ICCV 2017 Outline

Incorporating Skepticism in Visual Learning

Overview

In real-world situations, not all visual observations are optimized towards learning a task. A teacher demonstrating a skill may perform irrelevant, incomplete, or even incorrect actions, yet humans are capable of ignoring the mistakes.

To the best of our knowledge, we are the first to introduce the concept of *skepticism* in robot learning from demonstration. Our algorithm detects faulty human demonstrations to filter salient information from noise.

First, we learn a utility function from a small number of videos to rank preferences of states. Then, when a robot is learning a new instance of a skill through human demonstration, we consult the utility function to verify the utility-gain increases over time. Lastly, the robot interrupts human demonstration to ask if the action was essential for learning the skill.

We evaluate our experiments on a ground-truth dataset of 40 cloth-folding videos, half of which are faulty. We show how our algorithm correctly identifies mistakes using its learned utility function.

The utility function:

$$U(pg;\Lambda,F)=\sum_{\alpha=1}^{K}\lambda^{(\alpha)}(f^{(\alpha)}(pg)) \tag{1}$$

Strong ranking constraints:

$$pg^a \succ pg^b \Rightarrow U(pg^a) > U(pg^b) \tag{2}$$

$$\forall i \in \{1,...,N\} : U(pg_i^a) - U(pg_i^b) > 0 \tag{3}$$

Learning the utility function:

$$\begin{aligned} \underset{\Lambda,F}{\text{minimize}} \quad & Smoothness(\Lambda)+\sum_i \xi_i \\ \underset{\forall i\in\{1,...,N\}}{\text{subject to}} \quad & U(pg_i^a)-U(pg_i^b) > 1-\xi_i \\ & \xi_i \geq 0 \\ \text{where} \quad & Smoothness(\Lambda)=\sum_{\alpha=1}^{K}\int_x \lambda''^{(\alpha)}dx \end{aligned} \tag{4}$$

Given hundreds of fluents, pursue relevant fluents.