Dimensionality reduction with normalizing flows

Peter Bouss[1,2], Sandra Nestler[1,2], Alexandre René[1,3,4], Moritz Helias[1,4]

[1] Institute of Neuroscience and Medicine (INM-6) and Institute for Advanced Simulation (IAS-6) and JARA-Institute Brain Structure-Function Relationships (INM-10), Forschungszentrum Jülich, Jülich, Germany
[2] RWTH Aachen University, Germany
[3] Department of Physics, University of Ottawa, Ottawa, Canada
[4] Department of Physics, Faculty 1, RWTH Aachen University, Aachen, Germany

JÜLICH FORSCHUNGSZENTRUM

RTG 2416 MultiSenses Scales

RWTH AACHEN UNIVERSITY

Summary

- We demonstrate nonlinear dimensionality reduction using invertible neural networks (INNs).
- This is achieved by enforcing reconstruction with few (meaningful) dimensions.
- The method shows comparable results to PCA on close to Gaussian data.
- The network models require much fewer parameters on larger datasets than linear models.

Normalizing flows

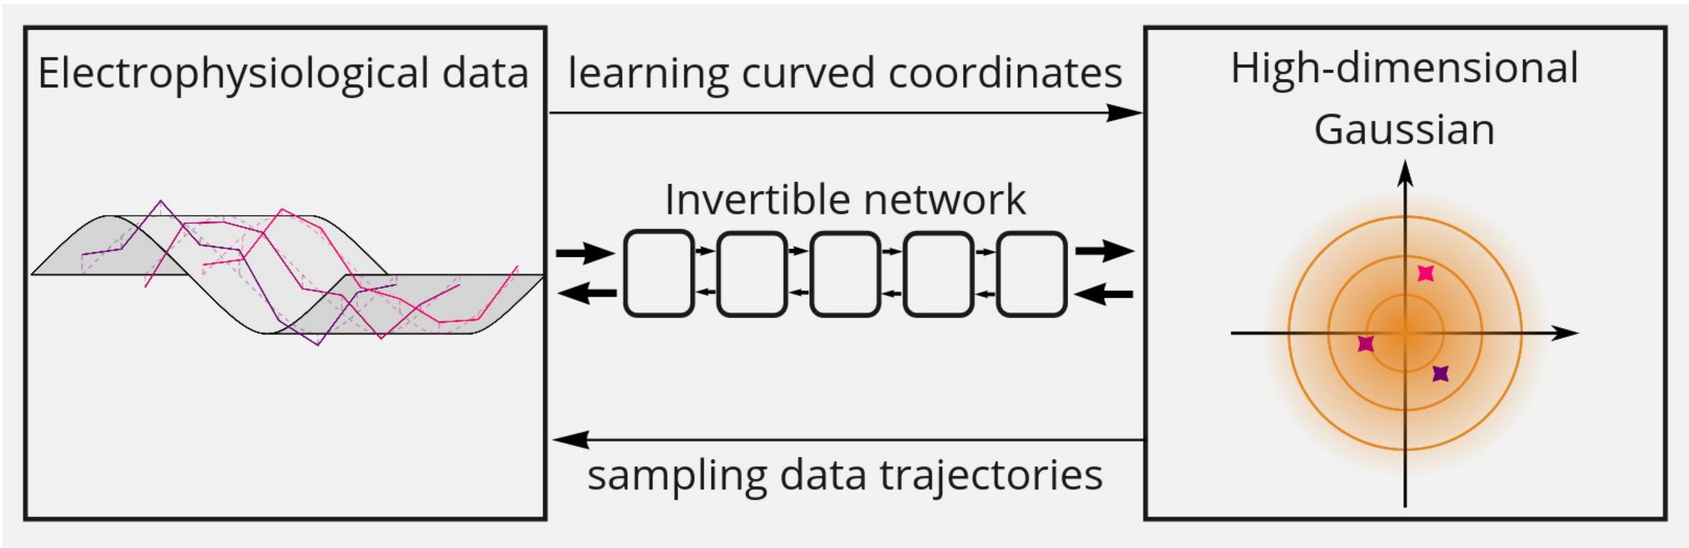

Enforce meaningful latent dimensions

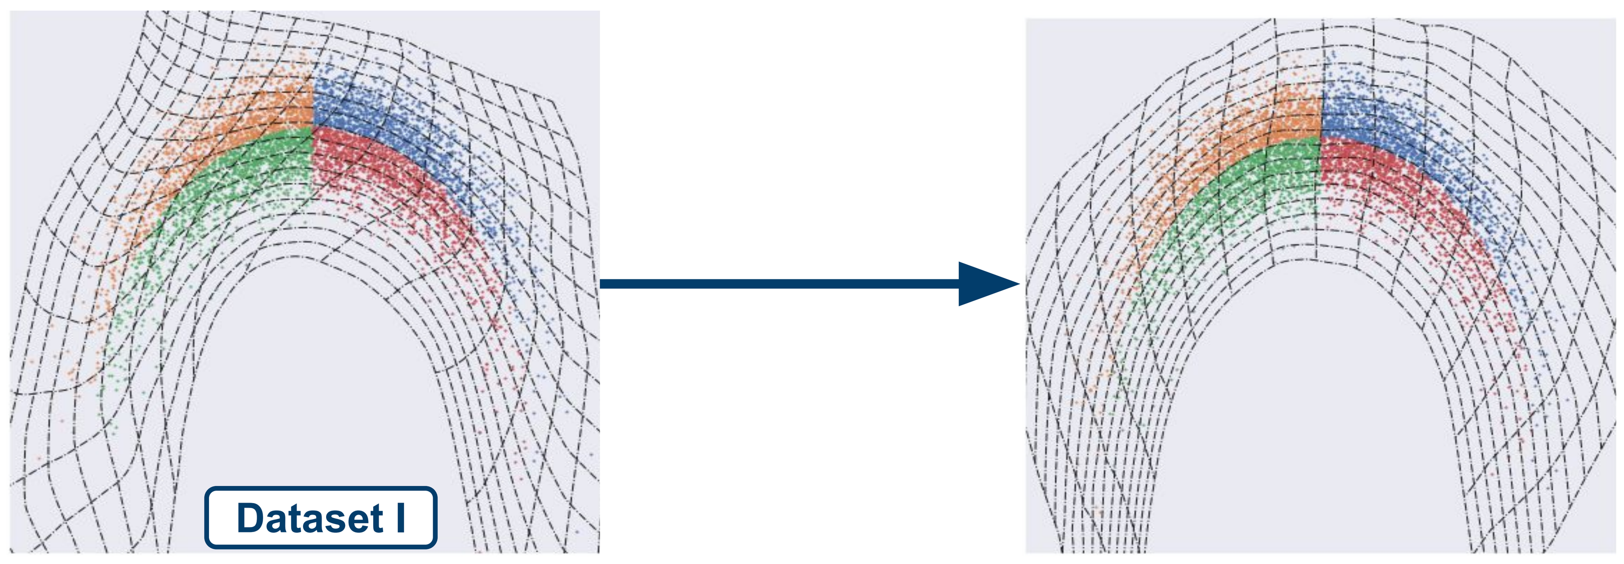

Multiscale Architecture

- At every scale: Invertible convolutions - additive / affine coupling layers are used [2-4]

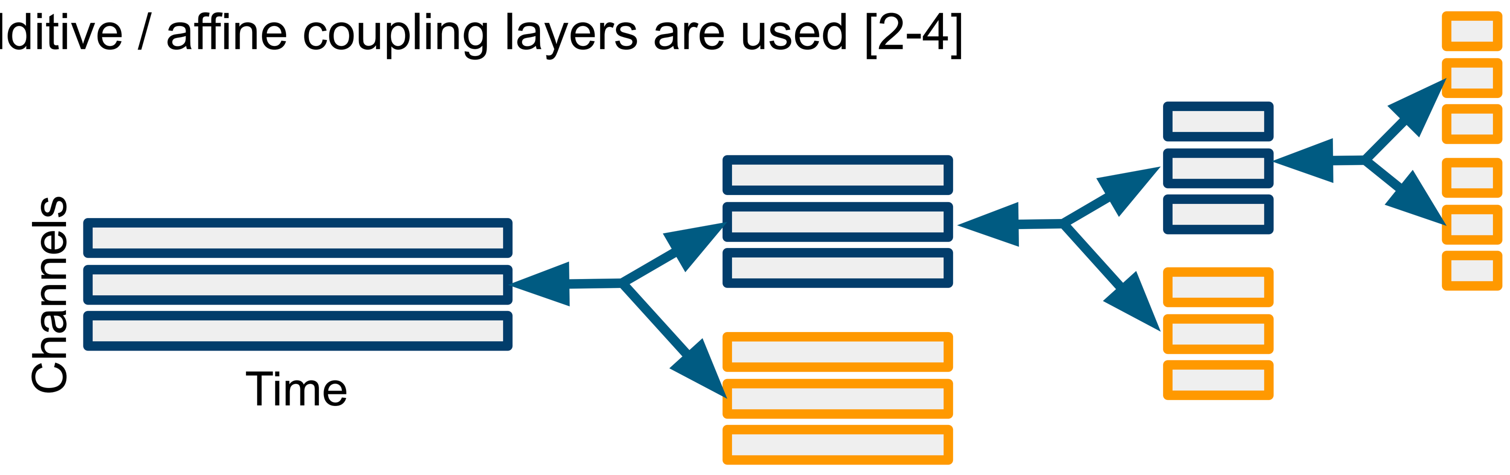

Reconstruct data with few dimensions

- In the latent space, data features are entangled and use all dimensions.
- To reduce dimensionality, we identify data-relative dimensions in the latent space.
- Reconstructed data with *i* dimensions
- Mean squared error for *i* dimensions
- Loss term added to standard loss enforcing low-dimensional reconstruction

$$\tilde{x}_i(x) = f^{-1}\left(\sum_{k=1}^{i} f_k(x)\,\hat{e}_k\right)$$
$$\tilde{x}_0(x) = f^{-1}(0)$$
$$R_i = \left\langle |\tilde{x}_i(x) - x|^2 \right\rangle_x$$
$$\mathcal{L}_{\text{dim}} = \gamma_{\text{dim}} \cdot \langle R_i \rangle_i$$

- For both data sets, INN and PCA have very similar performance → data is close to multivariate Gaussian
- Controlled data is clearly reconstructed by a few dimensions

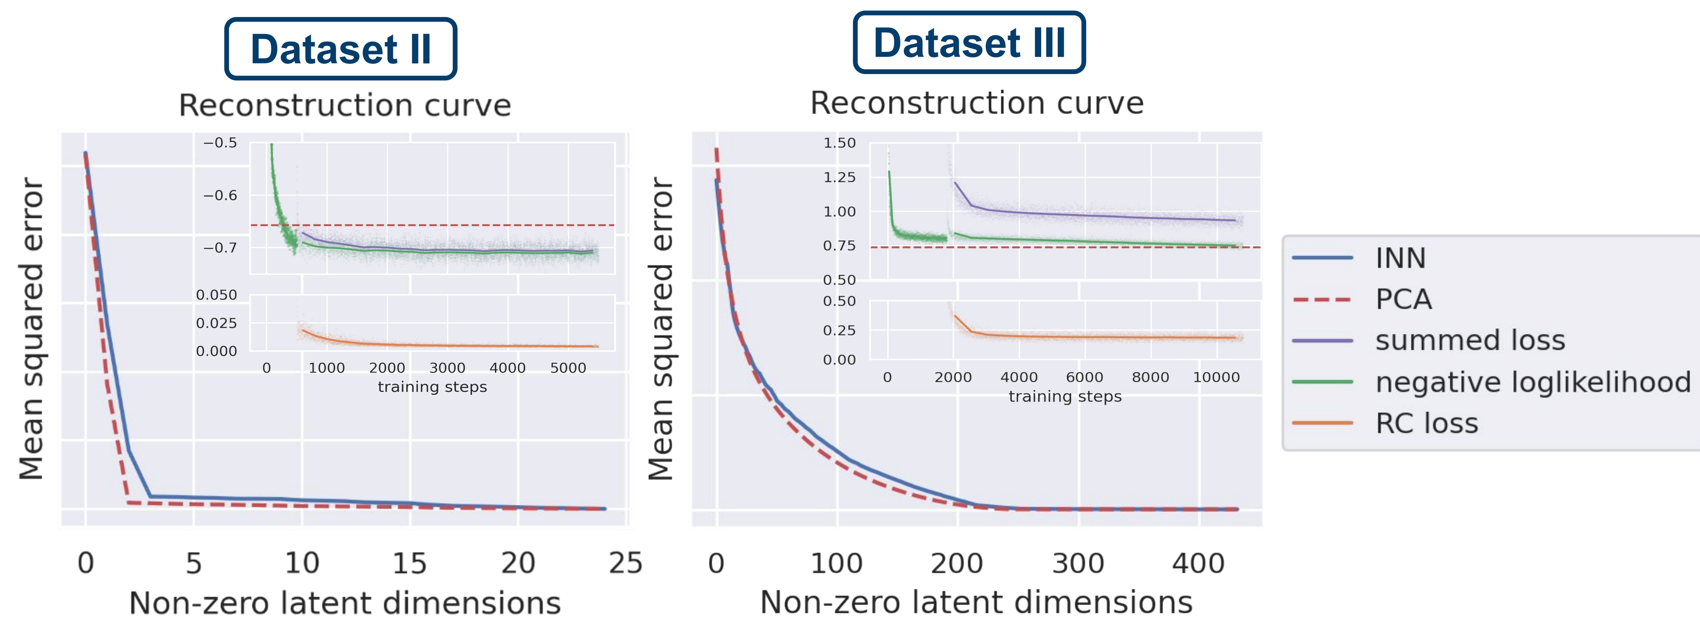

Data sets

I: Gaussian on Circle

- Phases and radii drawn from Gaussian distributions

II: Oscillations

- Oscillations with phases drawn from bimodal Gaussian
- Added white noise

III: EEG recordings

- High Gamma dataset from Tonio Ball's lab [1]
- Selected three channels subsampled to 32 Hz

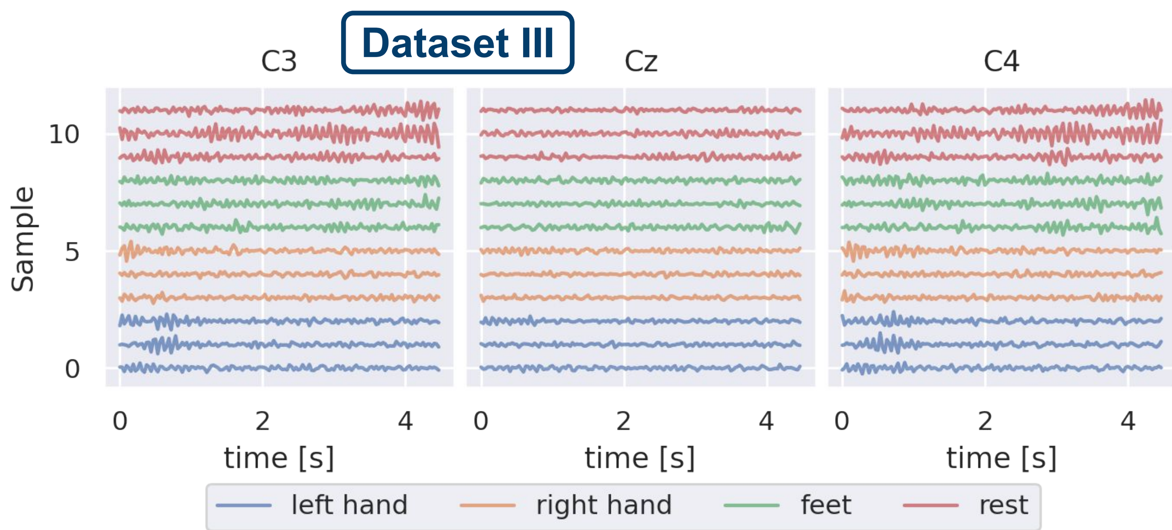

Identified latent dimensions enable effective reconstruction and sampling

- In controlled setting, only three latent dimensions modulate main information.
- Further dimensions provide the noise contribution.

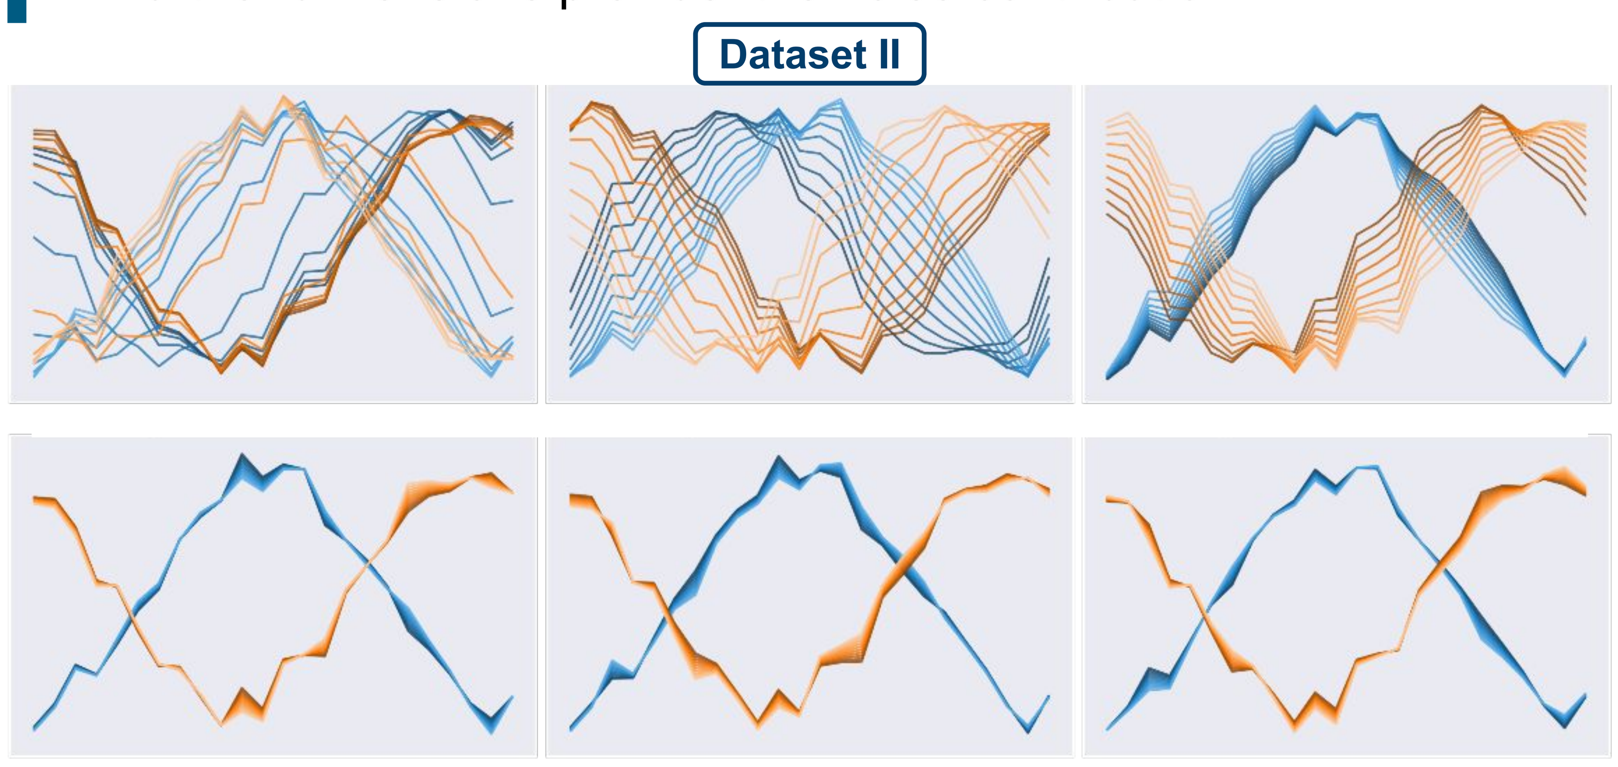

- Few dominant dimensions (<=10) contain prominent peak at 11 Hz.
- Dimensionality of < 30% sufficient to reproduce spectra with high fidelity.

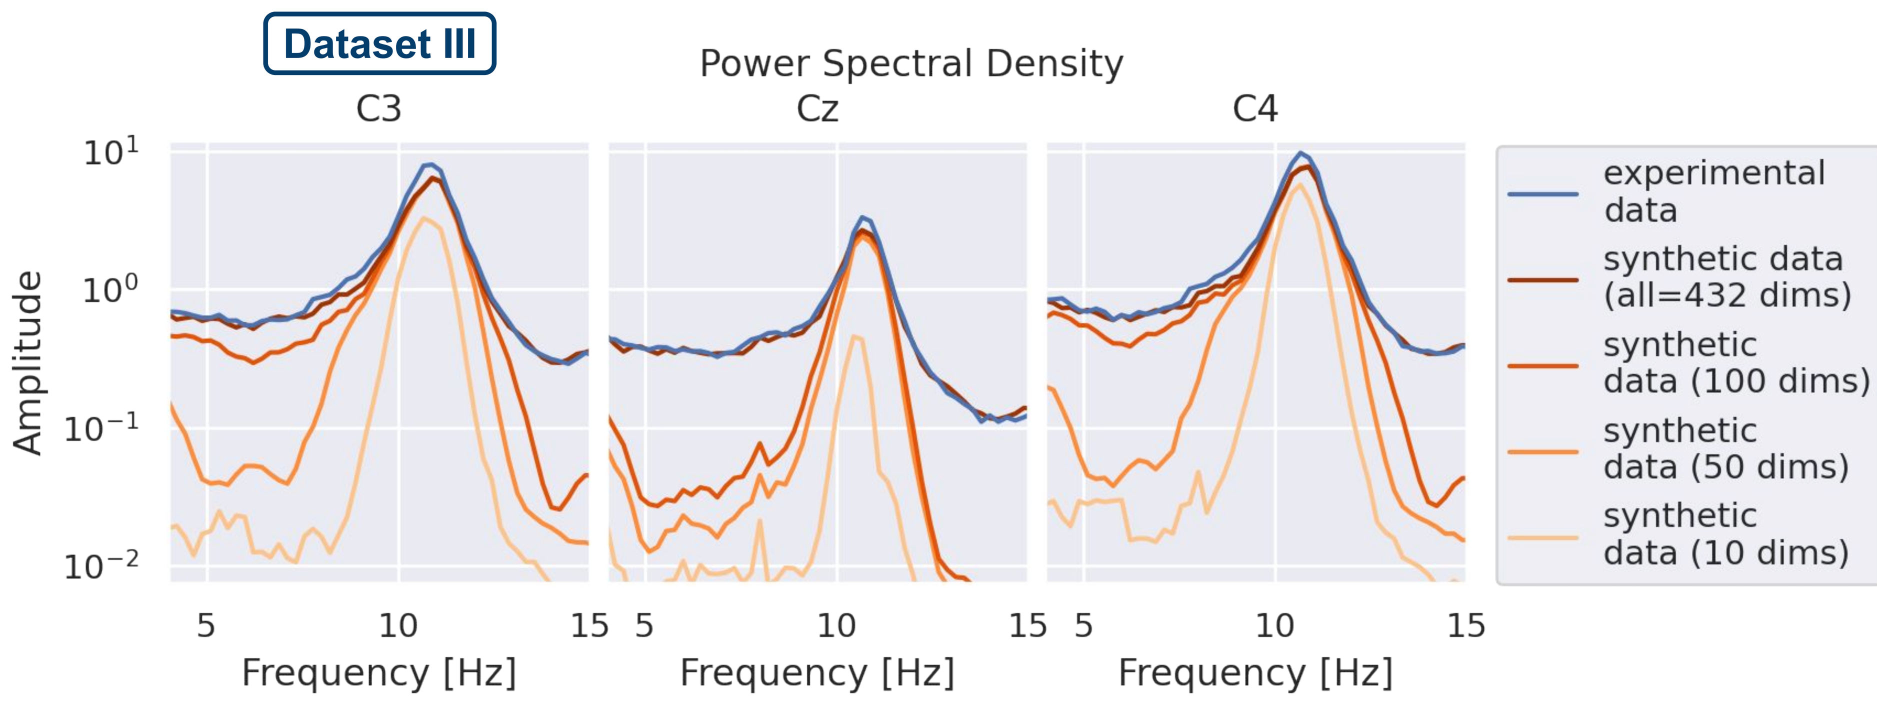

References

[1] Schirrmeister et al. (2017). Deep learning with convolutional neural networks for EEG decoding and visualization. Human brain mapping.
[2] Dinh et al. (2014). Nice: Non-linear independent components estimation. arXiv.
[3] Dinh et al. (2016). Density estimation using real nvp. arXiv preprint arXiv.
[4] Kingma & Dhariwal (2018). Glow: Generative flow with invertible 1x1 convolutions. Advances in neural information processing systems.