Fisheries New Zealand
Tini a Tangaroa

Feasibility of automating otolith ageing using CT scanning and machine learning

New Zealand Fisheries Assessment Report 2019/58

B.R. Moore
J. Maclaren
C. Peat
M. Anjomrouz
P.L. Horn
S. Hoyle

ISSN 1179-5352 (online)
ISBN 978-1-99-000866-5 (online)

October 2019

New Zealand Government

Requests for further copies should be directed to:

Publications Logistics Officer
Ministry for Primary Industries
PO Box 2526
WELLINGTON 6140

Email: brand@mpi.govt.nz
Telephone: 0800 00 83 33
Facsimile: 04-894 0300

This publication is also available on the Ministry for Primary Industries websites at:
http://www.mpi.govt.nz/news-and-resources/publications
http://fs.fish.govt.nz go to Document library/Research reports

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Moore, B.R.; Maclaren, J.; Peat, C.; Anjomrouz, M.; Horn, P.L.; Hoyle, S. (2019). Feasibility of automating otolith ageing using CT scanning and machine learning.

***New Zealand Fisheries Assessment Report 2019/58.* 23 p.**

Knowledge of the age of fish is an integral part of fisheries science, being a key requirement for estimating growth, age at recruitment and sexual maturity, longevity, mortality rates, population age structure, and age-dependent fishing gear selectivity, all of which are important components of age-based stock assessments. The current method for determining ages of most fish species relies on manually extracting, preparing (embedding, sectioning), and reading otoliths. This process is expensive, time-consuming, and can be subject to biases such as variations in age estimations between readers and within readers over time. Recent advances in imaging technologies and machine learning suggest that automation of at least some aspects of these processes may be possible. This study examined the feasibility of automating otolith ageing using CT scanning and machine learning, through a review of published work in these areas and trials on New Zealand fish species.

The utility of CT scanning technologies for imaging otolith annular structure was trialled using the MARS Bioimaging Ltd X-ray scanning machine housed at the University of Canterbury. Trials were undertaken on three species: snapper (*Pagrus auratus*); hoki (*Macruronus novaezelandiae*); and ling (*Genypterus blacodes*). Ages of individuals used in the trial were estimated prior to scanning through standard ageing techniques. The MARS CT scanner was able to resolve banding patterns, with best results generated from sections taken parallel to the distal surface, similar to the process used when otoliths are read whole. Spectral analyses revealed varying concentrations of calcium and other minerals across the otoliths. Transverse sections through the otolith core showed banding potentially indicative of annual growth bands, but the resulting images generally lacked sufficient resolution and contrast to detect outermost bands for most individuals. The first growth band was also often difficult to identify. Accordingly, ageing trials using systems capable of achieving higher resolutions are warranted.

To investigate the feasibility of using machine learning to estimate age in New Zealand fish species, we adapted a pre-trained convolutional neural network (CNN) designed for object recognition to estimate age using otolith images obtained via microscopy for snapper and hoki. For each species, the model was trained on a collection of images of fish previously aged by human readers (n = 687 and 882 for snapper and hoki, respectively). After training, the model gave the same age as the human reader for 47% of snapper in a test dataset, with a further 35% of ages estimated within 1 year of the human reader estimate of age. For hoki, the model gave the same age as the human reader for 41% of individuals.

Our preliminary examination suggests there is significant potential for imaging of otoliths for ageing purposes using a variant of CT scanning technology, for automating age estimation from otolith images, and for the potential to combine both of these techniques to form a fully automated ageing system. The key future research directions to the development of such a system are:

1. optimise and further evaluate the CNN-based approach used here to automate age estimation of snapper and hoki;
2. investigate the potential for the CNN-based model to automate age estimation of additional species from photographic images of sectioned otoliths;
3. further identify the potential to use otolith scans from CT scanning technologies and, in particular micro-CT systems, for imaging otolith annular structure; and
4. identify the potential to automate ageing from images generated from micro-CT scanning using deep learning approaches.

1. INTRODUCTION

Determining an accurate estimate of age of fish is an integral part of fisheries science. Knowledge of the age of fish is used to estimate growth, age at recruitment and sexual maturity, longevity, mortality rates, population age structure, and age-dependent fishing gear selectivity, all of which are important components of age-based stock assessments (Walsh et al. 2014a). For most bony fishes, age is determined by enumerating periodically-accreted growth marks in calcified structures, including scales, bones, fin rays, and otoliths (Welch et al. 1993, Campana & Thorrold 2001, Zhu et al. 2015). Of these, otoliths have received most focus in ageing studies, particularly in recent years. Otoliths are structures located in the inner ear cavities of all teleost (bony) fishes that assist in sound detection and are used for balance and orientation (Campana & Neilson 1985). Otoliths are composed mainly of calcium carbonate ($CaCO_3$), mostly in the form of aragonite. Teleosts possess three pairs of otoliths (sagittae, lapilli, and asterisci), with the sagittae generally the largest and the focus of most ageing studies. Throughout this report, the use of 'otolith' will be synonymous with sagitta, unless otherwise specified.

Ageing fish on the basis of periodically-accreted growth marks (laid daily, seasonally, or annually) is achieved using two successive processes: preparation and reading. In the preparation phase of annual ageing, otoliths are prepared so that annual bands can be accurately counted. The simplest approach is to immerse the whole otolith in a clear liquid, such as water or alcohol solution, illuminate it from above, and view it against a dark background, routinely performed using a stereo microscope. This method is suitable only if the otoliths are relatively thin and all annual bands can be seen. In many species, however, as the growth rate of the fish slows down, the outer growth bands become increasingly compressed and difficult to read in the whole otolith. To correctly determine the number of annual bands in these cases, alternative approaches, including transverse sectioning, or the 'break and burn' and 'burn and embed' methods currently utilised in New Zealand for snapper, *Pagrus auratus* (Walsh et al. 2014a), are required. Although several laboratory-specific nuances to sectioning exist, the general premise is that one or multiple thin (typically 200–400 μm) cross sections are taken encompassing the primordium (core) of the otolith. Sections are then fixed onto a microscope slide, covered with a coverslip, and then read using a microscope. Reading is then performed by one or multiple readers, and images are often taken for documentation, preservation, and quality assurance purposes.

Preparing and reading otoliths can be an expensive and time-consuming task. Recent estimates suggest New Zealand processes between 30 000–40 000 otoliths annually, at a cost of over $1,000,000 per year on preparation and reading. Moreover, otolith-based age interpretation is inherently subjective and uncertain. Individual readers may interpret the same otolith differently, and readers can also change their interpretations through time. Otolith collections often cover multi-decadal time periods, and differences between and within readers can result in long-term changes in interpretation, which can introduce bias into stock assessments. Improving the consistency of ageing through automated age determination would increase the replicability of age determination and the reliability of management advice.

Recent developments in imaging technology and machine learning may make it possible to improve the efficiency of many aspects of ageing, with potential to reduce both biases and long-term costs. For example, a range of alternative approaches to hard part (i.e., otoliths in fish, vertebrae in elasmobranchs) imaging have recently been undertaken to age commercially-exploited fishes and elasmobranchs, including X-ray imaging (Liu et al. 1999) and micro-computed topography (micro-CT) (Parsons et al. 2018). CT scanning technologies potentially offer a powerful approach to otolith imaging for ageing purposes. CT scanning uses X-ray technology to produce image slices through objects, which can be reconstructed into virtual, three-dimensional (3D) images that can be rotated and viewed in any orientation (Geraghty et al. 2012).

Automated image analysis using machine learning has the potential to provide more reliable, consistent, and rapid age estimates. Such methods have been explored for many years (e.g., Robertson & Morison 1999, Takashima et al. 2000, Fablet & Le Josse 2005, Fablet 2006); however, they have generally been found to be less precise than those obtained from experienced otolith readers (e.g., Robertson &

Morison 1999, Fablet & Le Josse 2005). In the last 2–3 years, there has been considerable progress in machine learning due to improved algorithms, greater computing power, and wider availability of labelled, digital training data. Many new applications have become possible, including automating otolith age estimation (e.g., Moen et al. 2018).

The current project aimed to explore the feasibility of using CT scanning and machine learning technologies to automate aspects of otolith ageing. The objectives of this project were as follows:

1. identify the potential to use otolith scans from CT-scanning technologies, such as the MARS CT scanner, for imaging otolith annular structure;
2. identify the potential for machine learning techniques to automatically age fish using available photographic images; and
3. develop a proposal that fully scopes and costs the further development of these methods.

For Objectives 1 and 2, we begin by presenting a review of the published literature on each technique, to assess whether such approaches have been conducted elsewhere, and, if so, whether lessons can be learned from those studies. We then follow this review with trials on New Zealand fish species.

2. REVIEW OF PREVIOUS STUDIES

2.1 Use of CT technology to examine the skeletal or non-skeletal properties of fishes

A number of studies have used CT technology to examine the skeletal or non-skeletal (e.g., soft tissue) properties of fishes, with direct implications for the utility of imaging otoliths. Metscher (2009) demonstrates the broad applicability of contrast stains for imaging various animal tissues using X-ray micro-CT approaches. Although the focus was on soft tissues, otoliths are clearly visible in scanned images of paddlefish (*Polyodon spathula*) and pike (*Esox lucius*) hatchlings (figures 2 & 5 in Metscher 2009). Otoliths *in situ* are also visible in the recent, high-resolution X-ray, CT-based re-description of the tuvirão, *Gymnotus inaequilabiatus* (Maxime & Albert 2014), and in a range of other studies employing CT imagery (e.g., Bignami et al. 2013, Edds-Walton et al. 2015, Felix et al. 2016, Fisher & Hunter 2018, Schulz-Mirbach et al. 2018). Although no attempt was made to age otoliths in any of these studies, these results demonstrate that *in situ* observation of otoliths using CT scans is possible.

Long & Snow (2018) used CT scanning in conjunction with oxytetracycline marking to examine the formation of otoliths in spotted gar, *Lepisosteus oculatus*. The three-dimensional rendered images generated were of sufficient resolution to reveal that otoliths (sagittae and lapilli) were not formed at hatching, but rather formed as loose collections of otoconia at 1-day post-hatch (dph), fusing and hardening into single crystals by 4 dph (Long & Snow 2018).

To our knowledge, there have been no published studies on the utility of imaging fish otoliths for annual age estimation using CT scanning technology. However, some work has been done on elasmobranchs. In one of the first applications, Geraghty et al. (2012) assessed the use of micro-CT as a valid and repeatable alternative approach for age determination in spinner shark, *Carcharhinus brevipinna*. They found that the quality and resolution of micro-CT outputs were sufficiently high such that growth bands were visible for each of four image types (whole vertebrae, radiograph, half vertebrae, and sagittal section), with comparable clarity in growth bands between methods. Moreover, reads from micro-CT-generated sections provided comparable and repeatable age counts relative to those from manually-produced sections across a wide age range (2–19 years), with no systematic bias in age counts between methods (Geraghty et al. 2012).

High-resolution X-ray CT was used by Parsons et al. (2018) to provide images of vertebrae to estimate age in spiny butterfly ray, *Gymnura altavela*, from the north Atlantic. Although annuli counts from reconstructed images were not tested against those from sectioned vertebrae, CT images revealed interpretable growth bands that were read with high precision between and within readers.

Francis et al. (2018a, 2018b) trialled the use of micro-CT scanning technology for imaging vertebrae to

estimate age in common electric ray (*Tetronarce nobiliana*), blind electric ray (*Typhlonarke aysoni*), carpet shark (*Cephaloscyllium isabellum*), and four species of deepwater shark, namely seal shark (*Dalatias licha*), Owston's dogfish (*Centroscymnus owstonii*), longnose velvet dogfish (*Centroselachus crepidater*), and Plunket's shark (*Scymnodon plunketi*). Growth bands were not visible in most specimens of deepwater shark, although some seal shark specimens exhibited lines in the corpus calcareum that may represent growth bands, as well as wavy lines of unknown significance in the intermedialia, and indistinct bands were sometimes visible on the central cone of Plunket's sharks (Francis et al. 2018b). Based on the observation that growth bands were evident in most specimens of carpet shark, common electric ray, and blind electric ray, these authors concluded that micro-CT imaging has utility in revealing the structure of hard parts and the growth bands, particularly as a comparative tool and aid for band interpretation (Francis et al. 2018a).

2.2 Machine-learning approaches to estimating fish age from otolith images

Machine-learning approaches to estimate fish age from otolith images have been trialled since the mid-1990s, beginning with the development of artificial neural networks (ANN). A neural network consists of layers of simple computational units (neurons), arranged so that the output of the units in one layer feeds into the inputs of the next layer's units. Each unit calculates a weighted sum of its inputs and applies a function that introduces non-linearity into the system (Malde et al. 2019). The weights of the inputs of each unit constitute the parameters to be learned. In artificial neural networks this is usually achieved using back propagation (Rumelhart et al. 1986) to calculate the gradient for a cost function, which is then minimised iteratively using some variant of gradient descent (Malde et al. 2019).

In one of the first applications of artificial neural networks to estimate fish age, Robertson & Morison (1999) tested ANN models as a means of objectively replicating the age estimates of an experienced human reader in three species of fish, including two species of sparid (black bream, *Acathopagrus butcheri*, and snapper) and one merlucciid (hoki, *Macruronus novaezelandiae*). Age estimates generated from the ANN structure correctly classified the age of fish for black bream and snapper at an accuracy level approaching that of an experienced reader, although there was evidence of underestimation of ages of the oldest fish, and their results were potentially biased by a lack of old fish (maximum ages examined were 9 years and 12 years for black bream and snapper, respectively). In contrast, error rates for hoki were high, with the model failing to classify more than 5% of fish correctly (Robertson & Morison 1999). Later, Robertson & Morison (2001) found that precision of age predictions of otoliths using neural networks from geometric features could be improved by using biological features, but the results obtained were still less precise than those obtained from experienced readers.

Multiple statistical learning techniques, including both neural networks and support vector machines, were used by Fablet & Le Josse (2005) to automate age estimation of plaice, *Pleuronectes platessa*). They observed correct classification rates that compared favourably with inter-reader agreement rates, although it should be noted that only 5 age classes (1 to +5) were used in their study.

Work on neural networks in the 1980s and 1990s was limited by low computational power, lack of large data sets for training, and limitations in machine-learning algorithms (Malde et al. 2019). In recent years, increased parallel processing abilities (e.g., by using GPUs), lowered costs of computing hardware, distributed computing, and advances in machine-learning algorithms have led to the construction of much larger and deeper neural networks than before, resulting in the emergence of *deep learning* approaches, where deep neural networks learn high-level abstractions in data by utilising hierarchical architectures (Schmidhuber 2015, Guo et al. 2016). Deep neural networks have been shown to outperform more conventional methods across a range of problems (Krizhevsky et al. 2012), increasing their applicability to complex tasks.

Recent studies using deep learning approaches have demonstrated significant utility for automating fish image analysis (e.g., Allken et al. 2019), including otolith images for age estimation. To our knowledge,

Moen et al. (2019) present the only study published to date to use deep learning of otolith images to estimate fish age in their exploration of the potential for a CNN to reliably estimate age from images of whole otoliths of Greenland halibut, *Reinhardtius hippoglossoides*. They observed precision in age estimations comparable with that obtained by human experts, with the age of 29% of otoliths estimated correctly, and another 38% estimated within one year of the read age. There was, however, a general tendency for the model to underestimate the age of old (more than 15 years) fish (Moen et al. 2019), potentially resulting from the relatively low abundances of these individuals in the sampled population, and thus the training dataset.

3. METHODS

3.1 Otolith imaging via CT scanning

The utility of CT scanning technologies for imaging otolith annular structure was performed using the MARS Bioimaging Ltd X-ray scanning machine housed at the University of Canterbury (see www.marsbioimaging.com). The system uses a photon-counting detector (Medipix3RX) to produce spectral image slices through objects, which can be reconstructed into three-dimensional images that can be rotated and viewed in any orientation.

Trials were undertaken on three species: snapper, hoki, and ling (*Genypterus blacodes*). Ages of all individuals used in the trial were estimated prior to scanning through standard ageing techniques (e.g., Walsh et al. 2014a, Horn & Sutton 2017).

Otoliths were adhered to the sample holder using a temporary adhesive to ensure they did not move during data acquisition (Figure 1). The scan protocol was optimised by varying the magnification, filtration, and the energy set-up. The final runs used a 120 kVp spectrum and a 1.96 mm aluminium filter to minimise beam hardening, while allowing the detector to capture the low energy photons. Four energy thresholds operating in the same mode along with the arbitration counter were set across the spectrum to ensure there are sufficient counts at each energy bin. This is particularly important in the first energy bin where the capture of low energy photons aids identification of calcium-like material, facilitating identification of areas of fast and slow growth (and therefore, potentially, growth bands). The scan protocol was calibrated using a set of calcium solutions, water, and lipid to capture a true signal in the calcium channel for each otolith sample. All optimisation tests were conducted using the otoliths from young (estimated 2–3 year-old) snapper.

The scanning time for each sample varied from approximately 40 mins to around one hour depending on the length of sample examined, with the first 8–10 minutes allocated to warming up the X-ray tube.

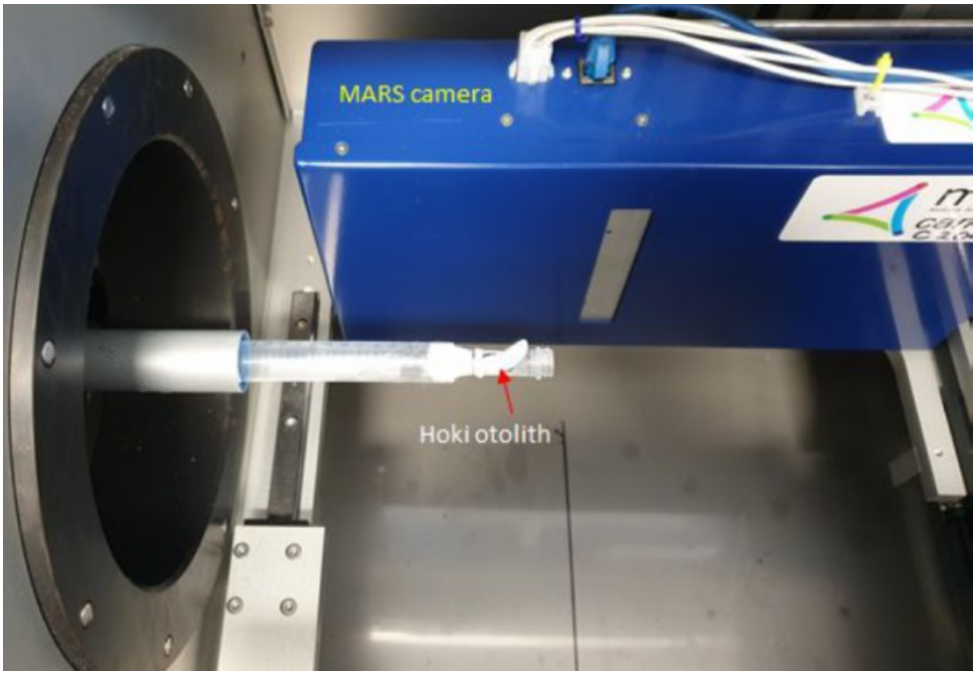

Figure 1: Hoki otolith on a small sample bed in the MARS scanner (credit: MARS Bioimaging Ltd).

Output data from the MARS system is in DICOM format. This includes the raw data (representing photon counts), three-dimensional attenuation maps (linear attenuation cm^{-1}), and coloured, three-dimensional material maps (density mg/mL). Data files were post-processed with MARS Vision software and Image J (Abramoff et al. 2004), using associated plugin routines that provided both serial sections and video animations for further analysis. MARS Vision software enables three-dimensional and two-dimensional visualisation of the specimen and facilitates extraction of virtual sections at chosen orientations through the specimen using digital clipping planes. The images presented herein are mostly extracted from the default mode of MARS Vision.

3.2 Age estimation via machine learning

3.2.1. Estimating age of snapper and hoki

We explored the utility of machine learning using CNNs to estimate fish age using two test species: snapper and hoki. All snapper and hoki otoliths used were previously aged by expert readers following the methods of Walsh et al. (2014a) and Horn & Sutton (2017), respectively (= observed age), and an image was taken of each otolith section using a dissecting microscope illuminated with transmitted light (Figure 2). The growth increments on the otoliths of these species vary in clarity, from quite clear in snapper to moderately difficult in hoki (Walsh et al. 2014a, Horn & Sutton 2017).

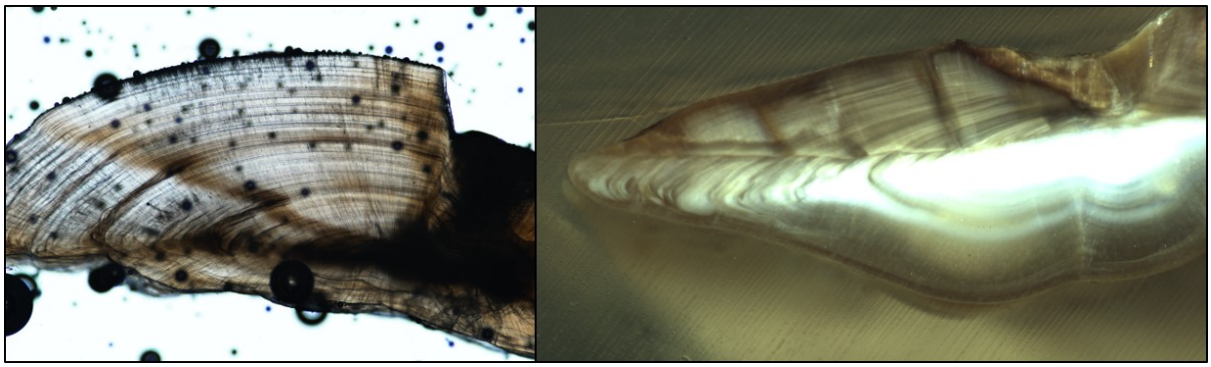

Figure 2: Examples of snapper and hoki otolith images used in the machine-learning trial.

We chose the TensorFlow (https://www.tensorflow.org/) and Keras (https://keras.io/) libraries to implement and train our models. TensorFlow is currently the largest and most popular library available for deep learning. Keras is a high-level API which runs on top of TensorFlow and simplifies implementation of TensorFlow models. We used a transfer-learning technique to develop a CNN for otolith age estimation. Transfer learning is the process of repurposing a machine learning model that has been pre-trained for another, related, task. Specifically, we started with the Inception V3 model from Google, pre-trained on the ImageNet database (http://www.image-net.org/). We removed the final classification layer, leaving what is effectively a spatial feature extractor, and attached a new regression layer. This design is shown in Figure 3 below. At this point, the neural network is trained to minimise the mean squared error (MSE) between predicted ages and human expert age estimates, using the otolith images as inputs. Various training metaparameters contribute substantially to final accuracy by using a stochastic gradient descent (SGD) optimiser and by leaving all network layers as trainable.

The CNN designed above was trained using 687 snapper otolith images and evaluated. This process was then repeated using 882 hoki otolith images.

Figure 3: Neural network architecture with the InceptionV3 convolutional layers connected to a neural network regressor.

We applied standard image augmentation approaches to the training dataset. Augmentation is an important technique for training deep learning CNNs on limited datasets. This process applies a set of random transformations that preserve class, whilst artificially inflating the training data set size. In doing so, it is unlikely that the classifier will encounter the exact same input twice and is thus less likely to overfit the data. Transformations applied included:

- rotations in the range -30 to 30 degrees;
- flipping images vertically and horizontally; and
- contrast adjustments by randomly scaling the RGB channels (Figure 4).

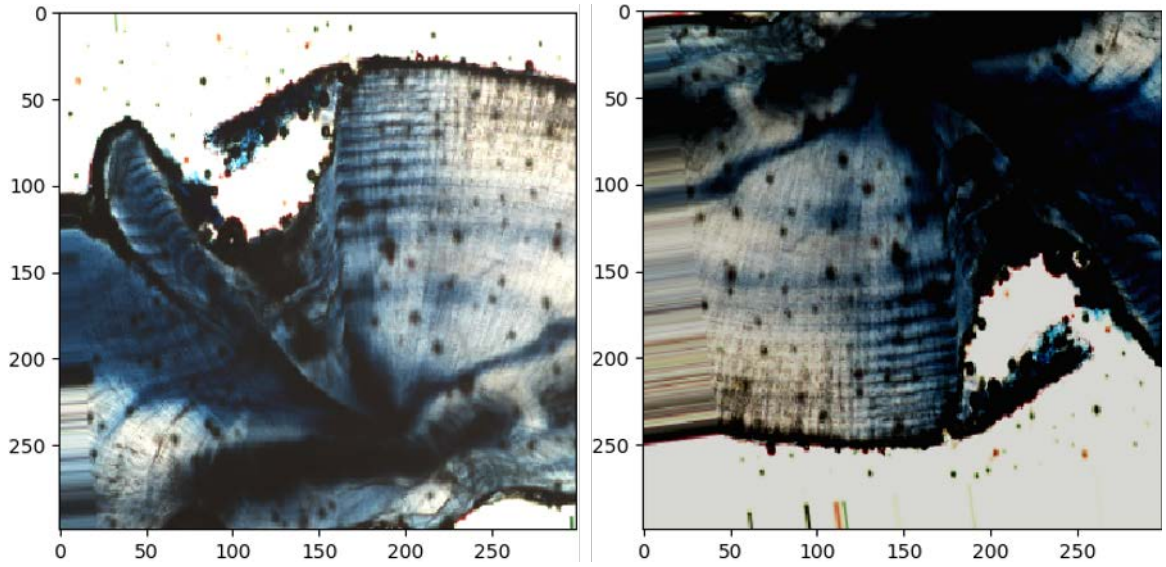

Figure 4: Two examples of random image augmentation on an image of a snapper otolith.

The training of the neural network for snapper otolith age prediction was done over 10 hours using a GTX 1080 Ti Nvidia graphics card. This model of graphics card is benchmarked to perform 10 TFLOPS on matrix multiplication workloads such as neural network training. In practice, performance is approximately 100x greater than a traditional central processing unit (CPU) making a graphics processing unit (GPU) a necessary part of training neural networks. A similar approach was undertaken for hoki, although due to time constraints the neural network for this species was trained for considerably less time (50 iterations, equating to approximately 50 minutes). The CNN was then evaluated with test datasets of 77 snapper otolith images and 99 hoki otolith images. Performance of the CNN to correctly assign ages was assessed via the mean squared error (MSE), root mean squared error (RMSE) and mean absolute error (MAE).

3.2.2. Image feature analysis

Experiments were performed to investigate which features in the otolith images are important to correctly predict the otolith age. In particular, we tested to what degree different aspects of image processing affected the resulting age estimation. Features trialled included the following:

- jpg exportation, to determine whether the artefacts introduced when the image is exported as jpg will impact the prediction;
- greyscale vs. colour, to determine whether colour information had an effect on the prediction;
- rescaling, to determine whether lower resolution images could be used; and
- masking, whereby various parts of the image are masked to determine the important parts of the image to the network's age estimation.

For these experiments, a snapper otolith image with a recorded age of 7 years was used. The predicted age of the unaltered otolith image was 7.03 years. It should be noted that these feature analysis experiments were done using one image and one neural network due to time constraints. Accordingly, these experiments are more indicative of potential future research, rather than proven results.

4. RESULTS

4.1 Otolith imaging via CT scanning

CT scans of otoliths from snapper, hoki, and ling are shown in Figures 5–11. Slices viewed from the distal surface (i.e., sulcus side down) sometimes exhibited growth bands, although the first growth band and outermost growth bands were often difficult to identify (Figures 5 & 6). Layers in the calcium channel were evident in most otoliths, but their presence was highly variable across different axes of the otolith (Figures 7–9) and are unlikely to represent annual growth bands.

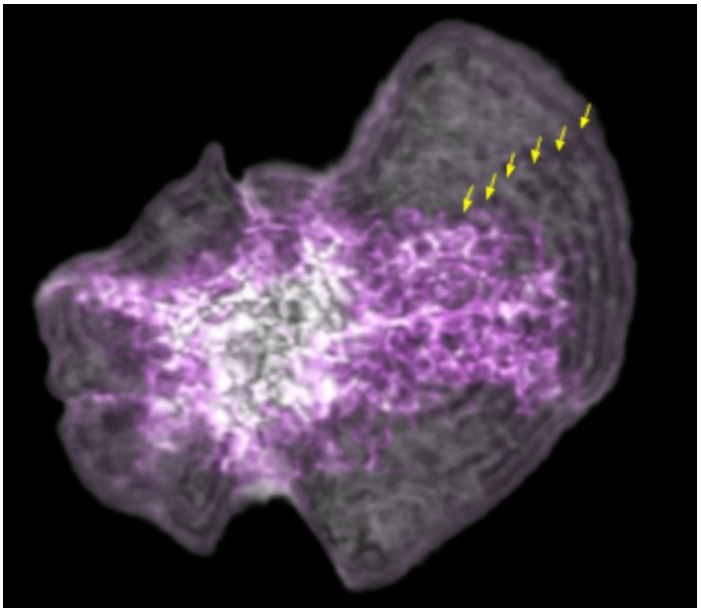

Figure 5: A slice parallel to the distal surface of a snapper otolith from a 47 cm fork length specimen with an estimated age of 7 years. Growth bands (estimated 6–7) evident through the CT scan are indicated by yellow arrows. Colours reflect different concentrations of calcium (credit: MARS Bioimaging Ltd).

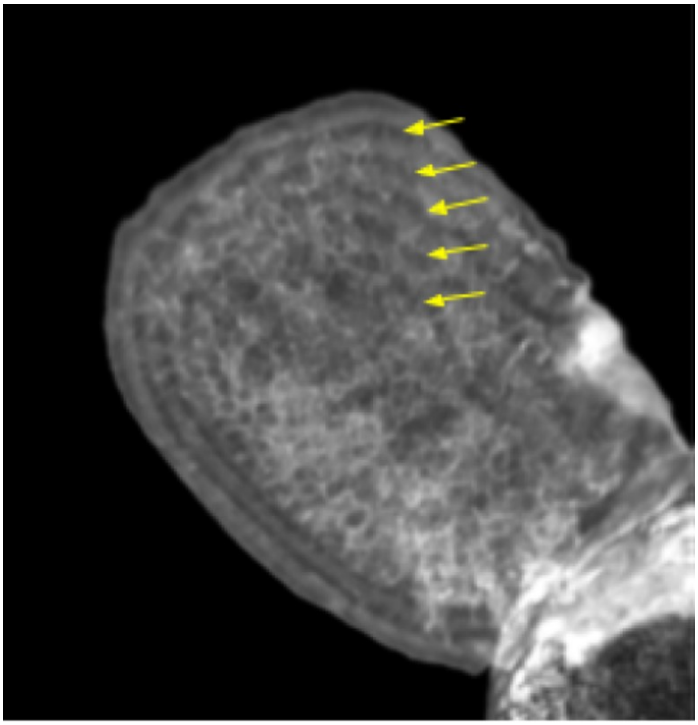

Figure 6: A slice parallel to the distal surface of a ling otolith from a 66 cm total length specimen with an estimated age of 5 years. Growth bands (5) evident through the CT scan are indicated by yellow arrows. The adhesive is evident in the bottom right-hand corner (credit: MARS Bioimaging Ltd).

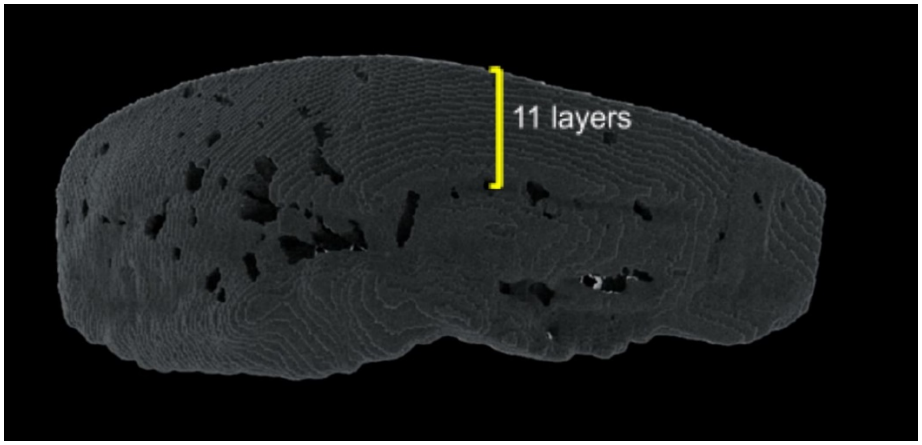

Figure 7: Image of a hoki otolith from a 93 cm total length specimen (estimated age = 11 years) in advanced 3D rendering mode showing layers in the calcium channel (credit: MARS Bioimaging Ltd).

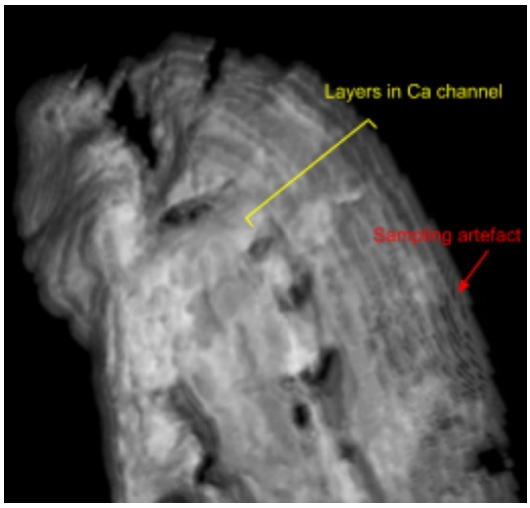

Figure 8: Image of a hoki otolith from a 93 cm total length specimen (estimated age = 11 years) in magnified view in simple rendering mode layers in the calcium channel (credit: MARS Bioimaging Ltd).

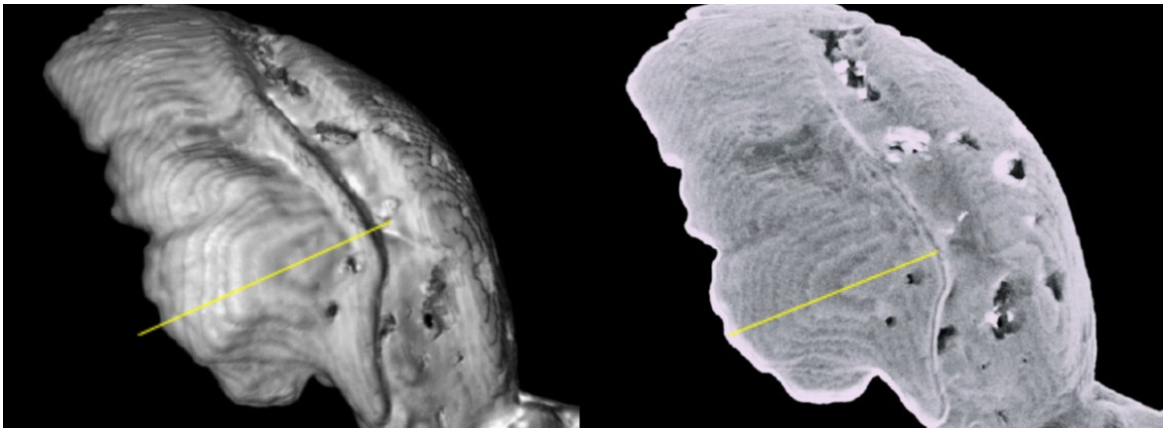

Figure 9: Image of a snapper otolith from a 44 cm fork length specimen (estimated age = 7 years) displaying layers in the calcium channel in simple rendering (left) and advanced rendering (right) modes (credit: MARS Bioimaging Ltd).

Transverse sections through the otolith core showed banding potentially indicative of annual growth bands; however, resulting images generally lacked sufficient resolution and contrast to detect outermost

bands for most individuals (Figures 10 & 11). Spectral analyses revealed different concentrations of calcium within the otoliths, but were generally of insufficient resolution to differentiate fast growth areas from slow growth areas (and thus detect annual growth bands), even in very young fish (Figures 12 & 13).

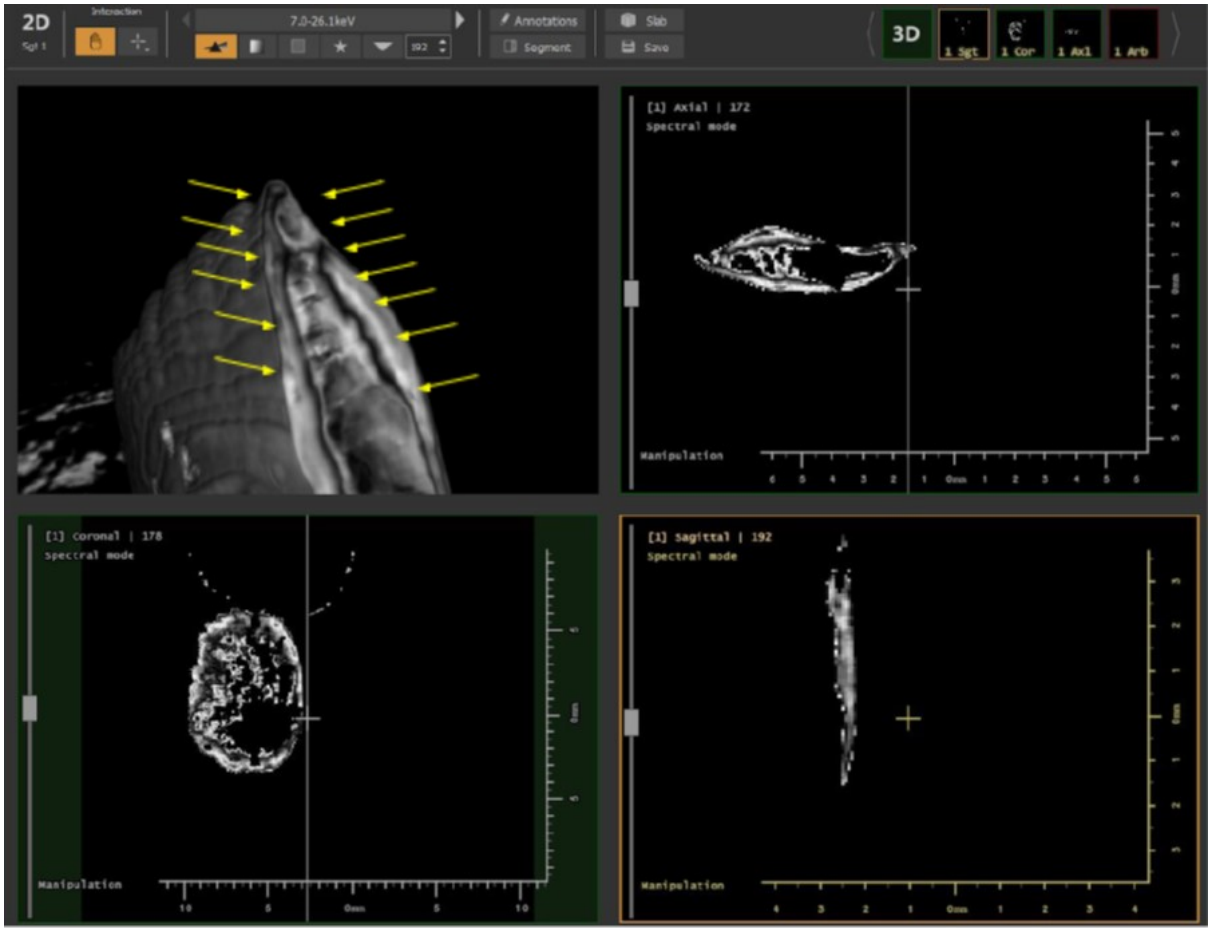

Figure 10: Image of a ling otolith from a 74 cm total length specimen (estimated age = 7 years) in different views of MARS Vision software. Top left: calcium layers formed through the otolith (yellow arrows). Top right and bottom left images show the axial and coronal views (credit: MARS Bioimaging Ltd).

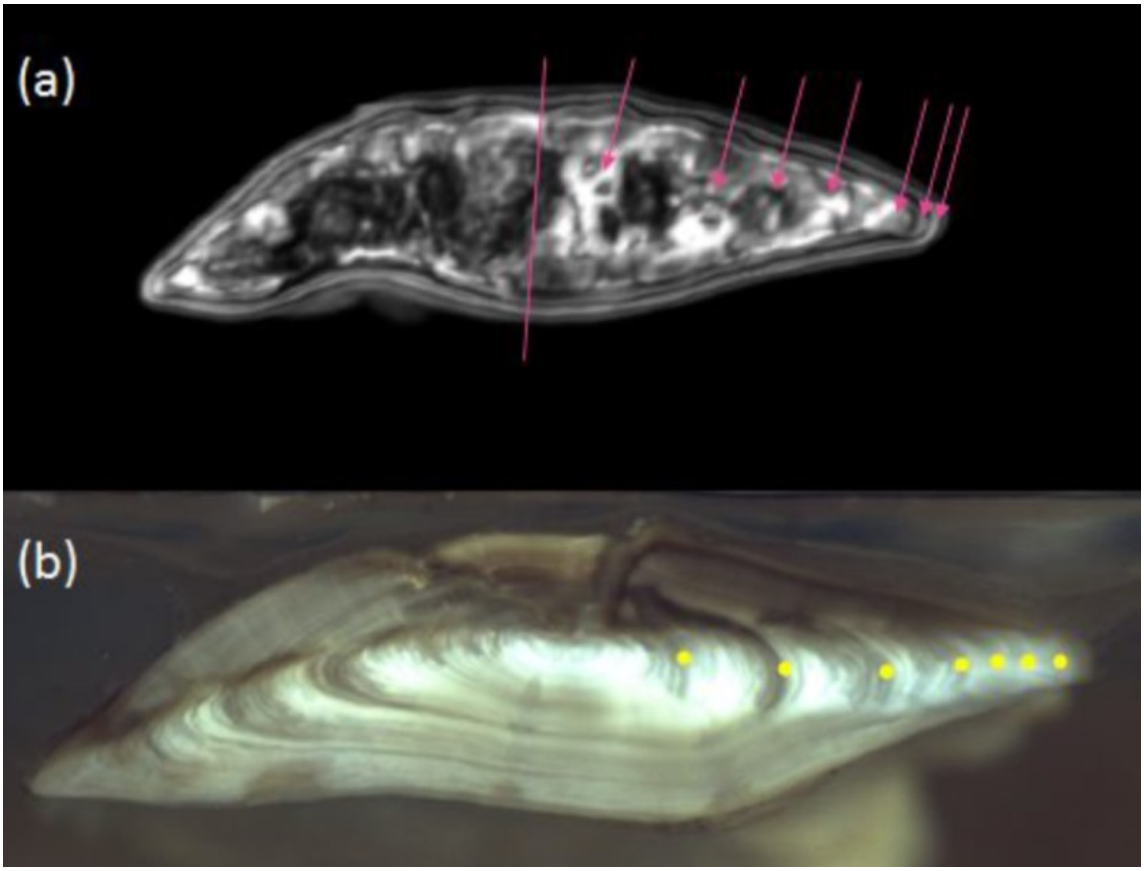

Figure 11: Images of a transverse section of a ling otolith from a 74 cm total length specimen (estimated age = 7 years) in MARS Vision software (top) and under a dissecting microscope (bottom). The expected location of growth bands in the CT-derived image are indicated by pink arrows (credits: a) MARS Bioimaging Ltd, b) Peter Horn).

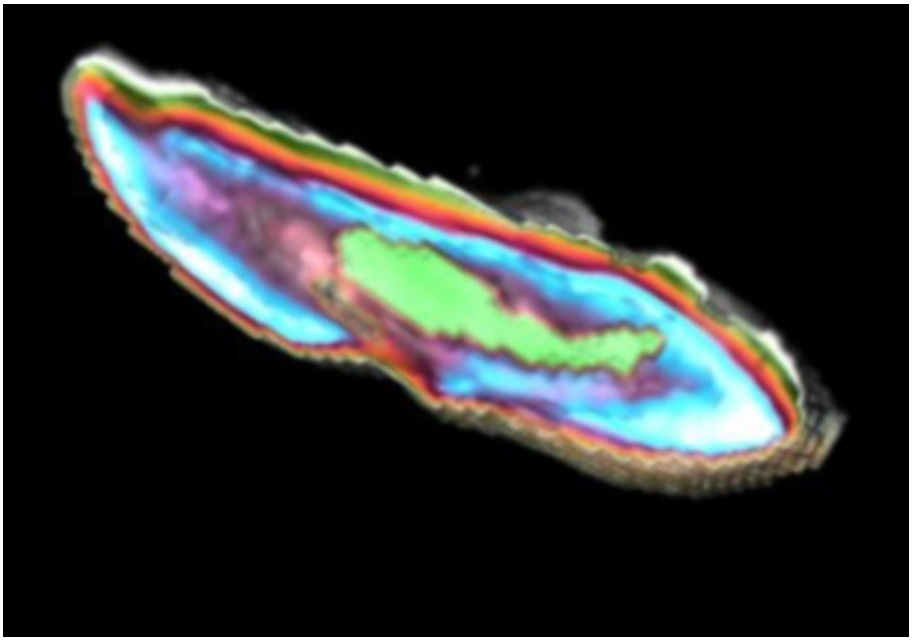

Figure 12: A transverse view sliced through the presumed otolith core of a snapper otolith (estimated age = 2–3 years old). Colours reflect different concentrations of calcium (credit: MARS Bioimaging Ltd).

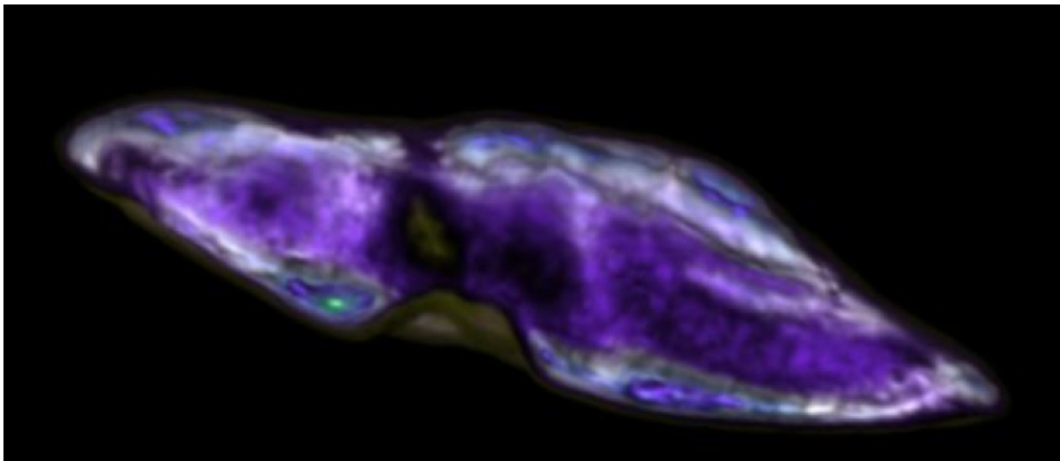

Figure 13: A transverse view sliced through the presumed otolith core of a hoki otolith from a 93 cm total length specimen (estimated age = 11 years). Colours reflect different concentrations of calcium (credit: MARS Bioimaging Ltd).

4.2 Age estimation via machine learning

4.2.1. Estimating age of snapper and hoki

The distributions of snapper predicted ages for both the test and training datasets generally approximated that of the observed ages (Figure 14, Table 1). In the test dataset, the ages of approximately 47% of snapper otoliths were estimated correctly using the CNN, and a further 35% of age estimates were within 1 year of the observed age. Age bias plots revealed the CNN tended to over-estimate the age of young snapper and underestimate the age of older snapper (Figure 14).

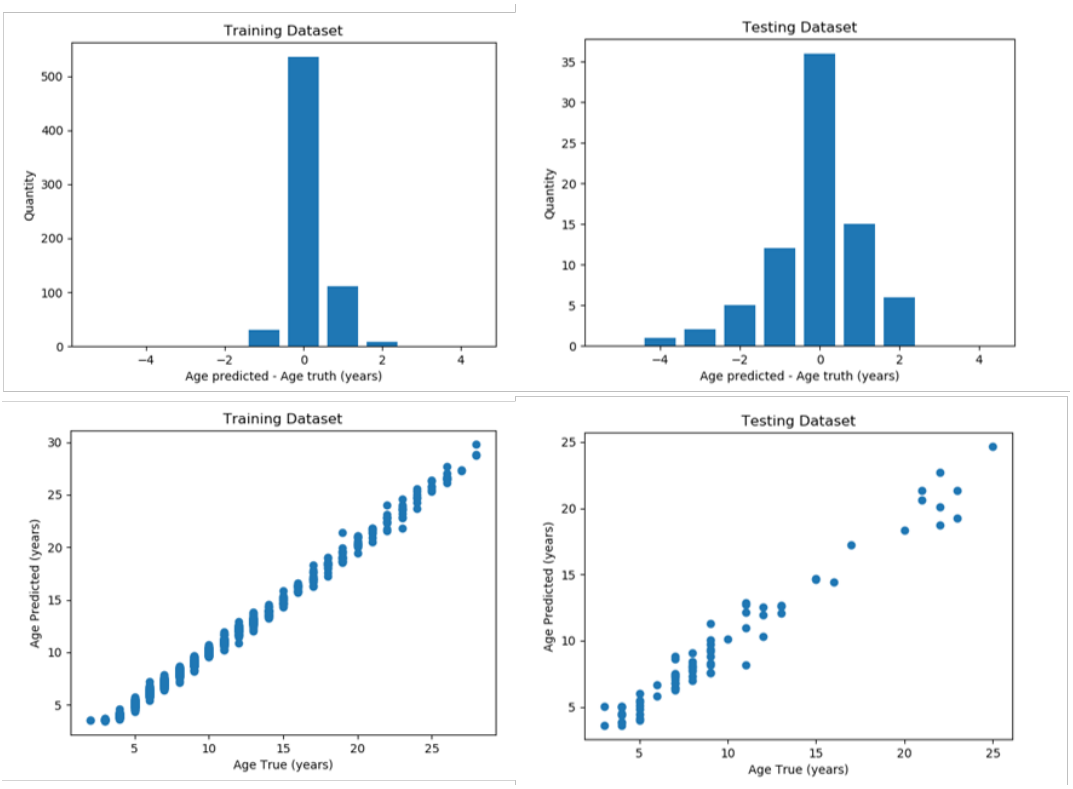

Figure 14: Histogram of differences in annuli counts (top) and age bias plots (bottom) between annuli estimations via machine learning and sectioned image interpretation for training (left) and testing (right) datasets for snapper.

Table 1: Neural network performance metrics for snapper, where MSE is mean squared error, RMSE is root mean squared error, and MAE is mean absolute error.

Metric	Training Dataset	Testing Dataset
MSE	0.197	1.203
RMSE	0.444	1.097
MAE	0.329	0.799
R^2	0.994	0.960
Accuracy	536/687 (78.0%)	36/77 (46.7%)

The training data accuracy for hoki was lower than snapper (40.6% *cf* 78.0%; see Table 1, Table 2). The difference in training data accuracy was likely due to the hoki neural network being trained for significantly less time (50 iterations, equating to about 50 minutes), and that growth bands in hoki were typically more diffuse than those in snapper training data (Horn & Sutton 2017). However, the accuracy of age estimations in the testing data were similar for snapper and hoki, with accuracy slightly less than 50% (see Tables 1 & 2). This likely results from the smaller sample sizes in the training dataset and the wider age range of snapper in both datasets, relative to hoki.

Table 2: Neural network performance metrics for hoki, where MSE is mean squared error, RMSE is root mean squared error, and MAE is mean absolute error.

Metric	Training Dataset	Testing Dataset
MSE	0.739	1.310
RMSE	0.875	1.144
MAE	0.935	0.850
R^2	0.923	0.84
Accuracy	358/882 (40.6%)	41/99 (41.4%)

4.2.2. Image feature analysis

Analysis of various aspects of image processing revealed contrasting results. Jpeg exportation had little effect on the predicted age of the single snapper otolith tested (7.02 years; Figure 15a). Changing the image from colour to greyscale resulted in an overestimation of age (9.16 years; Figure 15b), likely due to the neural network not being exposed to greyscale images during the training process.

Resizing the image revealed that some level of interpretation was possible down to 64x64 (6.33–7.59 years; Figure 16), breaking down at 32x32 (9.66 years; Figure 16d).

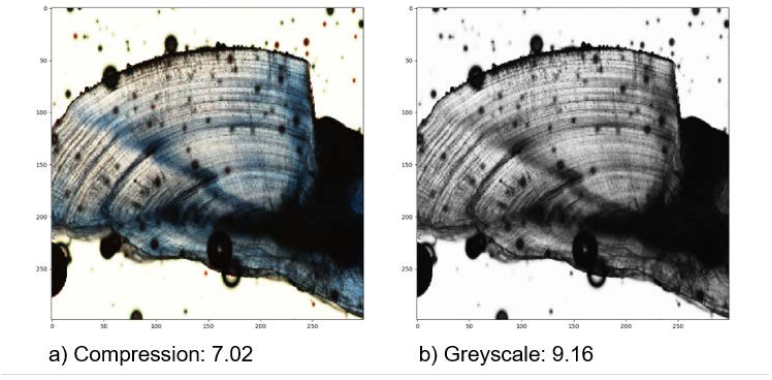

Figure 15: Image of a 7-year old snapper otolith after jpeg compression (a) and conversion to greyscale compression (b), with the predicted age from the neural network.

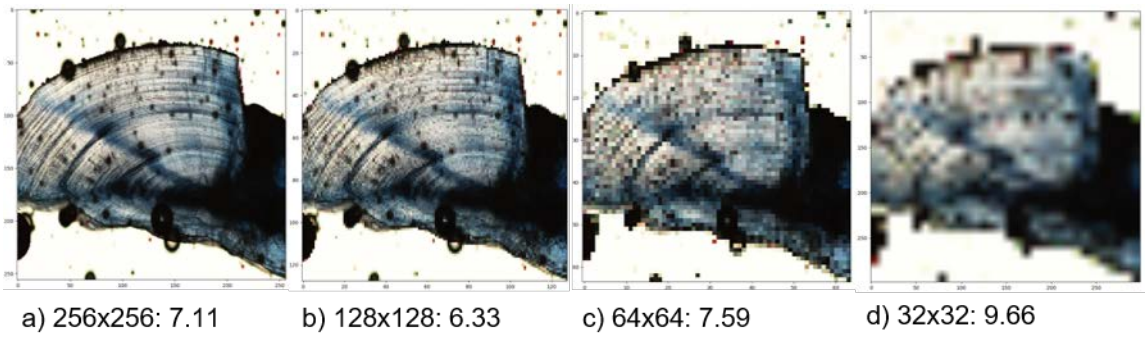

Figure 16: Image of a 7-year old snapper otolith after resizing, with the predicted age from the neural network.

Masking experiments revealed the model performed best when the image included as much information as possible (Figure 17). The critical requirement was for the image to include the axis from the primordium to the otolith margin on the ventral rim of the sulcus acusticus (Figures 17a & 17d).

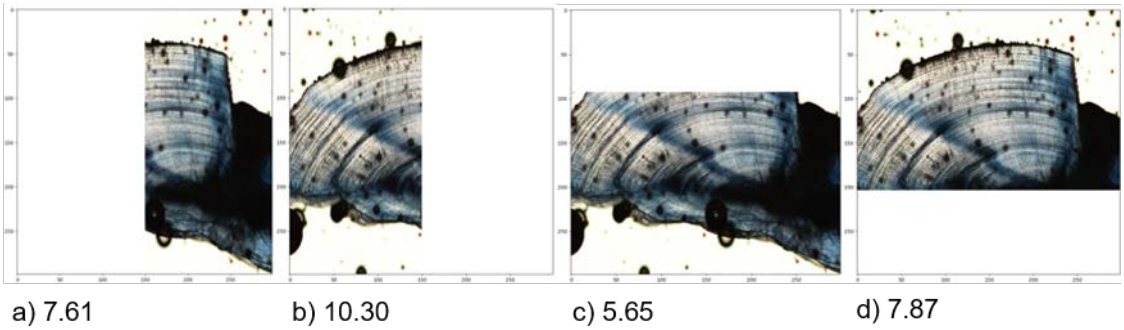

Figure 17: Image of a 7-year old snapper otolith after masking out sections, with the predicted age from the neural network.

5. DISCUSSION

5.1 Otolith imaging via CT scanning

Our trial using snapper, hoki, and ling otoliths revealed that the MARS CT scanner was able to detect and image otolith annular structure, although the resolution of the resulting images from this particular scanner was generally insufficient for ageing purposes across age classes. Best results were generated from sections taken parallel to the distal surface, similar to the process used when otoliths are read whole. In older samples, the outer growth bands become increasingly compressed and difficult to interpret. For example, successive paired growth bands (i.e., fast and slow growth bands) in snapper, considered to be a relatively 'easy' species to age because of the occurrence of distinct patterns of fast and slow growth bands in otoliths, are typically separated by less than 50 μm in 10+ year old individuals and can be as small as 20 μm in older individuals (NIWA, unpublished data). The nominal resolution of the MARS system used is 70 μm (M. Anjomrouz, pers. comm.)

Resolution issues may be resolved by using a micro-CT scanner to generate otolith images. Though not previously trialled for ageing fish otoliths, micro-CT scanners have shown potential for ageing elasmobranchs (Francis et al. 2018a), with comparable estimates of age to those by human readers (Geraghty et al. 2012, Parsons et al. 2018). Resolutions of available micro-CT scanners are as high as 0.35 μm (e.g., Bruker SkyScan 1272 desktop MicroCT held at the Auckland Bioengineering Institute, University of Auckland, www.auckland.ac.nz/en/abi.html), compared with medical CT scanners, which, at the time of writing, have best resolutions of around 50–70 μm. It should be noted that some

medical micro-CT scanners optimised for scanning small live animals are available (du Plessis et al. 2017).

For CT scanning to augment or replace human otolith preparation, the resulting images must be of appropriate resolution and quality to interpret growth bands at an equal or lower cost than current preparation techniques. Assuming image resolution issues can be resolved via micro-CT technology or further development of the current MARS system, the greatest limitation of this method is cost. In the current study, use of the CT scanner cost $1,500 per day (excl. GST), with optimised scan times for an individual otolith varying from approximately 40 mins to around one hour depending on the length of sample examined. At the time of writing, costs for the University of Auckland micro-CT scanners were somewhat lower at $85 per hour for instrument and operator, with image reconstruction and analysis additional to this. Further work is required to assess the scan time of micro-CT systems, whether efficiencies can be achieved such as processing multiple samples at the same time and/or reduced scan times and, if so, the cost-benefit implications.

Assuming comparable age estimates between CT-derived otolith images and human-generated otolith images can be obtained, CT scanning has several distinct advantages over manual sectioning and reading of otoliths. First, it is a non-destructive technique, allowing for unlimited multiple virtual sectioning from unlimited angles and perspectives. Moreover, otoliths are preserved whole and, therefore, are available for comparative studies as new methods are developed with advances in technology (Parsons et al. 2018) or are available for alternative applications (e.g., paired elemental and isotopic studies to discern movements and stock structure) to be conducted using the same specimens. In addition, variables inherent to manual processing, such as the width of sections or the location where they are taken, are eliminated because the digital sectioning of the virtual vertebrae can be precisely specified at the desired width or location. Finally, the low intrinsic X-ray contrast of non-mineralised tissues (Metscher 2009) means that otolith samples can be scanned in an uncleaned state, and potentially *in situ*, without affecting the quality or resolution of the output, substantially reducing sample processing time for embedding and sectioning. It is important, however, that researchers establish a standardised protocol to maintain consistency of angle and perspective for all otoliths when ageing individuals of a particular species, as with manual sectioning (Geraghty et al. 2012; Parsons et al. 2018). There is also a need to develop methods to optimise the speed and efficiency of CT scanning.

5.2 Age estimation via machine learning

Our preliminary trials on snapper and hoki, despite being limited by small sample sizes of the training datasets, revealed strong potential for using deep learning CNN-based approaches to predict ages of fish from otolith images. For snapper, and to a lesser extent hoki, the Inception-3 network estimated ages with relatively strong precision that were close to human expert age estimations despite minimal image or model optimisation, strongly supporting the potential of the method.

This study used fish aged up to 27 years for snapper and 18 years for hoki. Both species can live longer, with maximum ages in New Zealand waters estimated to be 67 years for snapper (Walsh 2008) and 25 years for hoki (Horn & Sutton 2017), although such old fish are rare in catches. The ability of the CNN to successfully estimate the age of older fish remains largely untested (and it would likely fail in its current form, due to the lack of any such images in the training dataset). The performance of the CNN in the test sample for snapper suggests a small but consistent downward bias for the oldest age classes, resulting in an underestimation of age. A similar bias was observed by Moen et al. (2019), where a clear tendency for their CNN-based model to predict lower ages for older individuals when compared with age estimation by human readers was evident. Future studies should ensure sufficient sample sizes of older age classes are included in both the training and testing datasets. Alternatively, one way to mitigate this is to implement a cost function that weights age classes evenly, i.e., each age class inflicts the same cost (Shen et al. 2015; Moen et al. 2019).

Although we did not undertake an analysis of those instances where the CNN failed to correctly predict age, a preliminary set of experiments revealed that image inconsistencies could impact the results. For

example, image colouration (or rather the removal of colour), scaling, and masking properties all had significant effects on the predicted age from the CNN for the single snapper otolith trialled. In this preliminary examination, we used available digital photographic images that often varied in their quality, lighting, and orientation. For hoki in particular, oversaturation of light (see Figure 2) often masked the otolith core, and potentially the innermost growth bands. Although some of these aspects (e.g., image orientation) were resolved during augmentation, our results suggest that if the image capture process could be standardised (by using consistent equipment, optical range, lighting conditions, etc.) then the ability for the model to predict ages could be improved. Increasing the sample sizes of the training datasets is also likely to improve the precision of the age estimates. Further experiments are warranted to examine these issues.

Recent estimates suggest that reading accounts for approximately half of the costs involved with age estimation from extracted otoliths for inshore fish species in New Zealand, with preparation accounting for the remaining half (Jeremy McKenzie, NIWA, pers. comm.). The CNN-based machine learning has the potential to considerably reduce costs involved with age estimation, even if otoliths need to be extracted, processed and imaged. Moreover, training otolith readers and achieving consistency among readers can be time-consuming and costly processes; much of the cost could be negated using an automated reading approach. Perhaps most significantly, the approach holds considerable promise for reducing problems associated with otolith reading by trained human readers, such as when an individual reader changes their interpretation over time and differences in interpretation between readers.

CNN-based approaches offer considerable potential for increasing the efficiency of age estimations whilst maintaining precision and human control over the ageing process. As the model provides a quantified level of confidence (the model output) for each age estimate, outputs can be screened to discriminate samples for which this confidence is low. Samples that produce low model outputs could then be examined by an expert reader or readers to determine the age class. This would eliminate the need for human experts to read ‘easy’ otoliths, whilst maintaining human-based decision control over more ‘difficult’ otoliths.

5.3 Potential for applying machine learning to otolith CT images

Advances in the fields of computer tomography and machine learning in recent years indicate that it may be possible to automate many aspects of fish ageing, with the potential for significant cost and time savings and improved consistency. Greater benefits may be obtained by combining both approaches. Future experiments are thus required to assess the performance of CNN-based age estimations (relative to human readers) from CT-derived images.

Ultimately, it may be possible to automate the entire process, by scanning a whole fish using CT technology (e.g., along a conveyor belt) and subsequently estimating its age using a machine-learning approach. Current technologies, however, are insufficient for this purpose. The chamber size for the MARS systems used in the current study is 280 mm by 100 mm, whereas the maximum object size for the Bruker SkyScan 1272 system is 75 mm by 80 mm. However, considerable development is occurring in this area of technology; e.g., MARS Bioimaging Ltd is currently developing an orthopaedic imaging system capable of scanning a human arm (albeit with an 80 μm voxel size; https://www.marsbioimaging.com/mars/wp-content/uploads/2018/07/MARS_Electronic.pdf).

To understand the required specifications for any imaging system designed for otolith ageing, it will be necessary to perform experiments to better understand which features the CNN relies on to perform ageing. When automating a process previously performed manually, it is tempting to imitate the same steps previously used by humans. In this case, the conventional manual approach is to view and count otolith growth bands. So, initially, it may seem reasonable to attempt to automate the imaging of an otolith with CT, such that the bands are resolved, and then automate the counting process. However, this may forego other, better, approaches which are now enabled by the use of machine learning. For example, it may not be necessary to image the otoliths at resolutions sufficient for human viewers to resolve, because the CNN may be able to arrive at an age estimate without directly counting bands.

Knowledge of this could potentially relax the imaging requirements enormously and may even allow fish ageing with lower resolution, high speed CT.

6. MANAGEMENT IMPLICATIONS / PROPOSED RESEARCH

Our preliminary examination, based on a review of the scientific literature and trials on New Zealand fish species, suggests there is significant potential to use CT scanning technology to image otoliths, to automate age estimation from otolith images, and to combine these techniques to form a fully automated ageing system. Below, in the form of a research proposal, we present the key next steps for improving and implementing these processes for ageing of New Zealand's exploited fish species.

6.1 Short-term research objectives

Objective 1: Optimise and further evaluate the convolutional neural network (CNN)-based model developed in the project 'Feasibility of automating otolith ageing using CT scanning and machine learning' to automate age estimation of snapper and hoki.

The first requirement of optimising the CNN approach trialled in the current study will be to improve the quality of images. For both snapper and hoki, digital images of otoliths (n = 2000–4000 per species) will be taken using a standardised protocol, with consistent equipment, lighting conditions, image size, and orientation. Care will be taken to ensure the images span a wide range of fish sizes, ages, sexes, and collection locations. Once images are optimised, further development of the model will be undertaken. In particular, images will be manipulated to determine which features affect ageing, to understand how the CNN uses information and to identify ways to improve the process.

The performance of the optimised CNN model to estimate age for both snapper and hoki will be evaluated by comparing estimates of age derived from the CNN model with previous age estimates achieved via conventional microscopy using current standardised approaches to assess within and between reader performance, including age-bias plots, indices of average percent error (IAPE), and the coefficient of variation (CV) (e.g., Horn & Sutton 2017). The performance evaluation will investigate potential covariates that may affect ageing success: e.g., fish size, otolith size, sex, sample site, reader-assigned age, reader-determined cohort, assigned readability score. The proposed budget includes staff time to generate a library of optimised otolith images for snapper and hoki (n = 4000 per species), optimisation of the CNN model, model performance evaluation, data analysis, and reporting.

Indicative cost: $140,000 NZD (excl. GST).

Note: Costs for these components could be considerably reduced should this work occur following the completion of Objective 1 of the proposed project SAM2019-02 'Development of imaging analysis techniques to determine ages from otoliths' (assuming snapper and hoki are used in SAM2019-02), subject to the availability of the resulting images.

Objective 2: Investigate the potential for the CNN-based model to automate age estimation of additional species from photographic images of sectioned otoliths.

Following optimisation of the approach for snapper and hoki using the methods outlined above, trials will be conducted on additional fish species. Trials will focus on the key species currently aged as part of routine monitoring and assessments, including ling, hake, Antarctic toothfish (*Dissostichus mawsoni*), tarakihi (*Nemadactylus macropterus*), and southern blue whiting (*Micromesistius australis*), with samples from the Fisheries New Zealand otolith collection, currently stored at NIWA, Wellington. For each species, a robust dataset of sectioned otolith images (n = 2000–4000, where available) will be generated, with care taken to ensure that a wide range of age classes and any spatial and temporal variation present is included, and that digitisation techniques are standardised as best as possible. Each image used by the CNN model will be accompanied by a final estimate of age from human reader(s)

obtained using standardised protocols. These estimates will then be compared with previous age estimates achieved via conventional microscopy (Walsh et al. 2014b, Horn & Sutton 2017) using current standardised approaches to assess within and between reader performance (including age-bias plots, IAPE, and CV). Because the testing set must be sufficiently large to generate statistically valid comparisons, we will use data from previous reader comparisons to determine the required sample sizes. Where possible, we will seek to include otolith images from previous reader comparisons in these trials.

Just as an experienced human reader can quickly learn to read a new species, the CNN is expected to require progressively fewer images to learn each new species. To optimise performance of the CNN it may be advantageous to pool data from all species prior to training the model. This may help mitigate the challenge posed by the relatively small training data sets. Alternatively, the model may then be less finely tuned to individual species. This pooled approach will be investigated and compared with the output where the model is trained separately for each species. As part of this investigation, we hope to gain a better understanding of the relationship between the number of training examples and final prediction accuracy of the model. The proposed budget includes staff time to generate a library of optimised otolith images for five additional species (n = 4000 per species), species-specific optimisation of the CNN model, model performance evaluation, data analysis, and reporting.

Indicative cost: $250,000 NZD (excl. GST).

Note: Costs for these components could be considerably reduced should this work occur following the completion of Objective 1 of the proposed project SAM2019-02 'Development of imaging analysis techniques to determine ages from otoliths', depending on the species used in this project and subject to the availability of the resulting images.

Objective 3: Further identify the potential to use otolith scans from CT scanning technologies for imaging otolith annular structure.

Investigations undertaken in the current study revealed that whereas some zonation was evident, particularly in younger fish, the resolution of the CT scanner used was insufficient for imaging fine-scale growth bands of older fish. Accordingly, this component will test the utility of micro-CT scanning technologies (e.g., the Bruker SkyScan 1272 desktop MicroCT held at the Auckland Bioengineering Institute, University of Auckland) for imaging otolith annular structure. Trials will be conducted on a range of fish species with different otolith morphologies: snapper, representing a relatively 'easy' species to age with relatively clear growth bands (Walsh et al. 2014a); hoki, representing a relatively difficult species to age because of the presence of diffuse growth bands (Horn & Sutton 2017); and tarakihi, representing a challenging species to age because of its relatively long lifespan (at least 45 years) and compressed growth bands towards the outer margins, particularly in older fish (i.e., those 10 years of age or older) (Walsh et al. 2014b). Approaches used to reduce scanning time, such as scanning a narrow band across an otolith or scanning several otoliths at once, will be trialled. Once three-dimensional reconstructions are generated, transverse slices will be taken through the otolith core using digital clipping planes. Growth bands will be counted from the micro-CT sections by two independent readers without prior knowledge of previous age estimations or the size of fish. These reads will then be compared with previous age estimates achieved via conventional microscopy (Walsh et al. 2014a, Horn & Sutton 2017) using current standardised approaches to assess within and between reader performance (including age-bias plots, IAPE, and CV). The proposed budget includes travel to use a micro-CT (budgeted on the use of the machine at the University of Auckland), time on the machine, data analysis, and reporting. At the time of writing, the daily rate for use of a micro-CT scanner, including operator costs, image reconstruction, and analysis, was approximately $1,000 NZD.

Indicative cost: $70,000 NZD (excl. GST).

Objective 4: Identify the potential to automate ageing from images generated from micro-CT scanning using deep learning approaches.

Should trials using micro-CT scanning technologies reveal annual growth bands and comparable age estimates to those of human experts be obtained in Objective 3 above, trials exploring the potential for estimating age via machine-learning approaches of micro-CT derived images will be conducted. Here, large numbers of otoliths (more than 1000) of the three species investigated in Objective 3 will be scanned with a micro-CT scanner using optimised approaches developed in Objective 3. Resulting images will be analysed using a CNN model. Age estimates from the model will be compared with those resulting from human interpretation of the images and previous age estimates achieved via conventional microscopy (Walsh et al. 2014, Horn & Sutton 2017) using current standardised approaches to assess within and between reader performance (including age-bias plots, IAPE, and CV).

Indicative cost: Costs for this component will be determined following the trial of Objective 3.

Note: should it be found in Objective 3 that micro-CT scans provide poor resolution of growth bands and/or unprecise age estimates, this component of the study will not be conducted.

6.2 Medium-term research objectives

The following components are additional objectives that could be investigated in the medium to longer-term, given their potential to provide substantial benefits and cost-savings. Objective 5 requires a longer timeframe to create a library of fish images with associated ages, and Objective 6 requires information from the above objectives to determine costs.

Objective 5: Exploring the potential for the CNN-based model to estimate fish age using non-annular information (including two-dimensional photographic images of whole fish and whole otoliths).

Given recent advances in CNN-based approaches, the potential to apply machine learning technologies to non-annular age estimations are worth exploring. In this objective, the potential for estimating age from two-dimensional images of whole fish, and whole otoliths, will be trialled. Trials will be conducted on a line-caught species, such as Antarctic toothfish, because these are expected to maintain their regular shape following harvest (compared with trawl caught species), and images will be able to be taken during routine surveys and biological sampling. During two successive Ross Sea shelf surveys, images will be taken of each Antarctic toothfish caught. Images will include the fish number for later reference and linkage of the photo to fish age. Fish will be processed as per standard approaches (i.e., each fish will be measured, sexed, and assigned its maturity status) and otoliths will be extracted. Otoliths will be imaged whole and aged using standard ageing procedures. Whole fish and whole otolith images will be added into a CNN-based model for estimating age. The model will be trained by pairing the images with their otolith-based ages, and then tested. Age estimates generated from the whole fish and whole otolith images will then be compared with previous age estimates achieved via conventional microscopy using current standardised approaches to assess within and between reader performance (including age-bias plots, IAPE, and CV).

Objective 6: Identify the potential for estimating age from otoliths *in situ*.

Should trials using micro-CT scanning technologies reveal annual growth bands and comparable age estimates to those of human experts (in Objective 3 above), further trials exploring the potential for *in situ* imaging of otoliths and subsequent ageing will be conducted. The long-term objective here is not laboratory-based CT-scanning, but scanning on industry production lines.

Whole fish, or portions of a fish containing the otoliths, will be placed in a micro-CT scanner with scans focussed on the head area to reduce scan time. Scanned otolith images will be constructed using relevant

software and aged by an experienced reader. Initial trials will focus on snapper, because it is locally available. Care will be taken to ensure a range of lengths (and thus ages) are included. Once scanned, otoliths from each scanned individual will be extracted and aged using standard protocols by an experienced reader (Walsh et al. 2014a). Estimated ages from scanned images will be compared with those of the experienced reader using age-bias plots, IAPE, and CV.

Note: should it be found in Objective 3 that micro-CT scans provide poor resolution of growth bands, this component of the study will not be conducted.

7. ACKNOWLEDGEMENTS

We thank Hannah Prebble, Professor Anthony Butler, and the team at MARS Bioimaging Ltd for advice, access, and expertise on the MARS CT scanner, and Jeremy McKenzie (NIWA) for reviewing the draft report. Ian Tuck (NIWA) kindly provided additional funding for undertaking the CT scans. This work was completed under Ministry for Primary Industries project SEA2018-30.

8. REFERENCES

Abramoff, M.D.; Magalhães, P.J.; Ram, S.J. (2004). Image processing with ImageJ. *Biophotonics International 11*: 36–42.

Allken, V.; Handegard, N.O.; Rosen, S.; Schreyeck, T.; Mahiout, T.; Malde, K. (2019). Fish species identification using a convolutional neural network trained on synthetic data. *ICES Journal of Marine Science 76*: 342–349.

Bignami, S.; Enochs, I.C.; Manzello, D.P.; Sponaugle, S.; Cowen, R.K. (2013). Ocean acidification alters the otoliths of a pantropical fish species with implications for sensory function. *Proceedings of the National Academy of Sciences of the United States of America 110*: 7366–7370.

Campana, S.; Neilson, J.D. (1985). Microstructure of fish otoliths. *Canadian Journal of Fisheries and Aquatic Sciences 42*: 1014–1032.

Campana, S; Thorrold, S. (2001). Otoliths, increments, and elements: keys to a comprehensive understanding of fish populations. *Canadian Journal of Fisheries and Aquatic Sciences 58*: 30–38.

du Plessis, A.; Broeckhoven, C.; Guelpa, A.; le Roux, S.G. (2017). Laboratory x-ray micro-computed tomography: a user guideline for biological samples. *GigaScience 6*: 1–11.

Edds-Walton, P.L.; Arruda, J.; Fay, R.R.; Ketten, D.R. (2015). Computerized tomography of the otic capsule and otoliths in the oyster toadfish, *Opsanus tau*. *Journal of Morphology 276*: 228–240.

Fablet, R. (2006). Statistical learning applied to computer-assisted fish age and growth estimation from otolith images. *Fisheries Research 81*: 219–228.

Fablet, R.; Le Josse, N. (2005). Automatic fish age estimation from otolith images using statistical learning. *Fisheries Research 72*: 279–290.

Felix, P.M.; Gonçalves, A.; Vicente, J.R.; Fonseca, P.J.; Amorim, M.C.P.; Costa, J.L.; Martins, G.G. (2016). Optical micro-tomography "OPenT" allows the study of large toadfish *Halobatrachus didactylus* embryos and larvae. *Mechanisms of Development 140*: 19–24.

Fisher, M.; Hunter E. (2018). Digital imaging techniques in otolith data capture, analysis and interpretation. *Marine Ecology Progress Series 598*: 213–231.

Francis, M.P.; Jones, E.G.; Ó Maolagáin C.; Lyon, W.S. (2018b). Growth and reproduction of four deepwater sharks in New Zealand waters. *New Zealand Aquatic Environment and Biodiversity Report No. 196*. 58 p.

Francis, M.P.; Ó Maolagáin, C.; Lyon, W.S. (2018a). Growth and reproduction of carpet shark, common electric ray and blind electric ray in New Zealand waters. *New Zealand Aquatic Environment and Biodiversity Report No. 195*. 39 p.

Geraghty, P.T.; Jones, A.S.; Stewart, J.; Macbeth, W.G. (2012). Micro-computed tomography: an alternative method for shark ageing. *Journal of Fish Biology 80*: 1292–1299.

Guo, Y.; Liu, Y.; Oerlemans, S.; Lao, S.; Wu, S.; Lew, M.S. (2016). Deep learning for visual understanding: a review. *Neurocomputing 187*: 27–48.

Horn, P.L.; Sutton, C.P. (2017). Age determination protocol for hoki (*Macruronus novaezelandiae*). *New Zealand Fisheries Assessment Report 2017/13*. 26 p.

Krizhevsky, A.; Sutskever, I.; Hinton, G.E. (2012). ImageNet classification with deep convolutional neural networks. *In*: Pereira, F.; Burges, C.J.C.; Bottou, L.; Weinberger, K.Q. (Eds.) *Advances in Neural Information Processing Systems 25*: 1097–1105.

Liu, K.M.; Chen, C.T.; Liao, T.H.; Joung, S.J. (1999). Age, growth, and reproduction of the pelagic thresher shark, *Alopias pelagicus*, in the northwestern Pacific. *Copeia 1999*: 68–74.

Long, J.M.; Snow, R.A. (2018). Posthatch development of otoliths and daily ring genesis in age-0 spotted gars. *Transactions of the American Fisheries Society 147*: 1146–1152.

Malde, K.; Handegard, N.O.; Eikvil, L.; Salberg, A.B. (2019). Machine intelligence and the data-driven future of marine science. *ICES Journal of Marine Science, fsz057*, https://doi:10.1093/icesjms/fsz057.

Maxime, E.L.; Albert, J.S. (2014). Redescription of the tuvirão, *Gymnotus inaequilabiatus* Valenciennes, 1839, using high-resolution X-ray computed tomography. *Copeia 2014*: 462–472.

Metscher, B.D. (2009). MicroCT for comparative morphology: simple staining methods allow high-contrast 3D imaging of diverse non-mineralised animal tissues. *BMC Physiology 9*: 11.

Moen, E.; Handegard, N.O.; Allken, V.; Albert, O.T.; Harbitz, A.; Malde, K. (2018). Automatic interpretation of otoliths using deep learning. *PLoS One 13*: e0204713.

Parsons, K.T.; Maisano, J.; Gregg, J.; Cotton, C.F.; Latour, R.J. (2018). Age and growth assessment of western North Atlantic spiny butterfly ray *Gymnura altavela* (L. 1758) using computed tomography of vertebral centra. *Environmental Biology of Fishes 101*: 137–151.

Robertson, S.; Morison, A. (2001). Development of an artificial neural network for automated age estimation. Final report to the Fisheries Research and Development Corporation Project No. 98/105. Department of Natural Resources and Environment, Queenscliff, Victoria, Australia. 289 p.

Robertson, S.G.; Morison, A.K. (1999). A trial of artificial neural networks for automatically estimating the age of fish. *Marine and Freshwater Research 50*: 73–82.

Rumelhart, D.E.; Hinton, G.E.; Williams, R.J. (1986). Learning representations by back-propagating errors. *Nature 323*: 533–536.

Schmidhuber, J. (2015). Deep learning in neural networks: an overview. *Neural Networks 61*: 85–117.

Schulz-Mirbach, T.; Olbinado, M.; Rack, A.; Mittone, A.; Bravin, A.; Melzer, R.R.; Ladich, F.; Hess, M. (2018). In-situ visualization of sound-induced otolith motion using hard X-ray phase contrast imaging. *Scientific Reports 8*: 3121.

Shen, W.; Wang, X.; Wang, Y.; Bai, X.; Zhang, Z. (2015). Deepcontour: a deep convolutional feature learned by positive-sharing loss for contour detection. *In*: Proceedings of the 28th IEEE Conference on Computer Vision and Pattern Recognition, pp. 3982–3991. 7–12 June 2015, Boston MA.

Takashima, Y.; Takada, T.; Matsuishi T.; Kanno, Y. (2000). Validation of auto-counting method by NIH image using otoliths of white-spotted char *Salvelinus leucomaenis*. *Fisheries Science 66*: 515–520.

Walsh, C. (2008). Swinging 60 Old man snapper caught. Seafood New Zealand, November 2008.

Walsh, C.; Horn, P.; McKenzie, J.; Ó Maolagáin, C.; Buckthought, D.; Stevenson, M.; Sutton, C. (2014b). Age determination protocol for tarakihi (*Nemadactylus macropterus*). *New Zealand Fisheries Assessment Report 2014/53*. 35 p.

Walsh, C.; Horn, P.; McKenzie, J.; Ó Maolagáin, C.; Buckthought, D.; Sutton, C.; Armiger, H. (2014a). Age determination protocol for snapper (*Pagrus auratus*). *New Zealand Fisheries Assessment Report 2014/51*. 37 p.

Welch, T.J.; van den Avyle, M.J.; Betsill, R.K.; Driebe, E.M. (1993). Precision and relative accuracy of striped bass age estimates from otoliths, scales, and anal fin rays and spines. *North American Journal of Fisheries Management 13*: 616–620.

Zhu, X.; Wastle, R.J.; Howland, K.L.; Leonard, D.J.; Mann, S.; Carmichael, T.J.; Tallman, R.F. (2015). A comparison of three anatomical structures for estimating age in a slow-growing subarctic population of Lake Whitefish. *North American Journal of Fisheries Management 35*: 262–270.