

Real-time error detection for a rowing training system

Leonard Johard* Alessandro Filippeschi* Emanuele Ruffaldi*

(*) PERCRO, Scuola Superiore Sant'Anna, TECIP, Pisa, Italy
E-mail: l.johard@sssup.it, a.filippeschi@sssup.it, e.ruffaldi@sssup.it

Abstract

The use of coaches in training provides several advantages to the skill transfer process. However, their use is limited because of cost reasons. Automatic recognition processes, on the other hand, have a vastly lower cost ratio and can replace or integrate many of the functions otherwise performed by a human. The paper presents an effective design of a real-time error recognition system for rowing in a simulator. It was designed in order to be easy to develop, be computationally efficient and offer an excellent classification ability.

1 Introduction

A coach or other human assistance in training is an important step in transferring motor skills between subjects. By collecting experience in analyzing the motion, the coach can be viewed to adapt the role of a trained critic in an actor-critic system, which vastly simplifies the learning task of the rower. However, human assistance is an expensive choice in training, which severely limits the possible extent of human assistance.

Modern machine learning techniques offer a plausible alternative and/or assistance to the human coach, at least for simple task evaluation and classification. The process is similar in that they learn from experience, but they are both cheaper and, in several comparisons with humans [6] [1], more efficient at generalizing the information gathered. On the other hand, machine learning systems lack a priori information assisting in the training process. This limits the usefulness of some complex data which the coach is able to use, such as raw visual information.

Previous studies in gesture recognition mainly focus on the use of markers or accelerometers to derive useful data for classifying the movement. In rowing specifically a previous study uses marker based capturing and commercial ergometers to test the ability to imitate human coaches [4], which achieved 71% - 95%

intra-subject error classification accuracy. The system used data from rowing on an ergometer recorded with VICON motion capturing system. In contrast, by taking advantage of modern rowing training systems, such as the SPRINT one developed in PERCRO [5], we get data of individual oar positions from the platform directly. This removes need for extrapolating occluded parts of the stroke. Also the rowing experience itself is more realistic, with two physical oar handles and a sophisticated virtual environment encouraging the rower to adapt a rowing pattern more similar to that used in a real boat.

Since the SPRINT system provides concurrent feedback to the user, delay is a critical issue. There are at least three reasons for minimizing delay. The first is that despite some studies showed some benefits in skills learning with delayed feedback [9] [7], recent research showed that when complex motor skills are involved, concurrent feedback is more effective both for performance improvement and retention [10].

2 Method

We have chosen to divide the rowing stroke into four main phases (entry, drive, finish and recovery) and perform error classification only once per stroke. The phase division is the central element to our design, allowing us to divide the recognition tasks evenly over the rowing cycle, limiting the tasks to a single error classification of each type per cycle and at the same time giving feedback delays much shorter than a cycle.

A large number of algorithms has been used in human motion recognition. There seems to be no clear distinction in performance between different methods in gait recognition [8], although there is plenty of theoretical considerations [2]. On the other hand, for classification tasks in general there are substantial comparisons available, as summarized by [3].

Artificial neural networks was chosen because of ease of implementation, good scaling characteristics

and low computing costs in classification tasks. In addition, the few available comparisons including Multi-Layer Perceptron (MLP) with regularization show excellent performance [3]. A Bayesian regularization was deemed suitable for the task for these reasons. As training is performed offline, such a slow training method will not affect the online simulation performance.

Expert (10 years of rowing practice), intermediate (4 ± 1 years of rowing practice) and novice rowers data were recorded on the rowing system and pre-processed. Then the network for phase classification was trained and test, in the end the same was done for the skying error detection network.

The classifiers were tested on a 2.4 GHz Intel Core 2 CPU, 2GB of memory, with Windows XP, and running all the computations as single-threaded. All computing times refer to the measured average running times on this system using real data.

2.1 Data acquisition

The data consisted of coordinates of the oars. Rowing oars are fixed by means of a spherical joint to the boat, allowing them to rotate with 3 DoF. Rotation of the oar about its axis is not currently considered, we call the remaining angles phi and alpha. Although data was recorded in a time resolution of 10 ms in the simulator, the actual sensors had a refresh rate of about 30 ms, resulting in extensive duplication of data points between sensor updates in an irregular manner. Five strokes each from a set of five rowers was classified manually, adding a phase value (1-4) to each data point. The rowers consisted of one expert, four intermediate and two novice rowers.

In an upcoming version of the rowing simulator the refresh rate of the sensors has been improved substantially and will match or be less than the resolution in the data. This will have no effect on running times, but will likely make the classification task easier.

3 Phase Identification

Dividing the rowing process into these four phases is a seemingly trivial problem. In fact, due to different lengths of the rowers and common mistakes in the calibration in rowing simulator before each session, the classifier has to be translation invariant to some degree.

Moreover, since the classification takes place in real-time, it is common for rowing errors such as wobbling to be temporarily almost indistinguishable from phase changes. It also has to handle a rather sparse description of the curve resulting from the low sensor refresh rate.

3.1 Feature extraction

A simple set of parameters was used: the mean around the classification point, the coordinates at the start of the time window and the coordinates at the end of the time window. To make the parameters translation invariant, they were all expressed as relative coordinates centered at the classification point.

3.2 Phase classifier

We used a regularized MLP with radial basis transfer function for phase classification, which gave superior performance in early tests compared to the sigmoid transfer function.

In the phase classification task we used a three-layer feed forward network with 24, 12 and 4 neurons respectively. The first and third layer use linear transfer functions while the second uses a radial basis function (RBF). In the last layer the output of each neuron represents a class. Training is performed using the Levenberg-Marquardt algorithm and bayesian regularization.

The addition of the first linear layer transforms the input to the radial basis layer as a linear combination of the inputs and does not extend the linear transform of the radial basis input layer. Its purpose is only to allow larger linear dependencies of the input layer under regularization by reducing the effect of weight decay.

Tests of our network classifier and feature extraction in a Simulink model showed that the system used about 0.06 ms of CPU time on average, which is certainly low enough to reach our goal of having several classifiers running simultaneously in real-time together with the simulation without affecting the overall performance.

3.3 State machine

For some rowers the classification could fluctuate between two phases for a brief period during the phase change. Although this only gives a few bad data points, it hinders us from locating the exact position of the phase changes as well as from deriving a simple cyclical transition between the phases. Our solution is a state machine, that changes phase if and only if the consecutive phase is classified in two steps in a row. After phase change occurred in the state machine, a minimum timer of 200 ms was added before another phase change was possible. This timer was set in order to guarantee only one classification change per phase boundary, but at the same time be small enough to stay clear of the next phase boundary in the cycle.

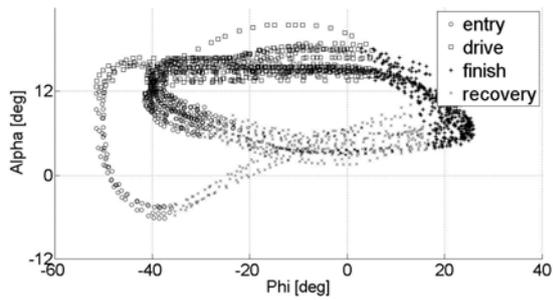


Figure 1: Manual classification of the four phases over five strokes.

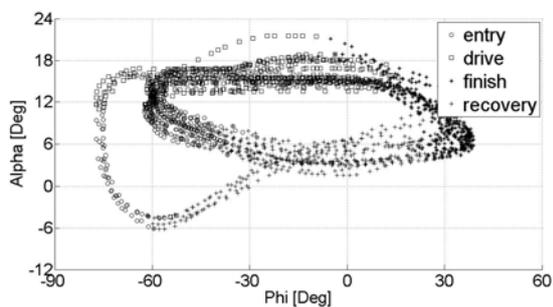


Figure 2: Merged out-of-set classifications from cross validation of automatic phase segmentation.

3.4 Validation of phase segmentation

An intra-subject cross-validation was performed and the result is shown in Fig. 1 and 2 in both a visual and analytical form. The visual result is important for two reasons: as the number of rowing samples is small, we have to be somewhat tolerant; in addition, the manual classification is somewhat arbitrary and we have to give some freedom of interpretation in the classification task.

Analytically the performance gave a correlation of 94% with the manual classification. The visual inspection of an experienced rower confirmed that the classification was of good quality overall. Most of the error comes from a small displacement of the classification boundary in comparison with the manual classification, which is acceptable. In some of the cross-validation runs we also noticed isolated points far from the appropriate boundary, which can probably be considered outliers due to measurement errors. This was successfully filtered by the state machine. Fig. 3 shows an example of final phase classification during a race simulation, illustrating the regularity and quality of the classification across strokes. This regularity is essential in providing consistent data for the skying classifier.

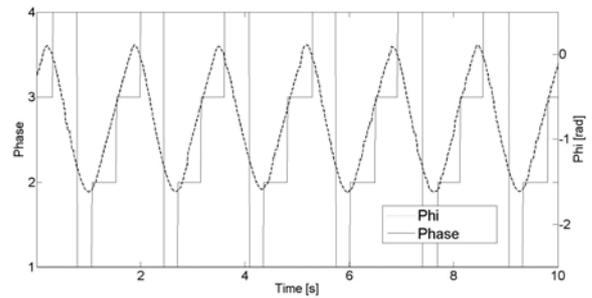


Figure 3: Phase classification against oar angles in a rowing session. Phi has been translated in order to ease visual comparison.

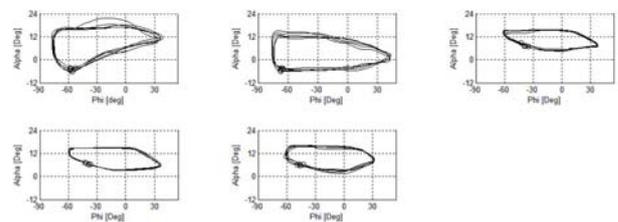


Figure 4: Rowing cycles with the phase change used as classification point marked by circles.

4 Skying classification

Low hands during the recovery is one of the most common errors in rowing. It occurs during the end of the recovery phase and consists of lifting the oars excessively over the water. Generally this can either be done by an overly high trajectory or by moving the oars upwards in the end of the recovery.

4.1 Training data

Data consisted of a time window (10 samples at 100Hz) of phi and alpha angles as described in 2.1. Five strokes from five subjects were used as training data. Of these two rowers were chosen because they naturally performed the skying error in their stroke, while the other three mostly performed the recovery in a correct way.

The above phase classification procedure was used to identify the phase change between recovery and catch. This is at the very end of the skying error, which can be visible from 0 to 500 ms before the phase boundary. This time window as well as a manual classification of each stroke as either skying or non-skying was used as training data in the skying classification task (Fig. 4).

4.2 Feature extraction

A similar set of features as in the phase classification was used, with the difference that the minimum absolute height in the time window was added as a parameter. This is needed to identify the first type of skying errors in which the trajectory has a normal shape but reaches excessively low heights. The features are extracted only once per cycle, at the phase boundary between recovery and catch.

4.3 Skying classifier

Since the classification of the low hands error is an easier task than phase classification, a linear network was chosen. The advantage of this is that one can easily derive not only if an error exists, but also the severity of the error on a continuous scale. One linear layer with five neurons was used and trained with the Levenberg-Marquardt algorithm Bayesian regularization. Training time was about 10 s. Classification time is less than 0.06 ms, but is hard to estimate exactly due to the small data set used and larger overheads from other parts of the application.

4.4 Validation of skying classification

Intra-subject out cross-validation was performed on the 5 subjects. Data was divided into five pairs of training-testing data, where the former data set contained four subjects and the testing set the remaining rower. This resulted in a 100% accurate intra-subject classification across the total of 25 test examples.

5 Conclusions

The design presented has proven to be successful in classification of phases and skying. It can easily be extended by a more extensive set of error classifiers without raising the computing requirements, as the new error classifiers could easily be distributed over the rowing cycle as demonstrated. The neural network approach is fast, which allows us to run several classifiers simultaneously.

More generally this proves that current classifiers are reliable enough to be put in interconnected systems, where one classifier identifies the data to be analyzed by another classifier. Although the current task is related to gait recognition, it could be applied to other fields. An example could be image recognition, where one classifier identifies the center of an object, which is then classified by a separate network.

References

- [1] A. Andler and M. E. Schuckers. Comparing human and automatic face recognition performance. *IEEE Transaction on Systems, Man, and Cybernetics, Part B: Cybernetics*, 37(5):1248–1255, 2007.
- [2] T. Chau. A review of analytical techniques for gait data, Part 2: Neural network and wavelet methods. *Gait & Posture*, 13(2):102 – 120, Oct 2000.
- [3] W. Duch. Benchmark datasets used for classification: comparison of results. Retrieved July 1st, 2010.
- [4] S. Fothergill, R. Harle, and S. Holden. Modeling the model athlete: Automatic coaching of rowing technique. In *Proceedings of the 2008 Joint IAPR international Workshop on Structural, Syntactic, and Statistical Pattern Recognition*, 2008.
- [5] A. Frisoli, E. Ruffaldi, L. Bagnoli, A. Filippeschi, C. Avizzano, F. Vanni, and M. Bergamasco. Preliminary design of a rowing simulator for in-door skill training. In *Haptics on Ambient System*, 2008.
- [6] W. M. Grove, D. H. Zald, B. S. Lebow, B. E. Snitz, and C. Nelson. Clinical versus mechanical prediction: A meta-analysis. *Psychological Assessment*, 12(1):19–30, 2000.
- [7] D. V. Linden, J. Cauraugh, and T. Greene. The effect of frequency of kinetic feedback on learning an isometric force production task in nondisabled subjects. *Physical Therapy*, 73:79–87, 1993.
- [8] A. Mannini and A. M. Sabatini. Machine learning methods for classifying human physical activity from on-body accelerometers. *Sensors*, 10:1154–1175, 2010.
- [9] C. Winstein, P. Pohl, C. Cardinale, A. Green, L. Scholtz, and C. Waters. Learning a partial-weight-bearing skill: effectiveness of two forms of feedback. *Physical Therapy*, 76:985–993, 1996.
- [10] G. Wulf, C. Shea, and S. Matschiner. Frequent feedback enhances complex motor skills learning. *Journal of Motor Behavior*, 30:180–192, 1998.