Artificial Intelligence for Human-Robot Interaction: Papers from the 2014 AAAI Fall Symposium

Automating Stroke Rehabilitation for Home-Based Therapy

Akanksha Saran
asaran@cmu.edu
The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, Pennsylvania 15213

Kris M. Kitani
kkitani@cs.cmu.edu
The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, Pennsylvania 15213

Thanassis Rikakiss
trikakis@andrew.cmu.edu
School of Design
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, Pennsylvania 15213

Abstract

In this work we present a conceptual design to automatically evaluate a subject's performance for a home-based stroke rehabilitation system. We propose to model a reaching task as a trajectory in the state space of hand part features and then use reward learning to automatically generate new ratings for subjects to track performance over time.

Neuro-muscular rehabilitation of the upper-extremity after a stroke requires dedicated hours of arm and hand exercises with a therapist. Often to improve the flexibility of the hands, therapists ask patients to manipulate differently shaped objects and move them around. We are working towards developing a home-based neuro-muscular rehabilitation system by doing away with markers and using 2D computer vision instead. As a first step, we are able to identify different grasp types using 2D hand part features. However such a system would also require basic automation in evaluation of a user's performance on grasping and reaching tasks. This can allow a user to self-monitor his/her performance at home and also allow remote monitoring by the therapist over time. The ultimate aim is that the user should be able to do therapy in the home in a convenient, productive and cost-effective manner.

In order to provide evaluation of a subject's grasps during therapy, we require robust and accurate predictions in the absence of a therapist. Expert therapists can rate the grasping or reaching performance of a subject just by observing the grasp. Such an approach has been widely used in other domains successfully (Akrour et al. 2013), (Fürnkranz and Park). Expert ratings give a coarse evaluation metric for the subject's performance. We propose that hand grasps be modeled in a feature space using 2D computer vision. The features incorporate relations between different hand parts. Then, we can apply optimal control (Ziebart et al. 2008), (Ratliff, Bagnell, and Zinkevich 2006), (Ratliff, Silver, and Bagnell 2009) to infer the reward function of the grasping policy that best justifies the ratings by expert therapists (Daniel et al. 2014).

For the proposed approach, certain limitations need to be kept in mind and a home-based solution can be made feasible by working around these limitations. A 2D camera provides a neat view in all kinds of lighting - indoors and natural light. However it does suffer from occlusions while capturing the hand. To overcome this, we always mount the camera at the same location on a table such that it always gets the same view and suffers similar occlusions, if any. This ensures that the home therapy system sees exactly what the therapist sees during the training and the subjects are also always asked to grasp objects starting from the same region of the image view.

Figure 1: An ideal home based therapy system using a single camera where a subject can do upper-extremity exercises by manipulating objects in the home on a simple kitchen table.

The ratings will be provided on a scale of 0 to 10. It has been proven that subjects perform better when providing ratings rather than classifying into categories (Hake and Garner 1951). Previously it has also been shown (Thomaz and Breazeal 2008), (Cakmak and Thomaz 2012) that humans have considerable noise in their ratings of actions. Though here we assume that a single therapist is fairly consistent in his/her ratings and with ample training with the therapist we can capture any possible uncertainty. The home-based therapy system is transferring the beliefs and understanding of a single therapist to rate new test scenarios.

Neuro-muscular therapy is often guided by a therapist

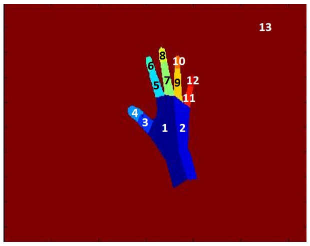

Figure 2: 12 different hand parts used for capturing the components of a grasp visible from an overhead 2D camera. The feature space of distance, mid-point, orientation and between-pair features are defined over the set of these hand parts.

with a highly individualized approach for each patient. However the ways of grasping an object can be clearly distinguished as bad or good grasps (Figure 3). Therapists distinguish between this broad category very well. Our approach assumes that given enough rating samples from therapists, the system could generalize for important classes of grasps (range of ratings) such as very bad, bad, not too bad, good, very good and so on.

2D Hand Part Features

We use a structured output random forest (Dollár and Zitnick 2013) with multiple levels of predictions capturing context to effectively label hand parts for a single 2D image (Figure 2). Once we obtain the hand parts, we can denote every grasp as a point in the space of pair-wise features proven to work for food recognition (Yang et al. 2010). We randomly select 1000 pixels from the segmented hand image (Figure 2) and generate the following set of features:

- Distance features - Euclidean distance between two pixels.
- Orientation features - The angle with the positive x-axis by the line joining the two pixels.
- Midpoint features - The probability distribution of the hand parts for the pixel which is the midpoint of the line joining the two pixels.
- Between pair features - The part labels for all the pixels lying on the line joining the two pixels.

Learning Reward

The reward and policy are specific to grasping a single object. We consider the contextual episodic policy search with continuous states $s \in S$ (representation in the pair-wise hand part features) and continuous control parameters $w \in \Omega$ (the hand-part distribution for an image). The distribution over states is denoted by $\eta^{\pi}(.)$. A trajectory τ results from a continuous set of images (hand part segmentation images). We can initialize the policy as a single Gaussian with random mean and update our policy to maximize the expected reward.

(a) On the left, the subject is unable to naturally open the clenched fingers though is able to make contact with the cup using the thumb and index finger. On the right, the grasp is looking more natural with all fingers grasping the mug in synergy.

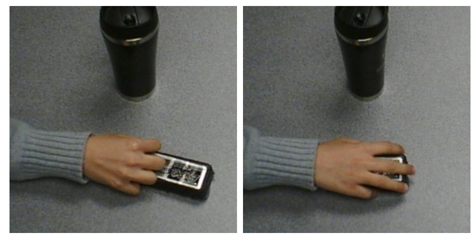

(b) On the left, the subject is unable to open the hand with a full aperture (inflexible) whereas on the right the subject has a more open grasp.

Figure 3: Illustrations of bad and good grasps which would achieve rating on the lower and higher side of the scale respectively by a therapist.

$$\mathbb{E}_{s,w,\tau}[R(f)] = \iiint \mathbb{R}(f = \phi(\tau))p(s,w,\tau)\mathrm{d}s\mathrm{d}w\mathrm{d}\tau \tag{1}$$

where $p(s,w,\tau) = p(\tau|s,w)\pi(w|s)\eta^{\pi}(s)$ and f is the feature space to which an image is transformed.

Our goal is to find a probabilistic model $p(R|o,D)$ that predicts the reward given an observed outcome and training data D, which is obtained from an expert as proven before(Daniel et al. 2014). Similarly, we couple policy search with reward learning (modeled as a Gaussian Process). The policy would vary for each reaching and grasping task (example grasping a mug, grasping a pen and so on).

Conclusion

We present this as a conceptual design paper as a first step in automation of neuro-muscular stroke rehabilitation for a home-based therapy system, guided by computer vision. Transforming a 2D image into a set of hand part features can allow us to model a reaching or grasping task during therapy as a MDP in the feature space. Ratings for these tasks from a therapist can help the system to learn optimal policy. Once the policy and rewards for state transitions have been learnt, when a new subject reaches an object, the system can predict the rating for his/her contact grasp by observing the trajectory in the vision based feature space.

Future Work

We will test our proposed approach with real data of non-impaired and impaired subjects with ratings available from therapists for training and testing.

References

Akrour, R.; Schoenauer, M.; Sebag, M.; et al. 2013. Interactive robot education. In *ECML/PKDD Workshop on Reinforcement Learning with Generalized Feedback: Beyond Numeric Rewards*.

Cakmak, M., and Thomaz, A. L. 2012. Designing robot learners that ask good questions. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*, 17–24. ACM.

Daniel, C.; Viering, M.; Metz, J.; Kroemer, O.; and Peters, J. 2014. Active reward learning. In *Proceedings of Robotics: Science and Systems*.

Dollár, P., and Zitnick, C. L. 2013. Structured forests for fast edge detection. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, 1841–1848. IEEE.

Fürnkranz, W. C. J., and Park, E. H. S.-H. Preference-based policy iteration.

Hake, H. W., and Garner, W. 1951. The effect of presenting various numbers of discrete steps on scale reading accuracy. *Journal of Experimental Psychology* 42(5):358.

Ratliff, N. D.; Bagnell, J. A.; and Zinkevich, M. A. 2006. Maximum margin planning. In *Proceedings of the 23rd international conference on Machine learning*, 729–736. ACM.

Ratliff, N. D.; Silver, D.; and Bagnell, J. A. 2009. Learning to search: Functional gradient techniques for imitation learning. *Autonomous Robots* 27(1):25–53.

Thomaz, A. L., and Breazeal, C. 2008. Teachable robots: Understanding human teaching behavior to build more effective robot learners. *Artificial Intelligence* 172(6):716–737.

Yang, S.; Chen, M.; Pomerleau, D.; and Sukthankar, R. 2010. Food recognition using statistics of pairwise local features. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, 2249–2256. IEEE.

Ziebart, B. D.; Maas, A. L.; Bagnell, J. A.; and Dey, A. K. 2008. Maximum entropy inverse reinforcement learning. In *AAAI*, 1433–1438.