Fish4Knowledge Deliverable D7.9

Final 6-monthly report to EC

Principal Author: UEDIN
Contributors: All partners
Dissemination: PU

Abstract: This document summarises the project activities in months 25-30.

Deliverable due: Month 30

1 Administrative Progress

The main administrative activities were 2 project consortium meetings December 12-13, 2012 (Luxembourg) and April 21-22, 2013 (Taiwan). Minutes and PDFs of all meeting presentations are available on the project "Members Only" pages.

The main administrative topics were the overview of the technical achievements and associated deliverables. Several technical working meetings were planned. The core technical goals of the December meeting were: integration planning, getting system fully working, performance issues, getting the user interface and workflow control working. The core technical goals of the April meeting were: performance issues, getting the user interface and workflow control fully working, data processing completion, data and code archival.

The project web site has been updated with recent publications and deliverables.

We have continued to promote the project. A poster summarising the project was presented April 15-17 at the 21st European Association of Fisheries Economists in Edinburgh. Papers will be presented at the 9th Indo-Pacific Fish Conference (Okinawa) and 48th Annual European Marine Biology Symposium (Ireland). We also held a full day meeting with the Taiwan Marine biologist community April 21, 2013 (see deliverable D6.5).

We are proposing to write a book describing the project, and are currently discussing oppertunities with publishers.

2 Technical Progress

This section summarises the progress made by each of the different teams during project months 25-30.

2.1 CWI team achievements

User Interface for ground-truth collection:

- Developed game-based labeling interface for fish species recognition/validation as image-matching game.
- Game requires selection of "most similar" textbook fish species image to image extracted from the video footage:

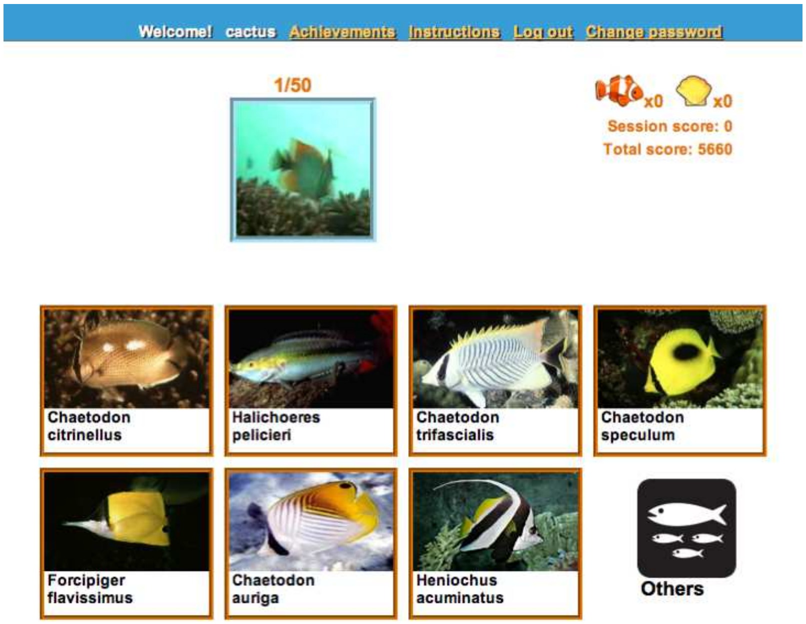

- Game includes competition elements to increase user engagement, and thus label more fish: top 10 scorers (in a single labeling session) and contributors (over all sessions).
- Study on quality of labels from lay-users compared with marine biology experts were carried out.

User Interface for data exploration

- Developed a first interface prototype to a complete end-to-end system, with access to all data. Prototype allowed users to select specific subsets of the data, and visualize fish counts over time.

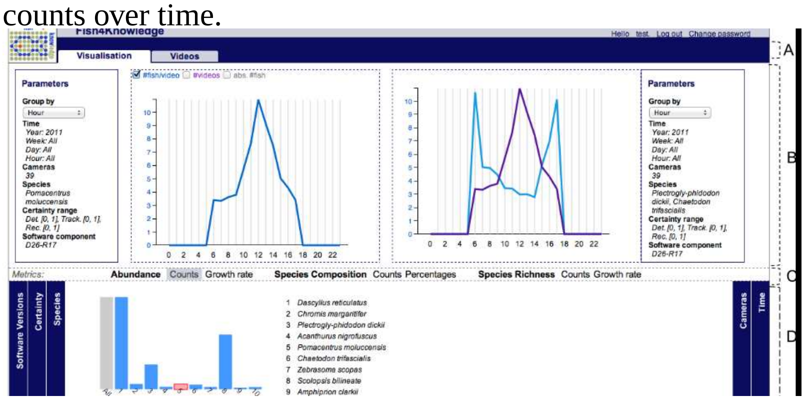

- Colected feedback on prototype, detailed in D2.4, lead to re-design of UI. Emphasis on improved UI on understanding the various levels of processing and their impact on the data being visualized.
- Conducted user study on user understanding of analysis results and submitted results to ACM MultiMedia 2013.
- Developed and evaluated a series of paper prototypes using heuristic evaluations with external experts.

- Implemented a second prototype and presented it to marine biologists in Taiwan. Prototype accessible at: `http://f4k.project.cwi.nl/demo/ui/`:

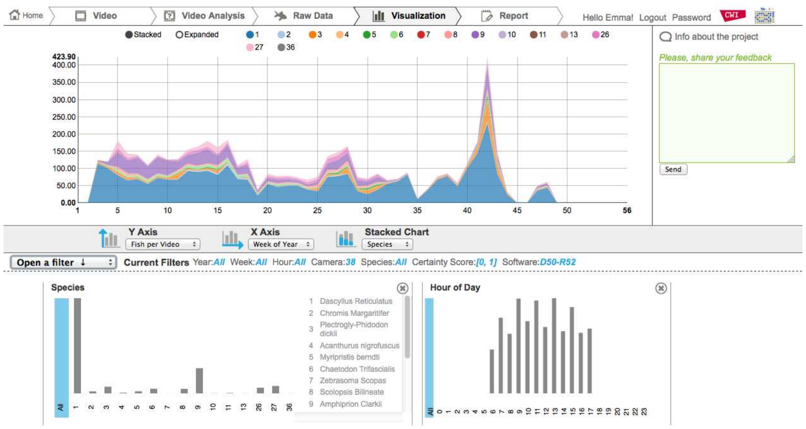

2.2 NARL team achievements

- Experimented with different formats and bitrates for storing videos
- About 677k videos in database, 33% of disk storage used
- Extended job control interface for all parallel machines. Model is:

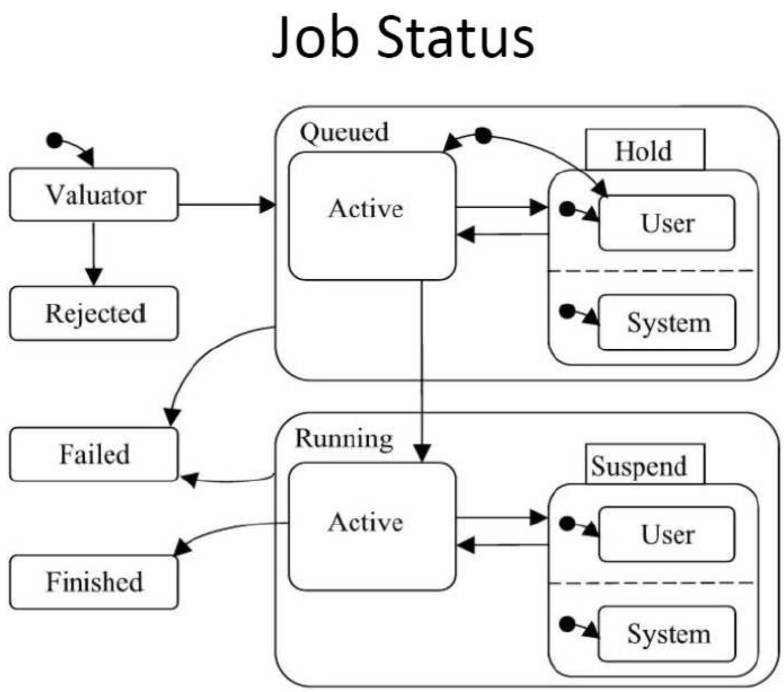

- Extended and analysed performance on SQL database
- Observed cost of communication overhead on individual database inserts, and are now exploring mechanisms for bundled inserts when many cores are processing videos
- Planning for extending access to WindRider (increased number of cores and core-hours)

2.3 UCATANIA team achievements

- Automatic Video Classification (Classified about 15000 videos).Examples of classified scenes are:

Video Classes (from top-left to bottom-right): 1) Algae, 2) Blurred, 3) Complex Scenes, 4) Encoding, 5) HighlyBlurred, 6) Normal

- Collection of new ground data on different underwater scenarios
- Extensive performance evaluation of the object detection algorithms on the different scenarios
- A powerful kernel density approach based on texture and colour features for fish detection
- Improving the performance of the post-processing module for false positives removal

2.4 UEDIN workflow team achievements

We have focused our work on integration. We have successfully implemented and tested the following cases on Windrider:

1. Processing on-demand and batch queries
2. Execution and monitoring of completed jobs. The schema is:

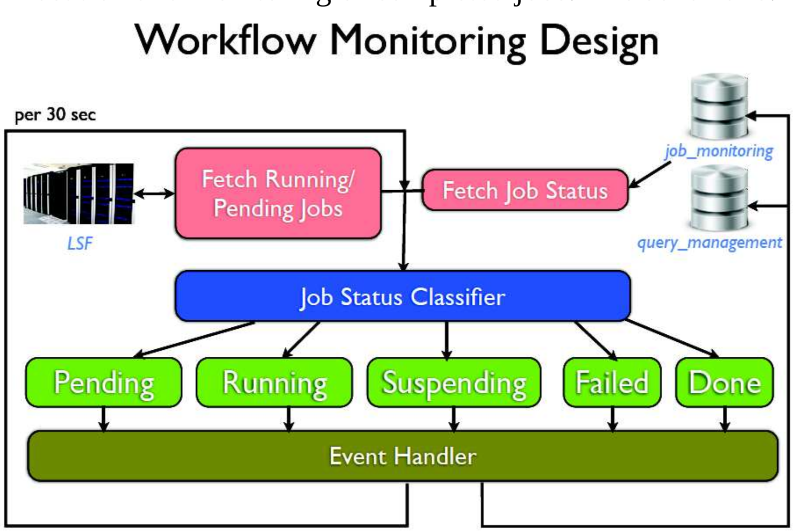

3. Rerunning of failed jobs up to 2 times
4. Dealing with job dependencies
5. Suspension of jobs
6. Resuming suspended jobs
7. Handling jobs with different priorities (low and high)
8. Detection of missing videos from queries

We have implemented the following on the VM group (consisting of 20 cores on 2 nodes at present):

1. Processing batch queries
2. Execution and monitoring of completed jobs
3. Rerunning of failed jobs up to 2 times

2.5 UEDIN vision team achievements

- Extended fish recognition to 23 species (from 35 observed):

- Both fish detection and recognition can run with $>$ 100 CPUs
- Multi-cores queues can be used by both VIP components
- Around 10% of the videos in the database are processed
- Abstract accepted in 9th Indo-Pacific Fish Conference (IPFC)
- Abstract and poster accepted at 2 other Marine Biology conferences

3 Responses to Reviewer's Comments

We have made progress with the comments/requests raised by the reviewers at the second annual review:

1. **The extension of the proposed 'objectness' algorithm that worked on still images to 'motion objectness' to distinguish between moving blobs produced by fish movement and blobs due to background object movements seems a good feature for object detection as well. Here they need to compare their approach with those which model the motion of the whole object by a complete first order affine model and not only mean motion.**:

   The activities related to motion objectness in the first six months of Year 3 have been devoted to feature selection, i.e. to identify the most suitable spatial and temporal features (and their combination) for discriminating objects of interest from background objects with the aim to improve the object detection performance. In detail, a systematic evaluation aiming at testing multiple combinations of features and classifiers in order to identify the most meaningful features and the most promising classifier for computing the motion objectness was performed. The results showed that the best performance was obtained by using the SVM classifier, whereas the other tested classifiers (e.g. Naive Bayes, Mahalanobis, etc.), though able to reach good accuracy, all suffer a 2-3% increase in the average misclassification rate with respect to the SVM. As far as it concerns the results of the feature selection, we noticed that features such as area and local contour contrast were excluded from the best selected set.

   In order to improve further the system's performance, we have also implemented an approach that models object motion by differentiating it from the scene dominant motion (computed via affine flow vector). Kinematic features, such as divergence, vorticity, etc. extracted from the estimated motion are then used to calculate motion objectness. Preliminary results shown a slight performance improvement, though computationally more expensive. However, in the last six months, a more extensive performance analysis will be carried out.

2. **The initial video analysis often takes an ad-hoc solution and the object description did not use a global model. We recommend using multi-resolution approaches (that will also speed up the process) or a particle filtering. In general detection and tracking algorithms are competent applications for this specific domain and show good results but they are not really novel or general. A research paper on particle filtering comprising fundamental references was supplied to the UC partner by the reviewers.** :

   Given the low spatial resolution of the underwater videos, multiscale tracking approaches can not be exploited. However, in order to improve the tracking performance, especially, in handling fish occlusions we have implemented a new tracking algorithm, based on the integration of the covariance model into a particle filter framework. Particle filters manage a set of weighted 'particles' information, (e.g. current speed and acceleration) representing the possible regions where an object can be found. A particle's weight is proportional to the similarity between the learned object model and the appearance model of the region described by that particle. At each new frame, all particles for a given tracked objects are 'moved' according to their own motion information and the adopted noisy motion model, to introd uce a fair amount of randomness) to the new candidate locations of the object. After particle weights are updated, the weighted average of all particles represents the object's new location. In our approach the particle weights are updated by using covariance information for its proven ability to homogeneously fuse different appearance features (e.g. color intensities in different color spaces, position, spatial derivatives). The main features of our approach are:

(a) the window search is intrinsically handled by the particle filter, through the independent and random movement of a large set of particles;

(b) the new candidate locations for an object is not constrained by the output of the background modeling algorithm; rather, this information is also integrated into the particle weighting model;

(c) multiple particles might overlap with a single moving region, thus being able to identify distinct objects within it;

(d) particle motion is based on a motion model whose order is related to the video frame rate (if high enough, first-order could be sufficient);

(e) in order to handle the cases when the covariance model cannot be used, we also keep a histogram-based model for each object, to be used as a fallback comparison method.

The approach is currently being tested and compared with state-of-the-art approaches by using both hand-labeled ground truth data and online methods and the results will be delivered in the final report.

3. **We recommend trying random forests/ferns within the BGOT (as it is being done). They might speed up the training process when lots of data are used, with the drawback that more memory is required to run them.**:

   We implemented the random forest method (Ho 1995)and compared it with our BGOT method. A random decision forest is made of a number of decision trees with binary splits for classification. It starts with single tree, and then moves onto the forest. It predicts responses for new data with the ensemble learned model. For each node of the tree, the feature on which to maximize the difference of entropy at that node is chosen. The new test sample is pushed down to every decision tree of the random forest. When the sample ends up in one leaf node, the label of the training sample of that node it is assigned to the test sample as tree decision. Then, a voting procedure is carried on to generate the final prediction. In our experiment using 15 species of fish, the random forest method is implemented with 50 decision trees. Each tree is constructed using 500 randomly selected features. In the later stage, we prune the maximum depth to 10 in case of over-fitting. The experiment result demonstrates that the performance of our BGOT method outperforms the random forest method:

Method	Average Recall (%)	Average Precision (%)	Accuracy by Count (%)
Random Forest	0.772	0.662	0.914
Ada-boost	0.753	0.769	0.923
Flat SVM	0.863	0.858	0.934
BGOT method	0.900	0.917	0.950

4. **The workflow process needs a better integration with the interface. The workflow computes lots of information (execution time, progress, etc.) that the user might want to see.** :

   The workflow engine is now fully integrated with its partner software via a relatively comprehensively design of database schemas and its corresponding database implementations. This integrated implementation is now successfully running on the NCHC WindRider machine. This means that the workflow machine can now communicate with CWI interface via the database, as well as with the underlying VIP modules via directly manipulating them.

   The UI reports on process progress are still under development.

5. **Some experiments running the computer vision algorithms on higher resolution of images could be shown for next year, and these higher resolution images can be shown to the experts for annotation to see if they still have the same problems.** :

   Experiments are planned

6. **Publications to more highly recognized computer vision conferences (ICCV, ECCV, and/or CVPR) are expected, and as the work during the third year will be more mature publications to journal papers such as IEEE Transactions on PAMI, Pattern Recognition, Computer Vision and Image Understanding are also encouraged.**:

   We have been submitting conference papers to CVPR, ICIP, BMVC and Computer Vision and Image Understanding in this period.

7. **We wonder if the interfaces design cover all use cases obtained from year 1 interviews with biologists, a paragraph explanation will help to answer this question.**:

   We synthesized the information needs into 5 general requirements. These requirements were drawn from year 1 interviews (Deliverables 2.1), but also from the Deliverable 2.3 which takes into account uncertainty issues that users were not able to envision. Some of the user needs expressed in year 1 interviews were excluded of the scope of the UI because of feasibility issues, as reported in the Deliverable 2.1. The implementation of the 5 general requirements into the UI is described in details in the Appendix I of D2.3.

8. **For the next period, it would be very helpful to report the replies to the above questions.**:

   Some are answered here; the remainder in the final report.

9. **The proposed scene detection algorithm is rather simple, it would be helpful to benchmark it against shot boundary detectors known from literature such as progressive changes in the distribution (histogram) of frames by R. Leonardi et al. or**

other methods which assume a progressive change of a scene and can be found in the proceedings of TRECVID workshops .:

Work in progress

10. **Recommendation for the future (stated last year and not followed): - report the percentage of actual versus budgeted costs for the period - report the percentage of actual versus budgeted resource usage to allow the assessment of spending of resources with regard to budget.** :

    Result will be given in the final report.

4 Plan for technical working meetings

We have identified a number of technical working meetings that are to take place.

- NARL + CWI + UCATANIA: HPC query interfacing
- All: second integration performance analysis
- UEDIN/CWI/UCAT (July): ROC to Counts + Multiple detections and classification result reconciliation and recomputations

5 Deliverable Summary

Below is a summary of the deliverables that are due since the second annual review.

We have caught up on the overdue deliverables. D6.5 has been completed early, but the D6.4 report is delayed due to the corresponding workshop happening in the summer. Reports D5.4 and D6.6 are in progress.

Each deliverable document, dataset and software component now has an identified person from another project partner who is responsible for reviewing the content and performance of the deliverable. The Quality Control person is identified in the table below.

Num	Team	Title	Mth	Done	QC
D1.3	UEDIN	Fish clustering and recognition	24	Y	Concetto
D2.4	CWI	First advanced UI prototypes available	27	Y	Daniella
D4.3	NARL	Process execution	24	Y	Bas
D5.6	UEDIN	Video Ground Truth Generation	24	Y	Jacco
D3.3	UEDIN	Process execution and control (merged with D4.3)	26	Y	Bas
D7.12	UEDIN	Third annual public report	26	Y	all
D5.4	UEDIN	Experimental evaluation report 1	30		CS
D6.4	CWI	Int. scientific workshop 3	30		BF
D6.6	CWI	Public query interface	30		DG
D7.9	UEDIN	Final 6-monthly report to EC	30	Y	all
D6.5	UEDIN	Int. joint biological and ICT workshop 4	36	Y	Jessica

6 Publications since Year 2 report

1. Ihunanya Martina Ugwuh, Exploiting Parallelism in Video and Image Processing Tasks: A-state-of-the-art survey, MSc Dissertation, School of Informatics, University of Edinburgh, 2012.

2. Yan Li, Fish Component Recognition, MSc Dissertation, School of Informatics, University of Edinburgh, 2012.

3. Konstantinos Vougioukas, Adaptive filters to remove blurring effects over time for underwater surveillance, MSc Dissertation, School of Informatics, University of Edinburgh, 2012.

4. Salvatore Roccella, Video Enhancement by Superresolution, MS thesis (in italian), Univ of Catania, 2013.

5. G.Nadarajan, Y.-H. Chen-Burger, R. B. Fisher. "Semantics and Planning Based Workflow Composition for Video Processing". Journal of Grid Computing, Special Issue on Scientific Workflows, 2013, to appear.

6. J. He, J. van Ossenbruggen, and A. P. de Vries. Do you need an expert in the crowd? A case study in image annotation for marine biology. In: Open Research Areas in Information Retrieval (OAIR 13), 10th International Conference in the RIAO series. To appear.

7. J. He, J. van Ossenbruggen, and A. P. de Vries. Fish4label: Accomplishing an Expert Task without Expert Knowledge (demo paper). In: Open Research Areas in Information Retrieval (OAIR 13), 10th International Conference in the RIAO series. To appear.

