

Frederik Warburg

<http://frederikwarburg.github.io>

frwa@dtu.dk

+45 42729858

Education

- 2020 - 2023 **PhD in Uncertainty Quantification in Deep Learning** [Technical University of Denmark](#)
I worked on 3D perception and uncertainty quantification. I was supervised by Søren Hauberg (Technical University of Denmark), Javier Civera (University of Zaragoza), Serge Belongie (University of Copenhagen), and Angjoo Kanazawa (UC Berkeley). I was awarded the Danish Elite Research Travel Stipend in 2022. Average grade 12 / 12. PhD defense 7 Dec 2023.
- 2018 - 2020 **MSc in Mathematical Modelling and Computing** [Technical University of Denmark](#)
I graduated from the honours program, which is an elite education with a more challenging course of study, with an average grade 11.9 / 12. During my studies, I found a special interest in machine learning and computer vision. I received 12 / 12 for my master thesis.
- Spring 2019 **MSc Computer Science** [University of California, Berkeley](#)
I received Sparnord Fonden's FinTech Entrepreneurial Scholarship to study at UC Berkeley in the Spring 2019. I received a GPA 4.0 / 4.0 Besides my studies, I followed and won an entrepreneurial track held by Innovation Center Denmark in Silicon Valley.
- 2015 - 2018 **BSc in Mathematics and Technology** [Technical University of Denmark](#)
I have obtained a solid mathematical foundation and advanced programming skills. I achieved an average grade of 10.3 / 12.0 I completed my BSc half a year faster than the standard time.
- 2014 - 2015 **BSc in Mathematics** [Lindenwood University](#)
After high school, I studied at Lindenwood University, MO, USA for one year. I was elected student senator, and as such I raised \$7700 for an outdoor study area. I achieved an average grade of 3.8 / 4.

Publications

- 2024 **Toon3D: Seeing Cartoons from a New Perspective** [preprint](#)
E. Weber, R. Peterlinz*, R. Mathur, F. Warburg, A. Efros, A. Kanazawa*
COLMAP cannot reconstruct non-geometric hand-drawn images even with perfect correspondences! Toon3D can recover camera poses and dense geometry with piecewise-rigid deformable optimization. <https://toon3d.studio/>
- 2023 **Nerfbusters: Removing Ghostly Artifact from NeRFs** [ECCV](#)
F. Warburg, E. Weber*, M. Tancik, A. Holynski, A. Kanazawa*
We propose a new dataset, evaluation procedure, and a data-driven, local 3D diffusion prior that can be applied with a novel density score distillation sampling loss. <https://ethanweber.me/nerfbusters/>
- 2023 **Learning to Taste: A Multimodal Wine Dataset** [NeurIPS](#)
Bender, Sørensen, Kashani, Eldjarn, Hyldig, Hauberg, Belongie, Warburg
We present a large multimodal flavor dataset and propose to add flavor as a modality for multimodal representation learning. https://thoranna.github.io/learning_to_taste/

- 2023 **K-planes: Explicit Representation of Space, Time, and Appearance** CVPR
S. Fridovich-Keil, G. Meanti*, F. Warburg, B. Recht, A. Kanazawa*
 We propose a simple planar factorization for radiance fields that naturally extends to arbitrary-dimensional spaces, and that scales gracefully with dimension in both optimization time and model size. We show the advantages of our approach on 3D static volumes, 3D photo collections with varying appearances, and 4D dynamic videos. <https://sarafridov.github.io/K-Planes>
- 2023 **Bayesian Metric Learning for Uncertainty Quantification in Image Retrieval** NeurIPS
F. Warburg, M. Miani*, S. Brack, S. Hauberg*
 Rather than relying on neural amortization as done in prior works, we learn a distribution over the network weights with the Laplace Approximation. <https://arxiv.org/abs/2302.01332>
- 2023 **DAC: Detector-Agnostic Spatial Covariances for Deep Local Features** 3DV
J. Tirado-Garin, F. Warburg, J. Civera
 We propose two post-hoc covariance estimates that can be plugged into any pretrained deep feature detector: a simple, isotropic covariance estimate that uses the predicted score at a given pixel location, and a full covariance estimate via the local structure tensor of the learned score maps.
- 2023 **Laplacian Segmentation Networks: Improve Epistemic Uncertainty from Spatial Aleatoric Uncertainty** MICCAI
K. Zepf, S. Wanna*, M. Miani, A. Feragen, S. Hauberg, F. Warburg*
 To ensure robustness to out-of-distribution segmentations, we propose Laplacian Segmentation Networks (LSN) that jointly model epistemic (model) and aleatoric (data) uncertainty in image segmentation. <https://arxiv.org/abs/2303.13123>
- 2022 **Searching for Structure in Unfalsifiable Claims** HCOMP
P. E. Christensen, F. Warburg, M. Jia, S. Belongie
 Social media platforms give rise to an abundance of posts and comments on every topic imaginable. Many of these posts express opinions on various aspects of society, but their unfalsifiable nature makes them ill-suited to fact-checking pipelines. In this work, we aim to distill such posts into a small set of narratives that capture the essential claims related to a given topic. <https://arxiv.org/abs/2209.00495>
- 2022 **Laplacian Autoencoders for Learning Stochastic Representations** NeurIPS
M. Miani, F. Warburg*, P. Moreno-Munoz, N. S. Detlefsen, S. Hauberg*
 We present a Bayesian autoencoder for unsupervised representation learning, which is trained using a novel variational lower-bound of the autoencoder evidence. This is maximized using Monte Carlo EM with a variational distribution that takes the shape of a Laplace approximation. We develop a new Hessian approximation that scales linearly with data size allowing us to model high-dimensional data. <https://arxiv.org/abs/2206.15078>
- 2022 **SparseFormer: Attention-based Depth Completion Network** CVCVR
F. Warburg, M. Ramamonjisoa, M. Antequera
 We introduce a transformer block that fuses 3D landmarks with deep visual features to complete a sparse depth map. <https://arxiv.org/abs/2206.04557>
- 2022 **Volumetric Disentanglement for 3D Scene Manipulation** WACV
S. Benaim, F. Warburg, P. Christensen, S. Belongie
 We propose a volumetric framework for (i) disentangling or separating, the volumetric representation of a given foreground object from the background, and (ii) semantically manipulating the foreground object, as well as the background. <https://arxiv.org/abs/2206.02776>

- 2022 **Danish Airs and Grounds: A Dataset for Aerial-to-Street-Level Place Recognition and Localization** RA-L (IROS)
A. Vallone, F. Warburg*, H. Hansen, S. Hauberg, J. Civera.*
 We present an aerial to street-level localization dataset covering both urban, suburban, and rural areas. It contains street-level images with accurate 6 DoF poses and associated aerial images from Denmark. <https://arxiv.org/abs/2202.01821>
- 2021 **ACDC-Net: Self-supervised Depth Completion for Active Stereo** RA-L (ICRA)
F. Warburg, D. Hernández, U. Bonde, A. Vakhitov, P. Alcantarilla.
 We present the first self-supervised depth completion method for active stereo that completes and refines the depth maps by closely integrating a visual-inertial SLAM system in the training and inference pipeline. <https://arxiv.org/abs/2110.03234>
- 2021 **Bayesian Triplet Loss: Uncertainty Quantification for Image Retrieval** ICCV
F. Warburg, M. Jørgensen, J. Civera, S. Hauberg.
 Instead of modelling image embeddings as points, we propose to model image embeddings as isotropic Gaussians. We develop a novel Bayesian triplet loss that extend the traditional triplet loss to enforce the triplet constraint for Gaussian embeddings. This allow us to quantify uncertainty in the retrieval system, which can be important for downstream decision-making. <https://arxiv.org/abs/2011.12663>
- 2020 **Probabilistic Spatial Transformer Network** UAI
P. Schwöbel, F. Warburg, M. Jørgensen, K. Madsen, S. Hauberg.
 We extended the Spatial Transformer Network to a probabilistic model enabling us to sample from a distribution of transformations. Thus, learning a suitable data augmentation for data of different modalities. <https://arxiv.org/abs/2004.03637>
- 2020 **Mapillary Street-Level Sequences: A Dataset for Lifelong place recognition** CVPR
F. Warburg, S. Hauberg, M. Antequera, P. Gargallo, Y. Kuang, J. Civera.
 We curated and validated a large dataset using Mapillary images. The substantial size and diversity of the dataset makes it relevant for training of deep neural networks for place recognition. Paper accepted with oral presentation. <https://research.mapillary.com/img/publications/CVPR20c.pdf>
- 2018 **Intensity Mapping for Mask Projection based Photopolymerization** ASPE, Berkeley
F. Warburg, M. Ribo, A. Luongo, A. Danielak, D. Pedersen.
 We presented a method for mapping the intensity field of the projected light in a photopolymerization system. We showed that the de-facto assumption about uniformly distributed light is invalid and we implemented a method for making the projection more uniform. https://www.researchgate.net/publication/339787809_Intensity_Mapping_for_Mask_Projection_based_Photopolymerization

Experience

- Fall 2023 - Now **Head of AI** [Teton](#)
We are building the AI care companion using computer vision and deep learning to help nurses and patients in hospitals and care homes. We work at the forefront of research in 3D reconstruction, human mesh recovery, action recognition, and visual-language models.
- Fall 2021 **Research Engineer Intern in Depth Completion** [Facebook](#)
Worked on depth completion as part of a four month research internship at Facebook. The internship resulted in a workshop paper submitted at CVARVR22.
- 2020 - 2021 **Research Intern in Depth Completion** [SLAMcore](#)
Under supervision of Pablo Alcantarilla, I developed the first self-supervised depth completion method for active stereo sensors that completes and refines the initial hardware optimized depth maps of the active depth sensor.
- 2019 - 2020 **Research Intern in Place Recognition** [Mapillary](#)
I comprised the largest dataset for lifelong place recognition using images from Mapillary's crowds source image database. I benchmarked state-of-the-art deep learning place recognition methods on this and competing datasets.
- Summer 2019 **Research Fellow in Deep Learning** [ETH Zurich](#)
I received ETH's Computer Science Summer Research Fellowship. I worked at Marc Pollefeys' Visual Computing lab at ETH. Under the supervision of Martin Oswald, Viktor Larsson and Mihai Dusmanu, I investigated a novel k-max pooling technique in several computer vision domains, including 3D reconstruction and super-resolution.
- 2018 - 2019 **Machine Learning Engineer** [Beep Analytics](#)
We use machine learning to create a predictive tool that delivers data-driven insights about repair parts for airplanes. The tool will provide improved maintenance and cost savings for airplane companies.
- Summer 2018 **Research Assistant in SLAM** [University of Zaragoza](#)
Under the supervision of professor Javier Civera, I worked with lifelong place recognition in SLAM. I comprised a large dataset for lifelong place recognition using images from Google Street View. I used state-of-the-art deep convolutional neural networks to post-process the data and to test the difficulty of the dataset.
- 2017 - 2018 **Data Scientist and App Developer** [Technical University of Denmark](#)
We scraped, cleaned, analyzed and presented data in an app that provides key-insights about the university's company collaborations. One feature of the app was an interactive graph representation where professors and companies were nodes and collaborations were edges.
- Summer 2017 **Software Developer Summer Intern** [AutoDesk](#)
I developed the data structure and the interface of a template selector that will radically change the work-flow of AutoDesk Fusion that has more than 100.000 users.
- 2016 - 2017 **Student Ambassador** [IBM](#)
I was responsible for the relationship between IBM and DTU. I facilitated guest lectures and hackathons while communicating technical content about IBM products to DTU students and professors.

Teaching Experience

- 2020 - Now **Project Supervision** [Technical University of Denmark](#)
I have supervised 9 Master, 3 Bachelor, and 4 special course projects on topics range from place recognition to applied robotics to generative adversarial networks, 3 of which resulted in publications in major ML conferences and 1 in a successful startup.
- 2018 - 2024 **Teaching Assistant** [Technical University of Denmark](#)
I taught DTU students in Introduction to Machine Learning, Advanced Machine Learning, Deep Learning, Software Startup Studio, Deep Learning for Computer Vision.

Honors, Awards & Certifications

- 2022 **Elite Research Travel Stipend (DKK 200.000)** [Danish Ministry of Higher Education and Science](#)
The grant is awarded to the top 20 Danish PhD students across all branches of science each year. The grant was presented by the Crown Princess and the science minister.
- 2018-2020 **Honors program** [Technical University of Denmark](#)
I graduated from the honours program, which is an elite education that offers a more challenging course of study, individual tutoring and ambitious research affiliation. The program is offered to the top 10 % students at the university.
- Nov. 2018 **Venture Cup Idea Hunt (SEK 5.000)** [Venture Cup](#)
We presented an innovative method for sowing wheat.
- Oct. 2016 **3'rd place winner of OI-X Big Data competition (DKK 10 000)** [DTU Skylab](#)
We developed a big data solution for wind turbine parks to optimizes the total energy production of the park.
- 2014 - 2015 **Dean Honours** [Lindenwood University](#)
Achieved a GPA above 3.5 both semesters at LU.