

ARTICLE

Automatic determination of heart rates from microscopy videos of early life stages of fish

R. Nepstad^a, E. Davies^a, D. Altin^b, T. Nordtug^a, and B. H. Hansen^a

^aSINTEF Ocean, N-7465 Trondheim, Norway

^bBioTrix, 7022 Trondheim, Norway

ABSTRACT

Toxic effects of organic hydrophobic contaminants include impacts on fish heart rate and cardiac functioning. Thus, in ecotoxicology as well as aquaculture and even medicine, fish heart functioning has important application areas. We here present a pipeline of image processing and statistical techniques which has been assembled to extract heart rate information from microscopy videos of the embryo and larval stages of three species of fish (Atlantic cod, haddock and Atlantic bluefin tuna). The method allows for automatic processing in parallel of a large number of individuals, saving a significant amount of time compared with manual processing, while simultaneously eliminating the type of errors such a manual process can incur.

KEYWORDS

Imageprocessing, ecotoxicology, cod, tuna, haddock, heart rate

1. Introduction and Background

The temporal and spatial overlap of the 2010 Deepwater Horizon (DWH) incident, in the Gulf of Mexico, with spawning of many commercially and ecologically important fish species resulted in a concern for potential effects of oil exposure on sensitive early life stages of large migratory and long-lived fish species, such as mahi-mahi (*Coryphaena hippurus*) (Incardona et al., 2014). Cardiotoxicity has been shown to follow from exposure of early life stages (ELS) of fish to polycyclic aromatic hydrocarbons (PAH) and other oil components (Incardona et al., 2006; Sørhus et al., 2015; Brette et al., 2014). Bradycardia (reduced heart rate) as well as other cardiotoxic effects (e.g. pericardial oedema, arrhythmia, cardiac looping and silent ventricle) have been observed in fish larvae exposed to PAHs, and such effects are linked to a reduction in the long-term survival and fitness of larvae (Incardona et al., 2014). Embryonic PAH exposure has been shown to cause bradycardia in hatchlings for several teleost species including cod, haddock (Sørhus et al., 2015), bluefin tuna and mahi-mahi (Incardona et al., 2014)). Embryonic bradycardia also appear to be a precursor for jaw, cranium and spine deformations in early larval development (Incardona et al., 2004). For a recent review of cardiac development in fish see (Incardona and Scholz, 2016).

Fast and robust methods to assess embryonic and larval health is of importance over a wide range of applications, including toxic exposure studies, aquaculture research, and environmental impact assessments. Video recording of embryos through a microscope may be used analytically to assess embryonic survival and health, and even used to predict delayed health effects. One endpoint which can be used for this purpose is heart rate, which can be extracted from videos. In medicine, zebrafish (*Danio rerio*) is used as an important *in vivo* vertebrate model to study the etiology of human cardiovascular diseases, for cardiovascular drug discovery, as well as for high-throughput screening of medicine (Kessler et al., 2015). Zebrafish bradycardia and atrioventricular block is used specifically as a model for screening the potential for drugs to induce repolarization abnormalities in humans (Milan et al., 2003).

Measurements of cardiac function in toxicity studies are generally done using laborious, manual, and to certain extent subjective analyses of videos of embryos and larvae, and there is a great potential for time saving and quality improvement by automating and standardizing these analyses. More generally in the biological sciences, and particularly in medicine, image processing techniques have been used and developed for several decades. There is a continuing trend towards faster and cheaper computational resources, and evolving imaging technology that provides higher resolution and faster acquisition rates from off the shelf products. Combined with development of powerful, high-level software toolboxes for image processing (e.g. imageJ (Schindelin et al., 2015) and scikit-image (van der Walt et al., 2014)), there is today great potential for fast development of automated analysis tailored to specific application. A recent highlevel review of image acquisition and processing in the biological sciences illustrate the main components involved in such undertakings, as well as challenges that must be overcome (Eliceiri et al., 2012).

In recent years, methods for automated image processing (Mikut et al., 2013; Pantazis and Supatto, 2014; Pylatiuk et al., 2014), and even machine learning techniques (Tharwat et al., 2015), have been developed to assist toxicology researchers in the analysis of video data. These recent efforts have focused mostly on zebrafish, a popular model species, which is particularly suited for such analysis due to very transparent tissue in early life stages. Many of these new advances in such automated analysis techniques have, however, been closely coupled to the hardware configuration used to obtain the images. This means that the analysis can become limited when applied to other microscope configurations or sampling procedures. There is therefore scope to widen the range of applications for automated analysis of heart function in microscope images, which may be obtained with sub-ideal conditions for computational analysis. It is often a challenge for automated image processing to handle unexpected features in images, which may appear relatively insignificant to a human. In the case of heart rate analysis, automated processing may become significantly more challenging if objects other than the heart itself move intermittently at the timescale of the heart rate, or if the frame rate, resolution or illumination conditions of the images is altered. Adapting image processing techniques to extract heart rates for multiple fish species with less optimal optical properties is possible, as we demonstrate here.

In the following sections we present an automated analytical approach for multispecies fish embryos and larvae to extract heart rates from video data, for use in toxicological applications to determine the potential for pollutants to cause cardiotoxic and developmental effects. The work-flow of the method is specifically constructed to

handle a relatively wide range of microscope configurations and magnifications so that it can be applied to video microscopy data which may not have initially been intended for the purpose of detailed automated heart rate analysis. Advances to existing video microscopy analysis methods are achieved by the application of signal processing and filtering techniques, combined with a robust identification of heart regions. We illustrate the method by extracting heart rates for three different species and two life stages. Comparisons with data previously obtained using traditional counting methods are made to assess the accuracy of the automated method. We also demonstrate the method by applying it to video data obtained from a toxicological exposure study of cod larvae, where a significant change in heart rate between control and exposure group is found.

2. Method

Before discussing the details of the automatic heart rate extraction method in section 2.2, we first briefly summarize the experimental setup used to collect the data which we have used to test our method.

2.1. Selection of fish species and maintenance

Three different species of fish were used for testing the method for automated heart rate measurements: Atlantic cod (*Gadus morhua*), haddock (*Melanogrammus aeglefinus*) and Atlantic Bluefin Tuna Fish (ABFT) (*Thunnus thynnus*). Examples of embryo and larval stages of these are shown in figure 1.

Newly fertilized eggs were transported by airfreight to SINTEF Sealab (Trondheim, Norway) within 12 hours from Spain (ABFT) or Bergen (cod and haddock) in temperature controlled transporter cases. Embryos and hatched larvae were kept at constant water flow (approx. 17 ml clean sea water/min) in 5L tanks at 5°C (cod and haddock) or 25°C (ABFT). The seawater used was natural seawater from Trondheimsfjord extracted at depth of 70 m, passed through a sand filter for removing larger particles before a final filtering step removing particles larger than approx 4µm by a cartridge filter (Cuno, US).

Microscope videos of embryos were obtained for ABFT approximately 32 hours post fertilization (hpf). Cod and haddock videos were obtained approximately 14 days post fertilization (dpf). Cod and haddock develop much slower than ABFT, so heart rates are visible only after organogenesis (>12 dpf at 5°C). As a positive control for bradycardia, cod embryos (11-15 days post fertilization, 7°) were treated with produced water containing approximately 30µg/L T-PAH, for 4 days of static exposure. Briefly, a dichloromethane (DCM) extract of a North Sea produced water was reconstituted into sea water by first removal of DCM by evaporation using nitrogen followed by addition of filtered sea water and ultrasonication for solubilizing produced water components in sea water. The PAH concentrations were measured in extracts of water samples using standard gas chromatography – mass spectrometry (GC-MS) as described elsewhere (US EPA, 1986; Faksness et al., 2012). Previous studies have shown that PAH exposures of fish embryos to this concentration causes bradycardia (Incardona et al., 2014).

Videos of hatched larvae were collected approximately 1 day post hatch (dph) for ABFT and 2 dph for cod and haddock. They were immobilized in methylcellulose (Incardona et al., 2014) and the cardiac area was imaged for 1 minute using a microscope.

Embryos and fish larvae were filmed using a dissecting microscope (Nikon eclipse 80i microscope with a 2xPlanApo objective) connected to a JVC-camera for ABFT and Leica MC170HD-camera for cod and haddock.

2.2. Automated heart rate extraction

The basis for the approach is founded on the hypothesis that the heart beat signal corresponds to variations in image intensity within a region of interest (ROI) that covers the heart of either fish embryos or larvae. For visual counting of heart beats to be possible, this must necessarily be the case. Thus, by extracting the time series of intensity for each ROI, heart rates, beat to beat variations, and other quantities of interest may be calculated. The main steps to achieve this are outlined below and described in detail in the subsequent sections:

- (1) Identify region of interest covering heart tissue,
- (2) Loop through video frames, extracting intensity signal from ROI,
- (3) Analyse signal from previous step to calculate quantity of interest.

Unfortunately, this procedure is complicated by several issues: video noise, motion of embryos/larvae, varying orientation of embryos in the same video, and ‘frame drift’, the latter being caused by motion of the whole imaging region relative to the camera. These can all be handled to some extent, but the most severe cases can cause the method to produce inaccurate results.

The presented analysis pipeline is written in Python, and makes use of the scikit-image package (van der Walt et al., 2014), as well as the standard Python scientific stack (numpy, scipy, matplotlib) (Hunter, 2007; Oliphant, 2007). To parse the video files, imageio is used (Klein, Klein).

The analysis runtime for a single video is typically a few minutes on a modern desktop computer, but does depend on video resolution, duration and framerate. For instance, in the case of haddock embryos (50s duration, 30 FPS, 1080x1920 pixel resolution), extracting heart rates for four identified embryos took approximately 40 seconds on a Intel Xeon 3.3 GHz machine. On a modern multi-core computer several videos can be analysed simultaneously (typically 4-8), while making use of a small computing cluster can bring this up to 100 videos or more.

2.2.1. Identifying heart tissue

A video frame (or image) at time t is a single-channel $n \times m$ integer array of intensity values in the range $[0, 255]$. We denote this as $I(t)$. For a discrete, given time t_k , we use I_k . Binary (black-and-white) images are denoted B . The intensity of a single pixel (i, j) in a frame at time t_k is denoted I_{ijk} .

To identify the heart region of interest (ROI) within a video, differences in intensity are calculated. Specifically, given a sequence of N video frame images I_1, I_2, \dots, I_N , the

absolute difference sum image D is calculated between images separated by a user-defined frame interval df as:

$$D = \sum_{k=1}^{N/df} |I_{(k+1) \times df} - I_{(k \times df)}| \quad (1)$$

The frame interval df is required in order to handle discrepancies between intensity difference and frame rate for differing species (i.e. a high frame rate and slowly changing intensity requires a larger df). Note that the difference must be computed with floating point representation (not integers) to achieve sufficient precision for ROI detection.

Calculation of D is performed on one image channel, typically either the V (HSV) or L (Luv) channels. An example of the resulting D image is shown in figure 2 (top). D is then further processed to identify candidate ROIs. First, filtering is applied to remove noise, using either a bilateral denoising filter (denoise bilateral function in skimage) or a Gaussian filter (embryos), producing \tilde{D} . Then thresholding is performed with the greatest value from the Otsu (Otsu, 1979) and Yen (Jui-Cheng Yen et al., 1995) methods, to produce a binary image B , from which small objects are removed, and any holes in the resulting image are filled. The cleaned binary image is then labeled, producing numbered objects, as shown in figure 2 for the case of embryos. For larvae, the thresholding is used to identify the single heart region, which is then enclosed by a convex hull, as shown in figure 4 (top).

2.2.2. Extracting the heart rate signal

With the approximate heart region established by the calculated ROI, an intensity signal for heart rate determination can be extracted. This is done by looping through the video frames, transforming each frame to the Luv color space and extracting the ROI in the L channel. Then, either the mean or standard deviation is calculated on the n masked pixel values,

$$\bar{s}_k = \frac{1}{n} \sum_{i,j \in ROI} I_{i,j,k} \quad (2)$$

$$\sigma_k^2 = \frac{1}{n} \sum_{i,j \in ROI} (I_{i,j,k} - \bar{s}_k)^2 \quad (3)$$

In a previous study on extracting heart rates from videos of zebrafish larvae, the standard deviation approach was used (Pylatiuk et al., 2014). In either case, an approximately periodic time series $h(t)$ is obtained, which should contain the heart beat signal.

The signal $h(t)$ is detrended and normalized before proceeding to the next step, yielding $h^0(t)$. Detrending is performed by subtracting the rolling mean with a window length which should cover several heart beat periods. This signal is then normalized by subtracting the signal mean μ_h , and dividing by the signal standard deviation σ_h ,

$$\bar{h}_k(t) = \frac{h(t) - \mu_h}{\sigma_h} . \quad (4)$$

The result is shown in figure 3 (left) for embryos, and in figure 4 (bottom) for larvae. The frequency content of the intensity signal is analysed using the Welch periodogram (figure 3, right). If a good quality and clean signal was obtained in the previous step, the periodogram will be dominated by a peak at the heart beat frequency, which is taken as the heart rate value. Higher harmonic peaks of the fundamental (heart beat) frequency may occur, which are caused by weaker periodic maxima in the intensity signal, phase shifted from the main heart rate signal. This is mostly an issue for the larval heart rate signals, because these resolve the complex heart contractions and motions to a much greater degree. The search for a heart rate value is restricted to a subset of the possible frequencies permitted by the signal, typically from near zero to a few multiples of the expected heart rate (red region in figure 3, right panel). This improves extraction of heart rates from videos with excessively high-frequency noise in the ROI signal.

In the set of videos considered here, motion of the larvae or embryos is a common occurrence. This will affect the extracted heart rate signal by introducing additional signal components, and in some cases the motion is such that the initially determined heart ROI is no longer valid. For these reasons, it is necessary to infer when motion occurs, and to discard the corresponding time intervals in the signal $h(t)$. One possible solution is to calculate the pixel-normalized frame-by-frame difference,

$$d_k = \frac{1}{m} \sum_{i,j} |I_{i,j,k} - I_{i,j,k-1}| \quad (5)$$

where m is the number of pixels in the image. This time series ($d(t)$) is then smoothed with a Savitzky-Golay filter (from `scipy`), see figure 5. Time intervals with excessive motion are identified by points (red markers) in $d(t)$ (thin blue full line) with greater than three times the median value of d (upper blue dashed line). The largest contiguous time interval without identified motion is then used to sub-sample the full heart rate signal. The candidate intervals are given by the black full lines. This approach was often successful in extracting a heart rate from noisy videos/signals where animal twitched, but does break down in the most severe cases, as discussed in the result section. It may be possible to further improve on the method by using Dense Optical Flow techniques (see e.g. (Uchida, 2013)) to quantify motion in the whole frame, and then use this to reposition the ROI, but we have not pursued this further in the present study.

3. Results and discussion

A number of videos for the three different fish species have been analyzed. An overview of the data and key performance metrics for the automated method is shown in table 1. The ROI accuracy indicates the percentage number of heart regions that were correctly identified, while the heart rate (HR) accuracy is the relative median difference between the values of the automatic (A) and manually counted (M) heart rate values,

$$\frac{\text{median}(|M - A|)}{\text{median}(M)} . \quad (6)$$

For haddock and tuna embryos, no manual data were available for comparison. We have two different datasets for cod embryos (A and B), with manual data only available for data set A. Thus, we made comparison with four different manual data sets (cod embryo A, cod larvae, haddock larvae and tuna larvae).

In figure 6 all extracted heart rates are shown for the three species and two stages (only dataset B for cod embryos). If outlier values are disregarded (indicated by black diamonds), differences in the median heart rate values both between species, and between embryo and larval stages for the same species can be seen. Notably, bluefin tuna has a significantly higher heart rate than cod and haddock. In all cases, larvae have higher (median) heart rates compared to the embryos. The outlier values correspond generally to cases where either the heart ROI was incorrectly identified, or where severe motion (twitching) drowned out the heart rate signal.

Comparing the automatically extracted heart rates with a dataset of manually counted values, we find overall good agreement, as shown in figure 7. For the cod embryos the agreement is particularly good, with a median difference of only 2%. The cod larvae results are also in close agreement, except for some outliers (9, red points, excluded), which we exclude from the median difference. Data points were flagged as outliers using two criteria: either the signal processing routine has flagged the heart rate signal or identified ROI as suspicious, or the video/heart beat was flagged as problematic during the manual counting procedure (either strongly arrhythmic or severe twitching). In the case of haddock and bluefin tuna there is more scatter, but the agreement is still fair. We also note that all tuna larval heart ROIs were correctly identified (table 1). The points far off the centreline in the haddock (4, red points, excluded) and tuna (3, red points, excluded) plots are caused by severe twitching in the videos, causing either ROI misidentification or heart rate signal drowning. The Pearson correlation coefficient has been calculated in all cases, with the red points excluded, and is given in the subplot headings (last number). Generally, outliers from the main cluster of heart rates can be automatically flagged for a quick manual inspection, using for instance the box plot method, cf. figure 6.

The method has also been tested on videos from a toxicology study, where cod embryos (11-15 days post fertilization) were exposed to a reconstituted produced water extract. Here we compare the control group with the group exposed to produced water, cf. figure 8, which were exposed to $30\mu\text{g/L}$ T-PAH for 4 days (static exposure). This T-PAH concentration has previously been shown to cause cardiotoxic responses (like bradycardia) in marine cold-water embryos (Sørhus et al., 2015; Sørhus et al., 2016). Discarding some outlier points, the exposed group had a lower median heart rate of approximately 10 bpm (26.2 vs 35.0 bpm). These differences are comparable to results from a study by Sørhus et al (2016) where HR reduction (bradycardia) was observed in haddock embryos exposed to $7\mu\text{g}$ T-PAH/L (HR: 20 ± 6 bpm) compared to controls (26 ± 3 bpm).

4. Conclusions

The potential for computer-assisted automation to greatly speed up analysis tasks in toxicology studies has been demonstrated and realized. We have developed an automated image processing method for extracting heart rates from video data of early life stages of fish, and applied it to the embryo and larval stages of three different species. The ability to analyse existing video data, recorded with standard techniques rather than a set-up optimized for image processing is useful. Given the premise of non-optimized videos, the method provides decent performance both in terms of heart rate accuracy and heart identification. Applicability to toxicology studies was shown by identifying a 25% decrease in median heart rate between a cod embryo control group and a group exposed to produced water.

5. Acknowledgements

The methodology presented was developed and tested with material from the projects "RESOLVE: Comprehensive chemical characterization of the unresolved petrogenic components of produced water" (243720) and "EGGTOX: Unraveling the mechanistic effects of crude oil toxicity during early life stages of cold-water marine teleosts" (267820) funded by the Research Council of Norway.

References

- Brette, F., Machado, B., Cros, C., Incardona, J. P., Scholz, N. L., and Block, B. A. 2014. Crude oil impairs cardiac excitation-contraction coupling in fish. *Science* 343: 772–776.
- Eliceiri, K. W., Berthold, M. R., Goldberg, I. G., Ibáñez, L., Manjunath, B. S., Martone, M. E., Murphy, R. F., Peng, H., Plant, A. L., Roysam, B., Stuurmann, N., Swedlow, J. R., Tomancak, P., and Carpenter, A. E. 2012. Biological imaging software tools. *Nat. Methods* 9: 697–710.
- Faksness, L.G., Hansen, B.H., Altin, D., and Brandvik, P.J. 2012. Chemical composition and acute toxicity in the water after in situ burning – A laboratory experiment. *Marine Pollution Bulletin* 64: 49-55.
- Hunter, J. D. 2007. Matplotlib: A 2D Graphics Environment. *Computing in Science & Engineering* 9: 90–95.
- Incardona, J. P., Collier, T. K., and Scholz, N. L. 2004. Defects in cardiac function precede morphological abnormalities in fish embryos exposed to polycyclic aromatic hydrocarbons. *Toxicology and applied pharmacology* 196: 191–205.
- Incardona, J. P., Day, H. L., Collier, T. K., and Scholz, N. L. 2006. Developmental toxicity of 4-ring polycyclic aromatic hydrocarbons in zebrafish is differentially dependent on ah receptor isoforms and hepatic cytochrome p4501a metabolism. *Toxicology and applied pharmacology* 217: 308–321.
- Incardona, J. P., Gardner, L. D., Linbo, T. L., Brown, T. L., Esbaugh, A. J., Mager, E. M., Stieglitz, J. D., French, B. L., Labenia, J. S., Laetz, C. A., et al. 2014. Deepwater horizon crude oil impacts the developing hearts of large predatory pelagic fish. *Proceedings of the National Academy of Sciences* 111: E1510–E1518.

- Incardona, J. P. and Scholz, N. L. 2016. The influence of heart developmental anatomy on cardiotoxicity-based adverse outcome pathways in fish. *Aquatic Toxicology* 177: 515–525.
- Jui-Cheng Yen, Fu-Juay Chang, and Shyang Chang 1995, mar). A new criterion for automatic multilevel thresholding. *IEEE Transactions on Image Processing* 4: 370–378.
- Kessler, M., Rottbauer, W., and Just, S. 2015. Recent progress in the use of zebrafish for novel cardiac drug discovery. *Expert opinion on drug discovery* 10: 1231–1241.
- Klein, A. imageio - a python library for reading and writing image data. [imageio.github.io](https://github.com/imageio/imageio).
- Mikut, R., Dickmeis, T., Driever, W., Geurts, P., Hamprecht, F. a., Kausler, B. X., Ledesma-Carbayo, M. J., Mar´ee, R., Mikula, K., Pantazis, P., Ronneberger, O., Santos, A., Stotzka, R., Str´ahle, U., and Peyri´eras, N. 2013. Automated processing of zebrafish imaging data: a survey. *Zebrafish* 10: 401–21.
- Milan, D. J., Peterson, T. A., Ruskin, J. N., Peterson, R. T., and MacRae, C. A. 2003. Drugs that induce repolarization abnormalities cause bradycardia in zebrafish. *Circulation* 107: 1355–1358.
- Oliphant, T. E. 2007. Python for Scientific Computing. *Computing in Science & Engineering* 9: 10–20.
- Otsu, N. 1979. A threshold selection method from Gray-level. *IEEE Transactions on Systems, Man, and Cybernetics SMC-9*: 62–66.
- Pantazis, P. and Supatto, W. 2014. Advances in whole-embryo imaging: a quantitative transition is underway. *Nature reviews. Molecular cell biology* 15: 327–39.
- Pylatiuk, C., Sanchez, D., Mikut, R., Alshut, R., Reischl, M., Hirth, S., Rottbauer, W., and Just, S. 2014. Automatic Zebrafish Heartbeat Detection and Analysis for Zebrafish Embryos. *Zebrafish* 11: 379–383.
- Schindelin, J., Rueden, C. T., Hiner, M. C., and Eliceiri, K. W. 2015. The ImageJ ecosystem: An open platform for biomedical image analysis. *Molecular Reproduction and Development* 82: 518–529.
- Sørhus, E., Incardona, J.P., Karlsen, Ø., Linbo, T., Sørensen, L., Nordtug, T., Van Der Meeren, T., Thorsen, A., Thorbjørnsen, M., Jentoft, S., Edvardsen, R.B., Meier, S. 2016. Crude oil exposures reveal roles for intracellular calcium cycling in haddock craniofacial and cardiac development. *Scientific Reports* 6: 31058.
- Sørhus, E., Edvardsen, R. B., Karlsen, Ø., Nordtug, T., van der Meeren, T., Thorsen, A., Harman, C., Jentoft, S., and Meier, S. 2015. Unexpected interaction with dispersed crude oil droplets drives severe toxicity in atlantic haddock embryos. *PloS one* 10: e0124376.
- Tharwat, A., Gaber, T., Fouad, M. M., Snasel, V., and Hassanien, A. E. 2015. Towards an Automated Zebrafish-based Toxicity Test Model Using Machine Learning. *Procedia Computer Science* 65: 643–651.
- US EPA, 1986. Method 8100, Polyaromatic hydrocarbons. 1986. <http://www.epa.gov/wastes/hazard/testmethods/sw846/pdfs/8100.pdf>.
- Uchida, S. 2013. Image processing and recognition for biological images. *Development Growth and Differentiation* 55: 523–549.
- van der Walt, S., Schönberger, J. L., Nunez-Iglesias, J., Boulogne, F., Warner, J. D., Yager, N., Guillard, E., Yu, T., and scikit-image contributors, T. 2014. scikit-image: image processing in Python. *PeerJ* 2: e453.

Table 1: A summary of key data set parameters: number of individuals per stage and species, number of correctly identified heart regions (ROIs), and the relative median deviation of automatic to manually extracted heart rates.

Data set	Num. individuals	ROIs detected	ROI accuracy	HR deviation
Cod embryos A	25	19	76%	2%
Cod embryos B	-	59	-	-
Haddock embryos	-	60	-	-
Tuna embryos	-	53	-	-
Cod larvae	41	40	98%	2%
Haddock larvae	24	22	92%	7%
Tuna larvae	36	36	100%	3%

Figure Legends




Figure 1: Example images of the three fish species used. Top from left to right: Cod embryo, haddock embryo and tuna embryo. Bottom from left to right: Cod larvae, haddock larvae and tuna larvae.

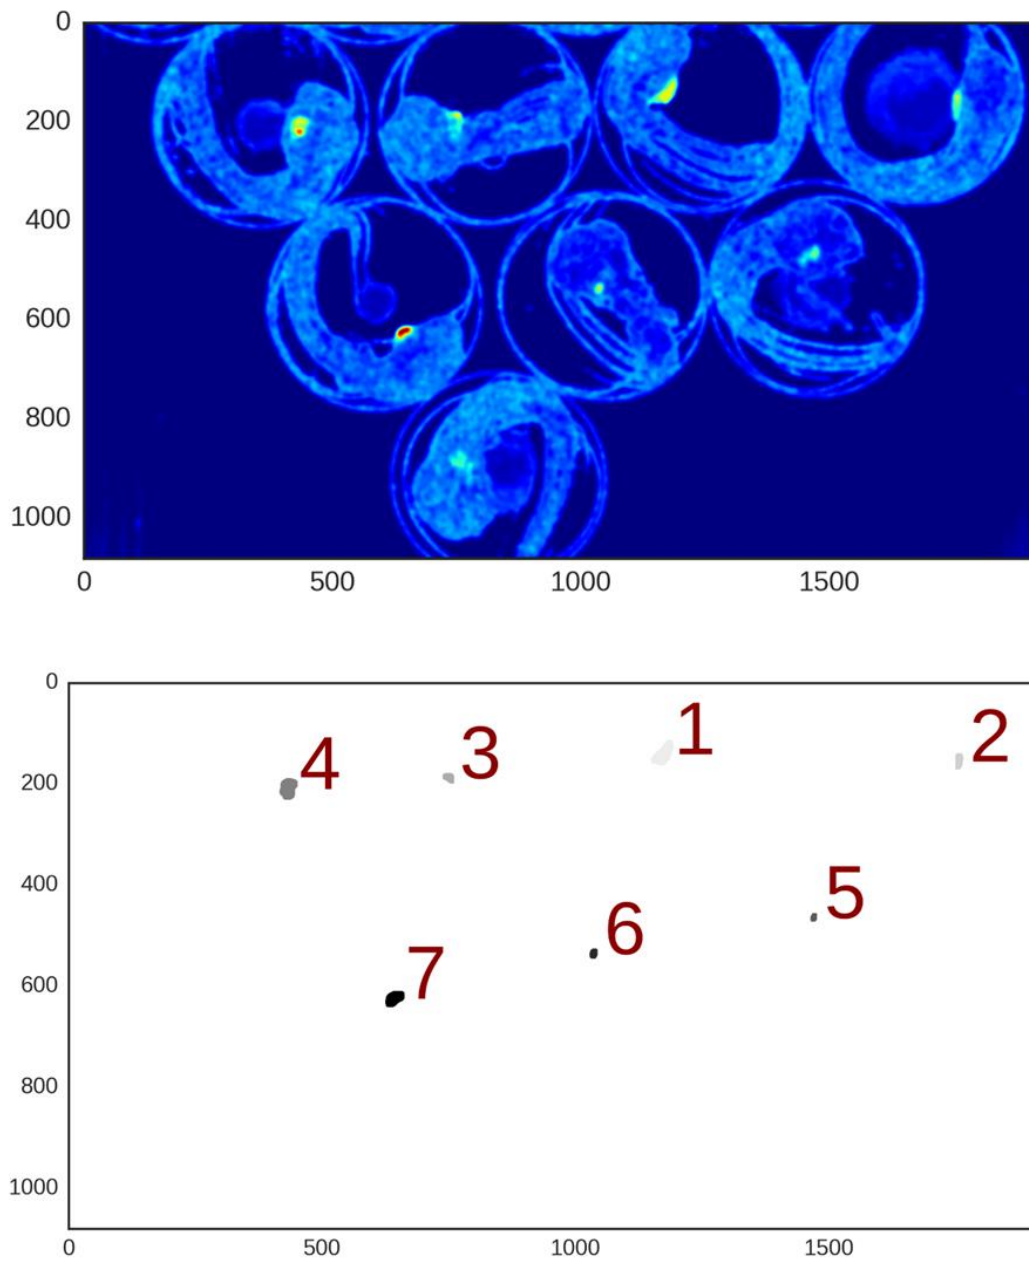


Figure 2: Top: Successive sum of frame differences (intensity), where red indicates areas of high intensity differences and blue for low intensity differences. Bottom: Automatically segmented and labelled heart ROIs.




Figure 3: Extracted signal for one identified embryo heart ROI, and the corresponding spectrum (Welch periodogram).

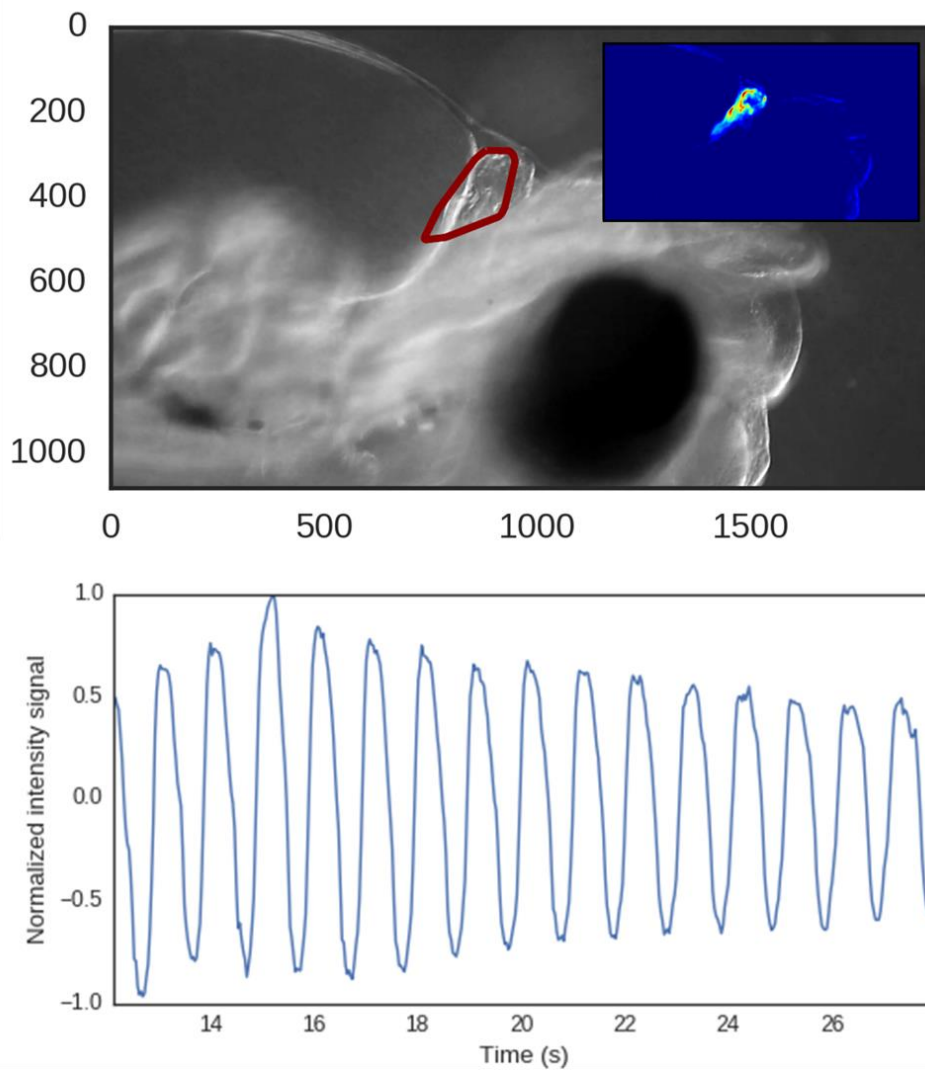


Figure 4: Heart rate region of interest identified from a video of a haddock larvae (upper figure), and the corresponding normalized intensity signal (lower figure). The inset in the upper figure shows the successive frame difference, highlighting the moving heart tissue.




Figure 5: Sum of frame-by-frame intensity difference, normalized to number of pixels (in the frame) and smoothed (thin blue line). This is used to remove time interval with too much general motion in the videos. See text for details.

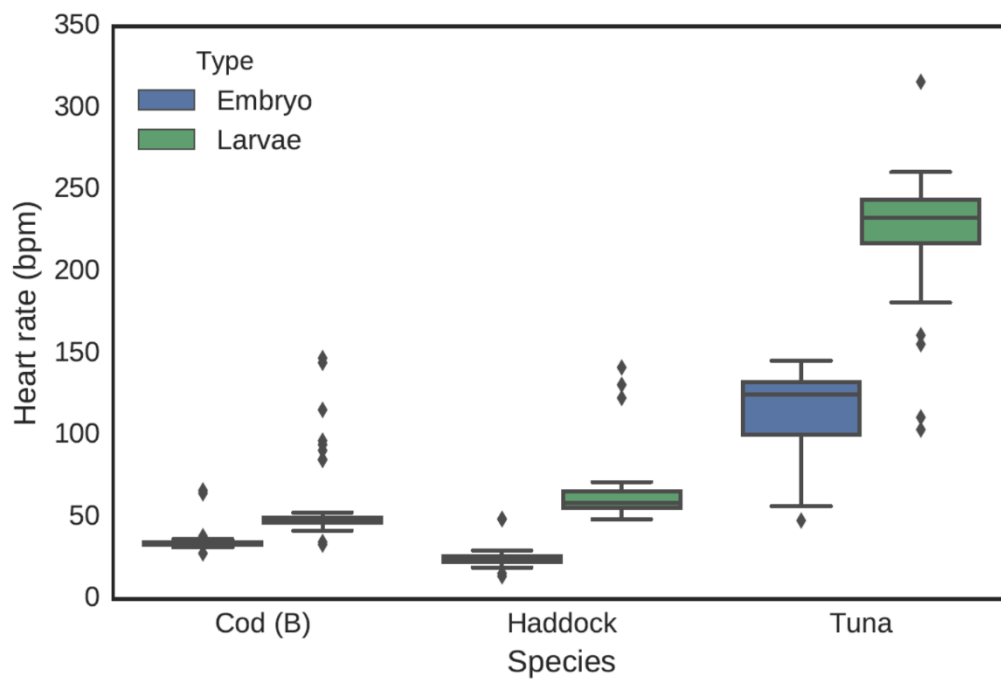


Figure 6: Distribution of heart rates for three species, two stages.

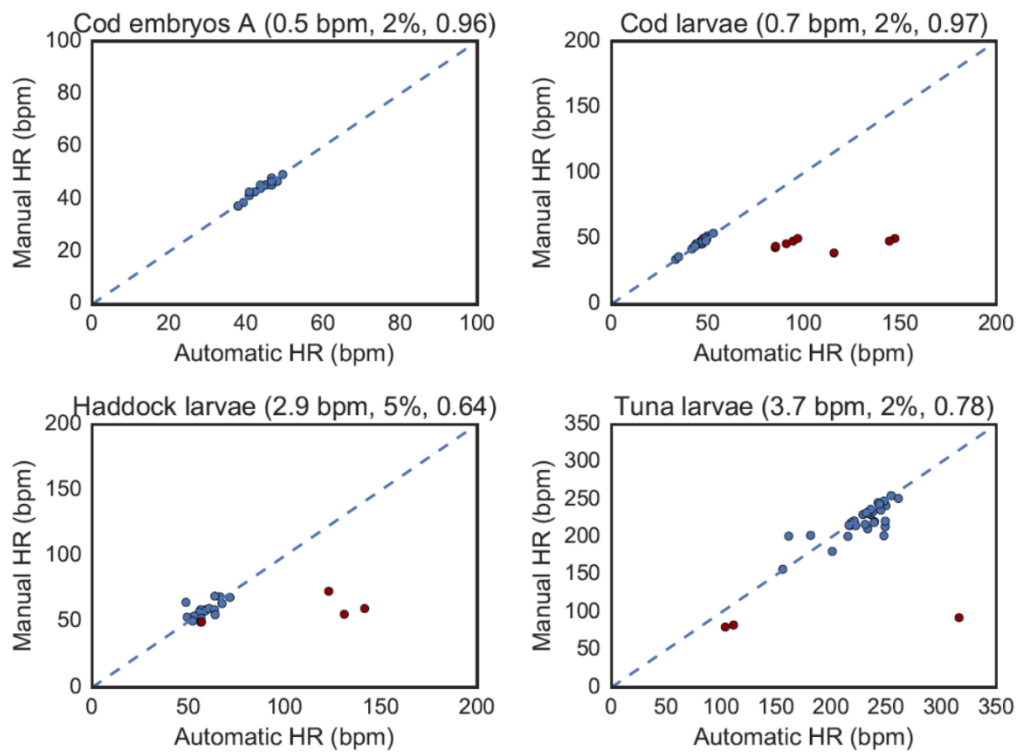


Figure 7: Correlation between automatic and manually determined heart rates. Median difference, relative median difference and Pearson correlation (subplot headings) are given for the blue points only. The dashed diagonal lines indicate perfect correlation, and have been added as a visual guide. See text for details.




Figure 8: Heart rates of cod embryos exposed to produced water vs. control.