the detected objects to distinguish between fish and non-fish objects. The features used for classification were object size, intensity, shape, and motion.
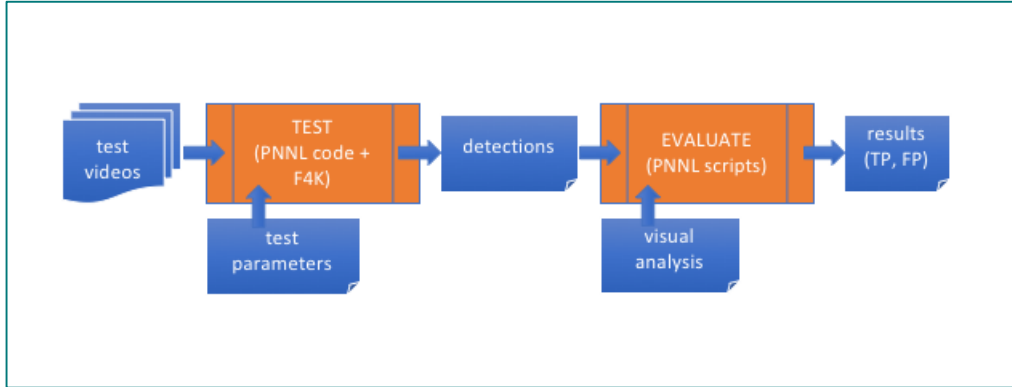
---

[3] http://opencv.org

Figure 4.3. The developed testbed pipeline for evaluating different detection algorithms.

Background subtraction is a computer vision technique that is used to separate an image (or video frame) into background and foreground, where foreground means objects or regions of interest and is application-dependent. In this study, foreground is defined as fish and everything else is considered background, even other objects that might be moving such as the turbine itself and floating debris. This is a challenging data set for background subtraction because of the low quality of the video and the highly dynamic background. RPCA, GMM, and ViBE algorithms were selected based on recommendations from researchers at the University of Washington and the Fish4Knowledge project (Boom et al 2014) as being robust relative to background motion. The recommended parameter values for each algorithm were used.

The foreground images resulting from the background subtraction were further processed to group connected pixels of similar intensity into "blobs". These objects were then classified as fish or non-fish. The blob size was highly variable ranging from 1 pixel to over 10,000 pixels, so the blobs were divided into five size groups and classification models for each group were developed separately.

The motivation for including optical flow is the hypothesis that fish motion is different than other motion in the scene, such as the motion of objects drifting with the current and the motion of the turbine foils turning. The researchers who performed the manual analysis said that one of the features they used to recognize fish was directed motion. Optical flow is the motion (spatial displacement) of light intensity from one video frame to the next. It is calculated for video by matching regions in one frame with regions in the subsequent frame, where the matching is based on edges and gradients of light intensity, and the flow is the displacement. There are several algorithms for calculating optical flow in the literature. For this application, the Farnebäck algorithm was chosen to calculate a dense optical flow over the entire image. Initial tests of both the sparse and the dense optical flow methods indicated that the sparse method was not effective when analyzing the raw video because of the lack of strong gradient features that could be tracked from one frame to the next. Dense methods are more robust relative to some changes in object intensity and shape because dense methods use more surrounding context for matching features.

In a parallel effort, a Deep Learning model was applied to the development data set. The open-source machine learning library TensorFlow was used to build a convolution neural network and train it on a portion of the data set. This type of neural network must be trained on labeled data to generate a model for classifying new data. For this video analysis, the inputs to the network were the individual video frames, labeled as "fish" or "no fish". A subset of the labeled data (video frames) was selected at random to train the network and the resulting model was tested on the remaining data. This process was repeated over multiple iterations, where a new subset of the data was selected for training at each iteration. This iterative process is necessary to find the subset of the data that produces the best model.

4.2   Results and Discussion

One project goal was to reduce human labor time, hence the performance objectives were a 90% detection rate and a 30% false positive rate. The detection rate is the percentage of actual fish present that are detected, and the false positive rate is the percentage of reported detections that are not fish. It was important to detect most of the fish at the cost of some false positives because the false positives could be sorted out by human analysts, but too many false positives would reduce the benefit of the automation.

For an initial comparison of the background subtraction algorithms, the recommended parameter values for each algorithm were used (Table 4.1). The algorithms were evaluated for how well they correctly identified which frames contained fish. A frame was classified as "fish" if it contained one or more detections (foreground objects). The algorithms all performed similarly on the test bed data set (Table 4.2). The best detection rate was 67.51% (ViBE), which was much lower than the goal of 90%. The false positive rate was high, but the best true negative rate was better than 57%, which means that over 57% of the frames containing no fish were correctly labeled as such.

Table 4.1. Background subtraction algorithm parameters.

RPCA	ViBE	GMM
Window size = 50 Interval = 10 Threshold = 50	History = 20 Learning = 50 Radius = 20 Match criteria = 2 Update probability = 1/8	alpha = 0.02 threshold = 0.7 upper difference = 220 lower difference = 30

Table 4.2. Background subtraction frame classification results.

Algorithm	Percent of Fish Frames Correctly Detected	Percent False Positives	Percent True Negatives
RPCA	57.45	92.18	**57.60**
ViBE	**67.51**	**91.51**	54.48
GMM	63.79	92.29	52.19

The figures in bold indicate the best performance between the algorithms.

The background subtraction alone is not sufficient to meet the performance objectives, but the results offer valuable insight into the effects of night and day, the use of lights, and camera placement (see Section 7.0 for specific recommendations). The individual videos were analyzed to better characterize the algorithm performance under different conditions (Figure 4.4). All algorithms performed best on the videos from camera 1 at night, where there was no turbine in view and lights were on but angled away from the camera's field of view. All algorithms performed poorly in terms of false positives on the video from cameras 3 and 4 at night. The turbine was in view in both these cameras and the lights were aimed at the turbine. All algorithms performed worst on the video from camera 1 during the day, when most of the reported detections were false positives and most of the actual fish present were not detected. During the day, the fish were in low contrast with the background, so they were more difficult to detect. The lights at night reflected off the fish, increasing the fish's contrast with the background, so they were easier to

detect. Floating debris that was similar in size to small fish also reflected the light causing false positives (Figure 4.5).


Figure 4.4. The detection rate and false positive rate by test video for RPCA (top) and ViBE (bottom). *The numbers on the scatter plots indicate the number of frames (out of 7500 per video) that contained fish and the number in parentheses is the number of Fish Events. A Fish Event is when a particular fish is in view; an event usually spans multiple consecutive frames.*

Figure 4.5. Example frame from camera 4 at night with small fish (a) and the detected objects from ViBE (b). The small fish are detected but there are also some false positives from the illuminated debris. These false positives were eliminated during post-processing.

Using the results of the initial evaluation as a baseline, preprocessing techniques were evaluated. The two techniques included with the Fish4Knowledge code were histogram equalization and contrast stretch. The preprocessing did not add significant computation time, but neither technique improved the performance and in some cases had a negative effect. A bilateral filter (Tomasi and Manduchi 1998) is often used in photographic applications, and this technique was tested on two of the videos, one from camera 3 and one from camera 4. Only two videos were processed because the computation was extremely slow on a desktop computer, but the results were promising (Figure 4.6).

Figure 4.6. The raw frames (a) are preprocessed with a bilateral filter to reduce the clutter from debris (b).

The parameters of the best performing algorithm, ViBE, were varied to find the optimal values for the Igiugig data set. An exhaustive search for the optimal values of all five parameters was beyond the scope of this project, so two parameters with a strong influence on performance, the radius and the match criteria, were varied. The radius is the relative difference in intensity between background and foreground; a higher value will reduce the sensitivity and a lower value will reduce the precision (more false positives). Because detecting all the fish was more important than eliminating false positives, the radius was reduced and values of 20, 18, and 16 were evaluated. The match criterion is the minimum number of historical values that must fall within a current pixel's radius to consider the pixel to be background. A lower value will reduce sensitivity and a higher value will increase false positives and processing time. Values of 2 and 4 were evaluated. The best combination of values was radius = 16 and match criterion = 4 (Table 4.3).

Table 4.3. ViBE parameter tuning results.[a]

	Match = 2	Match = 4
Radius = 18	52% / 39%	70% / 54%
Radius = 16	55% / 42%	**74% / 59%**

(a) The first number in each cell is the true positive rate and the second number is the false positive rate. The figures in bold indicate the best combination of values.

A classification model for the detected objects significantly improved the performance by reducing the number of false positives. The detected objects were classified as "fish" or "non-fish" based on human analysis, and were divided into five size categories. A random sample of approximately 50 of each class (~100 observations) in each size category was used to develop a linear discriminant model for each category. For each size class, forward-stepping linear discriminant analysis followed by canonical analysis was used to determine the best model. Variables considered for the model included the blob size, blob solidity, blob eccentricity, and blob intensity. The models were tested on the remaining blobs that were not used for model development and the results are shown in Table 4.4. The larger blobs were classified most accurately; the accuracy decreased with decreasing blob size. The size of fish that can be accurately classified is dependent on their distance from the camera (the same size blob could be a large fish further away vs. a small fish close to the camera), and having the capacity to judge distance more effectively is discussed further in the Recommendations.

Table 4.4. Detected object classification results.

	Fish		Non-Fish		
Object Size in Pixels (number of fish objects)	True Positive	False Negative	True Negative	False Positive	Percent Correct
200+ (549)	533	16	3753	63	**98.2**
100 – 200 (320)	290	30	12,715	22	**99.6**
5 – 100 (2805)	2281	524	61,356	10,369	85.4
2 – 5 (2114)	1485	629	109,995	23,365	66.3
Total	4589	1199	187,819	33,819	84.6

The optical flow was calculated to generate a displacement in the horizontal and vertical dimensions for each pixel in the video frames, dx and dy, respectively. The displacements were then used to calculate the direction and magnitude of the motion from frame to frame. Direction was defined with 0 being toward the right and 180 toward the left. Five points in different regions of the frame were selected, and the motion was characterized at those points for one of the videos from camera 4 at night (Figure **4**.**7**). The direction of the motion appeared to be uniformly distributed across all directions and the magnitude appeared to follow a Weibull distribution (Figure 4.8). Due to the flow of the current from left to right in the video, the direction of motion was expected to be concentrated around 0 degrees. The uniform distribution that was observed may be due to the small, random motion of the floating debris. It may also indicate that the optical flow algorithm was not accurately matching points from one frame to points in the subsequent frame, due to the lack of distinctive features in the scene and the algorithm's bias toward small motion. The distribution of the magnitude of the motion shows that most of the calculated motions were small, with some outliers. The small motions may be due to debris and noise in the images, and the

outliers may indicate darting fish. The results are inconclusive and indicate a more in-depth investigation is needed, including testing the use of the bilateral filter for preprocessing to reduce the clutter in the scene.


Figure 4.7. The detection rate and false positive rate by test video for RPCA (top) and ViBE (bottom). *The numbers on the scatter plots indicate the number of frames (out of 7500 per video) that contained fish and the number in parentheses is the number of Fish Events. A Fish Event is when a particular fish is in view; an event usually spans multiple consecutive frames.*
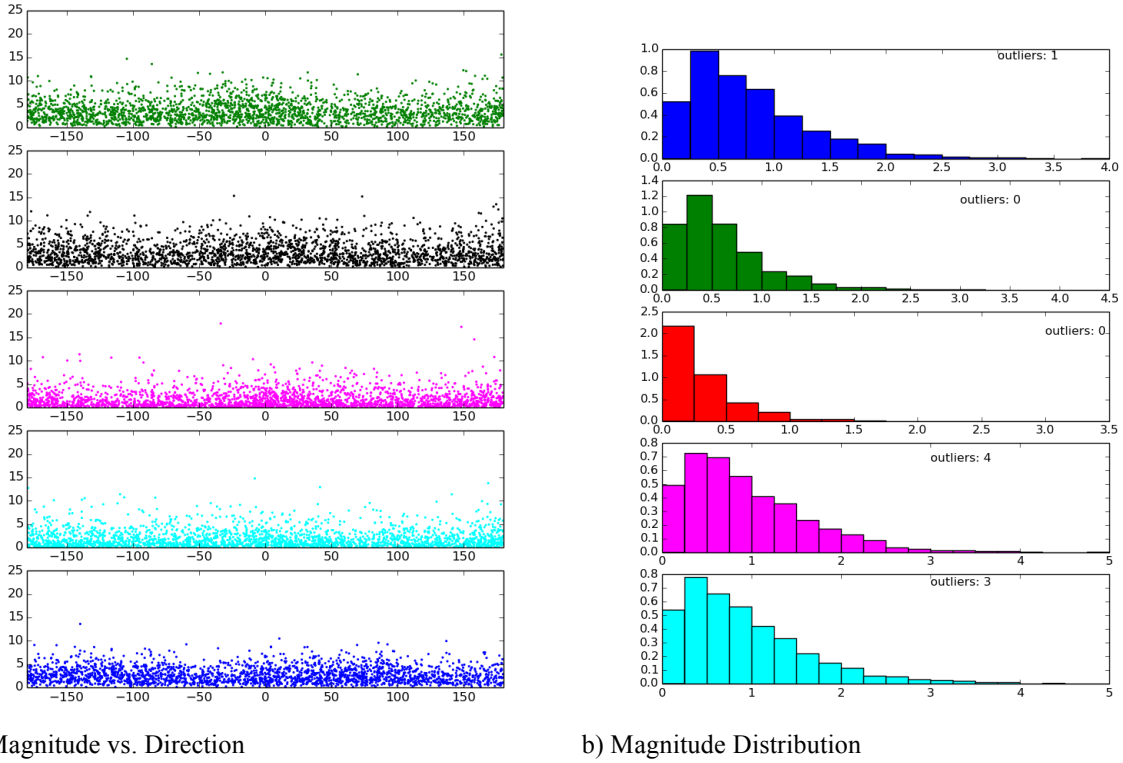
a) Magnitude vs. Direction　　　　　　　b) Magnitude Distribution

Figure 4.8.　The magnitude and direction of motion at five points in the video were characterized using optical flow. The direction appeared to be uniformly distributed across all directions (a), and the magnitude appeared to follow a Rayleigh distribution that implies the motion in the x and y directions are independent (b).

The best Deep Learning model developed over the iterative learning process achieved 79% correct classification of both "fish" and "no fish" frames. This result is promising and on par with the background subtraction/blob classification results. This approach could be suitable for batch processing a large volume of recorded data, like the Igiugig data set, especially if the processing can be done on a high-performance computing system of parallel nodes. However, this approach would not be suitable for a real-time system because of the computational intensity, and the requirement for training data that is specific to the location being monitored.

5.0　Software

5.1　Need and Requirements

EyeSea is a web application that was developed by PNNL to meet the need for a central repository for accessing and analyzing the terabytes of video data from the Igiugig project. During the manual analysis of the video data it was found that the proprietary format of the video data required the use of the vendor's specific software, which ultimately did not meet the requirements of the analysis team. Fortunately, an open-source video encoder (https://ffmpeg.org/) enabled the transcoding of the video data to the standard h264 format. This allowed the analysis team to use other more feature-rich software to perform the analysis of the video data. There was still the issue of how to make the h264 encoded video data available to the analysis team and also how to store results of the video analysis. To solve this issue a database-driven web site was designed, called "EyeSea". This web-based application was developed in

parallel with the algorithm development, and was envisioned as the framework for ultimately providing a user-friendly front-end to the automated analysis, combining all the analysis tools into one comprehensive "human-in-the-loop" system for video analysis.

5.2 Functionality

The database was designed to store video metadata (e.g., date, time, location, timezone), analysis results (e.g., fish detected, species of fish, location of fish in video frame), and site-specific data (e.g., log-in information, batch processing information). Figure 5.1 shows the schema of the database that was ultimately implemented in MySQL.

Once the database had been designed the next step was to implement the web application. Bottle (https://bottlepy.org), a web framework for Python, was used to implement the web site. An asynchronous architecture was designed to allow users to query video data and later return and view the transcoded results. This was necessary to enable users to query videos encompassing large amounts of time without causing a browser timeout. Although the website was designed to play back the video inside the web browser, an option was added to allow users to download the video to their local machine for later offline playback. The website was later extended to allow for in-browser analysis of video. Figure 5.2 shows a screen capture example of the in-browser video playback.

EyeSea was also designed to facilitate batch processing and analysis of video. A set of scripts were written that could be deployed on a cluster of servers for parallel processing of multiple videos. The servers query the database for jobs to process and communicate to the database the status of the job as they are completed. The batch processing feature of EyeSea was used to extract Fish Events for later use in the analysis algorithm development testbed.
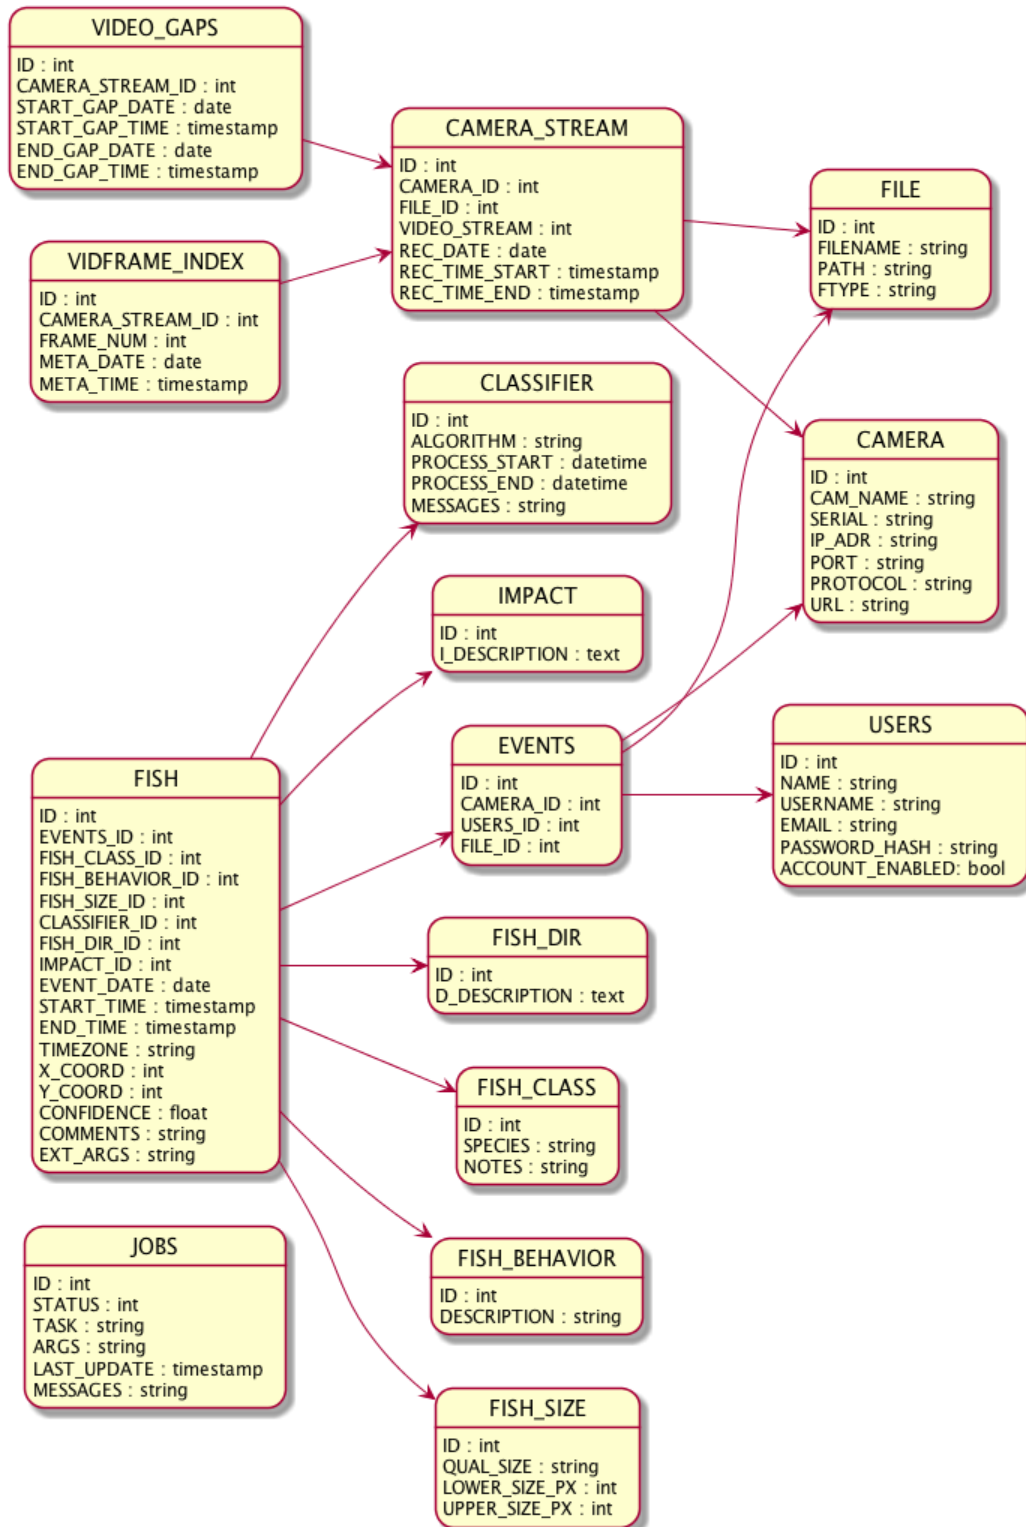
Figure 5.1. EyeSea database schema.

Figure 5.2. The EyeSea web application in playback mode.

6.0   Conclusions

6.1   Manual Analysis

The main points derived from the manual analysis of the data are as follows:

- Manual review of low-quality data is time-consuming. In this data set it took approximately 13–15 hours to manually review and annotate 1 hour of raw video data.
- Most interactions between the fish and the turbine occur at night.
- The frequency of fish interactions does not appear to be affected by whether the turbine is spinning or static.
- Processing subsamples (10 minutes) is likely effective for capturing unbiased event counts, but may not be effective for individual fish counts.
- Reviewer interpretation of Fish and Maybe Events in this data set is similar across two reviewers (qualitative analysis), but care is needed when assigning the designation of objects to introduced categories (quantitative analysis).
- Adult fish are qualitatively more likely to avoid collision or strike than juvenile fish.
- Adult fish are qualitatively more likely to show avoidance behavior as opposed to passive behavior relative to juveniles.

31

- More events occur when lights are on than when lights are off. Fish may be attracted to lights, lights may increase detection probability, or both.
- These data demonstrate the use of optical cameras for observing fish interactions with a deployed device in an underwater setting; however, improvements could be made with camera specifications and lighting parameters to increase the detection probability of fish, in both manual and automated review. Doing so may significantly decrease the manual and automated processing time.
- Observing "definite" vs. "probable" strike or collision is still extremely difficult and more research needs to be done to develop technologies or combine multiple technologies to gain confidence in determining actual contact with the device.

6.2 Automated Analysis

The main conclusions of the automated analysis effort were as follows:

- Tools available for detecting and tracking fish and other animals in underwater video are lacking. It was necessary to develop a new framework for semi-automated, human-in-the-loop analysis of underwater video. This framework can be used to test new algorithms and refine existing algorithms for automated fish detection and characterization, as well as support human expert analysis and standardized, reproducible information reporting.
- Reducing data volume is the first issue to address with automated processing. Large volumes of data are difficult to work with in terms of transferring, storing, and searching. A computationally simple background subtraction algorithm (ViBE) detected 74% of the human-identified Fish and Maybe fish, and is suitable for use in a real-time system to reduce data volume by saving only video that might contain fish.
- Reducing false positives is the second issue to address with automated processing. A statistical model can be used to classify detections as fish or not-fish, such as the one reported here that achieved a correct classification rate of 85% overall, and 92% for detections larger than 5 pixels. However, the statistical model required labeled training data that took time to assemble from the data and the model may not be transferable to other data sets. A classification model based on motion characteristics would potentially be more effective over a wider range of data.
- Underwater video recorded in energetic locations present challenges to automated processing that require algorithms specifically designed for this purpose. "Out of the box" algorithms such as those provided in the openCV library exhibited limited effectiveness, especially the optical flow techniques. Parameter tuning of the background subtraction algorithms did improve performance.

A combination of the automated detection developed under this project and human analysis could provide more accurate Fish Event information than the current practice of sampling, and with less labor time and cost than full analysis. Human analysis is currently the "gold standard" for accuracy, but it is very time-consuming so labor costs can be high and there may be long delays between collecting data and generating results. Sampling the video for analysis, e.g. analyzing 10 minutes of every hour reduces the labor time but sacrifices accuracy and increases the risk of missing rare events. The automated fish detection algorithms developed under this project can be completed quickly, but the resulting information is not as nuanced as that provided by human analysts and the detection accuracy is not yet sufficient so a "human-in-the-loop" approach is recommended. The automated detection software can be used to eliminate most of the video that does not contain fish, and the ensuing human analysis can be limited to those segments most likely to contain fish. With the developed processing system, this approach would reduce labor time by half over the full analysis, and would improve the reporting accuracy over sampling-based methods.

The performance of the automated processing can be further improved, based on the promising results demonstrated here. Future work should include incorporating computationally efficient bilateral filtering as a preprocessing stage, an intelligent scheme for parameter selection based on environmental conditions and video quality, and the integration of motion features. Further development of the EyeSea software should include a learning mode for tuning algorithm parameters using annotations provided by human experts.

7.0 Recommendations

The analysis of the Igiugig video data, both human/manual and automated, provided valuable insight into how to improve underwater video deployments in the future.

Since PNNL's review of the video data incorporated different approaches and anticipated outcomes than those of the original monitoring plan, recommendations arose regarding the monitoring process and methods for making project development and analysis more efficient in future studies. The water in the Kvichak River was described as being very clear compared to other rivers in the original monitoring report, yet the video data were described in PNNL's Quality Check Summary Report (Trostle 2016) as being "usable", in that the reviewers would be able to describe fish presence. The declaration of "usable" embodies the overall quality of the video data, including the following factors: resolution, frame rate, the placement of the cameras and light sources, the field of view, and the settings of the digital video recorder camera system. Careful consideration of the anticipated review and analysis objectives should be applied when making a camera selection. Those who will be reviewing the data should consider the questions they would like to answer and make sure that their camera selection, settings, and placement have the potential to address those questions.

To increase the quality of the video, future studies should use a low lux camera with a higher resolution and faster, even frame rate, but be aware that this will increase data accumulation because the files will be larger. Higher resolution video data would increase the likelihood that the manual reviewers could decipher between Maybe and Fish Events, identify taxonomic classification, and have more confidence regarding strike and collision. An increase in frame rate will improve the ability to detect actual strike because there would be more frames to describe the interaction around the turbine. It would also allow the reviewer, and possibly the algorithm, to use behavior as a qualifier, because sometimes the object would move significantly between frames, making behavior difficult to determine. Additionally, in some cases an object would only be in the field of view for one or two frames, making it difficult to determine if the object was a fish or not. In this study, objects that were only in one frame were not recorded as an event, because there was insufficient information to describe or categorize the object. With a more frequent frame rate and higher resolution, those objects could be included and give the study a broader picture, because the probability of missed events would be lower.

During manual video review, the reviewers realized that full manual analysis was too time-consuming. For this data set it took manual reviewers approximately 13–15 hours to process 1 hour of raw video data. A number of factors affect this approximation of time, including light operations, whether it was day or night, the number of fish, the behavior(s) of the fish, the amount of debris, the quality of the video, and whether or not the turbine was spinning. No 1-hour segment was identical to another in terms of time spent by manual reviewers due to the variability in the factors listed above, making the time spent extremely inconsistent. For this reason, future studies should be cautious when developing a timeframe estimate for manual processing, and reviewers should be wary that the anticipated estimation of time spent may change.

As described in the Quality Check Summary Report, a great deal of work and time was put into understanding the methods and settings implemented throughout the study, and converting the video from

a propriety format (.par), which was designed to be tamperproof, to a more appropriate format for the development of automated processing and analysis (.mp4). A more accessible format, such as .mp4 or .avi could be used for ease of use and to enable automated analysis before manual review.

As PNNL reviewers sifted through the video data, they noticed a variation in the light operations. In general, the lights seemed to be on at night and off during the day, with at least one exception on July 19, but a light operation record was not maintained during the study. Because there were many more Fish Events at night, it is important to get a better understanding of the effect artificial lights have on fish behavior (e.g., whether lights attract or repel fish). It is also important to quantify the difference the lights make to the physical parameters of detection (e.g., define more robust limits of detection, and define optimal placement and settings of light sources and cameras to increase manual and automated detection potential).

Recommendations for future underwater environmental monitoring are listed below:

- **General setup.** Include an indication of range within the field of view to help reviewers distinguish size and location in relation to the turbine. Also, aim the camera so the field of view is aligned with particular turbine components, possibly in combination with sensors on the turbine foils to increase the detectability potential and promote a higher level of confidence during potential strike and collision events. The aiming of each camera will likely require an iterative process of viewing early data and making adjustments to achieve ideal viewings for manual processors as well as ideal background for algorithm applications.
- **Video format**. A standard format (e.g., .avi, .mp4) should be used, rather than a proprietary format. When the video is in a standard format, researchers have a wide array of existing tools that they can use for analysis and processing. A proprietary format restricts researchers to using vendor-supplied software that often is not designed for the type of analysis required.
- **Camera type.** Choose a camera that has underwater capability with high pixel resolution in low light conditions, the capability to mount and adjust placement settings, appropriate data storage and transmission, and a suitable field of view range for the study area.
- **Camera resolution and placement.** The camera resolution will determine the size of objects in pixels at a given distance from the camera. Objects that are less than five pixels in total size are difficult to detect, both algorithmically and visually. Higher resolution will increase the volume of data, but low resolution will restrict the size of fish that can be reliably detected. The size of a fish in pixels, along the horizontal dimension, is

length of fish in meters / meters per pixel

The meters per pixel is

$$\frac{2r \tan \frac{\alpha}{2}}{n}$$

where r is the distance to the fish in meters, α is the horizontal camera field of view angle, and n is the number of pixels in the horizontal dimension. For example, a 10 cm (4 in.) fish would be 10 pixels long at 10 m from a 320 × 240 pixel camera with a 20 degree horizontal field of view. This calculation will also help determine how far from the turbine to locate the camera. Test the placement of the cameras and lights to optimize manual and automated detection probability. Note where the sun will be throughout the study and test different angles to avoid glare.

- **Frame rate.** Ideally, the frame rate should be constant, meaning that there is a fixed interval of time between frames. The Igiugig video had a variable frame rate that resulted in uneven motion of objects from frame to frame. Higher frame rates increase the volume of data, but if the frame rate is too low the number of frames in which a fish may be in the field of view is decreased, decreasing the

probability of detection. A rate of 30 frames per second is a reasonable choice to balance data volume with detection likelihood.

- **Lighting**. Fish specific to this study are typically more active at night, so some sort of illumination is needed if video is the only monitoring technique used. The light also generated more false positives from reflecting debris. An indirect light source, like the lighting viewed from camera 1 in the Igiugig video may be the best choice. If lights are to be used, the lights should be on throughout the study to maintain a more controlled environmental setup, and increase light sources with more angles of incidence to prevent fish from disappearing when they turn at an angle that does not reflect light from a single source. However, while improving the detection of fish and debris, this practice may also introduce possible bias because of the lights themselves increasing detection probability or attracting fish, both of which complicate comparisons when lights are turned off or confounded by diel differences.

- **Detailed record keeping**. The following aspects should be recorded:

  - all monitoring operations, including camera operation, light operation, power operation, turbine status, and any other introduced monitoring systems.
  - water flow, weather conditions and any significant events that occurred during the study.
  - any maintenance issues or disruptions throughout the study.
  - review efforts.

- **Other monitoring**. Consider adding other monitoring technologies to help determine whether actual collision or strike occurred and to have a backup technology for behavioral monitoring. Strain gauges or other devices physically attached to the blades of a turbine could be used to complement the video data for those times when a collision or strike is possibly seen. Having coincident data sets providing evidence of collision or strike would be better than just one. For instance, if a reviewer thinks a strike was seen on the video data, the same timestamp could be searched for blade-attached strain gauges to see if there was a spike. If there was an anomaly on the strain gauge, then that is more evidence of a strike. The absence of strain gauge data would be evidence that the interaction was more likely a near-miss.

To inform future studies, additional research into specific aspects would also be useful, in particular a study to assess the effects of lights on fish, in conjunction with an evaluation of light and camera placement and settings toward the optimization of detection probability to find an ideal experimental design for manual and automated detection. The study would record details of the placement, including heading, pitch and roll of both the lights and the cameras, light operations, intensity, wavelength, exact range in relation to the device and cameras, as well as camera operations, range, resolution, frame rate, and settings. For algorithm development, further research is recommended into different optical flow techniques, and the refinement of parameters for the background subtraction. Each of these aspects would improve the results derived from future monitoring. Each study site will have specific physical characteristics that will affect underwater video camera data collection. However, as research continues on future data sets, general application principles will arise that can be applied to most situations.

8.0   References

Avila A, Hull R, Matzner S, Staines G, Trostle C, and Harker- Klimeš G. 2016. *Igiugig Video Analysis – FY16 Progress Report* PNNL-25909 Pacific Northwest National Laboratory, Richland, WA.

Barnich O and Van Droogenbroeck M. 2009. ViBe: a powerful random technique to estimate the background in video sequences. In proceedings of *Acoustics, Speech and Signal Processing, 2009.*

ICASSP 2009. IEEE International Conference on Acoustics, Speech and Signal Processing (pp. 945–948). IEEE.

Boom BJ, He J, Palazzo S, Huang PX, Beyan C, Chou H-M, Lin F-P, Spampinato C, and Fisher RB. 2014. A research tool for long-term and continuous analysis of fish assemblage in coral-reefs using underwater camera footage. *Ecological Informatics*, 23:83–97.

Candès EJ, Li X, Ma Y, and Wright J. 2011. Robust principal component analysis? *Journal of the ACM (JACM) 58*(3), 11.

Farnebäck G. 2003. Two-frame motion estimation based on polynomial expansion. In proceedings of *Scandinavian Conference on Image Analysis* (pp. 363–370). Springer Berlin Heidelberg.

Lee DS. 2005. Effective Gaussian mixture learning for video background subtraction. *IEEE transactions on pattern analysis and machine intelligence 27*(5), 827–832.

LGL Alaska Research Associates, Inc. 2015. *Fish and Wildlife Monitoring Plan for RivGen® Testing on the Kvichak River, Alaska in 2015*.

Lotlikar R and Kothari R. 2000. Fractional-step dimensionality reduction. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22(6): 623–627.

Lucas BD and Kanade T. 1981. An iterative image registration technique with an application to stereo vision. *International Joint Conferences on Artificial Intelligence (IJCAI)* 81: 674–679.

Tomasi C and Manduchi R. 1998. Bilateral filtering for gray and color images. In *Computer Vision, 1998. Sixth International Conference on Computer Vision* (pp. 839–846). IEEE.

Trostle C. 2016. *Igiugig Fish Video Study– Quality Check, Summary Report*. Pacific Northwest National Laboratory, Richland, WA.

Appendix A

Manual Annotation

Appendix A

Manual Annotation

Annotation	Description of Annotation
Event	Reference number per event; restarting for each half-hour block of data
Date	Day video data was collected; yyyy/mm/dd
File	Filename given to each half-hour block of data; includes day and time
FileStartTime	Time that file starts on the given day it was collected
StartTime	Time that an event begins; begins 00:00:000 per half-hour block of data
EndTime	Time that an event ends; ends 29:59:999 per half-hour block of data
Lights	Either on or off; binary
Spinning	Either yes or no; binary
Camera	Designated number of camera; these data only include Camera 2
Fish?	Is the event triggering object a fish; yes, no, maybe
Number	How many objects or fish occur during an event
Size	Size of objects or fish seen during an event; measured as length on computer monitor; unidentifiable, small (<0.5 in), medium (0.5–3.0 in), large (>3.0 in.); was adjusted relative to monitor screen sizes
Species	Visually identifiable relative size designation or salmon; unidentifiable, juvenile, salmon, adult
VideoQuality	Relative anecdotal comparison of each event relative to others based on clarity of event triggering object in field of view; horrible, bad, okay, good, excellent
Notes	To clarify any previous annotation categories
Location	Where the event triggering object is in the water column; based on computer monitor divided into thirds; bottom, middle, top
Direction	All observed directions of the event triggering objects or fish; downstream, upstream, cross river toward, cross river away
Behavior	Reviewer description of all object or fish behaviors observed during an event. • straight across • against current • pause • avoid above • through turbine

- avoid below
- avoid reverse
- out of turbine
- milling
- toward static parts
- through wake
- avoid around
- unable to tell
- other

Impact	Reviewer determination if there was collision or strike during an event
Comments	To clarify any previous categories since "Notes"

A- straight across
B- against current
C- pause
D- avoid above
E- through turbine
F- avoid below
G- avoid reverse
H- out of turbine
I- milling
J- toward static parts
K- through wake
L- unable to tell (not shown)
M- other (not shown)

Appendix B

Video Data Set Used for Algorithm Development

Video Data Set Used for Algorithm Development

Video File (Date, time, camera)	Day | Appendix BLights	Turbine | Spinning | None	Fish Events	Fish Frames
`20150719_175830-1.mkv`	Day	None	7	308
`20150720_110030-1.mkv`	Day	None	9	307
`20150722_030200-1.mkv`	Lights	None	30	917
`20150722_030200-2.mkv`	Lights	Turbine	72	878
`20150722_030200-3.mkv`	Lights	Turbine	7	283
`20150723_000330-1.mkv`	Lights	None	2	97
`20150723_000330-2.mkv`	Lights	Spinning	68	519
`20150723_000330-3.mkv`	Lights	Spinning	7	268
`20150723_000330-4.mkv`	Lights	Spinning	8	423
`20150724_000000-1.mkv`	Lights	Turbine	5	499
`20150724_000000-2.mkv`	Lights	None	53	438
`20150724_000000-3.mkv`	Lights	Turbine	10	1007
`20150724_000000-4.mkv`	Lights	Turbine	9	311
`20150825_040330-1.mkv`	Lights	None	3	201
`20150825_040330-2.mkv`	Lights	Turbine	38	467
`20150825_040330-4.mkv`	Lights	Turbine	3	160
Total			**334**	**7569 (6%)**