

Progress Toward Automation of Salmon Escapement Counts

Mark Shortis

School of Mathematical and Geospatial Sciences, RMIT University

and

Ted Otis

Alaska Department of Fish and Game, Homer, Alaska.

Email: mark.shortis@rmit.edu.au

Abstract

The number of adult salmon that escape predators and fishing to successfully return to their natal streams to spawn, known as escapement counts, are vital to ensure sustainability of the salmon fisheries. Sufficient numbers of salmon must reach the spawning grounds each year to ensure the long term viability of the various species. Once a critical threshold is reached, the fishery is opened to commercial exploitation. To ensure the viability of salmon in Alaska, the Alaska Department of Fish and Game conducts escapement counts using a variety of techniques. One well established technique is to use remote video stations to record salmon passing upstream and periodically use human operators to view the video sequences to determine the escapement counts. This paper presents the initial stages of the development of a technique to automate the escapement counts and thereby relieve the demand for human operators. The technique uses a global segmentation and image morphology operations followed by a Kalman filter to identify and track the candidate salmon as they pass through the region of interest on the video recording. Candidates that complete the journey across the region of interest are counted. Finally, some of the future challenges posed by real time processing, water surface turbulence and lighting changes are discussed.

Mark Shortis

Mark R. Shortis is Professor of Measurement Science in the School of Mathematical and Geospatial Sciences and Director, Program Quality in the College of Science, Engineering and Health at RMIT University in Melbourne, Australia. Since the early 1980s he has been active in research in photogrammetry at close range, for example collaborative research with NASA Langley Research Center in Virginia on the structural dynamics of aerospace models, with the University of Western Australia on underwater assessment of marine fish populations and habitats, with ANU in Canberra on the characterisation of solar concentrators for energy generation and with University College London on high precision, industrial metrology.

Ted Otis

Edward O. Otis IV (Ted) has a Master of Science from the University of Arizona and is currently a Fishery Research Biologist with the Alaska Department of Fish and Game, based in Homer. Ted is responsible for assessment and forecasting of Kamishak Bay herring stock; directs salmon and herring catch/escapement-sampling programs; forecasts Lower Cook Inlet salmon returns; and develops new approaches to monitoring salmon escapement, for example remote video and time-lapse recording. He has previously been employed by the U.S. Fish and Wildlife Service in Kenai, Alaska managing projects to count salmon and track trout using radio-telemetry, and by the University of Arizona's Cooperative Fish and Wildlife Research Unit to assess fish stocks in tributaries of the Colorado River in the Grand Canyon.

Keywords: salmon, escapement, counts, video, image processing, identification, tracking

Introduction

Salmon escapement counts are vital to ensure sustainability of the fisheries in Alaska. Sufficient numbers of salmon must reach the spawning grounds each year to ensure the long term viability of the various species. In remote areas the salmon must escape their natural predators such as salmon sharks, seals and orca in open water, and then bears and eagles at the entrances to streams and rivers. In more populated areas the salmon are also subject to recreational and subsistence fishing, usually at or near stream and river mouths. Catches are limited by daily or seasonal limits per person/household. Only when sufficient numbers have passed through the predation zones and will reach the spawning grounds are the fisheries opened to commercial exploitation.

Commercial operators apply pressure to open the fisheries as early as possible to maximise the fishing season, consequently escapement counts have to be both accurate and reliable.

A number of techniques are used to monitor escapement. Traditionally, artificial weirs and manual counts by a resident fishery officer have been the mainstay for salmon escapement (Zimmerman and Zabkar, 2007). However, manual counts at every stream, especially in remote locations, are not practical or sustainable. Techniques such as hydroacoustics (Maxwell, 2007) and counting towers (Woody, 2007) are successful in some applications. Aerial surveys are widely used (Jones et al., 2007), but are a relative index rather than a count of individual fish because accurately counting thousands of salmon from the air is not feasible (Bue et al., 1998).

To bolster the accuracy and reliability of escapement counts and allow remote streams to be monitored efficiently, the Alaska Department of Fish and Game (ADF&G) has pioneered the use of remote video stations (Otis and Dickson 2002). Time lapse video is captured at a rate of 2-5 frames per second to a digital video recorder. The video recordings are transported to the ADF&G office in Homer for periodic manual counting. Whilst this approach is accurate and reliable, it is labour intensive and staff are subject to intense pressure as escapement counts approach the required threshold at each site. Automation of the process is a key aim to improve the efficiency of the counts.

This paper describes the progress made to develop an automatic video counting system that operates in real time or rapidly analyses pre-recorded videos. The developed algorithm has the capability detect, track and count the salmon passing through a defined region of interest in the video frame. The first step in the processing is to segment the fish objects in the current video frame using a global threshold technique. Morphological image processing (dilation, erosion and median filtering) is then used to eliminate speckle noise, eliminate small regions which are likely not related to fish objects, and connect parts of the same fish object separated by water surface turbulence. The third step of the system is to track the fish objects using a Kalman filter to overcome occlusions and gaps in identification. The Kalman filter predicts the track up to a certain number of frames, beyond which the object is eliminated as a false positive. Finally, fish objects with a full track across the region of interest are counted.

Remote Video Stations

In the 1990s the ADF&G at Homer began developing stand-alone remote video salmon escapement monitoring stations to supplement aerial surveys of remote, clear water streams (Otis and Dickson 2002). A typical system is based on an aluminium quadrapod sited next to a shallow area of the stream that affords a good view of the salmon migration (figure 1). A video camera in a weatherproof housing is mounted at the top of the quadrapod, viewing the stream obliquely from the south to minimise reflected surface glare from the sun. A light coloured substrate panel is used to provide sufficient contrast to make the salmon visible in most conditions of lighting and surface turbulence. The camera is outfitted with a circular polarizer to optimise the clarity of the view of the substrate panel area.

The video is recorded at 2-5 frames per second to a digital video recorder (DVR) inside a bear-proof strongbox. The DVR receives a NTSC composite (analogue) video signal from the camera but records digital video sequences onto a hot-swappable hard disk drive. Electrical power is provided by a bank of 12V lead-acid batteries connected in parallel. Wind and solar power generators are used to maintain the battery charge level. An automatic alarm system and satellite communications relay is used to alert ADF&G if equipment failures occur.

The field of view of the overhead cameras is optimised to match the panel area as far as possible. Figure 2 shows the fields of view from the systems at Chenik Creek and Mikfik Creek, both on the western side of Cook Inlet. Video sites are chosen based on a range of factors such as accessibility, stream morphology, minimisation of surface turbulence and reflections, availability of solar, wind or hydropower, and many others. As noted previously, the optimal direction for the camera view is to the north and this requirement can lead to a compromise between the field of view and the resolution of the view. The orientation of the format of the camera is physically limited by the weatherproof housing design, so in general the substrate panel is not aligned with the highest resolution axis of the camera. At some sites other requirements, for example a desire to monitor human traffic near the site, forces a wider than optimum field of view.

The DVR has four input channels and some sites also have a low-light camera mounted on the quadrapod or an underwater camera with a view across the stream bed. The underwater camera is used primarily to positively identify fish enumerated with the above-water camera and to apportion the total escapement among co-occurring species with overlapping run timing. The low light camera does not have a polarizing filter and is used primarily during twilight periods to minimize missing counts caused by the 4-5 hours of darkness during June to August. Previous research that compared video counts to those derived by a weir at Delight Creek in South-central

Alaska indicated that less than 2% of the salmon escapement passed the video counter during the hours of darkness (Otis and Dickson 2002).




Figure 1. Typical remote video station, located at the outlet of Chenik Lake into Chenik Creek.




Figure 2. Typical views from the overhead camera. Chenik Creek (top) and Mikfik Creek (bottom).

The hard disks have sufficient capacity to record video for 2-3 weeks. Regular visits to the remote sites are carried out to check the equipment and swap hard drives, with visits increasing in frequency as escapement counts approach the threshold necessary to open the commercial fishery. Maintenance visits are commonly pooled with aerial survey flights to efficiently use the float plane travel, as this is a significant component of the budget for escapement counts.




Figure 3. Example salmon escapement count profile from Chenik Lake, Alaska, 2009.

The hard disks are returned to the ADFG office for analysis. The disks are mounted in a multi-screen video station for replay. If multiple cameras are available then all captured videos can be shown on the screens with

synchronous replay. An operator manually records the numbers of salmon passing over the substrate panel using a hand counter. The operator can fast forward through periods of inactivity or step slowly through the individual frames of groups of fish passing over the panel. An operator can review a full day of video in as little as 10-20 minutes during periods of low fish count. However, 1-3 hours are often needed to review peak passage days which may involve 2000-4000 fish. Salmon runs span as many as 40 days and may require a total of 15 hours of analysis by an operator, dependent on the frequency profile and total count. A typical profile of daily salmon escapement counts is shown in figure 3.

Video Imagery

The quality of the video imagery recorded by the remote stations is, in most conditions, satisfactory for an operator to identify and count the salmon. The resolution of the video is limited to 720 by 480 pixels by the NTSC broadcast signal, which was adopted in the design of the remote stations for compatibility with the low cost surveillance cameras and security industry DVRs that are available off-the-shelf. This technology has many advantages, particularly that it has a proven track record, is very robust and there is a vast knowledge base of systems and technological solutions. In addition, the ability to readily inject a time code into the video image or the audio track allows a permanent record of the date and time, and allows accurate synchronisation of multiple cameras.

Nevertheless, the NTSC resolution is a limitation that is more evident for some stations, such as the Mikfik Creek site shown in figure 2. The field of view could be narrowed using a longer focal length lens, but in some cases the width of the stream is too wide to permit sufficiently high resolution images.

A further limitation on the quality of the video imagery is the video compression. The DVR units used in the remote video stations use a proprietary Motion JPEG compression to maximise the recording time by the units. The compression, combined with the relatively low number of frames per second, allows the units to be deployed for 2-3 weeks of 24 hour recording. This is particularly important during off-peak periods when the hard drives are not swapped out as frequently.

As is typically the case for surveillance DVRs, the operator cannot disable the video compression but does have some control over the level of compression. However the compression levels are high even within the allowable range. For example the lowest to the highest compression ratios per frame are 30:1 to 65:1. Visually there is little discernible difference between the extremes of quality recorded by the DVR. An apparent characteristic of the proprietary Motion JPEG algorithm is that adjacent odd and even horizontal lines in the image are adjusted to have similar RGB intensities, leading to an effective halving of the vertical resolution and further deterioration of the image quality through loss of resolution. Unfortunately testing with alternative compression algorithms, such as MPEG, produced even greater deterioration in image quality and also had a greater propensity for dropped frames when the video was reviewed at faster replay speeds.

Automation

As discussed above, operators can review an entire 24 hour video recording in a span of 10-180 minutes depending on the number of salmon to be counted. This level of efficiency is admirable for a human operator, however over a full season with many remote video stations the total time for review can be measured in hundreds of hours.

Computer-based image analysis has been used in marine and freshwater science for many tasks associated with detection and identification of fish from still images and video sequences. For example, Walther et al. (2004) detected candidates for estimates of biomass using motion analysis from a ROV-based video camera. Tillet et al. (2000) used image analysis for fitting to a species-specific template of the body shape of fish to extract specific metrics. There are many other examples of the use of image analysis for biological studies (Derry and Elliot 1997, Kocak et al. 1999, Rova et al. 2007) that use or build on these two fundamental processes of candidate detection and template matching.

More recently, Spampinato et al. (2010) reports on a system for detection, tracking, counting and classifying fish that is in continuous operation on a sub-tropical reef in Taiwan waters. Charalampidis et al. (2012) describe a NOAA research project investigating recognition and tracking of fish in video sequences. Both systems broadly use subtraction of the current image from a 'clean' background image; segmentation of the image using threshold or histogram techniques; morphological operations to eliminate noise and merge adjacent regions; and then edge detection and boundary definition to detect and identify candidate fish objects. Charalampidis et al. (2012) then uses a Kalman filter to predict the locations of each candidate in subsequent frames.

Full automation of the salmon counts is technically feasible based on a similar process of segmentation of the image to isolate the objects in motion, morphological operations to eliminate noise and merge regions, and then tracking of individual fish through a region of interest (RoI) to compile the count. The RoI can be placed within the video frame by an operator to ensure that all fish are detected and counted.

The behaviour of the salmon facilitates the RoI approach, rather than attempting the virtually impossible task of uniquely identifying each fish, because the salmon rarely deviate from a direct path across the substrate panels. The fish often pause at the threshold of the substrate panel because it is an unnatural feature, but then swim directly across the panel to the natural creek bed on the far side. The salmon are driven by instinct to swim upstream, so multiple factors militate against the fish repeatedly swimming or hesitating over the panel. Accordingly, a feasible approach is to track the fish objects across the RoI and count each detected object that completes the passage.

Video Processing

The video processing and tracking approach described in the previous section has been implemented using Matlab¹ as a development and testing platform. Once a mature, generalised solution is developed the intent is to implement the video processing algorithms as a stand-alone application that can be deployed on any PC. Each step in the process is described in more detail below.

The first step in the processing is to define the RoI within the video frame. This approach is feasible because of the fixed nature of the field of view of the remote video stations. The Matlab function *roipoly* allows the operator to specify a polygonal RoI within an image from the video sequence and creates a binary image that is used as a mask. The mask is then applied to every subsequent video frame to eliminate areas outside the RoI from further processing. The RoI reduces the computational burden on the system, increases the effectiveness of the segmentation performance and avoids the false detection of fish objects in the cluttered image areas outside of the RoI.



Figure 4. Definition of the vertices of a polygon to select the RoI within the substrate panel.




Figure 5. Polygonal mask created to exclude the area outside of the RoI.

An example of the operator selected polygon and the subsequent mask are shown in figures 4 and 5. The RoI should match the substrate panel, but should avoid the edges and shallow sections near the shore to provide uniform background intensity. If the salmon are hesitating, or entering and then re-entering the substrate panel area, the downstream edge of the RoI can be positioning further inside the substrate panel area to minimise tracking errors.

Segmentation is the next step in the video processing. This commences with converting the video frame from RGB colour into an average greyscale intensity for each pixel using the *ColorSpaceConverter* function. The *medfilt2* median filter function is applied next to remove noise and smooth the image. Under the simple assumption that the substrate panel will always be a relatively high intensity value and the fish objects will be a relatively low intensity value, a simple average of the minimum and maximum intensities in the RoI is used to globally segment the image and create a binary image of the candidate fish objects. This approach is based on the principle of a bi-modal distribution of intensity used in the Otsu method (Otsu, 1979). In mathematical terms:

$$thresh = \frac{m1+m2}{2} \quad \text{and} \quad dst(x,y) = \begin{cases} 1 & \text{if } src(x,y) \geq thresh \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

¹ <http://www.mathworks.com.au/products/matlab/>

where $m1$ and $m2$ are the maximum and minimum values of the intensities in an intensity image, $dst(x,y)$ is the pixel value of destination binary image and $src(x,y)$ is the pixel value of the source greyscale image. Two examples of original RGB images and the segmented, binary images are shown in figure 6. The binary images include fish candidate blobs, non-fish blobs and noise.




Figure 6. Original frames and corresponding segmentation. Test video on the left and quadrapod video on the right. Note that second salmon is not visible in the left segmented image because it is outside of the RoI.

The image morphology dilation and erosion functions *Dilate* and *Close*, using a 3 by 3 pixel square structuring element, are then applied to remove noise and merge adjacent blobs. The next step in the process is the removal of any remaining small blobs within the RoI. A threshold of size can be set empirically based on the apparent size of the salmon within the video images, which is in turn based on the set up at each remote video station. The difference in the image span of the salmon between Chenik Creek and Mikfik Creek is readily apparent in figure 2. An example of noise and small blob removal is shown in figure 7.




Figure 7. Segmented image (left), noise and small regions removed (right).

The Matlab *BlobAnalysis* function is then applied, with empirically determined minimum and maximum blob sizes, to analyse the RoI and identify the fish candidates. The blob analysis identifies the connected groups of pixels and computes characteristics such as the area, centroid and bounding box. The effectiveness of the blob analysis algorithm relies completely on the quality of the segmentation. A poor segmentation will result in no candidates being detected or many false positives due to lighting conditions or noise in the image. In turn this will lead to inaccurate results from the tracking and counting processes, described in the next section.

Tracking and Counting

In computer vision moving objects are tracked for variety of applications such as activity recognition, traffic monitoring and automotive safety (Hu et al., 2004; Moeslund et al., 2006). In this project, tracking is used for the purpose of counting the fish candidates. Each object is tracked from the point it enters the RoI to the point where it exits the RoI. The foundation of this approach is that migrating salmon are the only likely cause of an object transiting the RoI from the downstream edge to the upstream edge. If an object completes the journey across the RoI it is classified as a salmon and added to the count. Static objects and transitory candidates caused by surface ripples or reflections are automatically eliminated.

The multiple objects in the RoI are tracked with a Kalman filter (Kalman, 1960). The Kalman filter has been used for many applications in trajectory modelling, both within the 2D image space for computer vision and within 3D world space for positioning. The Kalman filter is a recursive solution to discrete-data linear filtering, used to estimate the state of a linear system assumed to be distributed as a Gaussian function. The filter supports estimation of past, present and future states, and can compute the estimates when the precise nature of the modelled system is unknown (Welch and Bishop, 2004).

In this project a Kalman filter has been used for object tracking by predicting the position from the previous information and verifying the existence of the object at the predicted location. In Matlab, the estimates are computed using the *vision.KalmanFilter* system object. The association of candidate detection to the same object is based solely on motion. The motion of each track is estimated by the Kalman filter which determines the likelihood of each candidate detection being assigned to each track.

Track maintenance becomes an important aspect of counting the salmon. In any given frame, some candidate detections may be assigned to tracks, while other candidates and tracks may remain unassigned. The assigned tracks are updated using the corresponding detections. An unassigned candidate detection begins a new track. Candidates must be tracked for a certain number of consecutive frames to be assigned to a track. Tracks that are unassigned for a certain number of consecutive frames are deleted. False positives caused by noise or surface turbulence tend to result in short-lived tracks. The limit is set empirically and in most circumstances a value of five consecutive frames is effective.

Assigning object detections in the current frame to existing tracks is done by minimizing a cost function. The cost is defined as the negative log-likelihood of a detection corresponding to a track. The algorithm involves two steps:

- Step 1:** Compute the cost of assigning every detection to each track using the distance method of the *vision.KalmanFilter* system object. The cost takes into account the Euclidean distance between the predicted centroid of the track and the centroid of the detection. It also includes the confidence of the prediction, which is maintained by the Kalman filter. The results are stored in an MxN matrix, where M is the number of tracks and N is the number of detections.
- Step 2:** Solve the assignment problem represented by the cost matrix using the *assignDetectionsToTracks* function. The function takes into account the cost matrix and the cost of not assigning any detections to a track.

The value for the cost of not assigning detection to a track depends on the range of values returned by the distance method of the *vision.KalmanFilter*. This value must be tuned experimentally. Setting it too low increases the likelihood of creating a new track, and may result in track fragmentation. Setting it too high may result in a single track corresponding to a series of separate moving candidates.

Figure 8 shows an example of fish candidates mapped onto the original frame, so that the results of segmentation and counting can be shown within the RoI. The top left annotation of the bounding box shows the fish counter and top right annotation shows the life of the candidate in number of frames. When the candidate is occluded or fragmented by surface turbulence, or is touching an edge of the RoI, the top right annotation indicates that the location is 'predicted'. The green cross marks show the centroids of the candidate blobs.

Candidates that complete an unbroken track across the RoI are added to an accumulation count of salmon.




Figure 8. An example of tracking and counting using the Kalman Filter.

In the limited testing of good quality video sequences that has been conducted to date, the counter has been 100% reliable. However this perfect reliability has to be strongly qualified by the quality of the video and the heuristic nature of the approach. To establish a baseline of effectiveness, the initial testing has concentrated on video sequences in which the salmon are clearly visible and counts can be easily verified by a human operator. Moreover, factors such as the limits on blob size, the upper limit for the elimination of small blobs and the selected number of frames used to eliminate false positives must all be chosen very carefully.

Conclusions

The approach to salmon escapement counts described in this paper has the potential to relieve the demand for manual counts by operators. The approach has been successfully applied to pre-recorded video sequences. However it is yet to be tested on live video and has been applied in a limited range of water surface conditions. Further, testing is needed to validate the results against reference data provided by experienced human observers.

The most difficult circumstances for salmon counts are twilight periods, wind-induced surface turbulence, sun-induced surface glare and discoloration of the substrate panel. Figure 9 illustrates the range of different conditions experienced over the course of a single day at Chenik Creek. It is evident that, as conditions deteriorate, identification and counts of salmon become commensurately more difficult. Despite the use of a polarizing filter, surface turbulence from the wind and upstream obstacles can exacerbate sun-induced surface glare to the point where the ability to see and count fish is severely compromised. The substrate panels can become discoloured by deposition of silt or algae growth, or may be interfered with by bears or floods. Fortunately, interference with the panels is infrequent and discoloration is rectified during regular maintenance visits.




Figure 9. Different water surface conditions at Chenik Creek remote video station. From left to right: overcast and still, ripples and reflections, blustery wind turbulence.

Further development and testing is necessary to extend the effectiveness of the approach to the full range of conditions that will be experienced at the remote video stations and decrease the reliance on empirically chosen factors to control the algorithms. The critical step in the video processing is the segmentation within the RoI, as this determines the success of the following steps. Different approaches, such as segmentation within the colour space and frame differencing from an averaged background image will be the subject of future research.

References

- Bue, B.G., S.M. Fried, S. Sharr, D.G. Sharp, J.A. Wilcock, and H.J. Geiger. 1998. Estimating salmon escapement using area-under-the-curve, aerial observer efficiency, and stream-life estimates: the Prince William Sound pink salmon example. *North Pacific Anadromous Fish Commission Bulletin* 1:240-250.
- Charalampidis, D., Gundam, M., Joginipelly, A., Quinteros, M., Ioup, G., Ioup, J., Yoerger, E. J., Thompson, C. H., 2012. Feature analysis for classification of fish in underwater video. Final Report, LA Board of Regents Contract NASA(2011)-STENNIS-02, December 2012. 19 pages.
- Derry, J. F. and Elliott, C. J. H. (1997). Automated 3-D tracking of video-captured movement using the example of an aquatic mollusk. *Behavior Research Methods, Instruments and Computers* 29(3): 353-357.
- Hu, W., Tan, T., Wang, L. and Maybank, S., 2004. A survey on visual surveillance of object motion and behaviors, *IEEE Transactions on Systems, Man, and Cybernetics*, 34(3):334-352.
- Jones III, E.L., S. Heintz, and K. Pahlke. 2007. Aerial Counts. Pages 399-410 in D.H. Johnson, B.M. Shrier, J.S. O'Neil, J.A. Knutzen, X. Augerot, T.A. O'Neil, and T.N. Pearsons, editors. *Salmonid Field Protocols Handbook: Techniques for assessing status and trends in salmon and trout populations*. American Fisheries Society, Bethesda, Maryland.
- Kalman, R. E., 1960. A new approach to linear filtering and prediction problems. *Transaction of the ASME Journal of Basic Engineering*, 35-45.
- Kocak, D., Vitoria-Lobo, N. and Widder, E., 1999. Computer Vision Techniques for Quantifying, Tracking, and Identifying Bioluminescent Plankton. *IEEE Journal of Oceanic Engineering* 24(1): 81-95.
- Maxwell, S.L. 2007. Hydroacoustics: Rivers. Pages 133-152 in D.H. Johnson, B.M. Shrier, J.S. O'Neil, J.A. Knutzen, X. Augerot, T.A. O'Neil, and T.N. Pearsons, editors. *Salmonid Field Protocols Handbook: Techniques for assessing status and trends in salmon and trout populations*. American Fisheries Society, Bethesda, Maryland.
- Moeslund, T. B., Hilton, A. and Krüger, V., 2006. A survey of advances in vision-based human motion capture and analysis. *Computer Vision and Image Understanding*, 104(2-3): 90-126.
- Otis, E.O. and M. Dickson. 2002. Improved salmon escapement enumeration using remote video and time-lapse recording technology, Exxon Valdez Restoration Final Report (Restoration Project 00366), Alaska Department of Fish and Game, Division of Commercial Fisheries, Homer, Alaska.
- Otsu, N., 1979. A threshold selection method from gray-level histograms. *IEEE Trans. Systems, Man, and Cybernetics* 9(1): 62-66.
- Rova, A., Mori, G., Dill, L. M., 2007. One fish, two fish, butterfly, trumpeter: Recognizing fish in underwater video. Proceedings, *IAPR Conference on Machine Vision Applications*, May 16-18, 2007, Tokyo, Japan. ISBN 978-4-901122-07-8 pp 404-407.
- Spampinato, C., Giordano, D., Di Salvo, R., Chen-Burger, J., Fisher, R., Nadarajan, G., 2010. Automatic fish classification for underwater species behaviour understanding. Proceedings of *First ACM International Workshop on Analysis and Retrieval of Tracked Events and Motion in Imagery Streams*. Firenze, Italy, ACM: 45-50.
- Tillett, R., McFarlane, N. and Lines, J. 2000. Estimating dimensions of free-swimming fish using 3D point distribution models. *Computer Vision and Image Understanding*, 79, 123-141.
- Walther, D., Edgington, D. R. and Koch, C. 2004. Detection and tracking of objects in underwater video. In *IEEE International Conference on Computer Vision and Pattern Recognition*, F.M. Titsworth (ed.). New York: Institute of Electrical and Electronics Engineers 1, 544-549.

Welch, G. and Bishop, G., 2004. An introduction to the Kalman Filter, In University of North Carolina at Chapel Hill, Department of Computer Science. Tech. Rep. 95-041, available at http://www.cs.unc.edu/~welch/media/pdf/kalman_intro.pdf

Woody, C.A. 2007. Tower counts. Pages 363-384 in D.H. Johnson, B.M. Shrier, J.S. O'Neil, J.A. Knutzen, X. Augerot, T.A. O'Neil, and T.N. Pearsons, editors. *Salmonid Field Protocols Handbook: Techniques for assessing status and trends in salmon and trout populations*. American Fisheries Society, Bethesda, Maryland.

Zimmerman, C.E. and L.M. Zabkar. 2007. Weirs. Pages 385-398 in D.H. Johnson, B.M. Shrier, J.S. O'Neil, J.A. Knutzen, X. Augerot, T.A. O'Neil, and T.N. Pearsons, editors. *Salmonid Field Protocols Handbook: Techniques for assessing status and trends in salmon and trout populations*. American Fisheries Society, Bethesda, Maryland.