

Simple Clustering for an Effective Visual Similarity Model

Adaimi, George
VITA, EPFL

Kreiss, Sven
VITA, EPFL

Alahi, Alexandre
VITA, EPFL

Abstract. From Tesla to Google to Uber, the race to developing fully autonomous vehicles is slowly growing with big companies and startups investing in self-driving cars and delivery robots, which eventually will populate our already crowded public spaces. These autonomous vehicles will have to learn to navigate complex social situations in their interaction with humans. They will have to anticipate people's behaviors and plan their own actions accordingly. Re-identifying people is one of the building blocks for tracking systems that need to be robust to occlusions and towards systems that can anticipate person-specific behavior from previous observations. Several previous works have tackled this problem by hand-crafting systems specific to the human body and utilizing this information to improve the discrimination process. In this paper, we show that similar, and in some cases, better results can be achieved using a generic but properly trained model. We focus our attention on the clustering method, the loss function and optimization procedure rather than on hand-crafted human characteristics. We investigate the person re-identification (re-id) problem in the context of the more general metric learning, clustering and image-retrieval domains. The code will be made available.