Samuel Webster
Discussion Leader Cover Letter

Paper: Distillation-guided Representation Learning for Unconstrained Gait Recognition

This paper presents GAit DEtection and Recognition (GADER), a gait recognition framework. GADER addresses limitations in existing gait recognition methods, which typically perform well on controlled indoor environments but struggle to generalize to real world implementations. They note certain factors which negatively impact existing techniques, including varying viewpoints, low resolution, changes in appearance, and identification at a long distance.

GADER consists of two main parts: a gait detection module and a gait recognition module (GAR). The gait detection module addresses challenging gait sequences that contain walking and standing segments as well as obstructed or incomplete body views, both of which can negatively impact gait recognition performance. Walking segments are merged to construct a Double Helical Signature, a 2D pattern based on knee movement over time. Body size is normalized to keep the signature proportional over time. In creating this, moving and still segments can be distinguished. The double-helical pattern and the body normalization ratio are the input to GAR. GAR infers from the silhouette modality, but is trained using cross-modality distillation from a model trained on RGB sequences, improving robustness to appearance changes as well as making GAR more lightweight.

GADER is evaluated on several gait detection datasets: BRAIR, an outdoor dataset; CASIA-B, a controlled indoor dataset; Gait3D, another outdoor dataset. They demonstrate that when GADER is applied to SOTA gait recognition backbones, Rank-1 accuracy is consistently improved. They also perform an ablation study to confirm the effectiveness of the ratio attention and cross-modality distillation components. They demonstrate the gait detector module's individual use on identifying gait segments across datasets, identifying possible non- and part-gait segments in curated datasets.