

GMDS-PCA: Project Plan

Dr. Roy Schestowitz

August 5, 2011

Summary: Estimated milestone towards comparable results to state of the art

After about 300 hours of work, including work on technical documentation and learning of the associated methods, the target design is generally in place. Its form is still crude at places and therefore it is worth considering what was previously done better, eying potential reuse of external or peer code, even though there are implementational overlaps. One can identify the following as items requiring completion in order for results to improve further, based on observation of improper identifications (more of those appear, especially as false pairs are added and get misclassified, reducing recognition rate down to ~93%):

1. consistently ordering PCA for sampling of distances/spatial points, maybe finding the right balance between both
2. cheek sampling/completion (a slider for that added)
3. mustache as an issue motivates heightening the binary mask to never account for facial hair
4. partial matching needed for cases where hair hides the eye region
5. cropping more consistently to further assist GMDS

6. less smoothing, more hole filling instead, especially around the eyes
7. better ICP for initialisation of points

The above requires little actual work and more testing, including the building of large models that depend on choice of parameters and need to be rebuilt upon changes to the algorithm. A lot of this would be passive, meaning that it requires letting the computational servers (now 2 are readily available, not one) run overnight, whereas actual work – that which is being measured – should last just dozens of hours. Face-to-face consultation with people at the lab has helped considerably when it comes to eliminating poorer solutions and exploring what previously worked, first in 2005 when GMDS was applied to and tested on face data, then around 2009 when recognition rate as measured on large datasets reached 97%.

The problem of face recognition is a thoroughly explored one for reasons that are simple to grasp. The effectiveness of some methods in this already-crowded space is due to use of the rich structural information that is also accurate at the photometric side, permitting accurate measurements to be taken (both in 2- and 3-D). The strength of GMDS lies within the fact that many geodesic distances can be measured quickly, e.g. using FMM with friendly hardware architecture. This, in turn, gives many sample points on the surface. PCA can autonomously assist the weighting of the different distances, just as GMDS autonomously finds approximations of analogous points. By forcing GMDS to latch onto features that are more easily identifiable, competitive performance can probably be attained.

One suggested path to explore would be to try the finer quality database from Texas. Generally speaking, one of the profound advantages (or 'selling points') of GMDS is that it is generic, thus it need not depend on *a priori* knowledge

or markup about the problem domain. To measure the merit of GMDS based on face data is like judging a Swiss army knife based solely on the sharpness of its blades. Its versatility and problem-agnosticism ought to accentuate its real power, hinged upon sophistication and not *ad hoc* or brute force methods. While perfectly acceptable given the ubiquity of the application of face recognition, this might not be the ideal domain on which to use GMDS. The intrinsic appearance of faces is simple enough for a human observer to interpret accurately, whereas it's flexible tissue and underlying morphology where humans lose the ability to discern one from another.