Final projects

Raquel Urtasun
TTI Chicago

April 1, 2010

1 Logistics

Each student student should select a topic, either from this document or from their own research (if approved). At the end of the semester, you should submit and prepare the following:

- An extensive report on the project you have done. This report should include theoretical and practical aspects of the project you have chosen.
- There will be a 20 minutes presentations for each project (15 talk + 5 questions). Please prepare slides for this 15 minutes, not more!
- There will be an oral exam on the theoretical aspects of the class. This will last for 15-20 minutes.

IMPORTANT: Email me as soon as you select a project, so that the distribution of projects is balanced.

2 Available Projects

In this section a list of available projects is described in detail. If you have any questions, do not hesitate to contact me.

2.1 Inverse kinematics with multiple priorities

During the second lecture of the class we covered least-squares and extensions to include multiple priorities. In this project you will investigate two different strategies to incorporate priorities

- Weighting strategy: give the constraint a higher weight, but yields composite solution.
- Task priority strategy: find a posture that satisfies the joint limits

For detailed information about the optimization, please refer to the second lecture or to [1]

The optimization framework would be tested to solve one of the following problems

- Character animation from user specified constraints:
- Human pose estimation from images

Choose the problem that you are more interested in, code the optimization and experiment with it.

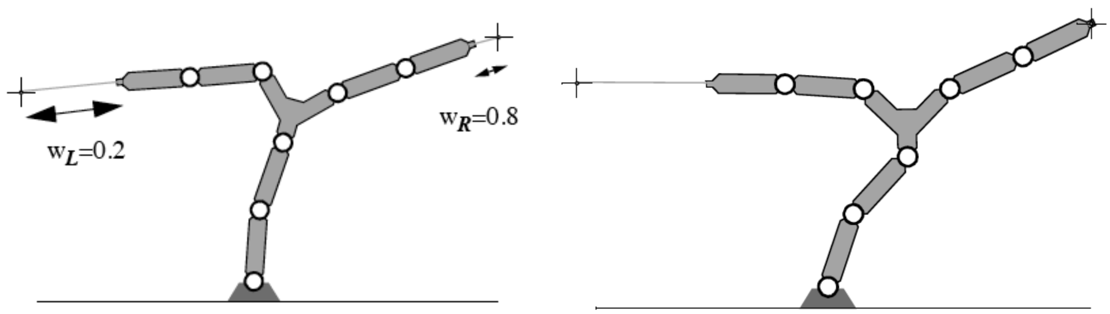

Figure 1: (left) Weighting strategy. (right) Priority strategy (Baerlocher01)

2.2 Evaluation of particle filtering techniques

This project consists in the implementation of a particle filter, and evaluate different sampling strategies in order to solve a simple human body tracking problem, where the image likelihood is simply the reprojection error of 2D to 3D correspondences. In order to generate the data, dowload sequences from the CMU mocap dataset [1] , generate a virtual camera by defining a projection matrix, and project the 3D location of the markers into 2D to create 2D correspondances. Test the influence of adding Gaussian noise.

2.3 Evaluation of likelihood models

This project consists in the evaluation of different likelihood models. You can do this either in an optimization framework that seeks the MAP estimate, or in a particle filter framework. Note that for the optimization framework you need to compute derivatives of the image likelihood with respect to the state. This reduces the number of possible image likelihoods to differentiable ones, but makes the implementation of each likelihood more difficult since you need to also compute the gradient. Download the Humaneva dataset [2] in order to get the video sequences to test.

2.4 Character animation via motion graphs

This project consists in the implementation of the well know motion graphs [2]. This is a Nearest-Neighbord (NN) approach to character animation. Download lots data from the CMU dataset and compute possible transition points. Implement also an inference program that given a begin point and an end point fills out the animation in between.

2.5 Discriminative approaches to human pose estimation and classification

In this project you have to investigate discriminative approaches to human pose estimation. In particular, you will use the Humaneva dataset , where images and 3D poses are given at the same time. Use the background substraction code that is provided to get a bounding box around the person. Compute object-recognition type features (i.e., Bag of words of SIFT, PHOG, etc) on the corresponding bounding box, and use these features as your image representation. Investigate different regression methods, e.g., NN, Gaussian process [3] [3], etc.

[1] CMU mocap, mocap.cs.cmu.edu

[2] Humaneva dataset, http://vision.cs.brown.edu/humaneva/

[3] Code available at http://www.gaussianprocess.org/gpml/

2.6 Activity recognition

This project consists in performing activity recognition for human motion. Perform experiments in the Weizmann [4] and KTH datasets [4]. Compare the performance of using approaches based on object recognition (e.g., Bag of Words of spatio-temporal features, temporal histograms of HOG features, etc) to HMM type approaches. Note that implementations of these features exist in the web.

References

[1] Baerlocher, P.: Inverse kinematics techniques of the interactive posture control of articulated figures. PhD thesis, EPFL, Lausanne, Switzerland (2001)

[2] Kovar, L., Gleicher, M., Pighin, F.: Motion graphs. In: SIGGRAPH. (2002)

[3] Rasmussen, C.E., Williams, C.K.I.: Gaussian Processes for Machine Learning. MIT Press (2006)

[4] Gorelick, L., Blank, M., Shechtman, E., Irani, M., Basri, R.: Actions as space-time shapes. PAMI (2007) 2247–2253

[4] KTH action dataset, http://www.nada.kth.se/cvap/actions/