Categorizing Polyhedral Dice Using Image Data

Presenter Thomas Whaley, Computer Science major
Mentored by Dr. Ana Stanescu

Recognizing the number of sides that a dice has is a type of object recognition problem, and it is arguably a very difficult computer vision task mainly due to the complexity produced by an object's position in the image, its size, lighting, relativity to viewer, background, etc. Moreover, as the number of sides increases, dice categorization, either automated or by human eye, becomes even harder. In this work, we conduct an empirical analysis of several supervised machine learning algorithms applied to the problem of categorizing polyhedral dice. From a labeled dataset of 16,384 images, 15% of them is set aside for evaluating the final performance of the selected model and the rest is used in training using k-fold cross-validation. Finally, the best algorithm is selected and evaluated against the previously held-out test data. We found that an ensemble-based model predicts the number of sides the dice has in an image with 94% accuracy. This type of modeling can be especially useful in identifying small parts and sub-components from image data in production facilities. For example, the model could verify that parts on an assembly line are built without errors.