modus vivendi

Project Fishface

Final Report

Project Fishface research, fieldwork and writing by:
Dr Malcolm MacGarvin,
Director: modus vivendi ltd
Ballantruan, Ballindalloch, AB37 9AQ
macgarvin@modus-vivendi.co

Fishface Final Report. Version 1.0 10th November 2015

Cover photograph John Dory, *Zeus faber* Falmouth Bay

The Lighthouse Foundation Hamburg part-funded this pilot for Fishface, a project that explores the voluntary use of video cameras on inshore fishing boats to better document their catches for various purposes. This Final Report sums up the results and possible next steps.

modus vivendi is grateful for this funding, and also to Dr Carl O'Brien for offering Defra, London, as hosts for the Workshop; to Cheryl Sykes (Cefas) for organisational assistance, and to all the attendees of the workshop for self-funding and attending a highly speculative event!

download link to low resolution version, ca. 1.8 MB

high resolution version (**suitable for assessing resolution of video camera images**), ca. 26 MB:

modus vivendi

Project Fishface

November 2015

Final Report

Summary

Video documentation can and has benefitted inshore fishers and others with an interest in the fisheries. **This *Fishface* pilot demonstrated that affordable consumer HD video cameras with GPS, mounted at a fixed location, and so not requiring an observer operator, have the potential to be—and in some cases already are—a valuable tool.** The work is also relevant to larger vessels. It is easy to identify most fish to species, to quantify numbers, catch rates and location, to evaluate discard rates and how the catch is handled. It is also possible to estimate size: work is required to establish achievable accuracy.

***modus vivendi* organised a workshop to discuss how *Fishface* might be scaled up, and yet cope with quantifying the large amounts of video data generated.** It included people working in fisheries, from government (Cefas/ Defra) and from the regional inshore management IFCAs. It also included technical experts in citizen / volunteer science and in automated computer vision—two routes for dealing affordably with Big Data.

On citizen science *Zooniverse* noted that ***Fishface* covers two areas of current interest to mass volunteer platforms: the use of video rather than still images, and the production of training sets for computers.** There is contingent interest in *Fishface* as a *Zooniverse* project. But also the interface and other coding is open source and *Zooniverse* welcomes third party usage. There is a simple *Zooniverse* 'build your own' project kit, although this currently wouldn't meet *Fishface* needs. There were various ways of dealing with location confidentiality, and if a rational case for volunteer involvement could be made to volunteers, neither the commercial interest in the results, nor the fact that wildlife—fish—were being killed to be eaten were necessarily an impediment.

Pioneering fish work involving computer vision, and the challenges to be overcome, were discussed at the workshop. The essential message was that computer vision was worth pursuing, but also a caution that it could take time to develop practical tools for routine use. **If computer vision is to be developed the most important starting requirement was for as many (i.e. tens of thousands) of identified training images as possible**—hence the citizen science linkage, and the interest in *Fishface*.

The workshop conclusion was that Fishface responded to various needs and the meeting discussed how this might be carried forward and what funding sources might be available in the UK. **Carl O'Brien (*Defra* Chief Fisheries Science Adviser) summed up best for the meeting when suggesting that the balance of emphasis was perhaps to prioritise getting video raw material, and then exploring the many potential ways it could be used.** The costs seemed relatively small compared to alternative ways of gathering the data (people, money, research ship time). The meeting also considered next steps. For the UK at least, *Seafish*, with the industry levy on catches seemed one obvious port of call when considering a funding bid.

Since the workshop was held in September **two important developments have occurred:** First, a **Garmin *VIRB XE* video camera is now available which essentially turns it into a 'plug in and forget' device with regard to**

the crew time required, the greatest concern of *modus vivendi* arising from the pilot.

Second **Zooniverse's *Chimp and See* has successfully used video clips** (rather than photos)**, drawing from a video bank even greater than that anticipated for the next stage of *Fishface*.** This includes developing a user interface essentially identical to that that which would be required for non-expert volunteers to enumerate *Fishface* videos.

A compact next stage, assessed in this report for practicalities of data handling etc., could be to **capture, archive and back up video for one entire year from up to ten fishing vessels, at a cost of £60,000**. It is suggested that, to avoid a hiatus, data, including training sets of identified fish are produced from this material and made widely available, at an additional cost of £20,000. These are illustrative examples: proposals would be discussed with partners, not least with fishers. There may be merits in going further faster and bigger, in which case *modus vivendi's* core interest is in delivering video capture, reception and storage, while also participating in characterisation and wider coordination.

Contents

The Pilot

Background

Pisces

Fishface grew out of a fishing industry and restaurant initiative *Pisces* (2004-14) which created premium markets and direct supply chains between the two. *Pisces* did on-board evaluations of fishing boats to provide independent backing to assertions of sustainability and handling. The evaluations included the species caught, their numbers and size, and the level of discarding. An on-board assessment was also made of wider environmental impacts and fish handling for food quality (important information for premium market buyers). **This required going out to sea on multiple occasions and documenting the fishery, including photography and video.** The video were primarily long uncut sequences rather than unedited highlights, intended to ensure that they were representative, and to help viewers to draw their own conclusions about the fishery as well as seeing the *Pisces* assessment. Discussion onboard with the crew was generally avoided (to keep the fishing trips as typical as possible). Nor is there any commentary. Nevertheless, viewing figures of the videos, placed on YouTube, was successful beyond the immediate needs of the fishers and restaurants. The first were added in June 2012, with most in 2013-14). By November 2015 they had generated over a quarter of a million views and 4.5 million minutes watched. Interestingly, although the videos hold nothing back, the ratio of likes to dislikes run at a ratio of 4.6:1. Although no new videos have been added for a year, the viewing rate has steadily increased to over 10,000 views per week. One reason for the relatively high viewing may be that fishers have little opportunity to go out on vessels other than their own, and unedited extensive footage is more useful as a resource than 'highlight' moments.

Fisheries Science Partnership Large Mesh Gill-net Project

The *Pisces* videos (as well of those created as part of a *Fishface* pilot) in turn proved useful as part of a Cefas *Fisheries Science Partnership* project FSP44 on the Cornish inshore large mesh gill-net fishery for haddock and other gadoids. This assessed the relative environmental impact (size of fish caught, discard levels, and wider environmental impacts) and relative economic returns against other métiers, and so provide guidance on the merits of awarding greater quota. Interestingly the skippers involved, many of whom had been involved in *Pisces*, asked for funding for full video coverage of the season. This was not possible, so they used the *Pisces* and *Fishface* video material instead to help argue the case. The FSP report concluded that the fishery had low discards, relatively little wider environmental impact, and that the boats supplying premium markets were gaining a substantial price premium per kg. but that they were severely limited by quota availability. This work strongly influenced (Chris Bean, *pers com.*) a decision by uplift 300 tonnes of additional gadoid quota towards the inshore sector—a significant result.

Hauling in 'Big Data'—without being swamped

Putting observers onboard boats is a time consuming and so an expensive process, even for a few trips per year that are unlikely to characterise all important aspects of the fishery. Much of the time and expense was absorbed as part of modus vivendi's *pro bono* work. Some vessels conditions are cramped for an observer, and ideal viewing positions difficult to obtain. Inshore day boats are vulnerable to weather disruption and are more likely to operate multiple fishing methods, both increasing the required observer time to gather data and costs. Little commonality in environmental effects, species composition and size, or discards rates can be assumed (or likely to be accepted by critics without verification) between different métiers operated by the same vessel e.g. large mesh monk tangle nets for monkfish, trammel nets targeting sole, large mesh gillnets for gadoids or smaller mesh targeting red mullet—even thought to casual observers one net may look much the same to another.

However, technology is rapidly developing, and costs are falling, for video cameras, video editing, storage capacity and data transfer speeds. Storage capacity potentially allows all fishing activity to be recorded. In such situations mounted video cameras, self-operated by the fishers, further reduced the costs. Potentially this is more credible that a few observer trips every year. Moreover—if challenged—the skipper can e.g. invite comparison of the video against *Registered Seller* landing data for the dates concerned. It makes it feasible to work with larger groups and gather more data per unit cost for métiers that provide a clearly identifiable view of (mostly) individual fish and size as they are hauled. Capturing GPS data adds value. Potential uses include stock assessments; evaluating discard rates; assessing the size of fish caught, discarded and retained. This could support the development of analytical tools for fishers, such as (with soak times and net characteristics recorded and individual nets identified) determining optimal soak periods, optimal life time of nets, data analysis of catch rates in different areas and (with linkage to other data) different conditions. It also provides promotional materials and concrete evidence of practices for both fishers and those buying the fish.

A trajectory going from specialist use for premium markets to widespread use can be envisaged, (see table below) with the evaluation of increasing amounts of data going from expert observers, through crowd sourcing/citizen science, to automated machine recognition. **Extensive *Fishface* video data already has commercial value, where a buyer agrees to buy at a favourable price from a vessel in return for access to the video material, allowing them to assure themselves, and respond to third party queries and criticism.** Citizen science is good for ground breaking research, but arguably may be less suitable for routine (but still essential) data analysis. Nevertheless, repetitive data collection is a feature of important volunteer projects such as wildlife atlases. As the number of identified fish in the database increases, this becomes a valuable training resource for developing AI (artificial intelligence) computer (aka robotic) identification of fish species, size, date, location, and other parameters from video. Ultimately the AI identification rate may become sufficient to do away with the need for routine video, with onboard realtime logging of the data created by companies such as e.g Garmin, who then have the prospect of integrating into their other marine recorders and analytical tools used by fishers.

Small scale premium market	medium scale, increasingly used by vessels with 'a good story to tell'	medium/ large Scale	Large Scale
Expert evaluation of video: can include species, size, location, CPUE, discard rate, wider environmental impact, handling & food quality	*Expert/Citizen Science evaluation of video*	*Computer vision evaluation of video. Multiple uses. Depends on training sets from experts/citizen science*	*Onboard real-time recognition & recording without video. Integration w. sounder & plotter data, gear characteristics & performance, & automated export e.g. for stock assessments, fisher Apps*
Already happening	**Current Fishface proposal, initial one year(?) project to gather video & trial evaluation**	**Current Fishface proposal to create a research resource** **5-10 years widespread?**	**10 years?**

Possible route-map, left to right, for use of video data. Some niche uses, such as broad assessments of discard rates, species caught and food quality premium have already utilised video and can be expanded, reducing the need and costs of on-board observers (left). Other uses, involving more routine identification and quantification, or the creation of training sets for computer recognition, require greater effort, either by experts or (more data) citizen science. This might be done as the next step. If computer vision could displace human effort this further reduces costs and increases value. A modest first step would be to gather video in a form so that this becomes a valued resource for subsequent development. Timescales are partly dependent on resources.

Leveraging 'Big Data' through Citizen Science

A project like *Fishface* would generate large amounts of video data – potentially 80 TB for ten vessels hauling 3 hr/day in one year (see later calculations). By way of everyday comparison, an *Apple* Time Capsule, used for back up, has a 2 TB drive. There are two interesting and novel ways (for fisheries science) to deal with the *analysis* of Big Data volumes. The first is to use Citizen (or Volunteer) Science. The second is to use the burgeoning field of Computer Vision, where computers self-learn how to identify items from photographs and video.

Citizen Science Many commercial fishers (and anglers) collect and analyse their own data. In this they are on a par with bird watchers and other naturalists. But the latter have been far more active in collaborating to create big data sets and make use of this data, producing population surveys and atlases in the British Isle and beyond since the 1970s. These have been amongst the largest Citizen Science collaborations before anyone had coined the term. Anxieties about data quality have generally receded and the benefits to policy makers, planners and others have become apparent. During the survey period of the latest British Isle 2007-11 Bird Atlas, maps were updated overnight every night. Broad results were evident almost in real time, rather than the years to publication for earlier Atlases. **This atlas is reckoned the 'single most important' bird publication in the British Isle in 20 years, and the inspiring BTO video shows the potential for *Fishface*.** Continuing monitoring of bird abundance and distribution is maintained e.g. via BTO's *Bird Track*, (and other projects) while recording of all species (i.e beyond birds, and including marine species) is increasingly co-ordinated via the *National Biodiversity Network Gateway* (NBN). This includes citizen scientist initiatives increasingly integrated into the traditional formal system, including *iRecord*, (to record species, validate records and get them into the national system) and *iSpot* (linking volunteers, and also experts, to identify species) as well as many regional and local initiatives. Much is based on available open-source coding toolkits, such as *Indicia*.

Fishers may be concerned about naturalists trawling through their data (or rivals identifying fishing marks). But biodiversity recorders' and database managers' concerns are not so different from those of fishers regarding confidentiality, data access, and sensitive locations becoming public knowledge. As a result, **there is generally a well beaten path governing confidentiality amongst biodiversity recorders, including voluntary observer projects**. In the UK the presumption has been that location should be recorded as precisely as practical in the underlying database, but both e.g. BTO and NBN essentially control the map resolution available to the public. In long term projects, such as the *National Plant Monitoring Scheme*, landowners' permission is also required by survey organisers. Public access to map data has traditionally been at the 10 km grid square level, with BTO and NBN vetting requests for data access at higher resolution, and a legal agreement regarding the purpose for which the data is to be used. 10 km has proved sufficient for many scientific purposes. For sensitive species the geographic resolution for public access may be 'fuzzed' even further. Delaying access to time sensitive data (by days, months, years) can also desensitise records without significantly harming scientific value. Releasing or restricting data access has benefits and disadvantages to all interests. More recently, bottom up citizen science platforms such as *iSpot* and *iRecord* have allowed entry and viewing of records down to 1m resolution with little or no control by the provider. This is a disruptive innovation: nevertheless the data provider still determines locational accuracy of data collected, and revealed.

Fishers, the data providers here, want their fishing marks to be confidential. Data on species abundance located within 2 or 10 km squares, and/or perhaps time delays in release of data, is likely sufficient to develop a better understanding of changing distributions and abundance, and so the better stock assessments and other developments that fishers desire. Subsequent data pooling in an area (i.e vessel identity unspecified), which increases anonymity, may actually increase the value of the data (including to fishermen, who see e.g. how their catch rate fits within local trends). Some vessels, certain they operate best practice, will want to be identifiable. Overall, going with the established framework (c.f. curated BTO, NBN), rather than an 'open access' route, would seem to be a sensible starting point, and as already noted fishers have control by virtue of being the data providers. There will likely be some hard discussions, reflections, and adjustments of position, but this is not an insuperable challenge. **Such improvements are needed for inshore fish distributions, both for stock assessment insights and biodiversity records**, the latter currently being aseasonal, and biased by accessibility of habitats to typical data collectors–see for example the scarce records for the NBN map for haddock.

On a different tack, *Zooniverse* is reputedly the largest citizen science initiative in the world. The historic *Zooniverse* project *Snapshot Serengeti* and the current *WildCam Gorongosa* feature a user-friendly identification interface for non-expert citizen scientists to identify and record various aspects of wildlife which is already very close to that which could be used to identify fish. Users click on a range of characteristics to narrow down the species choice, can check 'easily confused species' and access a field guide. They also record other aspects, such as numbers, species' activity and whether they are present with young. This interface could readily be adapted to include other aspects (such as size estimates—the code is open source, available for anyone to use and modify) These use trail camera photos rather than video. However a third *Zooniverse* project, *Chimp and See* is based on 15 second video clips. *Chimp and See* aims to better understand chimpanzee behaviour, as well as identify other species and document their activity, at 15 sites across Africa. The viewer first sees 9 still images to quickly judge whether there is anything in the video clip. If there is, the scorer proceeds to the clip, and identifies species (with the aid of a side-bar guide if necessary), their numbers and their activity. *Chimp and See* has ca. 7,000 hours of video (by comparison, 10 inshore vessels are anticipated to generate ca 6,000 hours of video; see later discussion). Between April and October 2015 over 1.5 million classifications had been made by *Chimp and See* volunteers. Each image will have been viewed and evaluated multiple times, which allows a statistical approach to be taken to flagging challenging images.

There is no reason why citizen scientists cannot do the same service identifying fish species and other data. This also creates the necessary large training set of identified images required for computers to work out how to identify fish …

Leveraging 'Big Data' through Computer Vision

In principle the features that people use to identify fish appears amenable to computer identification. They include clear (if sometimes subtle) colour combinations (varying even for closely related fish such as gadoids (cod and relatives)); from various body markings and shape; and from the ratios between fins, eyes, gill covers and/or other parameters. If the observers can judge perspective, they can also learn to determine length or weight with some accurately. In UK fisheries there are rarely more than ten species making up the bulk of the catch, which also keeps identification manageable for computer recognition. The East Anglian long line fishery for cod and thornback ray are dominated by two visually very different species. Unidentified fish can be flagged for later identification. These are less likely to be commercially significant, although they may be significant for other reasons. Identification from video, is easier—potentially for computers as well as humans—than from a single photograph, where the image may be distorted or hide a key feature. Video, unlike life, can be paused and replayed.

What is often not appreciated, and a barrier to the use of computers, is how much subconscious processing is required to interpret images that computers have to learn to replicate. However, as is apparent from e.g. the issues covered by recent IEEE (Institute of Electrical and Electronics Engineers) Conferences, ever more challenging and significant issues and enhancements are being addressed, many of which are relevant to the recognition of fish in fishing environments. This includes research relevant to handling distortions and perspective as fish are hauled over the side[1] ; distinguishing true object boundaries behind, e.g a

[1] Papandreou, G., Kokkinos, I., and Savalle, P. -A. 2015. *Modelling Local and Global Deformations in Deep Learning: Epitomic Convolution, Multiple Instance Learning, and Sliding Window Detection. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition:* 390-399. http://www.cv-foundation.org/openaccess/content_cvpr_2015/papers/Papandreou_Modeling_Local_and_2015_CVPR_paper.pdf.

mesh fence[2]; and otherwise sorting out from background clutter (in a paper that recognises the state-of-the-art importance of Convolutional Neural Network methodology, but also CNN's need for large labelled datasets, over 50,000 images in this case)[3]; practical issues, such as automated computer labelling of images—i.e. automatically collating and attaching created parameters (fish species, size, estimates of statistical probability, but also other data in file, such as GPS position and time) into the image captions[4] and/or into databases; and also the perplexing misclassifications[5] that still occur between successfully and unsuccessfully identified images, when only minute differences are apparent to humans (although 'acceptable', not 100% accuracy, is required for *Fishface*).

So the characteristics required to identify fish appears amenable to resolution by computers, at least as judged by research effort and direction. Whether 'appears amenable' becomes 'is amenable' has yet to be demonstrated.

Pilot Results

Detailed results from the pilot project are available in the *Initial Report* and are not repeated here. This includes detailed timeline documentation of videos that have been placed on a YouTube playlist, with a log of individual fish caught (species identification, size estimate, time into video, and other notes).

The results are summarised below under two headings. The first, *Does it Work?* deals with the immediate questions regarding the quality of the video for fish identification and recording other parameters. The second, *Practicalities*, deals with operational issues, such as the time and storage capacity required to save large amounts of video, and so whether it is sensible to develop the project. An assessment is also made of the amount of crew time that be required to operate the Garmin *VIRB Elite*, and whether this is too demanding. The equipment used, new in 2014, has already been supplanted by a significant update (the Garmin *VIRB XE*), which has positive implications on e.g. crew time required, so it is also important to read the later section dealing with this.

This deals with the pilot results: fieldwork already conducted. Potential future steps, such as citizen science and computer vision, are discussed later.

---

[2] Uijlings, J., and Ferrari, V. 2015. *Situational Object Boundary Detection*. arXiv preprint arXiv: 1504.06434. http://www.cv-foundation.org/openaccess/content_cvpr_2015/papers/Uijlings_Situational_Object_Boundary_2015_CVPR_paper.pdf.

[3] Cimpoi, M., Maji, S., and Vedaldi, A. 2015. *Deep filter banks for texture recognition and segmentation.* Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition: 3828-3836. http://www.cv-foundation.org/openaccess/content_cvpr_2015/papers/Cimpoi_Deep_Filter_Banks_2015_CVPR_paper.pdf.

[4] Vinyals, O., Toshev, A., Bengio, S., and Erhan, D. 2014. *Show and tell: A neural image caption generator.* arXiv preprint arXiv:1411.4555. http://www.cv-foundation.org/openaccess/content_cvpr_2015/papers/Vinyals_Show_and_Tell_2015_CVPR_paper.pdf.

[5] Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. 2013. *Intriguing properties of neural networks.* arXiv preprint arXiv:1312.6199. http://static.googleusercontent.com/media/research.google.com/en//pubs/archive/42503.pdf.

Does it work? Video quality for species ID & more

Ease of fish species identification The quality of the Garmin *VIRB Elite* 1080 HD video was excellent. It was as easy to identify fish from the video as in real time vessel-based identification. Identification was based on a combination of overall shape, colour and colour combinations, and characteristic specific markings and ratios. Subtile colour variations important for within genera identification in real life monitoring were also evident in the video material. Indeed, the identification success rate (and other parameters such as size estimation) is likely higher that real time on-board census data, because video can be reviewed, for example for species pairs, such as red and tub gurnard (with blue edge to pectoral fins) which can be difficult to see in a brief real life viewing. The typical flatfish species whether viewed from the top or underside could be distinguished by an expert observer by e.g. body shape and by subtile colour variations. Gadoids can similarly be distinguished by subtile but consistent differences in colour, and consistent markings, even when shape is contorted. Prior knowledge/ previous history/ likely accelerates the identification—even if subconsciously—for example where some species that cannot easily be distinguished from others have rarely occur at that location, date and or fishing method. Computer identification might use similar procedures. Errors in identification are more likely to be picked up if the species concerned is a significant part of the catch volume, so errors are most likely between similar species where some make up a small part of the catch (although this doesn't mean that capture has no conservation significance).

The Garmin *VIRB Elite* used in the pilot, plus waterproof housing, articulation arm and adhesive flat mounting pad. One euro coin for scale. The camera is inserted from the housing through the hinged front facing away from the camera. Note the video on-off slide on the top of the housing. The fixed joint articulation was not as easy to orientate as a ball and socket but was manageable.

Counting of numbers and evaluation of size/weight: suitable métiers The boats assessed in the pilot were all netters, which meant that fish would come over a hauler individually or a few at a time, and so at a fixed location on the vessel. In addition earlier video and/or still photography and/or earlier video was available for long lining, various potting/trapping techniques for crab, cuttle fish or other shellfish, and fixed nets for salmon. For all of these identification of species and counting of individuals would be highly feasible. For other documented methods, including commercial bass angling and squid and mackerel jigging, the fish come over the side at less precisely defined locations.

Top: Freeze-frame jpg extracted from the VIRB video. "Ultra-wide angle" setting. This is in a near ideal position both for giving a close view of fish as they come over the *Spencer Carter NHO-03* hauler, and also of other activities on the vessel allowing e.g. discarding and icing practices to be demonstrated. The position could be slightly higher, as fish are *marginally* out of focus at nearest point; their is slight crew blocking of deck activity, slight hauler blocking of discarding on left side. See *Initial Report* for a more detailed discussion. Note optional data overlay. **Bottom:** An example where freezing the image was required to distinguish (here) turbot from brill, as the fish passed rapidly through the hauler. The fish on the left is a male phase cuckoo wrasse.

Fixed location recovery via a hauler makes size estimation easier in a 2D video. The ease of size estimation depends on the hauler used. The UK market is dominated by Spencer Carter with two main haulers used by under 10s, the smaller NHO-01, and the larger NHO-03 which has a guide channel and tray, upon which the fish are particularly well displayed, with the fish leaving rearward. This channel and tray design provides multiple fixed reference points, assisting the judgement of size for an experienced assessor. Including graduated scales/ scored lines and/or distinctive reference points assessment might assist human and automated size assessment (colour reference charts might also assist computed species identification). Paint wear on the haulers indicates safer areas to locate scales, or a grid could be embossed/recessed. Notwithstanding the different orientations of fish coming over a hauler enmeshed in net, and the wide angle distortions of the camera lens, it is anticipated that an accuracy of ±5 cm could be achieved by an experienced observer. This may be sufficient to statistically resolve year classes of major commercial species to a acceptable and useful degree, but would need to be confirmed.

Screenshot from playback of an earlier , lower resolution, Sony compact video camera. The hauler is the smaller *Spencer Carter NHO-01*. The transmission route is 'up and over', and there is no tray to provide more reference points for estimating fish size. The distinguishing characteristics of the fish being hauled (here the dark 'thumbprint' on the side of the haddock) is sufficient to allow identification despite lower resolution of the video camera. Link to video etc. in the *Initial Report.*

Applicability to vessels over 10m Although this work was carried out on under 10m inshore boats, these conclusions are relevant to larger vessels operating these métiers. Indeed, sight should not be lost of the point that the evaluation cost/value of catch ratio may be better for some larger vessels, possibly making these candidates for trialling methodologies and for early adopters.

Métiers for which method might require modification Methods where large numbers of fish are discharged at once, and/or at multiple location such as rod and line, trawling or dredging might need a different approach. For example video taken on conveyor belts below decks of a large trawler, and multiple counting points for dredgers, may be viable approaches. For rod and line a different vantage point may be sufficient. The compact size and low relative cost of e.g. the Garmin VIRB may be an advantage over previous equipment that has been trialled. Opportunities and challenges may differ for expert, citizen science and computer assessment, and for the various parameters measured.

Demonstrating other aspects of vessel operation It was possible to capture a good view of the deck, allowing discarding rates and handling of fish subsequent to capture (for example handling time to storage in e.g ice slurry) to be demonstrated, while still having more than adequate resolution of fish for species identification and size estimation. One location point was was from the mast above the cabin. The quality would have been sufficient to allow identification and size estimation of fish, as well as evaluating discard rates and handling of the catch. However this would be less accessible to the crew, for example to clean the lens (and know this was necessary); to remove the camera for security; and exchanged SD cards. Instead the best location was judged to be on the rear wall of the cabin, close to the vessel side, above the hauler and above head hight, looking obliquely across the deck with a wide angle view.

Top: Alternative viewing point, 'super wide angle' setting, more vulnerable to obscuration during normal working operations. **Bottom:** crop from the full image of sole. See *Initial Report* for links to video and stills. Note the *Spencer Carter NHO-03* hauler transmission route to the rear of the vessel, and the tray that restrains the track of the fish.

Métiers other than netting vary in layout and ease of assessment. Fish from **long-lining**, like netting, come onboard at a fixed position. They may have fewer species, and be less obscured, that netting. Salmon from **fixed nets** are recovered at the same point on the vessel. **Shellfish potting** and **trapping**, here for cuttlefish, are recovered with a hauler, but swung onboard so may have a more variable recovery position. **Hand loop** for mackerel and **squid jigging** involves more mobility and there may be more than one fisher for squid, a fishery that takes place at night. Commercial **rod and line**, here for bass, may also have more than one fisher. **Trawling**, here beam trawling lands all the catch at once. **Dredging** lands the catch at once and at multiple locations. (all pictures Malcolm MacGarvin).

rod - both on same vessel

Long-lining

squid-jigging

mackerel hand loop-line

salmon fixed net

shellfish potting

cuttlefish traps

beam trawl

scallop dredge

GPS location tracking was good; selectable data such as vessel track, distance covered, time was overlay-able on the Garmin VIRB Edit software and could be burnt into the video for viewing on any video playback app. The GPS data (a generic .gpx file) was exportable to third party GIS mapping software (e.g. QGIS). This provides a bridge to analytical linkage with other marine databases.

Practicalities: data capture, transfer, storage and analysis

For any large scale use the ability to identify species and other parameters, as described above, is just the start of the process. It also needs to be practical to store, process and back up the large amounts of video data produced. This section makes an assessment of the practicality and the resource implications of the equipment used in the pilot. **Higher specifications that push the limiting boundaries are now available**; see the later discussion on the *VIRB XE*.

Selection of video camera The Garmin *VIRB Elite* was selected, as this included HD 1080 video, GPS, a waterproof housing, and multiple means of mounting. GPS is crucial for many potential uses, ranging from stock assessments, through connection to other data sets, and the development of future tools for fishers to explore catch histories for better future catches. The then top of range *Elite* (around £350 with waterproof housing and mounting equipment) used could store just under 7 hours of HD video on up to a 64 GB micro SDHC card (c.a. £30 for a fast card). The alternative, the GoPro range, does not include GPS. Car dash video include GPS, and are cheaper, but lack of weatherproofing would restrict them to within cabin use, which is too limiting.

Battery life Shooting continuous video drains battery life fastest. One battery, used to record HD video continuously at ≈ 15°C, and also recording GPS data, lasted 108 minutes. The *VIRB* automatically starts a new file every 23’57” and ≈3.67 GB, with the remainder in a fifth file. These can be run together without loss of video. One of two of the lithium ion batteries failed, no longer take a full charge after a couple of cycles. Batteries cost £15-£20.

Card to computer transfer speed Extracting the video from the camera to a 2009 then high-end Apple MacBook Pro 17 inch took 1.3 minutes per GB [108 mins, 16.82 GB, in 22 minutes, i.e a transfer rate of 12.74 MB/sec]. Transferring ca. 3 hrs of a boat’ s net hauling and clearing activity during a day was easily done in 2014, taking 36 minutes from camera card to computer. Every hour of video required ≈9.34 GB of SD card storage. As files on the Garmin VIRB *Elite* are automatically broken down into segments of 23’57”, copying from cards to disk could be done in the background while working on earlier video. Current high end desk top machines, such as the Mac Pro, used with fast SD cards and USB3 transfer can be expected to achieve a real world transfer speed ten fold faster (128 MB/s) and have claimed transfer speeds of up to 500 MB/s. *These high end machines would be required to ensure that card > disk transfer did not become a bottleneck when processing larger numbers of vessels.* Note that current maximum flash card transfer rates, of 160 MB/s, are not (yet) available for micro SD cards—as of October 2015 the maximum was 90 MB/s, so card to computer transfer rate is the current rate-limiting factor but still approaches 10x faster than the 12.74 MB/s achieved in the pilot. Multiple parallel transfers can be performed on high end machines.

Storage capacity required for one boat for one year Backing up a video set to a second hard drive is much faster than taking data off a SD card, typically taking a few seconds per GB on high end machines. Assuming a vessel hauls for 3 hours per trip, works 5 days a week for 40 weeks of the year (allowing for bad weather, tidal conditions and boat maintenance) that is 6 TB of video per year, doubled to allow for a backup copy. Put another way, 10 under 10m vessels hauling 3 hours per day might be expected to generate 2 TB of data per month. Consumer (i.e lower cost, lower data transfer speed, nominally lower reliability/ life-time) hard drives up to 8 TB, the largest amount of storage likely required if they fish longer each day/more weeks, are available. Thus storage requirements are significant, but not unfeasible, at around £600 per boat for two 8 TB drives (i.e. including back-up). Using typical high speed RAID arrays (typical for professional video editing) with faster read/write speeds and higher levels of data corruption security would be around £1,000 per boat including back up.

Video editing and viewing Scrubbing (rapidly scrolling forward or backwards, changing position in video) using the dedicated Garmin software produced stuttering and pauses that slowed down the enumeration of the data using 2009 equipment. This is something that might improve with state of the art hardware. Nevertheless, viewing the original video files in other software, e.g. Quicktime (i.e no conversion or exporting of files required) eliminated these problems even on the older machine. This is significant because it means that third party enumerators (experts or citizen scientists counting fish!) who simply view video do not need state of the art computers to participate.

Video division into short clips It may be necessary to divide the video into clips, for ease of scoring by third parties using internet connections, or as part of a randomised evaluation/ or anonymisation process for those counting the fish. Sufficiently short clips may also mean that time and location of fish can be recorded just once for the clip (for association with fish density, time of year etc.), the significance being that the amount of data capture required by human hand is a rate limiting factor for the number of boats that can be handled. Marking up (i.e indexing) and then (if required/ considered desirable) breaking up and saving clips in different files, is best done in a professional video editor. Adobe Premier has been tested, and this allow batch exporting of the video clips from the original without further involvement. Five minute clips (perhaps suitable for expert enumeration) would require 36 clips to be marked up in 3 hours total video. Zooniverse's *Chimp and See* uses 15 second clips, which has been suggested as more appropriate for mass citizen scientist involvement. This would require 720 clips, which would require an automated process.

Hard baking data overlays into video Exporting the video from within *Garmin VIRB Edit* software allows, if required, the visual 'hard-baking' of chosen data overlays onto a new 1080 HD MP4 video file viewable by Quicktime (and other standard video viewers/editors). This took 5x the length of the file on the 2009 laptop, and 2x the length on a 2009 desktop MacPro. This would be faster on state of the art equipment. The conversion resulted a small but acceptable loss of resolution video. This would be useful if observers needed to observe data such as time and location of capture from generic video viewing applications (although *VIRB Edit* is a free app) to overcome *VIRB Edit* stutter (see above).

Upload speed to e.g. YouTube Using channels on services such as YouTube are one means of distributing video to observers, and incidentally currently a free means of backing up data and making the archive accessible without funding or maintenance. Uploading 1080 HD video to YouTube, with a transfer rate of ca 0.5-1 MB/second (perhaps a UK average upload speed) took one and a half times the running time of the video clip. So this would be a rate limiting factor for distribution. So if internet distribution to viewers was used, this would likely require batch uploading at a high speed location, or the mass transfer of data on disk to cloud storage (such as the Amazon S3/Snowball option described later).

Practicality of viewing high definition (HD) video over internet Viewing 1080 HD video on YouTube required a transfer rate varying between 0.3-0.5MB per second to be viewed and to build up a buffer to avoid pauses. YouTube allows video to be played in slow motion or accelerated by variable amounts up to 2x life speed, useful for reducing handling times during periods where no fish are caught (it is assumed that non-hauling time is not included in e.g. YouTube video produced for counting).

Practicalities: crew time

Another important point is amount of crew involvement required. The *VIRB Elite* used in the pilot can be wired into a boat's electrical system, but it is an unofficial bodge, and would also require drilling a hole in the housing that might compromise waterproofing. Battery operation does not compromise sea-proofing, but would require at least one battery change to record three hours of hauling. This would require working gloves to be removed, and damp hands in some operating conditions may overwhelming the silica gel within the enclosure that prevent lens misting. It would also require two batteries to be charged overnight, and also require silica gel to be placed into the waterproof housing overnight (the internal desiccant pads cannot be removed easily once installed).

The crew would then need to remember to turn on the VIRB at the start of each hauling episode, and to be alert to the red flashing recording light stopping, which indicates that the SD card is full or the battery is flat. It is likely that these will be overlooked on occasion: the aim would be to keep this within bounds that do not undermine credibility, for those aspects of data collection where this is relevant. They may also face criticism that e.g. discarding was done during periods when the camera was turned off (although they would be able to call landing data in aid). Lack of continual recording also means that the skipper will need to note e.g. gear deployment times and start/end locations if catch per unit effort/gear deterioration with age/gear drift and drag are documented. Once on land the *VIRB Elit*e would be removed from its mount (easy) the SD card usage checked and returned every couple of days (with a 64 GB card) in a postal shuttle with sufficient cards to ensure that a replacement is always available.

Such levels of involvement are not impossible. Crews do spend equivalent amounts of time on other research, for example measuring fish. The initial boats recruited could be expected to be highly motivated, believing they have a good story to tell. *However* uptake may be higher if there is financial compensation - direct payments and/or relaxation of quota or other restrictions. **Nevertheless the crew time required was the greatest concern identified by *modus***

***vivendi* in the the pilot**: data recording should not interfere with normal operation, even if the vessel stands to benefit financially from their installation.

Significant practical developments since the pilot

Garmin VIRB XE

However, a new model, the *VIRB XE*, released in Summer 2015, resolves virtually all of these potential boat operation issues. The *XE* can be connected directly into the boat's electric system with a Garmin bare wire connector kit at the boat-end of the connection. It no longer requires an enclosing waterproof housing for operation at depth (and hence sea conditions), and has a hydrophobic lens coating (i.e water droplets are more likely to run off). The maximum SD card capacity is increased to 128 GB. There are also some subtile but important features regarding management of video parameters. For example white balance can now be manually set and be invariable, which means that colour can be better characterised.

Overall, this means that, with the 14 hours HD video capacity on a 128 GB card, and the VIRB cradle wired to the boat electrics, the skipper simply plugs in the camera at the start of the day, the camera comes on when the boat is powered up, and switches off with final power down. The skipper then removes the camera (to avoid theft) and on land each day swaps over the card with a replacement and puts it in a stamped addressed envelope for return. The most likely crew involvement during the sea trip is reduced to checking that the lens remains sufficiently clean. Moreover, additional data, such as gear setting times and location can then also be taken off the VIRB video upon return of the card, rather than requiring the skipper to document this. Complete coverage (beyond hauling) means that boat operators could demonstrate fish handling, and counter accusations of discarding at points other than hauling. The value of video storage of the entire trip is debatable (although the potential for unexpected uses should not be discarded - for example verification of satellite estimates of cloud cover or wave state. Possibly a reduced resolution video may be sufficient for such uses without making storage impractical).

Storing data off-site

Cloud storage, such as Amazon's S3, at the time of writing costs $0.03 per GB, so accumulating an *additional* $60 per month for 2TB maximum video per month. Amazon also provide a physical hard drive (Snowball) to deliver data to their data centres, at $200 plus standard courier costs. Each unit has capacity of 50 TB, compared to the possible 80 TB annual storage capacity demand for 10 vessels. There is also a reverse process for returning data from Amazon's cloud to onsite or other storage. There are other cloud storage capacity providers. Physical transfer for uploading to cloud storage helps resolves the issue of uploading video for (private or public) remote viewing.

Other options may be available, for example if working with a university or institutional partner.

Workshop Discussion

The workshop included representatives from the fishing industry, the Inshore Fishing and Conservation Authorities, Defra/Cefas, from a global citizen science initiative and experts in AI/ Robotic Vision research (see table below).

name	organisation	interest
Carl O'Brien	Cefas/Defra - Chief Advisor, Fisheries	multifaceted, but the appeal of better data in multiple ares; and the potential cost savings / freeing up of resources currently used for stock assessments, research vessel time.
Simon Pengelly	Southern IFCA	application to inshore research; stimulating economic returns of local fisheries
Jerry Percy	EU LIFE project, New Under Tens Fishermens' Association	Utility for fishers
Grant Miller, Ali Swanson (not able to attend but followed up with call and emails)	Zooniverse	Citizen / Volunteer Science
Sam Devlin	University of York (also representing Simon Hickinbotham and Adrian Bors)	Working with games developers for application development. interfaces for engaging citizen scientists. Automated vision. Linkage to other fisheries science
Krystian Mikolojczyk	Department of Electrical and Electronic Engineering, Imperial College, London	Computer vision
Mark Fisher, Michal Mackiewicz & Geoffrey French	University of East Anglia Computing Department	Computer vision. Already working with Marine Scotland on demersal trawlers.

The results of the pilot work were presented. As per this report, it was suggested that the technology had advanced to the point where voluntary use of video cameras on inshore day boats to quantify fish numbers, species and size, as well as other features such as discard rates among others, was a practical proposition. It was suggested that the technology was near 'ready to go' for replacing the routine and expensive use of on-board observers for specialist and premium fisheries. It was also suggested that citizen science initiatives could increase coverage, as well as help produce the large number of 'training sets' of identified fish images for training computers to recognise and size fish species; that the use of computers for identification was becoming a practical proposition.

Potential User Reactions

Of the potential users of the data, **Jeremy Percy** thought that the use of video could be controversial with fishers, but if it was dealing with issues such as demonstrating low discard rates that would be in its favour, as there was currently very little documentation of such issues. Issues such as confidentiality / anonymity needed to be clear, and a working balance found with making the video available, when, and to whom. Malcolm MacGarvin (noting *Pisces* and the recent *FSP* haddock under-10 gillnet project) suggested this would likely depend on whether the fishers thought they would be likely to benefit, and that some fishers saw did see this as an opportunity and were requesting video coverage. Beyond perhaps some minimum ground rules for project participation, agreed with fisheries leaders and others beforehand (and participation being the fisher's choice), fishers would be in control of their own data—for example they might choose to make it available to to potential fish buyers. **Simon Pengelly** thought that *Fishface* was an interesting development and that IFCAs could find this very useful. Southern IFCA were willing to assist in future development and encourage boat participation. **Carl O'Brien** noted the general lack of knowledge for the under 10m fleet, and of changes in fishing patterns. Generally, the resources demanded for issues such as gathering data for stock assessments, both of people and of research vessel time, were significant. If *Fishface* both reduced these costs and provided valuable new data (and/or increased data plus a means of handling it without increasing the burden of work), that would be welcome. **If the costs for a year long pilot were, as suggested, around £60,000** (see later)**, that was not a huge amount of money for the potential benefits.** Carl suggested prioritising gathering the data first, and work out what to do with it later.

Citizen Science: What works?

Regarding the possible use of citizen science, *Zooniverse* is the world's "largest online platform for collaborative volunteer research", involving hundreds of thousands of people around the world. **Grant Miller** was due to attend the workshop, but in the event was unable to. Instead *Fishface* was discussed before and after the workshop with Grant and **Ali Swanson**. *Zooniverse* has an interest in *Fishface*, as it deals with two aspects – the use of video, and the production of training sets for computer recognition – that are areas that are of active interest. As already been noted *Zooniverse* has recently started using video as well as still images in the projects with which it is associated. For example *Chimp and See* uses fifteen second video clips, using an interface that is very close to that required for *Fishface*. **It was suggested that fifteen second video clips was likely the optimum length to engage and retain e.g. *Zooniverse* volunteers.** Experience suggested that identification guides and other features should be available, but not dominate the site. Also lengthy training sessions may be counterproductive; rather people like to dive into classification, and issues such as accuracy can be dealt with in the background.

One could say that *Fishface* consists of hours of video of animals being killed, which raises the question of how might the *Zooniverse* volunteer

community react? It was speculated that this would depend on the purpose. If it was evident that the purpose was to increase the sustainability of fisheries; to see more sustainable fisheries prosper; and to develop future techniques (such as computer recognition), and that this could not be done without mass volunteer effort to produce training sets, then that could be appealing to participants. It might be that the volunteer community would be most engaged by an important development phase, dependent on mass volunteer involvement, rather than routine monitoring. On the other hand you could say that volunteers are attracted in large numbers to the production of Bird Atlases, and monitoring work that depended on constant recording effort (c.f. http://app.bto.org/birdtrack, the *Cornell Lab of Ornithology* – Cornell already includes automated photo ID. The same is true for other taxonomic groups. The project might also continue by including new parts of the world.

Confidentiality and anonymity. The amount of locational and date information provided can be determined for each project: in various *Zooniverse* projects there is another partner that takes the lead in 'facing' the project, for example updates and blogs, and also taking the lead on decisions regarding anonymity. Here individual fishers and coastal scenery will appear in the video, so the material is not strictly anonymous even if the vessel and location are retained. This might be a sticking point for some fishers: it would depend on how confident they were on the relative merits of their fishery, and how they anticipated the video might be used by others, for good or ill. On the other hand, if *Fishface* was conducted in different parts of the world, the videos provided could be regionalised by *Zooniverse* so that volunteers only saw videos *outside* their region, should this be considered important.

So, to sum up, *Zooniverse* would consider *Fishface* as a potential project, and it might also consider putting in some development work with regard to features such as video and producing training sets for computer recognition, though funding would need to be discussed. However, much of the coding for *Zooniverse*, including *Snapshot Serengeti*/ *Wildcam Gorongosa* and at least some of the *Chimp and See* code, is freely accessible and adaptable on GitHub. So it was also suggested that it was not necessary for *Fishface* to be dependent on direct *Zooniverse* involvement.

Sam Devlin noted that there was diverse cross-discipline expertise at the University of York *Ron Cooke Hub*, including fisheries scientists, and people with a background in app and game development. This includes those working on Complex Systems Analysis and at the Creative Technology Centre. Sam had an interest in creating interfaces that work well, drawing on game development skills. There is already a project underway that seeks to transfer skills and learning gained from games development. Simon Hickenbotham and Adrian Bors had experience in fisheries research and in computer vision. *Fishface* fits well with these interests and they would like to be kept in the loop.

Teaching Computers to Count Fish: What needs?

Krystian Mikolojczyk provided a summary of how Machine Learning, alternatively Computer or Robot Vision had developed. He noted that a prominent **technique known as Convolutional Neural Networks (CNNs) was the most promising for fish recognition**. This essentially lets the machines work out how to identify things, rather than humans attempting to teach them. However they require 'training sets' of thousands of pre-identified images, (see earlier discussion): the more the better.

Mark Fisher, Michal Mackiewicz and Geoffey French agreed that CNNs were the best way forward. They then explained their work on Computer Vision onboard Scottish demersal trawlers with Marine Scotland. Some members of the fleet have been gathering video footage for some years as part of a compliance programme, in return for enhanced fishing access. Video cameras are in various locations. The UEA Computing Department worked on footage below decks, as the catch comes down, reasonably well spread out, on a conveyor belt. Access is cramped, limiting camera location points. **Each boat has a unique layout and installation, and the footage is relatively primitive low resolution analog video. In the event, only one vessel had footage really suitable for attempting Computer Vision.** The fish are dead (i.e immobile). The area of the conveyor is isolated in the video analysis, although e.g. hands and arms intrude into the recognition area. Analysis is effectively that of individual frames in the video. **The robotic vision learnt to distinguish fish from non-fish, and to distinguish between some species. Distinguishing other species was more problematic, for example flatfish presenting white underside up.** Colour balance, with analog and the under-deck available lighting for video, was problematic, restricting the reliability of colour as a distinguishing feature. In this iteration the algorithms for identifying fish were unique to each vessel surveyed.

The general discussion then concentrated on the applicability of Computer Vision to fishing vessels that operate métiers where fish came over the top side of the vessel one or a few at a time—specifically haulers operating nets and long-lines. Other such methods, such as haulers hauling pots or traps, and commercial rod and line angling (no fixed position on deck) were noted but not pursued. The discussion was focused on under-10s, but it should be noted that over 10s also operate these methods. It was noted by the workshop that the video quality from the *Fishface* pilot was good. Fish were readily identifiable to the human eye. It was noted that the design of haulers was relatively consistent, which enhanced the prospects of generic recognition across vessels—although there appeared to be a non-standard design of hauler for the long liner catching cod and thornbacks featured in the presentation (also shown in the metier photo compilation). Long-lining appealed to the computer vision experts because of the simplicity of the species composition, lay-out, and marked differences in fish shape and colour.

The computer vision experts thought that this was a worthwhile project. However they also noted various cautions: the camera location would not be exactly identical between vessels, and (fixed) colour guide charts may deteriorate in operating conditions. Moreover the colour balance settings of video camera were often uncontrollable and vary with lighting conditions [note that the 2015 *VIRB XE* does allow control and custom setting of colour balance and other image parameters]. The distortion of fish body shape in nets, and the colour and presence of nets themselves were both complicating factors. It was also cautioned that the total storage and handling capacity for video needed to be taken into account [these appear manageable for the pilot and beyond, see earlier discussion]

It was suggested that if possible **the visually simplest fisheries should be included, such as the East Anglian winter long-line fishery**, that predominantly catches cod and thornback rays, be included in the pilot.

The importance of sufficient human identified training sets was also emphasised. It was noted that **Marine Scotland had found expert fatigue / boredom was a factor in the production of training sets**, so had switched to many experts doing relatively short identification sessions.

Next Steps

There was a consensus at the workshop that the initiative was worthwhile, and a wish for participants to be kept in the loop with future developments.

The next step would be to pursue funding for gathering video from fishing boats in a proposal that the fishers involved see as beneficial. The indicated costs in the original *Fishface* proposal, of up to £60,000 to gather a year's video from 4-10 vessels (depending on how effort was divided between securing data and identifying fish in that data) appears reasonable given the potential benefits. It was suggested that the balance of effort was perhaps best spent more towards securing the data and developing the practicalities, over actual identification at this stage.

Post workshop, the cost/time implications of minimising significant crew time/ monitoring (and potentially cost/ increased errors) by using Garmin *VIRB XE*s was explored, with a daily shuttle of 128 GB SD cards recording from engine start to stop, and then extracting fishing time video. Maxing out computer and RAID array hard drives to cope with a maximum of 10 vessels delivering full video adds £5-7.5k, There would be extra daily handling time rendering the video down to fishing time (or otherwise have ≈4-fold data storage requirements). A maximum of ten vessels can still be handled for £60k, for one year's fishing activity, requiring three quarters of a working year equivalent for one person, but this comes at the cost of no identification or counting of fish, or production of training sets ready for future steps. The presumed (significant) benefit is a larger pool of fishers willing to participate, and fewer dropping out (both due to crew time required) and fewer losses of data. Putting identification and counting back in, by increasing expert time up to one person/year equivalent, and other costs of this aspect, could increase total cost to £80,000.

In the UK it was suggested that *Seafish* was the obvious body to approach regarding funding, and that both fishers representatives and those of Cefas/ Defra would make them aware of this work and its potential.

modus vivendi will continue to explore and work up funding possibilities, develop proposals and keep those with an interest informed, and is open to all available collaboration and assistance in this process! *modus vivendi's* interest could be either developing a self contained small scale project of up to 10 vessels for one year (as outlined in this report) or—if there was enthusiasm for a larger multi-partner project from the outset—to offer to handle the on-vessel video gathering aspects, along with reception of the video and initial processing; and to be involved in the onward passage of data into storage, involvement in fish identification and the creation of training sets, and generally participate in project coordination and development with partners.

For practical reasons this project is discussed in the context of the UK and the inshore fleet, but in principle the concepts have wider utility and relevance: interest has already been expressed by potential North American partners and by FAO staff.

modus vivendi