STRUCTURES
CLUSTER OF
EXCELLENCE

UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

STRUCTURES JOUR FIXE

Ullrich Köthe

IWR, Universität Heidelberg

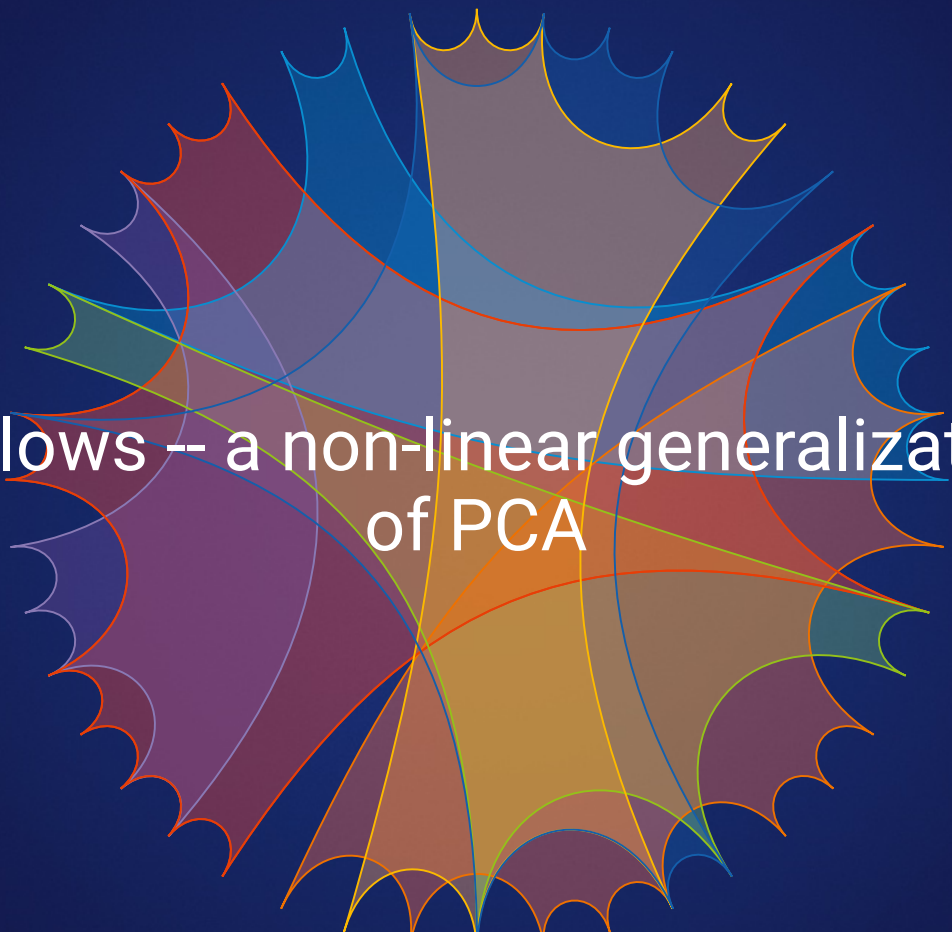

EOFlows -- a non-linear generalization of PCA

Pretalk by Armand Rousselot will start at 1 pm

May 08, 2026, 1:30 PM, Phil 12 GHs

COFFEE & SNACKS IN ROOM 106

ZOOM: Meeting ID: 935 6549 3662, Code: 928036

CONTACT: office@structures.uni-heidelberg.de

STRUCTURES
CLUSTER OF
EXCELLENCE

UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

ABSTRACT

Principal Component Analysis (PCA) is a standard tool to identify the important factors of variation of a dataset in an unsupervised manner. From a machine learning perspective, PCA can be interpreted as an encoder-decoder pair restricted to linear transformations. Normalizing flows (NF) are a natural non-linear generalization of encoder-decoder architectures, but lack the interpretability of PCA.

Entropy-ordered flows (EOFlows) overcome this limitation by using orthogonality regularization during NF training, such that the converged decoder induces an approximately orthogonal curvilinear coordinate system that is aligned with the data geometry. As a result, each latent dimension has a distinct semantic effect, and different dimensions can be sorted by importance according to their "explained entropy", analogous to "explained variance" in PCA. We show how EOFlows are trained, what factors they find on the portrait dataset CelebA, and how stable their outputs are under repeated training and in comparison to existing methods. Ideally, the colloquium would identify promising EOFlow applications in physics to pursue in the future.

ZOOM: Meeting ID: 935 6549 3662, Code: 928036
CONTACT: office@structures.uni-heidelberg.de