Two computer-vision papers co-authored by BGU CS researchers have been accepted to CVPR 2018

Two computer-vision papers co-authored by researchers from the Department of Computer Science at Ben-Gurion University have been accepted to the IEEE conference on Computer Vision and Pattern Recognition (CVPR 2018).

The first paper, **"Curve Reconstruction via the Global Statistics of Natural Curves"**, was co-authored by PhD Student Mr. **Ehud Barnea** and Prof. **Ohad Ben-Shahar**, both from BGU CS.

The second paper, **"Deep Diffeomorphic Transformer Networks"**, was co-authored by Mr. **Nicki Skafte Detlefsen**, Dr. **Oren Freifeld**, and Dr. **Soren Hauberg**. Dr. Freifeld is with BGU CS while his two co-authors are with the Technical University of Denmark. During Mr. Skafte Detlefsen's MSc under the joint supervision of Dr. Hauberg and Dr. Freifeld, he also visited at BGU CS as part of this international collaboration.

The work described in [Barnea & Ben-Shahar], seeking to reconstruct missing image parts, focuses on the veridical reconstruction of curve gaps and presents **a data-driven approach for generating curves that are both visually pleasing and close to the true underlying structure** by employing the global shape statistics of natural curves. An example can be seen in the figure of the flower below, where the completed curve obtained by the proposed method (green) outperforms the one obtained by a competing method (purple). This novel method may be used, e.g., to improve images by removing distracting or unwanted elements by reconstructing the hidden structures.

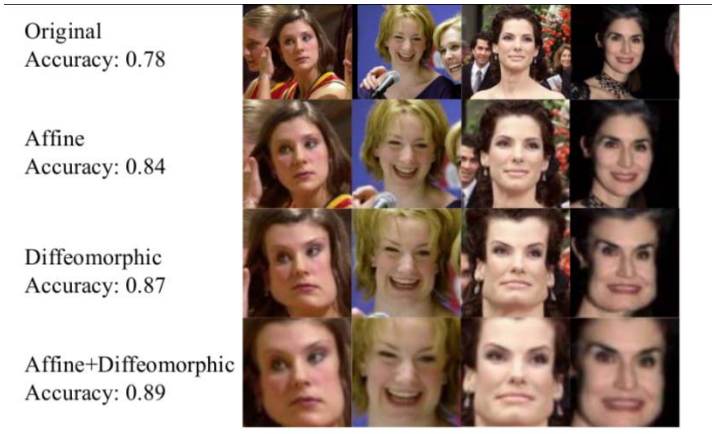

Figure 1: The spatial transformer layer improves performance of deep neural networks for face verification. By learning an affine transformation, the network can "zoom in" on the subjects face; when learning a flexible transformation (proposed), the network here stretches an oval face to become square. This significantly improves performance.

The second paper, by Skafte Detlefsen et al., **combines ideas from deep learning and geometry** and is the first to incorporate highly-expressive diffeomorphisms (namely, well-behaved spatial transformations) explicitly within deep-learning computer-vision architectures. This results in deep nets with better expressive power and improved task-dependent invariance to spatial deformations. For example, the figure below shows how, in a face-verification task, the system automatically gets rid of most of the uninformative pixels by applying a learned input-dependent nonrigid deformation of a face of a person such that it almost fills the entire image. As verified empirically on real-world datasets, this results in significantly-increased accuracy.