KESTRELEYE – Phoenix

Automated training of neural networks

Version v2.2.4 | *Released* August 2026

Fully automated end-to-end training software for neural networks, particularly for image processing. Offers user-friendly selection of models and training recipes with only basic requirements for understanding of the underlying technology. Designed to train networks for later evaluation using the KESTRELEYE VEGA software and offers a broad range of networks covering both low-latency and high-throughput as well as slower, high-accuracy state-of-the-art architectures.

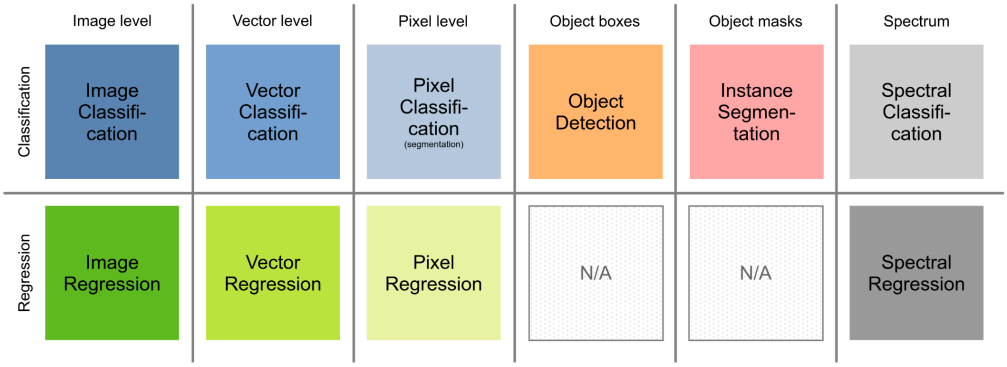

Models available for training in Phoenix. All models can be made available for monochrome, RGB, or hyperspectral data.

Functionalities offered

Data handling: Automated dataset generation (from prepared data), including validation and report summary

Data pre-processing: Various ways of data pre-processing are available, including pre-scaling, sampling, weighting, augmentation and (limited) synthetic data generation

Network training: Automated training of neural networks, including

- Wide selection of state-of-the-art and tailored, low-latency and high-throughput network architectures
- Pre-trained weights
- Training procedure on CUDA-capable GPUs or CPU
- Testing of networks using standard performance metrics, export to machine-readable files
- Comprehensive, easy-to-read report in PDF format
- Export of models to encrypted ONNX
- Logging throughout the entire training loop

Data evaluation: Tools to evaluate data using trained models and export the results

Models and use case examples

Image Classification: e.g. product categorization, sorting applications

Pixel Classification (segmentation)**:** e.g. defect detection / quality control, change detection

Image Regression: e.g. quality or age estimation, scoring, angle estimation

Pixel Regression: e.g. density or depth estimation

Vector Classification (one class per image row / column)**:** e.g. audio (spectrogram) analysis, 1D quality control

Vector Regression (one measurement per image row / column)**:** e.g. line height estimation, deflection measurements

Object Detection and Instance Segmentation (localization + class (+ mask))**:** e.g. part counting, defect detection, pedestrian / car detection, monitoring

Spectral processing: available with all models above for hyperspectral data, e.g. water quality assessment or contamination estimation (also for single spectra)

Key benefits

Managed pipeline: Pre-defined and tested models, packaged with full definition of model architecture and training configuration for specific use cases

Available models: State-of-the-art, highly accurate as well as lightweight, low-latency, high-throughput models (developed by us) available for challenging industrial environments

Hyperspectral data: Handling of hyperspectral data in all models, not typically available in off-the-shelf architectures

Customization: Model library, for tailored solutions for specific use cases

Reporting: Detailed reporting with curated, useful metrics

Software packaging

Phoenix is delivered as a Docker image via container registry, with easy updates through the internet. It offers a command-line interface, API and graphical front-end.

System requirements (recommended)

OS	Ubuntu 22.04+ LTS or another Linux variant
CPU	8+ cores, 4.0 GHz or better
GPU	NVIDIA CUDA-capable GPU, 16+ GB VRAM
RAM	32+ GB (more is better for large datasets)

Contact

For product demos and enterprise enquiries:
heinz.fleischhacker@kestrel-eye.com | www.kestrel-eye.com