



Goal

“The goal of the Janus program is to enable dramatic improvements in unconstrained face recognition...to develop novel representations to encode the shape, texture, and dynamics of a face for the purpose of improving face recognition performance from video and still images, and then allowing higher performance search and retrieval based upon these more efficient representations”

- IARPA BAA-13-07



Overview

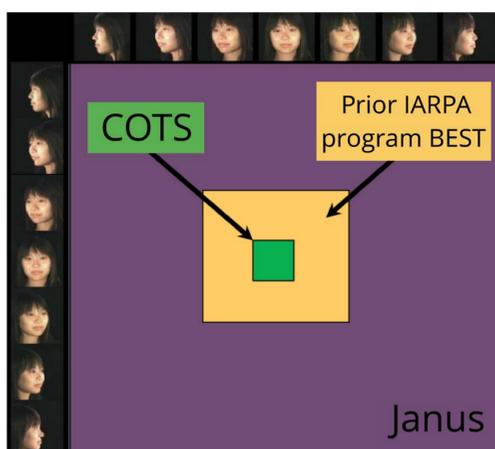
Dramatically improve face recognition performance in massive video collections through novel approaches capable of leveraging the rich spatial and temporal information available within the multiple views captured in unconstrained video.



Intelligence analysts often rely on facial images to assist in establishing the identity of an individual, but too often, just examining the sheer volume of possibly relevant images and videos can be daunting.

A Hard Problem

- Difficult to automatically identify several unique people in a set of videos and photos
- The current state of the art is limited to face photos from persons staring at the camera
- Searching faces is increasingly slow and unwieldy as the database grows to millions
- Recognizing faces in lots of videos is time-consuming and requires vast computational resources
- Analysts need to quickly find the unique faces in a set of videos and how often each unique face occurs across multiple videos
- Partial or incomplete faces are impossible to recognize
- Current approaches severely under-utilize and under-exploit all available face information in a video



Program Structure

Phase 1 – 18 months

- Four multi-disciplinary teams
- Develop new algorithms for unconstrained video
- Develop the most challenging, real-world datasets
- Provide to teams a dataset every six months
- Teams deliver software systems with a Janus API for testing
- Test algorithms against 500+ subjects and tens of hours of video under extreme pose, illumination and expression (PIE) variations
- Objective accuracy metric: 0.85 TAR @ 0.01 FAR
- Objective query time: linear

Phase 2 – 18 months

- Provide to teams datasets challenging for face detection
- Provide to teams datasets challenging for occlusion
- Provide to teams datasets challenging for aging
- Test algorithms against 2000+ subjects and hundreds of hours of video
- Objective accuracy metric: 0.85 TAR @ 0.001 FAR
- Objective query time: sublinear

Phase 3 – 36 months

- Test algorithms against 10,000+ subjects and thousands of hours of video
- Objective accuracy metric: 0.85 TAR @ 0.0001 FAR
- Objective query time: logarithmic