Object detection in historical portraits

Master research project (5 ECTS)

Object detection

- Art workshops of the 16th century often reused their motifs directly or to some extend
- To compare visually striking image patches we need to detect their location as a bounding box or as landmarks

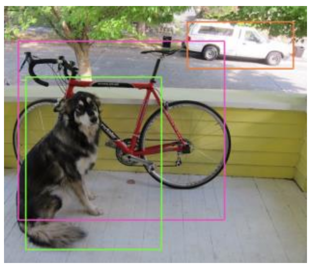

Object detector (e.g. YOLO[1])

Transfer to paintings and prints

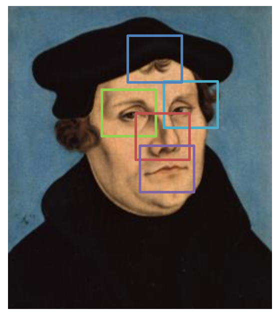

IT_GdU_1160

DE_LmKKO_15-572

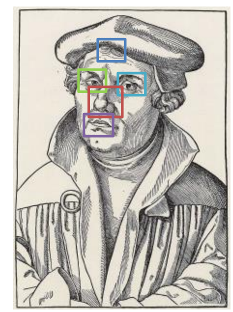

Ehemals Sammlung Liechtenstein

Bounding box detection

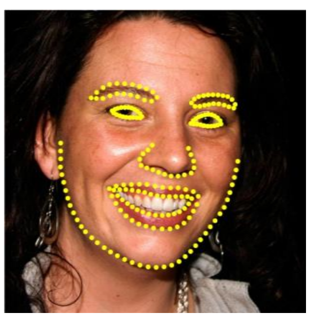

Facial landmark detector (e.g. Dlib[2])

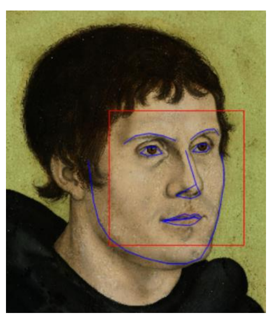

DE_GNMN_Gm1570

Transfer to prints

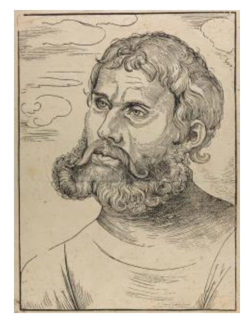

DE_GNM_Mp_14637_a

Facial landmark detection

[1] J. Redmon, S. Divvala, R. Girshick and A. Farhadi, "You Only Look Once: Unified, Real-Time Object Detection," 2016 IEEE CVPR, 2016, pp. 779-788.
[2] V. Kazemi and J. Sullivan, "One millisecond face alignment with an ensemble of regression trees," 2014 IEEE CVPR, 2014, pp. 1867-1874.
Image sources of paintings and prints: Lucas Cranach, Portrait of Martin Luther, Cranach Digital Archive (CDA) and Germanisches Nationalmuseum (GNM) Nürnberg

Task

- Adapt and extend existing machine learning / deep learning methods for:
  - Bounding box detection of visual striking elements (such as eyes, mouth, hands) in paintings and prints
  - Facial landmark detection for prints
- Implementation in Python

- Master research project (5 ECTS)
- In cooperation with the Germanisches Nationalmuseum

Contact:
Aline Sindel
Room 10.138
aline.sindel@fau.de

Tide Water Glacier Front Segmentation in Radar Images Using Deep Neural Networks

5, 10, 15, 30 ECTS: (Combined) Research Project or Master Thesis

AmirAbbas Davari: amir.davari@fau.de

Pattern Recognition Lab

University of Erlangen-Nürnberg

Motivation

- Glacier and ice sheets are currently contributing 2/3 of the observed global sea level rise of about 3.2 mm/a.
- Many glacier on glaciated regions, e.g., Antarctica, show already considerable ice mass loss in the last decade.
- The continuous and precise extraction of glacier calving fronts is hence of paramount importance for monitoring the rapid glacier changes

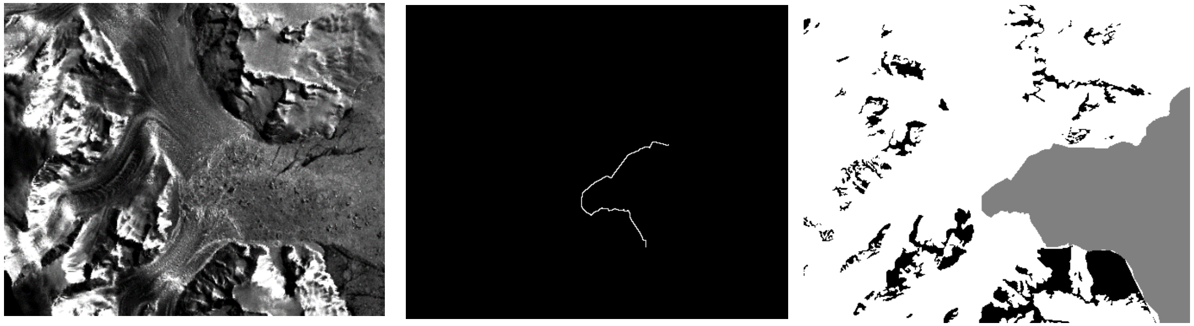

a b c

a) A sample SAR image of a glacier, b) its calving front line, and c) the corresponding region-based mask. White, grey and black regions represent ice (glacier), water (sea) and rock, respectively.

Project Description

- Fully automatic deep learning-based glacier front segmentation using time series synthetic aperture radar (SAR) imagery.

- U-net has performed extremely well in image segmentation, specifically in medical image processing community.

- The main tasks are to:
  - experiment with different U-net architecture variants,
  - investigate the optimality of each architecture for this task,
  - try newer state-of-the-art deep learning-based segmentation

- The tasks in this project are suitable for 5, 10, 15 or 30 ECTS.

- Interested? Please contact AmirAbbas Davari (amir.davari@fau.de) if:
  - you rate yourself B+ or higher in python programming language,
  - you have solid knowledge on deep learning basics.

Topics

Compressing PHOC-like representations

- PHOC-like: An vector representation of a string that can be generated from word-images
- Can we compress them? Will standard compression techniques work?
- Do they preserve their their joint image-string searchability

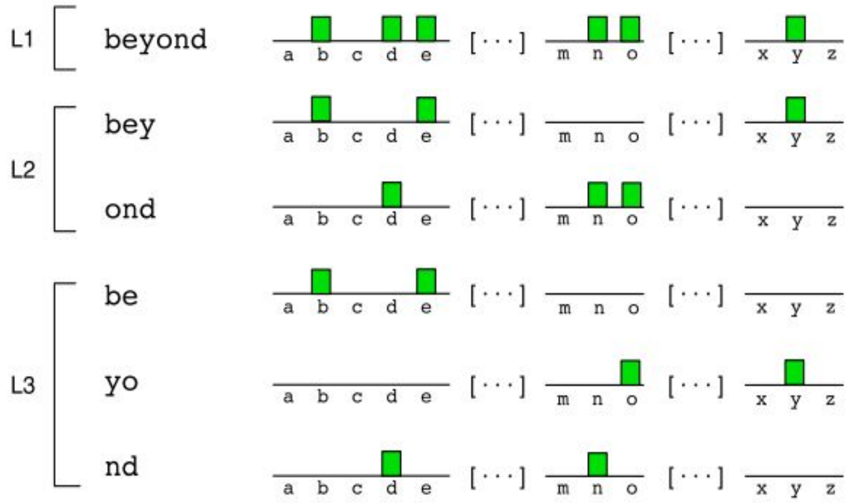

Contact: ***anguelos.nikolaou@fau.de***

ECTS: 5/10/MT (depends)

Deep regular expressions

- Regular expressions are easily compiled to NFA (Non-discrete Finite-state Automata)
- Typically regular expression engines are implemented by compiling NFA to larger DFA (Discrete Finite-state Automata)
- Can we work directly on NFA?
- Can we use it on top of a Deep Neural Network?
- What are the benefits?

Contact: ***anguelos.nikolaou@fau.de***

ECTS: 5/10/MT (depends)

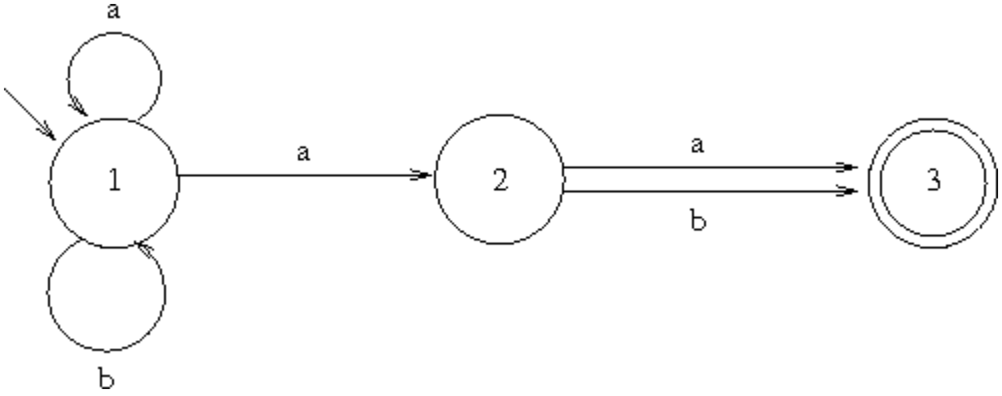

Decomposing 2D Convolutions

- 2D Convolutions complexity: N^2
- Two consecutive 1D Convolutions complexity N*2
- How much do we lose if we train on 2D and do inference on 2x1D?
- How much do we lose if we train on 2x1D and do inference on 2x1D?

Contact: ***anguelos.nikolaou@fau.de***

ECTS: 5

2D Finder

- LR Finder (fastAi) is a nice algorithm to find an adequate starting LR (s. plot right)
- Task: expand to 2D incorporating weight decay and evaluation
- 5 ECTS
- Contact: vincent.christlein@fau.de

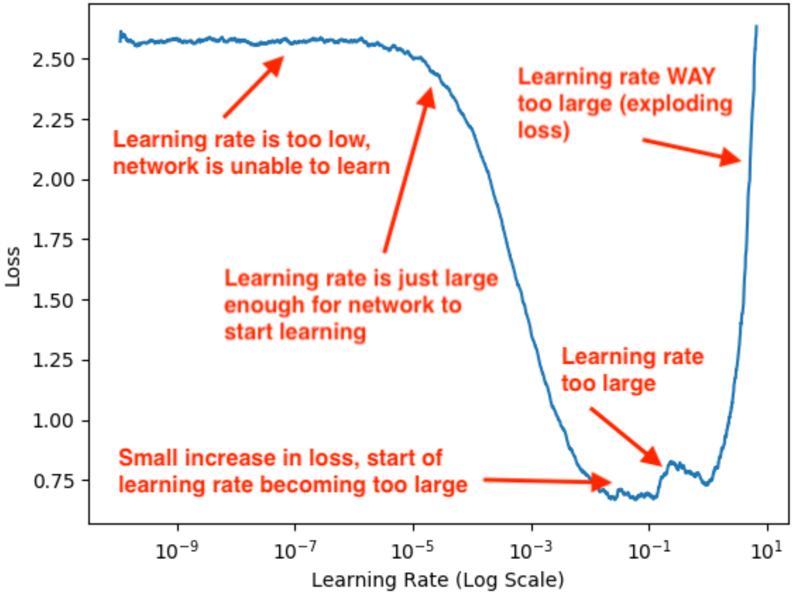

Writer Verification

- Highlight similar/dissimilar regions in writings online
- Local feature extraction + local naive Bayes NN
- Requires knowledge in Web-technologies
- 10 ECTS

Contact: vincent.christlein@fau.de

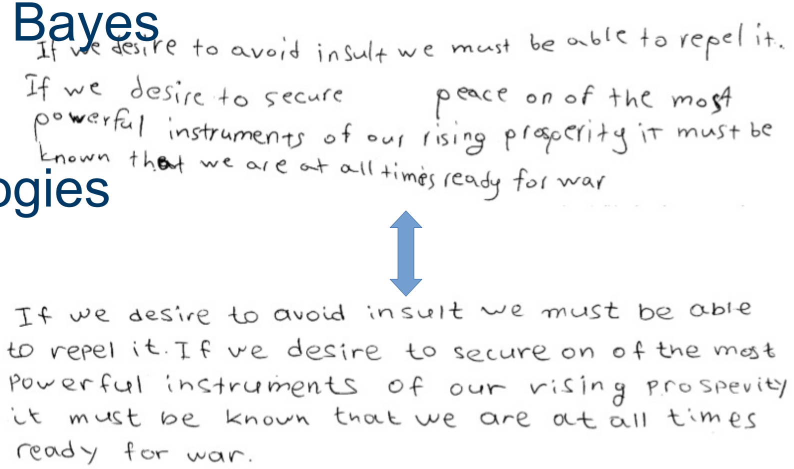

Online vs. Offline Writer Identification

- Comparison between online and offline writer identification
- Implementation of a DL-based online identification system
- Comparison with existing offline system
- 5 ECTS
- Contact: vincent.christlein@fau.de

Style Classification in Posters

- Style classification using WikiArt
  - Crawl WikiArt (images+styles)
  - Train DL-based network w. WikiArt data
  - Apply to poster data

- 5/10 ECTS Project
- Contact: vincent.christlein@fau.de

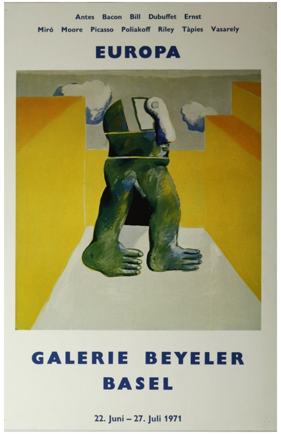

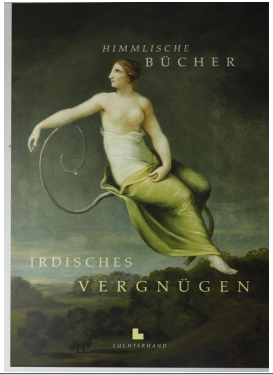

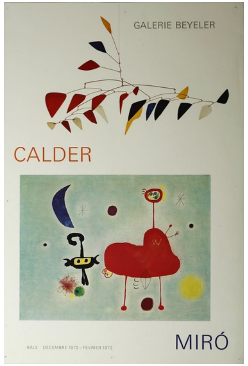

Super-resolution

- Reconstruction-based superresolution
  - Implementation in PyTorch
  - Matlab Code exists

- 5 ECTS Project
- Contact: vincent.christlein@fau.de

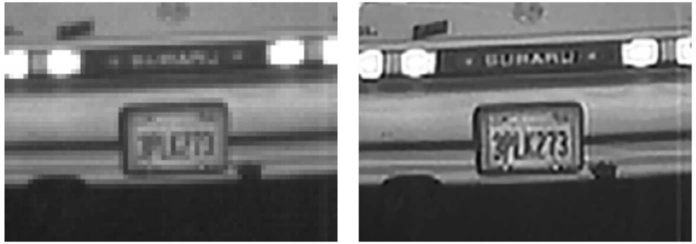

Jigsaw Puzzling of Historical Fragments

- Implementation of deep learning-based jigsaw puzzle solver
- Evaluation on historical fragment dataset
- 5/10 ECTS Project / BT
- Contact: vincent.christlein@fau.de

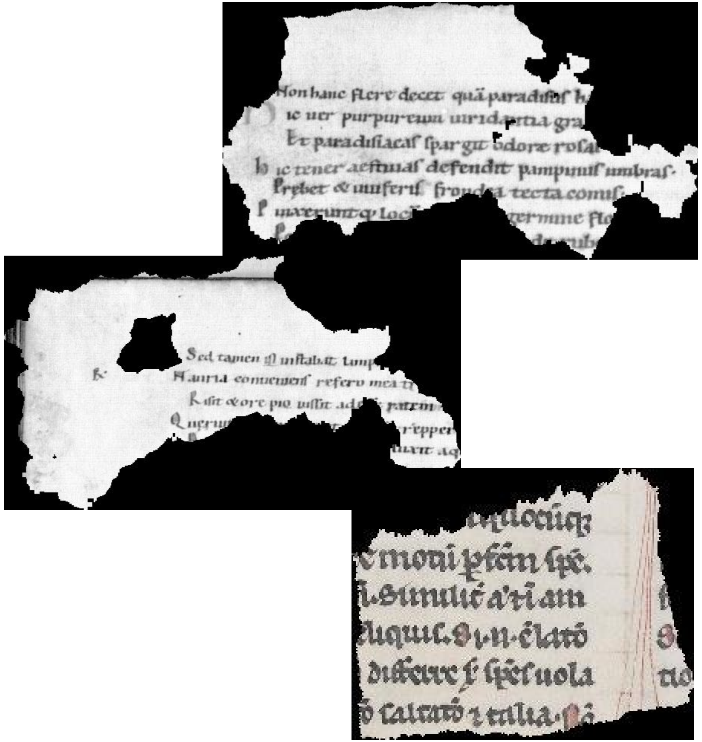

Line-based Binarization

- Comparison of two different line segmentation algorithms for binarization
- Evaluation using different DIBCO challenge datasets

- 5/10 ECTS Project
- Contact: vincent.christlein@fau.de

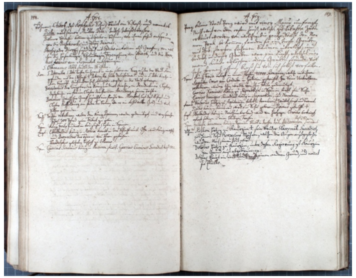

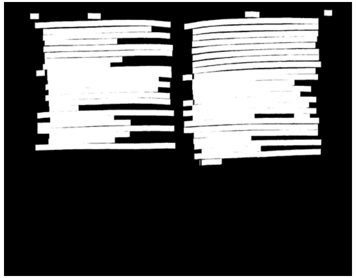

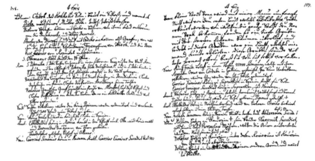

DNN Optimization in Audio

Axel Plinge

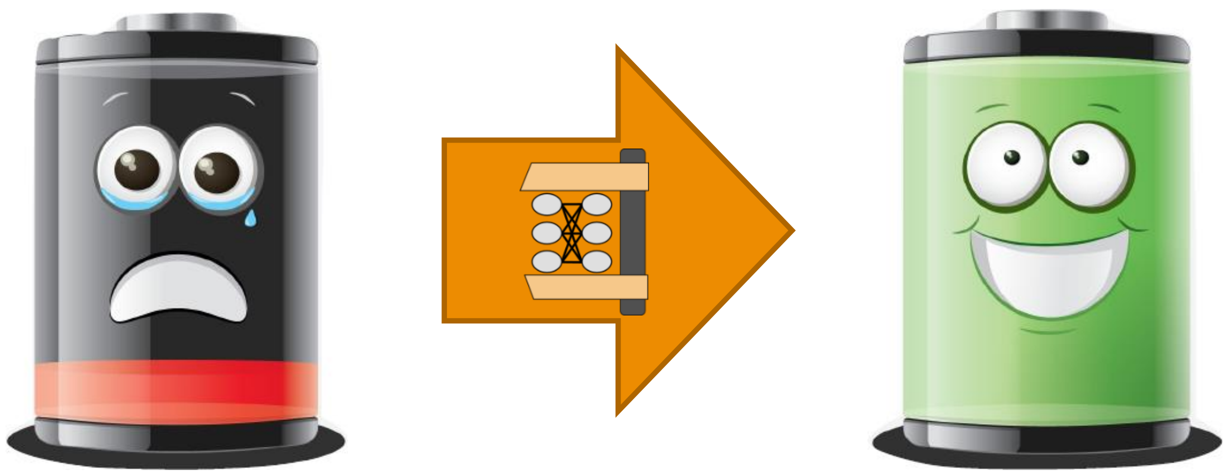

DNN Optimization in Audio

Motivation

- Training DNNs requires Graphical Processing Units (GPUs)
- They still need considerable resources (energy) at run-time
- Applications should run on embedded devices in real-time!
- It can be done: AlexNet (244MB) ➔ SqueezeNet (5MB)

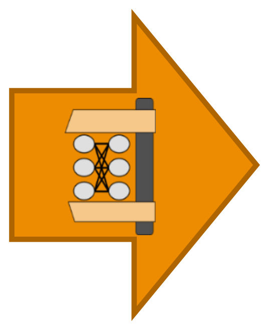

[I16] Iandola, F. N., Moskewicz, M. W., et al. "SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <1MB model size" arXiv 1602.07360

DNN Optimization in Audio

Fraunhofer IIS

- Fraunhofer IIS in Erlangen is the “home of mp3”
- 250+ employees working on audio, video, multimedia, virtual reality and more

DNN Optimization in Audio

Master Thesis Topics

We want YOU to optimize our Applications!

Apply, investigate and develop state-of-the-art deep compression methods to one of the following:

i. Speaker localization with microphone arrays & CNN
ii. Speech separation using (B-)LSTM
iii. Language modelling by RNN for natural language interfaces
iv. Speaker verification with ResNet-like architecture

[H15] S. Han, H. Mao, et al., 2015, "Deep Compression: Compressing Deep Neural Networks with Pruning, trained Quantization and Huffman coding." ArXiv:1510.00149

Multi-frame superresolution for defect detection on solar cells

Master project

Idea

- **Low**-resolution-images of solar modules → **high** resolution images
- Enables detection of more defects

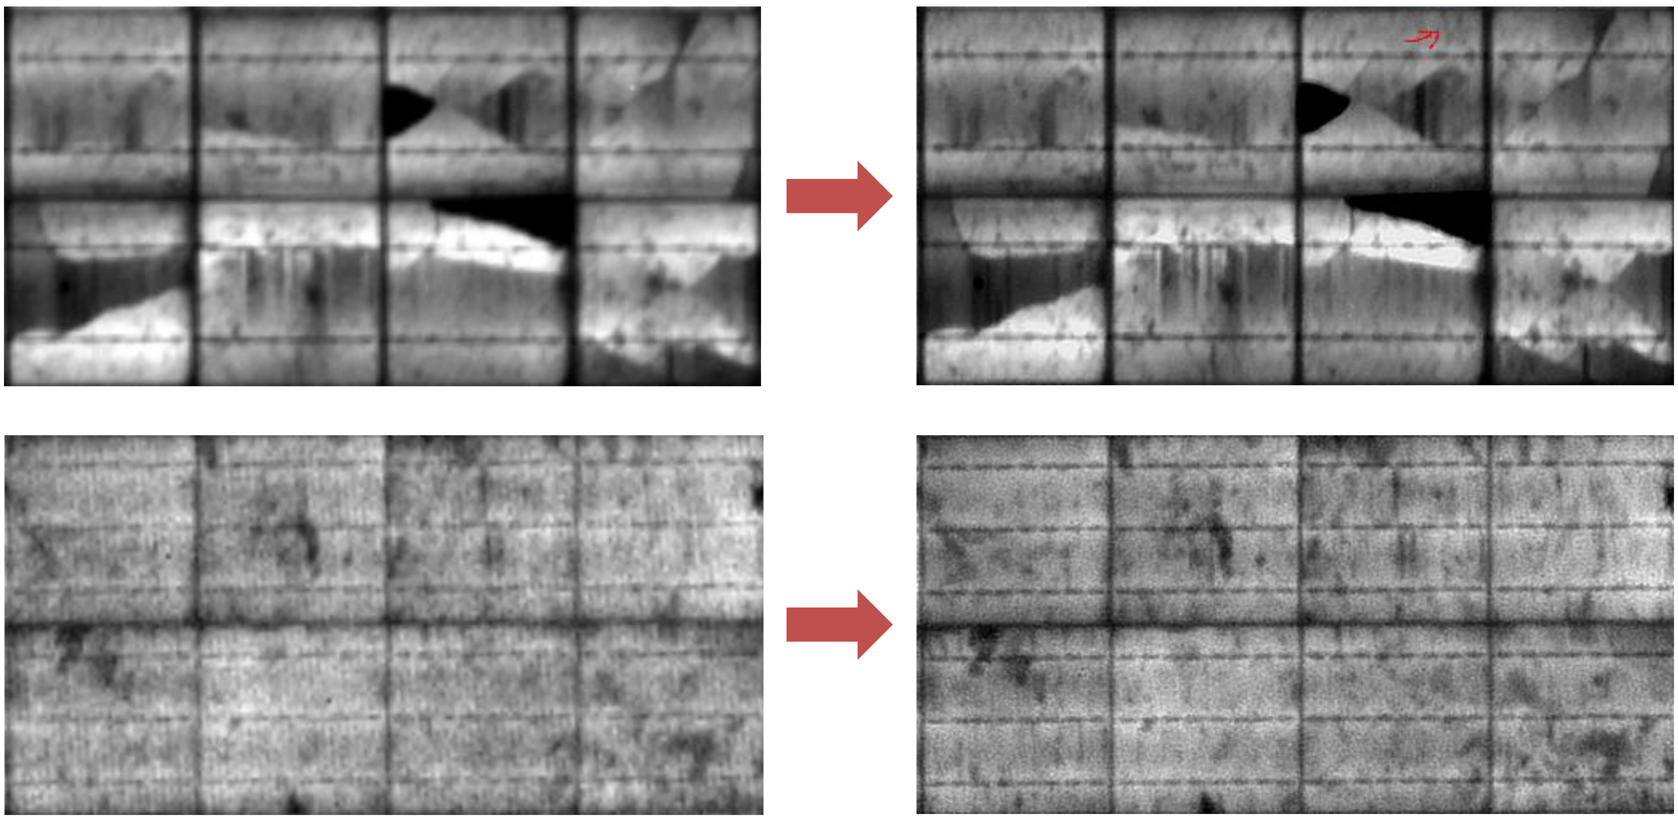

Principle

Registration

- Estimate parameters of pinhole camera
- Warp pixels into common grid

Interpolation

- Interpolate with higher resolution
- Well suited for **CUDA** implementation

Restoration

- Invert camera PSF by deconvolution

Caught you attention?

- Implement a classic CV pipeline
- Have **fun coding** C++ and CUDA
- Method is **known to work**

Contact:
Mathis Hoffmann (09.153)
mathis.hoffmann@fau.de

HI ERN
Helmholtz Institute
Erlangen-Nürnberg

Quality Control of Solarparks - Failure detection and analysis using statistical methods

Thema – Projektarbeit – Bachelorarbeit - Masterarbeit

Mai 2019 ||| Dr.-Ing. Claudia Buerhop ||| High Throughput Methods in Photovoltaics

part of
JÜLICH Forschungszentrum

in cooperation with
FAU Friedrich-Alexander Universität Erlangen-Nürnberg
HZB Helmholtz Zentrum Berlin
ZAE BAYERN

Quality control of solarparks

Inspection using imaging techniques, e.g. thermography

benefit:

Fast

Contactless – without operation interruption

During real operating conditions – during sunshine

Quality check on module level

Predicting the power the basing on IR-images is advantageous because time-consuming electrical measurements are avoided and no operating interruption is necessary.

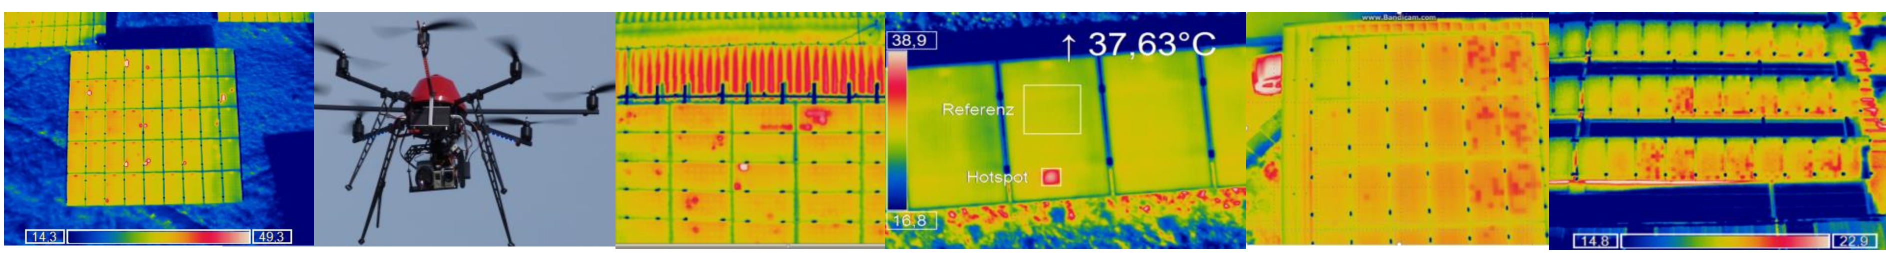

HI ERN Helmholtz Institute Erlangen-Nürnberg

Task:

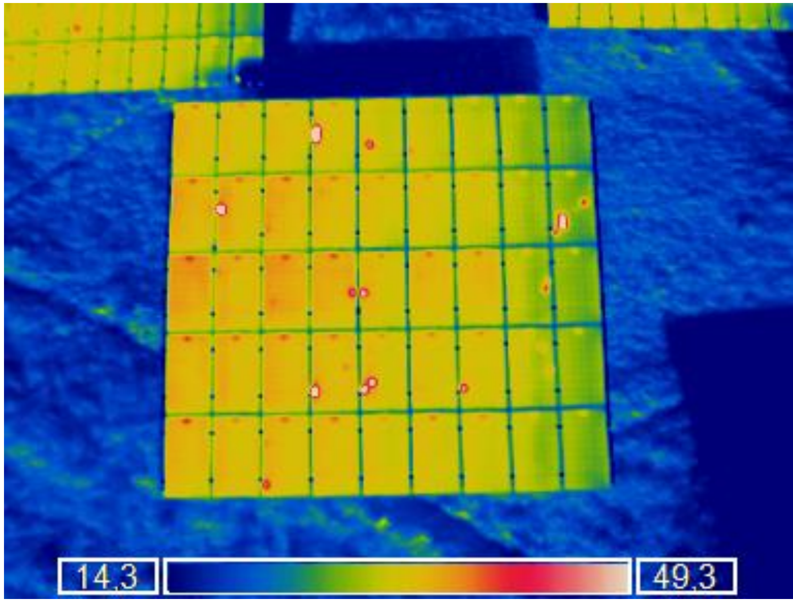

- detection of thermal anomalies
- Identification of malperforming PV-modules inbetween mostly well-performing PV-modules
- Prediction of the module power basing on IR-images

TODOs:

- Machine learning techniques for power prediction, deep learning
- Processing the recorded IR-movies and -images of PV-systems recorded at field conditions
- Training a deep learning model on modules with known power
- Ensuring that it generalizes to unknown data under varying conditions

Stitching of solar modules

Master project

Idea

- **Partial views** of a solar module → image of the complete module
- Obtain a higher resolution per cell

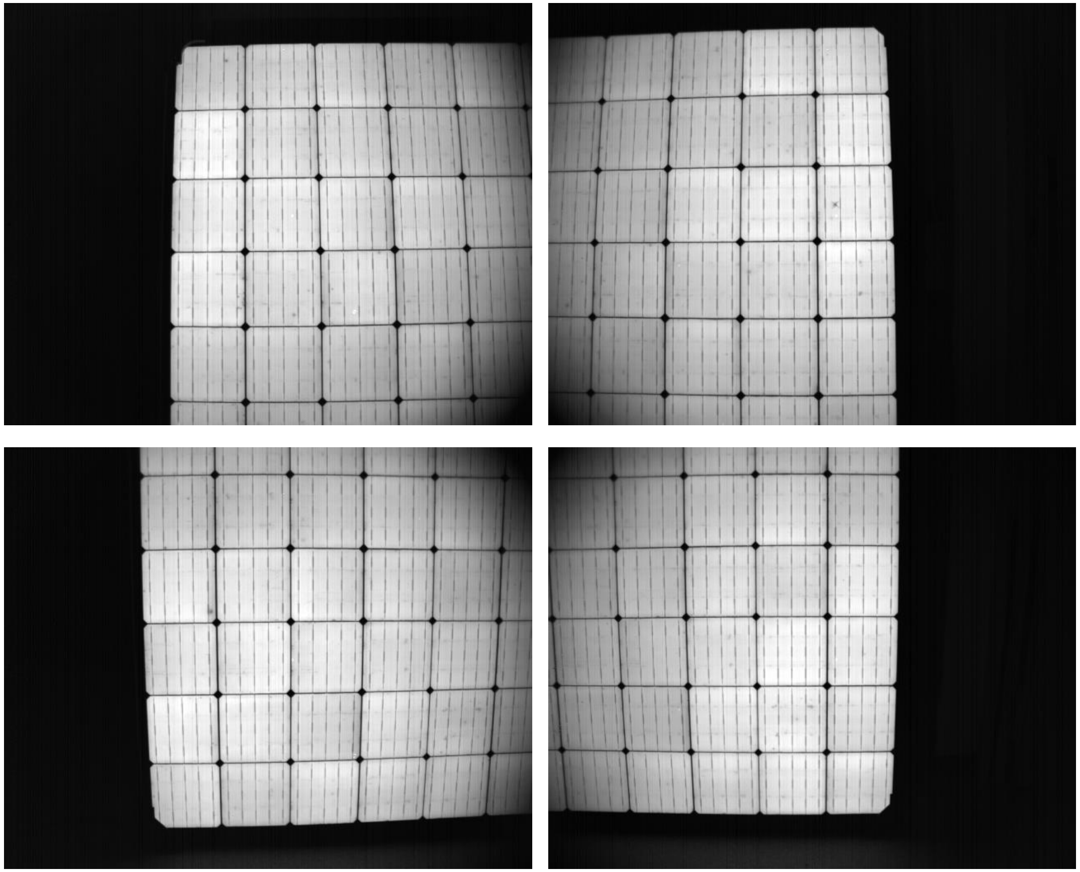

→ ?

Steps

- Detection:
  - Extend existing module detection pipeline
  - Alternatively: Code your own
- Match keypoints between images
- Compute stitched image

Caught you attention?

- Find a creative solution
- Code in whatever language you prefer
- Get 5-10 ECTS

Contact:
Mathis Hoffmann (09.153)
mathis.hoffmann@fau.de

Understand and Improve Handwriting Imitation Pipeline

Idea:

- Evaluate different parts of the existing pipeline
- Improve Writer Style Transfer (and Image Style Transfer)

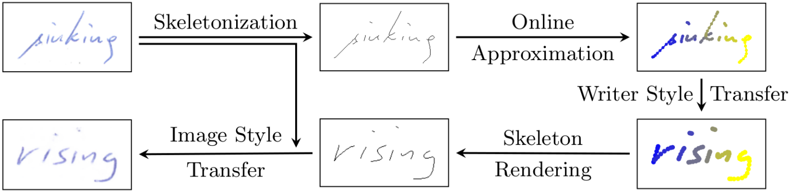

Figure: Offline-to-Offline Handwriting Style Transfer Pipeline

Understand and Improve Handwriting Imitation Pipeline

Concrete Tasks:

- Retrain all models of the pipeline
- Compute baseline
- Evaluate: Classical skeletonzation vs. pix2pix skeletonization
- Improve Writer Style Transfer
- (Optional) Implement artificial pen to simulate drawing along the points of online handwriting data while preserving the style of the input pen

General Information:

- Bachelor's Thesis / Research project (10 ECTS)
- Contact: martin.mayr@fau.de
- Implementation in PyTorch

Historical HTR using Transformers

Idea: Use methods of neural machine translation, like transformers, to build an HTR system for historical documents. For evaluation and training the Nuremberg Letters of Correspondence should be used.

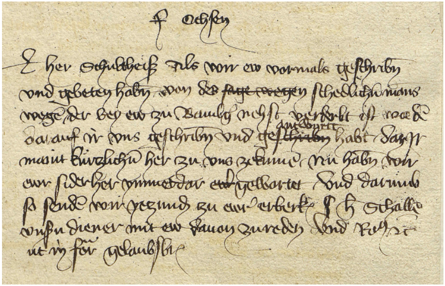

Figure: Example of Nuremberg Letters of Correspondence

Historical HTR using Transformers

Concrete Tasks:

- Literature review
- Implement pipeline to download/update data from Transkribus API
- Compute baseline with CITlab HTR method[1]
- Develop approach using transformers[2] [3] [4] (Implementation in PyTorch)
- Evaluate results
- (Optional) Evaluate performance decrease using imperfect GT for training

General Information:

- Master's Thesis
- Contact: martin.mayr@fau.de

[1] CITlab ARGUS for historical handwritten documents (https://arxiv.org/pdf/1605.08412.pdf)

[2] Attention is all you need (https://arxiv.org/pdf/1706.03762.pdf)

[3] Reformer (https://arxiv.org/pdf/2001.04451.pdf)

[4] ImageBERT (https://arxiv.org/pdf/2001.07966.pdf)

Pattern Recognition Lab

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
FACULTY OF ENGINEERING

Composition diagrams in art history

Bachelor’s Thesis/Forschungspraktikum

Image Composition Canvas (Current Work)

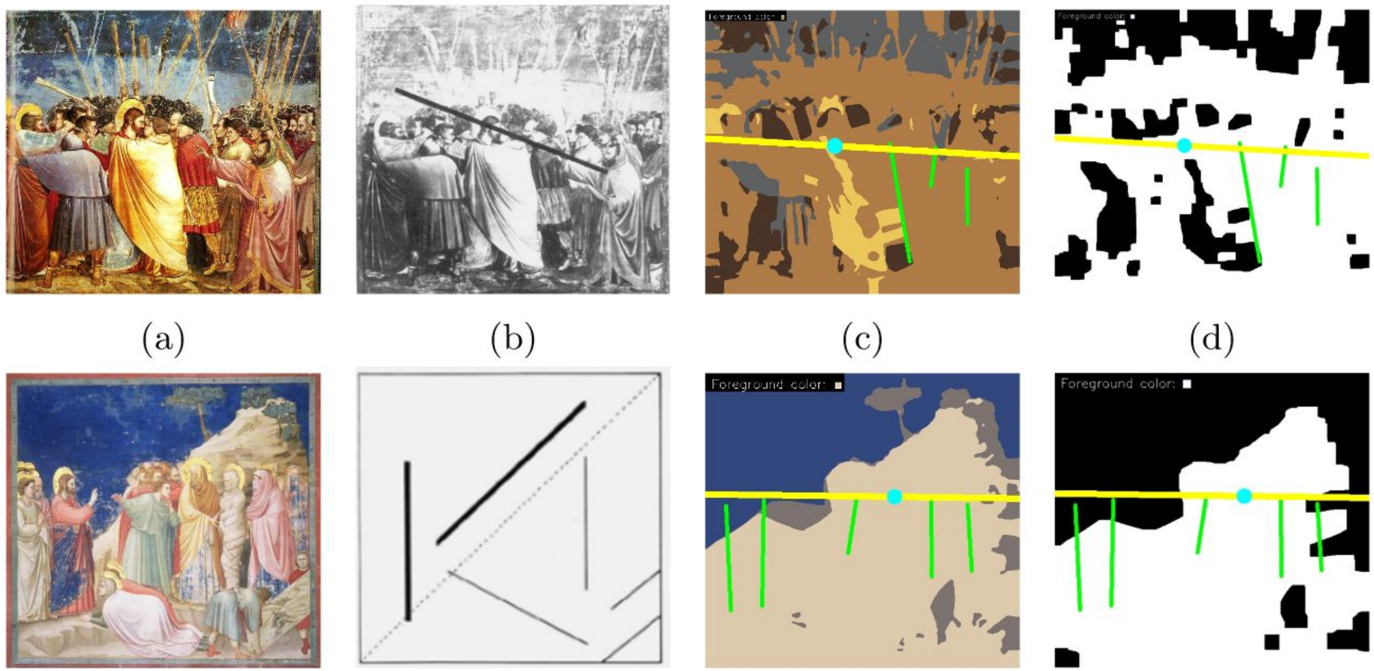

(a) (b) (c) (d)

(e) (f) (g) (h)

Next steps

- **Modify the algorithm**
  - Introduce “Gaze Follow [1]” pipeline
  - FineTune “OpenPose [2]” on art-history data
  - Improve the global action line(s) (yellow line in previous slide)
  - Quantitative evaluation on various datasets

- If interested, please drop us a line:
  - Prathmesh Madhu : prathmesh.madhu@fau.de
  - Ronak Kosti : ronak.kosti@fau.de
  - Project page : ICONOGRAPHICS

[1] Recasens, Adria, et al. "Where are they looking?." Advances in Neural Information Processing Systems. 2015.
[2] Cao, Zhe, et al. "OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields." arXiv preprint arXiv:1812.08008 (2018).

Pose based image retrieval in greek vase paintings

Bachelor’s / Master’s Thesis

Tasks

- Develop a simple pose based ***image retrieval*** tool
- Enhance existing ***pose estimation*** models for vase painting data
- *Style-transfer* and *transfer learning* based models
- Existing state of the art methods fail miserably
  (Try it out yourself, image provided on the right)

- **Expectations**
  - Python
  - Reading papers
  - (any one) : keras, fastai, pytorch, tensorflow

Interested?

- If interested, please drop us a line:
    - Prathmesh Madhu : prathmesh.madhu@fau.de
    - Ronak Kosti : ronak.kosti@fau.de
    - If you're interested to work / collaborate with us on any problem that you can define for our data, here's our project page.

      Project page : ICONOGRAPHICS

Head and legs orientation in greek vase paintings

Bachelor’s / Master’s Thesis/ Forschungspraktikum

Tasks

- Develop an algorithm to detect head directions
- Detect the leg orientations for protagonists (characters) in vase paintings (example on the right)

- **Expectations**
  - Python
  - Reading papers
  - (any one) : keras, fastai, pytorch, tensorflow

Interested?

- If interested, please drop us a line:
  - Prathmesh Madhu : prathmesh.madhu@fau.de
  - Ronak Kosti : ronak.kosti@fau.de
  - If you're interested to work / collaborate with us on any problem that you can define for our data, here's our project page.

    Project page : ICONOGRAPHICS

Pattern Recognition Lab
FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
FACULTY OF ENGINEERING

Super-Resolve real world images

Bachelor’s / Master’s Thesis

Tasks

- Implementation of the paper "Unsupervised Learning for Real-World Super-Resolution [1]"

- Use "Deep Image Prior [2]" to improve the pipeline

- Requirements:
  - Python : keras, fastai, pytorch, tensorflow (any one of these)
  - Keen interest in programming

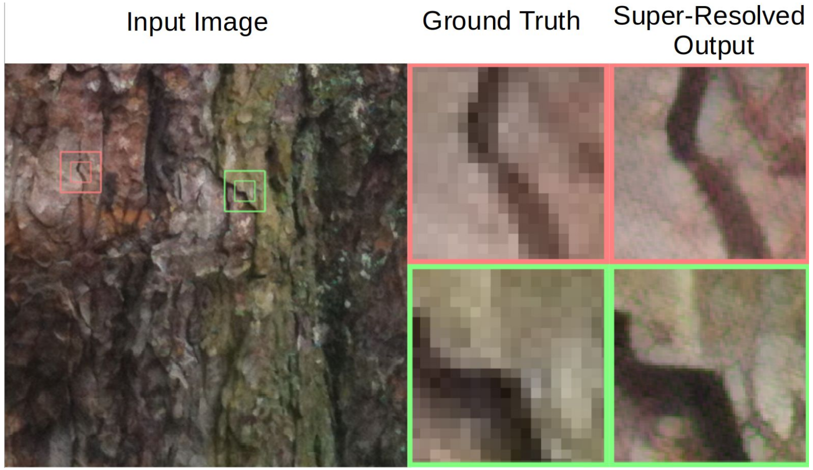

[1] Lugmayr, Andreas, Martin Danelljan, and Radu Timofte. "Unsupervised learning for real-world super-resolution." arXiv preprint arXiv:1909.09629 (2019).
[2] Ulyanov, Dmitry, Andrea Vedaldi, and Victor Lempitsky. "Deep image prior." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2018.

Interested?

- If interested, please drop us a line:
  - Prathmesh Madhu : prathmesh.madhu@fau.de
  - If you're interested to work / collaborate with us on any problem that you can define for our data, here's our project page.

    Project page : ICONOGRAPHICS

Emotion detection in Art

Research Project (Bachelor/Master) 5/10 ECTS

Computer Vision Group, Pattern Recognition Lab, Friedrich-Alexander University Erlangen-Nürnberg

Motivation

It is challenging to detect emotions of people in art paintings:

(a) This person is in the state of *Anger*

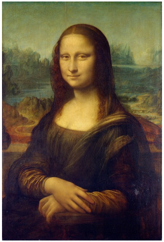

(b) What can be said about the emotion of this person?

Outline

Using current emotion recognition pipelines, modify various deep networks to detect emotions in Art images (or paintings in digital format).

1. **Current Research** Reviewing current state-of-art methods for emotion detection of people in images
2. **Data** Choosing an appropriate dataset for training (or already chosen!?)
3. **Implementation** Evaluate the performance of different models on the collected data
4. **A**nalysis and **C**onclusion

Interested?
Contact for further information/discussion:

Ronak Kosti (Room: 10.136)
ronak.kosti@fau.de

Saliency detection for Emotions

Master Thesis

Computer Vision Group, Pattern Recognition Lab, Friedrich-Alexander University Erlangen-Nürnberg

Motivation

Detecting the regions of image that are salient for emotion recognition **AND/OR** sentiment elicitation [1]

(a) Source Image

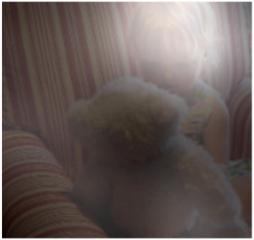

(b) *Expected Salient Region*

Figure: An image has lot of information. Which regions have more significance for emotion analysis?

[1] Fan, Shaojing, et al. "Emotional attention: A study of image sentiment and visual attention." Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. 2018.

Outline

Using the background research/models, find the salient regions (objects, people, stuff) which elicits emotions - *Saliency as a bridge between low and high level vision.*

1. **Literature review** Emotion Recognition *AND/OR* Sentiment Analysis
2. **Data** Mining and building Datasets/Resources
3. **Methods** Attention Models, Context Analysis, etc
4. **I**mplementation, **A**nalysis and **C**onclusion

Interested?
Contact for further information/discussion:

Ronak Kosti (Room: 10.136)
ronak.kosti@fau.de

Deep Learning based Noise Reduction for Hearing Aids

Hendrik Schröter
Speech Processing Group, Friedrich-Alexander University of Erlangen-Nürnberg
July 22th 2019

Hearing Aid Pipeline

Replace conventional noise reduction algorithms with deep learning based approach:

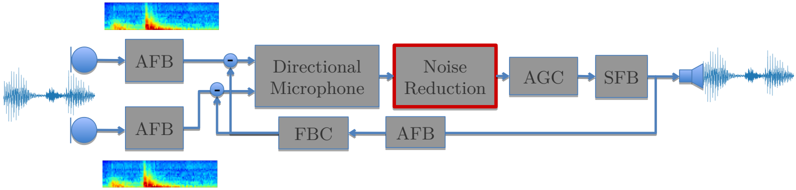

AFB: Analysis Filterbank
SFB: Synthesis Filterbank
AGC: Automatic Gain Control
FBC: Feedback Canceler

Figure: Typical hearing aid pipeline[1].

[1] Figure from: Ehrensperger, Kai, "Deep Learning-based Noise Reduction for Hearing Instrument Applications", MA thesis (Friedrich-Alexander University Erlangen-Nürnberg, 2018)

Denoising using Deep Learning

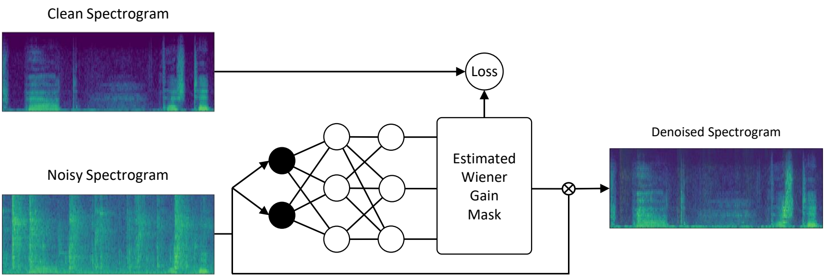

Figure: Simplified schematic figure of the neural network training.

Example: Denoising using Deep Learning

(▶ Noisy)

(▶ Conventional)

(▶ DL based)

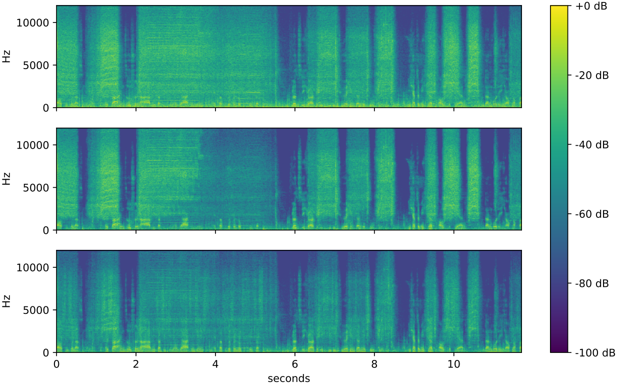

sivantos
the hearing company

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
FACULTY OF ENGINEERING

Distillation Learning for Noise Reduction

Research Project Master (10 ECTS) / Master Thesis

Hendrik Schröter
Speech Processing Group, Friedrich-Alexander University of Erlangen-Nürnberg
WS 2019/20

Distillation Learning

Improve an already existing deep-learning based noise reduction and reduce the number of parameters using distillation learning.

Concept distillation learning (or student/teacher networks):

- A powerful teacher network is trained on the data with hard labels.
- The student is trained to model the teacher's output distribution.
- I.e. the student does not try to predict the hard labels, but rather should learn to imitate the output of the teacher.

Distillation Learning

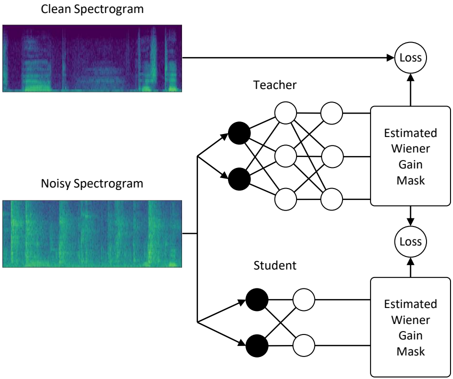

Teacher network:

- Deeper network, more parameters
- “Easier” input, i.e. higher SNR
- Relaxed real-time constraints

Requirements:

- Deep learning basics
- Signal processing basics (complex numbers, Fourier transform)

Contact:

Hendrik Schröter (Room 10.138)
☎ +49 9131 85 27882
✉ hendrik.m.schroeter@fau.de

sivantos the hearing company

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
FACULTY OF ENGINEERING

Deep Learning based Beamforming for Hearing Aids

Master Thesis

Hendrik Schröter
Speech Processing Group, Friedrich-Alexander University of Erlangen-Nürnberg
WS 2019/20

Hearing Aid Pipeline

Improve an already existing deep-learning based noise reduction using multi-channel signals, which enables to exploit directional information.

Using this, we want to replace traditional directional signal processing and noise reduction with deep learning based approach:

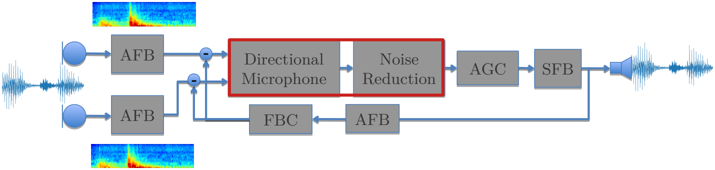

AFB: Analysis Filterbank
SFB: Synthesis Filterbank
AGC: Automatic Gain Control
FBC: Feedback Canceler

Figure: Typical hearing aid pipeline[1].

[1] Figure from: Ehrensperger, Kai, “Deep Learning-based Noise Reduction for Hearing Instrument Applications”, MA thesis (Friedrich-Alexander University Erlangen-Nürnberg, 2018)

Data

- Multi-channel noise signals from hearing aids
- Clean speech signals, transformed with HRTFs (Head-related transfer function)

Beamforming

a) Use multiple channels to estimate a multi-channel Wiener filter

b) Use multiple channels and positional information of the microphones to estimate beamforming coefficients

Requirements:

- Deep learning basics
- Signal processing basics (complex numbers, Fourier transform)

Contact:

Hendrik Schröter (Room 10.138)
☎ +49 9131 85 27882
✉ hendrik.m.schroeter@fau.de

The Lottery Ticket Hypothesis
Finding Sparse, Trainable Neural Networks

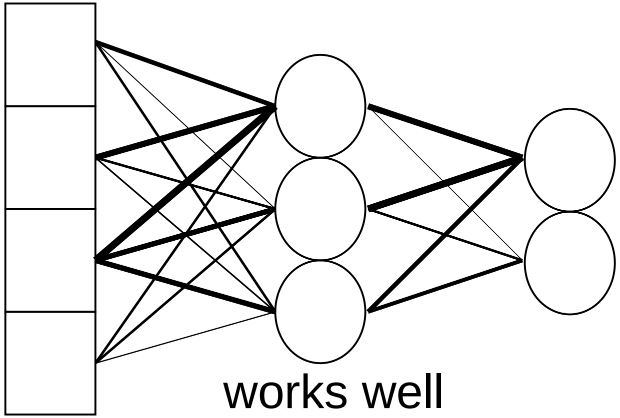
works well

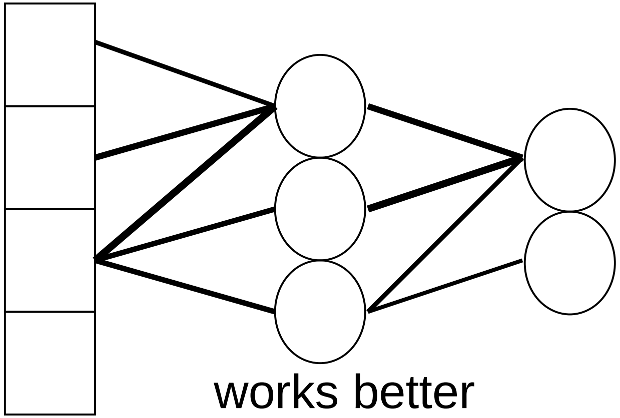
works better

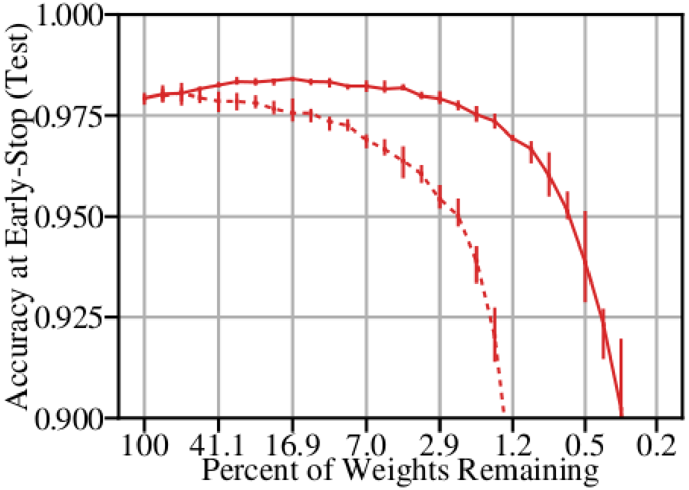

- New network pruning approach
- Removing up to 80 % of weights
- Produces networks as good/better

- Are resulting structures consistent?
- Is transfer learning possible?
- Can layers be resized?

… many other open questions

5-10 ECTS project – contact: mathias.seuret@fau.de

Based on a paper by Jonathan Frankle, Michael Carbin

Gradient-domain Data Augmentation : Degradation Model

Color domain

Gradient domain

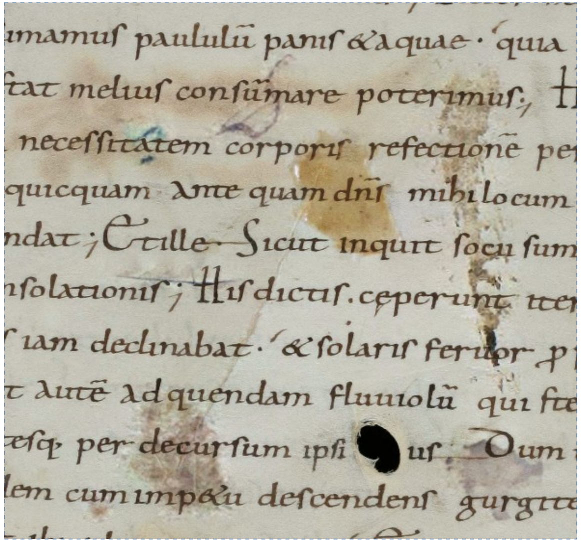

- Data augmentation method
- Paste gradients of stains
- Pixels reconstructed from gradients
- "Fools" human experts

- Noise location model needed (e.g., Fingerprints in margins, or water stains at top/bottom)
- GAN degradations generator, document as parameter
- Train/eval analysis method on augmented data

5-10 ECTS project – contact: mathias.seuret@fau.de

Mimicking Typesetting & Printing

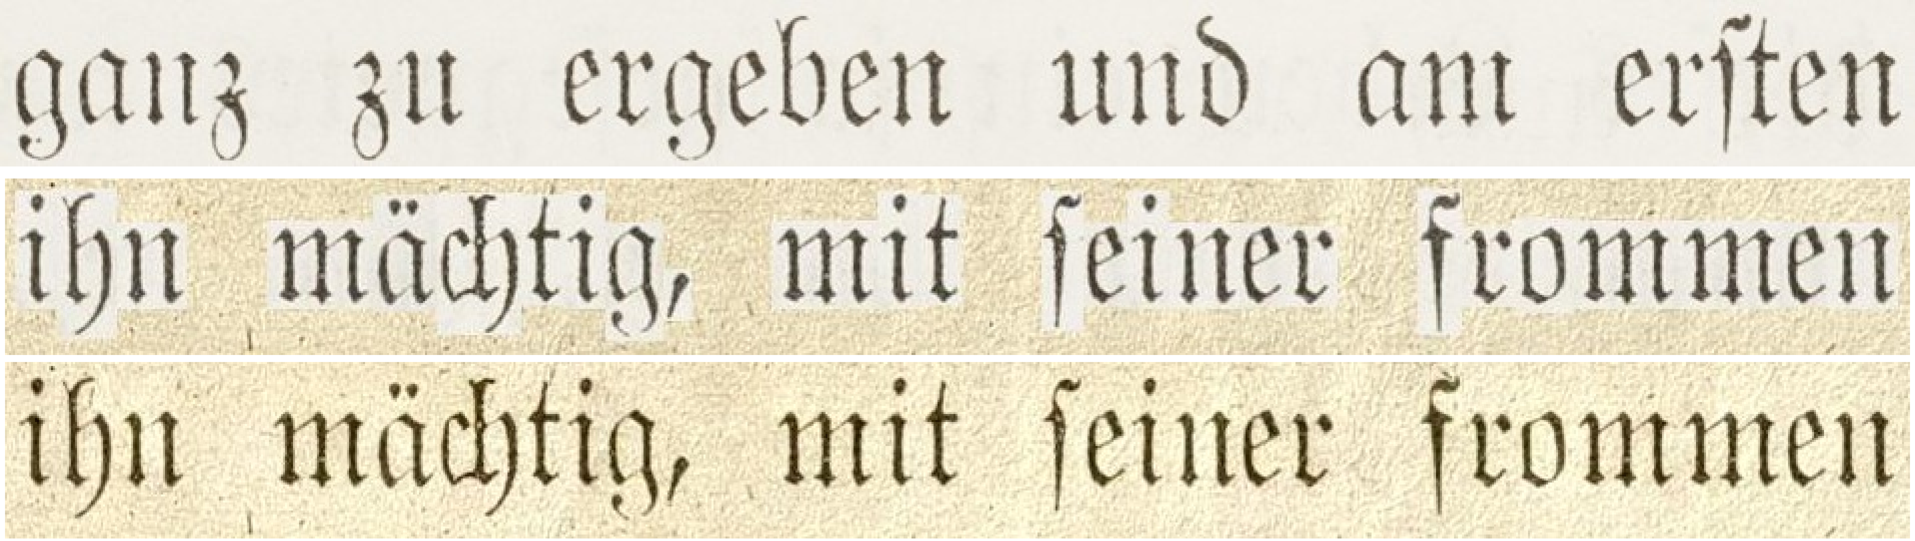

- OCR for ancient documents: open problem
- Synthetic data needed
- Gradient-domain approach
- Toy-example proof of concept

- Automatic character & baseline extraction
- Character- & forme-level augmentation (GAN?)
- ‘’Print’’ pages with multiple fonts
- Evaluation through OCR

From 5 ECTS to Master project – contact: mathias.seuret@fau.de

Font-Based Segmentation of Historical Documents

- Fonts (& other content) known

- Localization unknown

- Paragraphs, words, parts of words in different fonts

- More than 35k pages

- Goal: use a classifier to localize the content

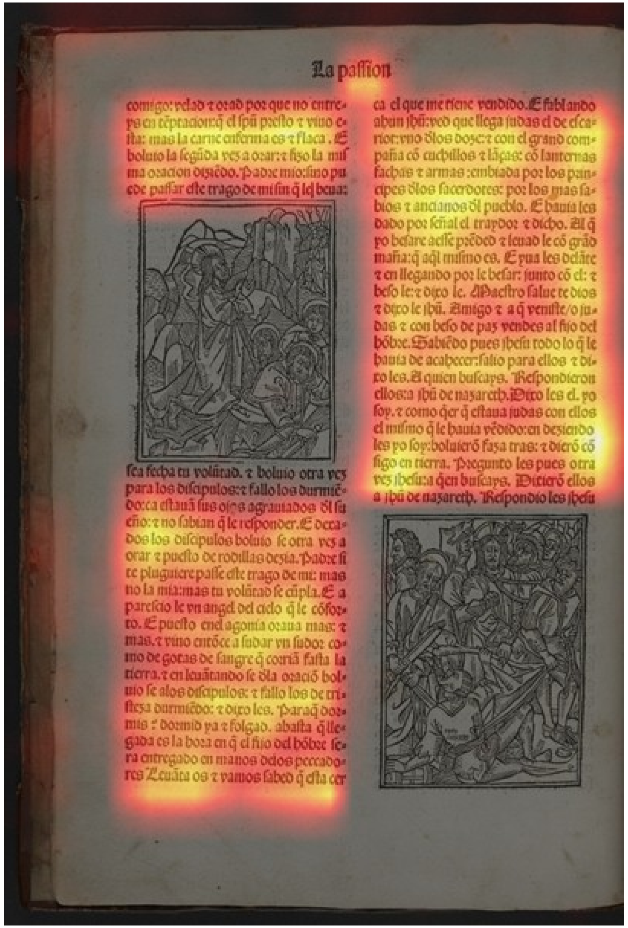

5-10 ECTS project – contact: mathias.seuret@fau.de

Weakly supervised multimodel lesion detection and classification in mammogram & ultrasound

Master's Project (10 ECTS)

Motivation

- Mulitmodel breast image analysis for malignancy **detection** and **classification**

Mammogram (2D)

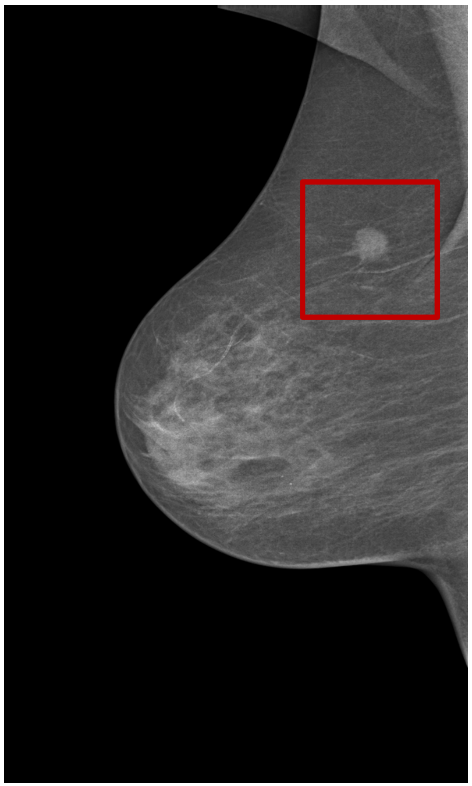

Tomosynthesis (3D)

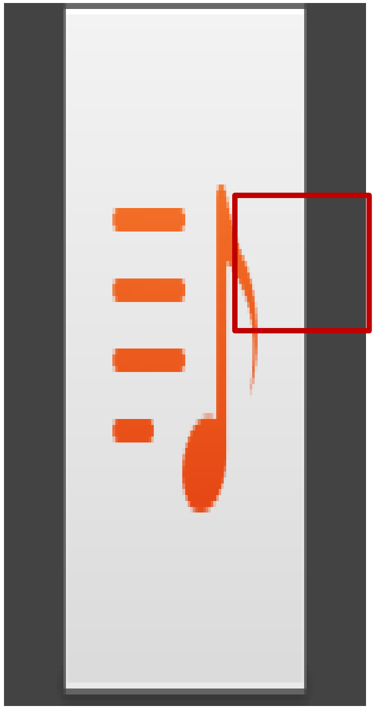

Ultrasound (3D)

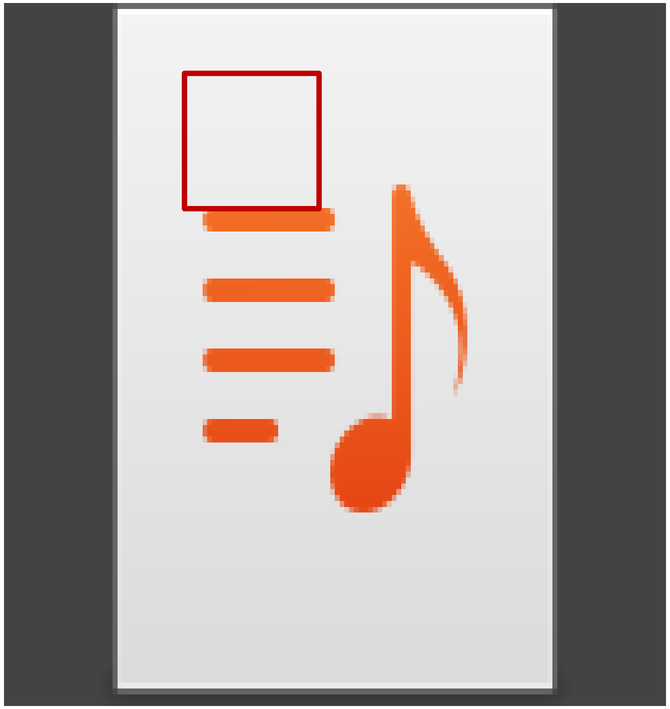

Figure 1: A 55 years old patient with a malignant lesion in left-side breast, diagnosed with BI-RADS 5

Steps

- Design weakly supervised **multimodal learning** method using cross-modality fusion
  - Feature learning level
  - Classifier/decision-making level
  - For learning: no manual annotation, but pathology label

- **Requirements:**
  - Programming skills: Python + Keras/TensorFlow
  - Deep understanding of volumetric/high-dimensional data

Contact for further information/discussion:
Sulaiman Vesal M.Sc. (Room: 10.136)
Sulaiman.vesal@fau.de

Left ventricle quantification using spatiotemporal feature learning

Master’s Project (10 ECTS)

Motivation

- Assessing the **heart's function**, the **left ventricle (LV) function**, morphology and temporal dynamics is of clinical interest
  - Cavity and myocardium size
  - Cavity dimension
  - Regional wall thicknesses
  - Heart phase (systole or diastole)

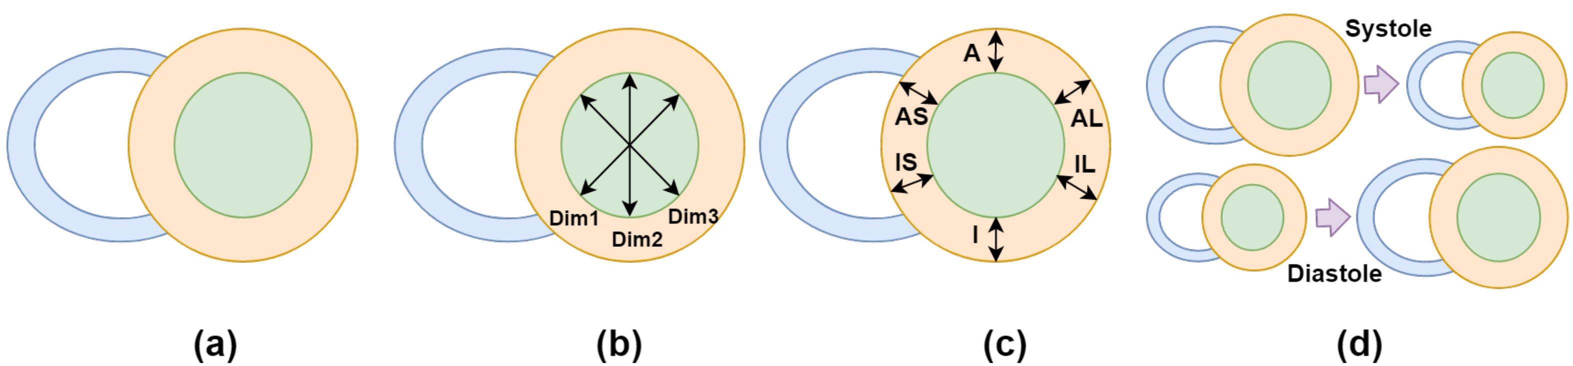

Figure 1: Illustration of LV indices to be quantified for short-axis cardiac image.(a) Cavity (green) and myocardium (yellow) areas. (b) directional dimensions of cavity (black arrows). (c) Regional wall thicknesses (black arrows). (d) Phase(systole or diastole)

Steps

- Develop effective machine learning models that can estimate a set of clinically significant LV indices
  - Supervised localization of LVs in short-axis cine MR images
  - Investigate the use of **spatiotemporal convolutions**
  - **Multi-task learning** for both cardiac phase detection and LV indices estimation

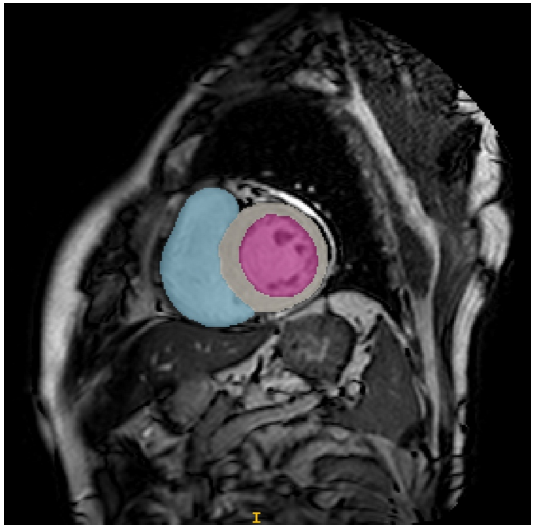

Figure 2: Cine-MR image with segmented left ventricle and myocardium

- **Requirements:**
  - Programming skills: Python + Keras/TensorFlow
  - Deep understanding of volumetric/high-dimensional data

Contact for further information/discussion:
Sulaiman Vesal M.Sc. (Room: 10.136)
Sulaiman.vesal@fau.de

Multimedia Security Group

- **Image enhancement:**
  - Superresolution of compressed data
- **Image/video forensics:**
  - Has an image been retouched?
  - Is part of a video computer-generated?

How "dangerous" is GAN-generated CGI?

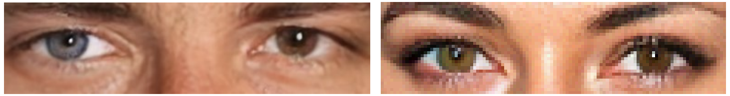

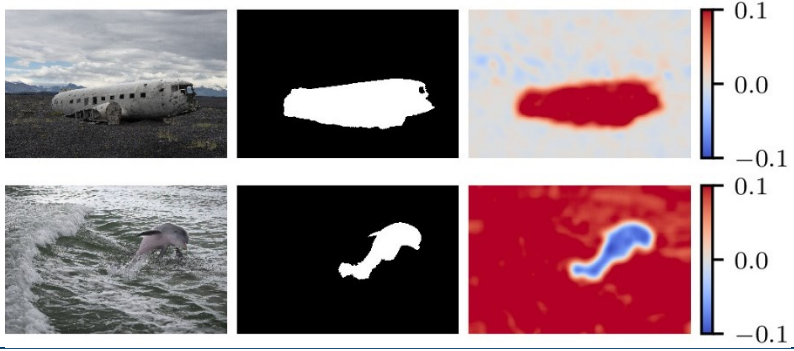

What traces leave manipulations in the compression container?

How can we learn to detect manipulated faces from few training examples?

Example Open Projects or Theses

Guess characters on unreadable licence plates	Statistical video manipulation detection	Physics-based image manipulation detection	How easily can DL-based forgery detectors be fooled?
-> CNN to deal with strongly compressed video frames of licencse plates	-> Deep anomaly detector / device parameter regressor	-> Learning-based methods for classical vision tasks, e.g., shadow segmentation	-> Can we construct a counter-forensics adversarial example image laundry just from “innocent” JPEG settings?

Who to talk to

- We run projects between the Pattern Recognition Lab, the Computer Graphics Lab, and the IT Security Infrastructures Lab
- Group Members


Amir Davari


Benjamin Hadwiger


Benedikt Lorch


Patrick Mullan

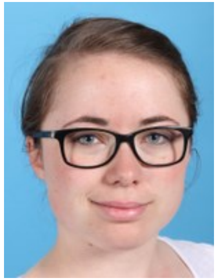
Franziska Schirrmacher

- For concrete Projects or Theses:
  Contact Franziska Schirrmacher, franziska.schirrmacher@fau.de