Institute for Digital Communications
Prof. Dr.-Ing. Ralf Müller
Friedrich-Alexander-Universität Erlangen-Nürnberg

Bachlor Thesis / Research Project / Master Thesis

Empirical Analysis of Convex Separability in Modern Deep Learning Frameworks

Background Deep learning and specifically neural networks (NNs) [3] have emerged as some of the most powerful algorithms applicable to a wide range of problems [2, 7, 4]. Recent results have investigated the convexity of high-dimensional datasets, both in the input domain [8], as well as in the latent space of feature extractors [6]. It is found that image datasets such as CIFAR-10 and MNIST form convex hulls in the input domain and admit (partially) convex decision boundaries in their respective latent representation. Convexity is, furthermore, associated with few-shot learning and generalization [6]. This finding motivates us to study the emergence of convex decision boundaries in modern deep learning architectures and datasets.

Convex Separability and Nearest Convex Hull Classification The emergence of convex hulls in image datasets intuitively motivates the application of nearest convex hull classification (NCHC) [5]. NCHC of n classes requires to solve n constrained quadratic programs, and moves the complexity of gradient descent from training to inference. The concept of convex separability is defined analogously to linear separability but with respect to NCHC: If there exists a training set $\mathcal{Z}_{\mathrm{tr}} \subseteq \mathcal{Z}$ which admits an NCHC generalization error $\epsilon_{\mathrm{g}} = 0$ on a classification problem $\mathcal{Z} = \{\mathcal{Z}_1, \mathcal{Z}_2\}$, then $\mathcal{Z}$ is convex separable [1]. See Fig. 1 for examples. It is postulated that the generalization error ϵ_{g} of NCHC is a measure for the convexity of decision regions, similarly to the Euclidean measure introduced by Tětková et al. [6].

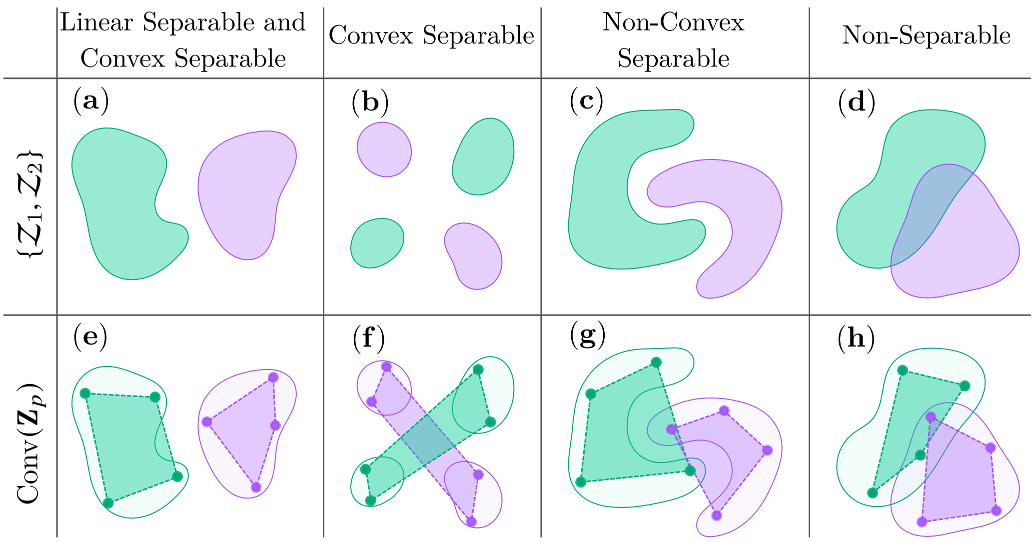

Figure 1: Examples of convex separable and non-convex separable classification problems. (**a**) A learning task which is both linear and convex separable, a possible training set configuration and the corresponding convex hulls which admit $\epsilon_{\mathrm{g}|\mathcal{Z}_{\mathrm{tr}}} = 0$ are depicted in (**e**). (**b**) An XOR style classification problem which is convex separable (see (**f**)), but not linear separable. (**c**) A non-convex separable dataset, where (**g**) plots possible corresponding convex hulls. (**d**) A non-separable problem, as the sets $\mathcal{Z}_1$ and $\mathcal{Z}_2$ overlap. (**h**) Two convex hulls of overlapping training sets.

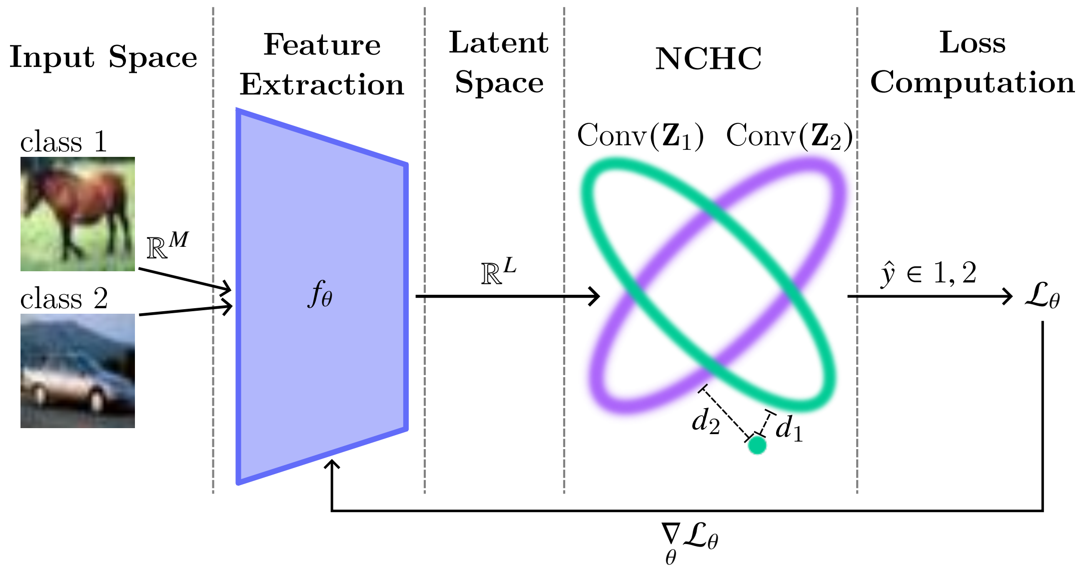

Figure 2: Illustration of an end-2-end training procedure of a feature extractor followed by NCHC read-out layer.

Recent Results The theoretical and empiric study [1] found that there are convex separable problems on which multi-layer perceptrons (MLPs) struggle, whereas NCHC excels with significantly reduced error floors at lower number of training samples. This effect, however, was only observed on synthetic data which is asymptotically convex separable in the noise-free regime. As the noise-power is increased, the MLP exhibits lowered generalization error when compared to NCHC. We postulate that this is due to the convex decision regions breaking down systematically which hurt NCHC performance while the MLP can learn non-convex decision boundaries.

Overall Project Aim The project has two main aims. Firstly, we want to analyze if NCHC can translate the significantly reduced generalization errors on convexly separable synthetic data [1] to latent space representations of real-world data. This includes the study of NCHC on latent representations of NN-based feature extractors, end-2-end training/fine-tuning of a NN followed by NCHC read-out layer as illustrated in Fig. 2, and the study of further datasets of different modalities, i.e. video, audio or text. Secondly, we need to refine our understanding of how convex separability relates to convex decision boundaries. To this end the synthetic dataset of [1] shall be analyzed for convex decision boundaries using the Euclidean convexity measure of [6]. Additionally, we aim to investigate if modern deep-learning architectures handle convex separable tasks better than standard MLPs.

Project Objectives The main objectives of this thesis can be summarized as follows:

- Identification of state-of-the-art (SOTA) feature extractors on image datasets.
- Evaluation of NCHC on latent space representation of feature extractors on 1-3 datasets.
- Implementation of end-2-end fine-tuning of the feature extractor and NCHC.
- Comparison of NCHC generalization error and Euclidean convexity measures.
- Evaluation of modern deep learning architectures (transformers, convex neural networks, modern Hopfield models, ...) on the synthetic, convex separable dataset.

The objectives may be adapted based on the project type (Bachelor Thesis, Research Project, Master Thesis).

Prerequisites The following qualifications are required in order to pursue this project:

- Good knowledge of linear algebra.
- Basic knowledge of (convex) optimization.
- Good programming skills in Python (JAX, PyTorch or Haiku, numpy).
- Beneficial: experience with Git.
- Beneficial: experience with convex solvers (e.g. cvxpy or jaxopt).
- Beneficial: experience with bash and slurm.

Bibliography

[1] Levi-Pascal V. G. Bohnacker and Ralf R. Müller. "Learning Convex Decision Regions is Hard: When Neural Networks Fail on Convexity". Unpublished manuscript.

[2] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. "ImageNet classification with deep convolutional neural networks". en. In: *Communications of the ACM* 60.6 (May 2017), pp. 84–90. DOI: 10.1145/3065386.

[3] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. "Deep learning". en. In: *Nature* 521.7553 (May 2015), pp. 436–444. DOI: 10.1038/nature14539.

[4] James Murphy, John E Ward, and Brian Mac Namee. "An Overview of Machine Learning Techniques for Onboard Anomaly Detection in Satellite Telemetry". en. In: *2023 European Data Handling & Data Processing Conference (EDHPC)*. Juan Les Pins, France, Oct. 2023, pp. 1–6. DOI: 10.23919/EDHPC59100.2023.10396403.

[5] Georgi Nalbantov, Patrick Groenen, and Cor Bioch. *Nearest convex hull classification*. Tech. rep. Econometric Institute, Erasmus University Rotterdam, 2006.

[6] Lenka Tětková et al. "On convex decision regions in deep network representations". en. In: *Nature Communications* 16.1 (July 2025), p. 5419. DOI: 10.1038/s41467-025-60809-y. URL: https://www.nature.com/articles/s41467-025-60809-y.

[7] Ashish Vaswani et al. "Attention is All you Need". en. In: *Advances in Neural Information Processing Systems*. Vol. 30. Curran Associates, Inc., 2017.

[8] Roozbeh Yousefzadeh. "Deep learning generalization and the convex hull of training sets". en. In: *Information Geometry* (Oct. 2025). DOI: 10.1007/s41884-025-00178-z.

Supervisor Levi-Pascal Bohnacker, levi.bohnacker@fau.de, Room 04.037

Start date: TBD
End date: TBD

Prof. Dr.-Ing. Ralf R. Müller