Ambarella
AI envisioned™

VisLab internship opportunities

Maggio 2026

VisLab Introduction

A **high-tech company** based in Parma, VisLab is focused on **all aspects of ADAS and autonomous driving i**ncluding machine learning, deep neural networks (DNNs), stereovision, robotics, sensor fusion, planning, system calibration, and more.

Founded in the late ‘90s as a research laboratory at the University of Parma, VisLab quickly established itself as a **pioneer in vehicle autonomy** before spinning off from the University in 2009.

In 2015, VisLab was **acquired by Ambarella**, a publicly traded semiconductor Company based in **Silicon Valley** headquartered in Santa Clara (California, USA) and with offices in Detroit (Michigan, USA), Munich (Germany), Hong Kong, Hsinchu (Taiwan), SeongNam City (South Korea), Shanghai (China), Shenzhen (China), and Yokohama (Japan).

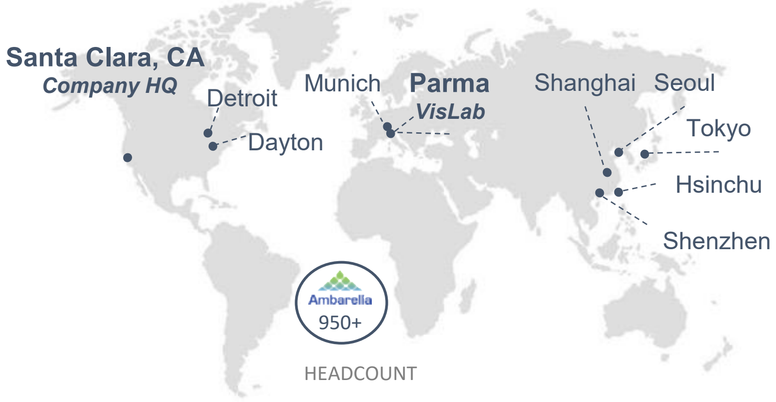

Specializing in the development of advanced computer vision hardware processors, **Ambarella relies on VisLab for core research and development** across all the markets with a **particular focus on automotive**.

VisLab provides top level research in the most contemporary high-tech fields such as autonomous driving, robotics, and deep learning. VisLab equips and operates a **fleet of autonomous vehicles** and received the **first authorization from the Italian Ministry of Transportation** to run autonomous driving tests on open roads.

Check www.vislab.it for further information.

VisLab: 25 Years of CV and Autonomous Driving

VisLab Today

- ***Algorithm first* approach** applied to Computer Vision
- Continuous research benefits **future Ambarella chips**

1998: Mille Miglia 2000+km on Italian highways 94% autonomous steering

2005, 2007: DARPA Grand and Urban Challenge. 100% AD

2010: VIAC 15K km cross-continent drive AD following

2013: PROUD: Braive drove 13km in Parma, Level 4 demo 100% driverless AD

2014: DEEVA, 13 stereo camera systems; CV integration; 100% AD

2015: Ambarella acquires VisLab

CES 2018: camera perception stack demonstrated on CV1

2018: AD demonstration in Santa Clara office

CES 2019: AD driving on the roof track of Lingotto building in Turin

CES 2020: L4 autonomous driving demo in Las Vegas with automatic parking. CV2-based

CES 2022: L4 autonomous driving demo and new Oculii-based sensing suite.

2023: Full L4 autonomous driving stack demo on CV3

Internship 1

Title: Testing deep-learning-based architectures for heterogeneous data fusion

Description: The goal of the internship is to analyze and test different network architectures for the fusion of heterogeneous data, coming from radars, cameras, lidar and stereo cameras.
The intern will start from the study of the state of the art, to move on to the analysis and the tests.
The intern will have to evaluate the different architecture and to select the best one.
More details will be provided during the interview.

Requirements:

- The intern should have at least basic knowledge of neural networks.
- Proficiency in the Python programming language is also required.

Tutor: Alessandro Pacchioni

Duration: 6 months

Internship 2

Title: Testing deep-learning-based architectures for heterogeneous data fusion, using publicly available datasets

Description: The goal of the internship is to analyze and test different network architectures for the fusion of heterogeneous data, coming from radars, cameras, lidar and stereo cameras.
The intern will start from the study of the state of the art, to move on to the analysis and the tests.
The intern will have to evaluate the different architecture using publicly available datasets.
More details will be provided during the interview.

Requirements:

- The intern should have at least basic knowledge of neural networks.
- Proficiency in the Python programming language is also required.

Tutor: Alessandro Pacchioni

Duration: 6 months

Internship 3

Title: Real data based imitation learning

Description: The goal of the internship is to develop/tune a network and train it using different datasets.
The network should learn how to imitate the action of the human driver.
The intern will begin by studying the state of the art, then proceed to develop the network and conduct tests using public datasets based on data acquired on real world (e.g. NuPlan).
The network will be tested and validated in simulation environments, with the possibility for the deployment on a real-world vehicle.

Requirements:

- A good understanding of neural networks.
- Proficiency in the Python programming language is also required.

Tutor: Alessandro Capasso

Duration: 6 months

Internship 4

Title: Automatic MPC Weight Tuning via Learning-Based Methods

Description: This thesis proposes a learning-based framework to automatically estimate the cost matrices and key tuning parameters of a continuous-time MPC controller based on the PRONTO (Projection Operator-based Newton method for Trajectory Optimization) algorithm. The approach combines reinforcement learning and deep neural models to adapt the controller objectives from data and task performance, while preserving the analytical derivatives and real-time optimization structure of PRONTO. The goal is to improve robustness, adaptability, and control performance of in-house simulator.

Requirements:
- A good understanding of neural networks.
- Proficiency in the Python programming language.
- Good understanding of non-linear optimization

Tutor: Alessandro Capasso

Duration: 6 months

Internship 5

Title: From Imitation to Offline RL for planning policy improvement

Description: The goal of the internship is to study and implement offline reinforcement learning strategies to improve neural driving policies in a traffic simulator.
Currently, a deterministic omniscient policy selects the actions to perform, and a neural network is trained through supervision to imitate those actions. The thesis will investigate how this pretrained imitation-learning network can be fine-tuned or retrained using offline RL, starting from previously collected simulated trajectories.
The intern will study state-of-the-art offline RL methods, with particular focus on approaches that require little or no handcrafted reward function.
Different strategies will be implemented and tested to determine whether the learned policy can improve beyond pure imitation by exploiting long-term trajectory information.
The final network will be evaluated in simulation and compared against both the supervised baseline and the original expert policy.

Requirements:
- A solid understanding of neural networks, especially reinforcement learning.
- Proficiency in the Python programming language is also required.

Tutor: Alessandro Capasso

Duration: 6 months

Internship 6

Title: Automatic generation of high-definition maps with neural networks from satellite views

Description: The goal of the internship is to develop a network to generate high-definition maps starting from satellites views (aka google earth images). The intern will start from an in-depth study of the state of the art, to move on to network development and testing using real and simulator data provided by the company. The intern will have to try different network architectures to evaluate them and select the best one.

Requirements:
• At least basic knowledge of neural networks;
• Knowledge of the Python programming language.

Tutor: Dafne Molin

Duration: 6 months

Internship 7

Title: Automatic generation of high-definition maps with neural networks using diffusion models

Description: The goal of the internship is to develop a diffusion neural network to generate high definition/topology maps or refine existing ones in an offline setting. The intern will start from an in-depth study of the state of the art, to then move on to network development and testing using real or simulated data provided by the company. The intern will have to try out different network architectures to evaluate them and select the best one.

Requirements:
• Basic knowledge of neural networks
• Knowledge of the Python programming language

Tutor: Dafne Molin

Duration: 6 months

References:

P. Jia et al: DiffMap: Enhancing Map Segmentation with Map Prior Using Diffusion Model
T. Monninger: MapDiffusion: Generative Diffusion for Vectorized Online HD Map Construction and Uncertainty Estimation in Autonomous Driving

Internship 8

Title: A MLLM-based framework to detect and correct multi-sequence HD maps

Description: Large-scale map construction plays a pivotal role in applications such as autonomous driving. When constructing large-scale lane-level maps, two essential challenges are inevitable: (1) occlusions caused by obstacles, which can lead to missing or inaccurate road predictions; and (2) incomplete annotations resulting from poor reconstruction or from the methods used to generate them.
For these reasons, the goal of the internship is to develop a multimodal large language model (MLLM)-based framework to detect and correct errors in HD maps (e.g., missing connections, incorrect numbers of lanes, or misclassifications) in an offline setting. The intern will begin with an in-depth study of the state of the art and then move on to network development and testing using real and simulated data provided by the company, as well as public datasets. The intern will explore different network architectures, evaluate their performance, and select the most suitable one.

Requirements:
- Basic knowledge of neural networks
- Knowledge of the Python programming language

Tutor: Dafne Molin

Duration: 6 months

References:

Y. Yuan et al.: UniMapGen: A Generative Framework for Large-Scale Map Construction from Multi-modal Data

Internship 9

Title: Large-scale visual localization and mapping

Description: The goal of the internship is to develop and test novel algorithms for large-scale visual localization and mapping (SLAM). After an initial phase of literature review, the intern would then focus on one of the following directions:

- Benchmarking state-of-the-art algorithms on public real-world data.
- Extending a vision-only system to support LiDAR and radar data.
- Extending a pinhole-only system to natively support feature extraction and matching across multiple camera models (e.g. fisheye).

The project could then be extended by exploring the other directions.

Requirements:

- Good understanding of classical computer vision, especially multi-view geometry and 3D reconstruction.
- At least basic understanding of neural networks.
- Knowledge of the C++ programming language.

Tutor: Marco Orsingher, Pietro Azzari

Duration: 6 months

Internship 10

Title: 3D reconstruction and novel view synthesis with Gaussian Splatting

Description: The goal of the internship is to develop and test novel algorithms based on Gaussian Splatting (GS) for urban-scale 3D reconstruction. An outline of the possible activities include:

- Reviewing existing literature on GS in urban scenarios.
- Testing and benchmarking state-of-the-art systems on public real-world automotive datasets.
- Improve such systems in terms of efficiency and scalability.

The project could then be extended further by considering dynamic objects and exploring the creation of self-driving simulators with GS.

Requirements:

- Good understanding of classical computer vision, especially multi-view geometry and 3D reconstruction.
- At least basic understanding of neural networks and computer graphics.
- Knowledge of the Python programming language. Some previous CUDA exposure would be appreciated.

Tutor: Marco Orsingher, Paolo Medici

Duration: 6 months

Internship 11

Title: 3D reconstruction and novel view synthesis with Transformer-based 3D foundation models

Description: The goal of the internship is to develop and test novel algorithms using Transformer-based 3D foundation models for urban-scale 3D reconstruction. An outline of the possible activities include:

- Reviewing existing literature on Transformer-based 3D foundation models in urban scenarios.
- Testing and benchmarking state-of-the-art systems on public real-world automotive datasets.
- Improve such systems in terms of efficiency and scalability.

The project could then be extended further by considering dynamic objects.

Requirements:

- Good understanding of classical computer vision, especially multi-view geometry and 3D reconstruction.
- At least basic understanding of neural networks and computer graphics.
- Knowledge of the Python programming language. Some previous CUDA exposure would be appreciated.

Tutor: Marco Orsingher, Paolo Medici

Duration: 6 months

Internship 12

Title: Category-level dense 3D reconstruction

Description: The goal of the internship is to develop and test novel algorithms for dense 3D reconstruction of specific objects that frequently arise in urban scenarios, such as vehicles, human and animals. Some category-level approaches have been proposed in literature in order to simplify and improve 3D reconstruction, assuming to know which kind of object is being reconstructed. The intern would need to:

- Review state-of-the-art methods for each category.
- Focus on one category and implement a state-of-the-art method on public real-world urban data.
- Scale up the method to potentially reconstruct tens or hundreds of objects in parallel, across a whole sequence.

The project could then be extended naturally to other categories.

Requirements:

- Good understanding of classical computer vision, especially multi-view geometry and 3D reconstruction.
- At least basic understanding of neural networks and computer graphics.
- Knowledge of the Python programming language. Some previous CUDA exposure would be appreciated.

Tutor: Marco Orsingher, Paolo Medici

Duration: 6 months

Algo3.0 Team Internships

Internship 13

Title: Deep Learning-based Visual-LiDAR Odometry (VLO)

Description: Odometry is a key task for autonomous vehicles, correctly estimating the pose is crucial for subsequent perception and map generation algorithms. The goal of the internship is to develop a deep learning-based algorithm to estimate the vehicle position between consecutive frames, given the RGB camera and 3D (Radar and/or LiDAR) information. The intern would need to:

- Review state-of-the-art methods on the VLO task.
- Implement a state-of-the-art method, reproducing papers' results on a public real-world urban dataset.
- Investigate how this method generalize on other datasets, comparing model's performance when fine-tuning, training from scratch and zero-shot.
- Comparison with non-deep odometry algorithms (e.g. KISS-ICP)

Requirements:

- Basic understanding of classical computer vision.
- Good understanding of neural networks.
- Knowledge of the Python programming language.

Tutor: Fabjan Kallasi, Lorenzo Cipelli

Duration: 6 months (From September '26)

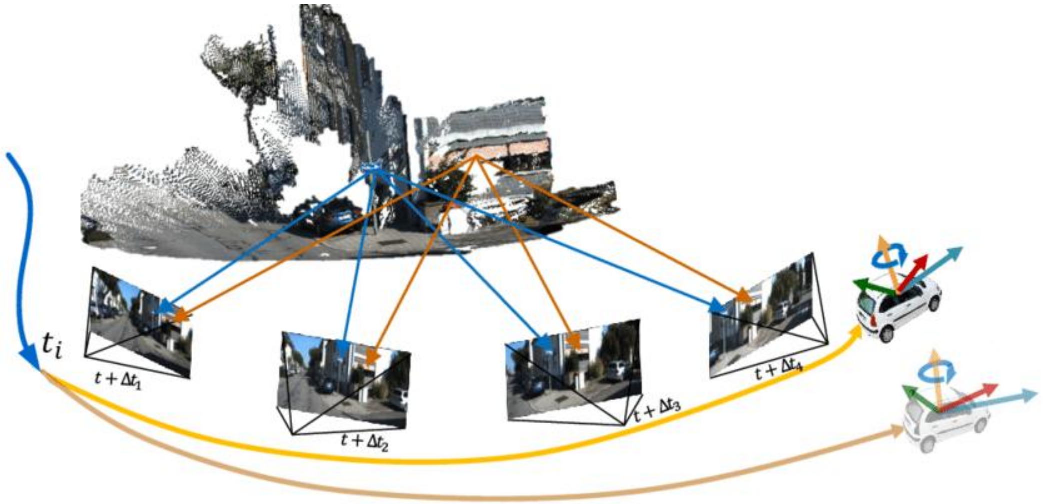

References:

1. Liu, Jiuming, et al. "Dvlo: Deep visual-lidar odometry with local-to-global feature fusion and bi-directional structure alignment." *European Conference on Computer Vision*. 2024.
2. Vizzo, Ignacio, et al. "Kiss-icp: In defense of point-to-point icp–simple, accurate, and robust registration if done the right way." *IEEE Robotics and Automation Letters*. 2023.

Internship 14

Title: Multi-modal Deep Learning-based 3D Semantic Segmentation

Description: 3D semantic segmentation (3DSS) is a fundamental task in the auto-labeling pipeline of autonomous driving systems. It involves assigning a semantic label to each point acquired by a 3D sensor (LiDAR or radar) in the 3D space, thereby enabling a detailed understanding of the environment surrounding the vehicle. While most of the 3DSS approaches are based only on lidar pointclouds, the tri-dimesional information can be enhanced with the RGB information coming from the cameras. To improve the quality of the 3DSS, the intern would need to:

1. *Preliminary study*: review the current state-of-the-art (SOTA) approaches for 3DSS and image backbones.
2. *Baseline implementation*: Focus on one SOTA method, akin Pointnet [1] or PTv3 [2], and implement their own solution using public datasets.
3. *Contribution*: Propose and implement a solution that improves the baseline with the usage of images.

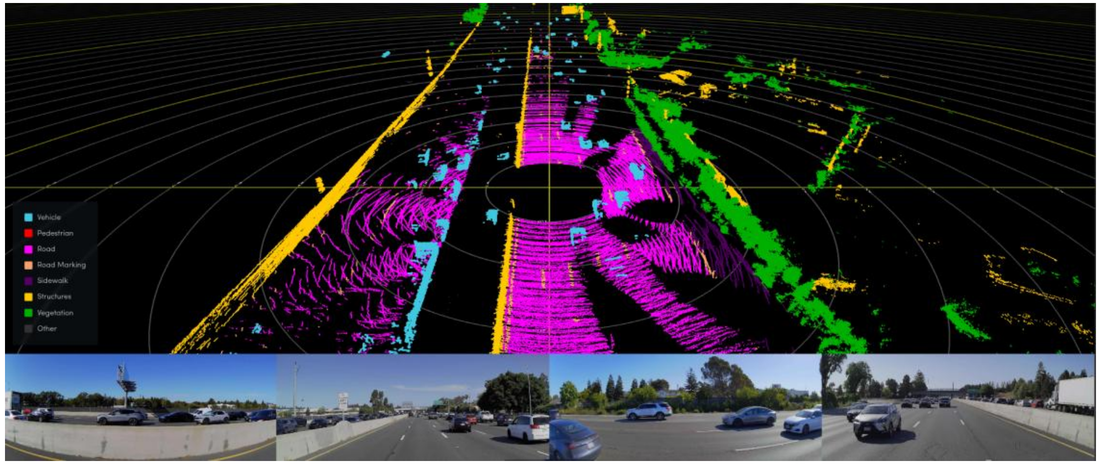

Requirements:

1. Basic understanding of classical computer vision and neural networks.
2. Knowledge of the Python programming language.
3. [+] prior experience with DL frameworks (pytorch, tensorflow)

Tutor: Fabjan Kallasi, Emanuele Ghelfi

Duration: 6 months

References:

1. Qi, Charles R., et al. "Pointnet: Deep learning on point sets for 3d classification and segmentation." Proceedings of the IEEE conference on computer vision and pattern recognition. 2017.
2. Wu, Xiaoyang, et al. "Point Transformer V3: Simpler, Faster, Stronger." CVPR (2024).

Internship 15

Title: Self-Supervised Learning (SSL) for Autonomous Driving Perception

Description: Self-Supervised learning (SSL) methods have recently gained popularity with 2D images. However, the application of these methods for 3D point cloud is still in early stage, especially in autonomous driving domain. The goal of this internship is to study and improve the existing SSL methods for 3D data, using a multi-modal approach (images and 3D point cloud) to boost representational power in automotive domain.

- Review state-of-the-art SSL for images and point cloud.
- Focus on SSL methods on a public dataset and evaluate its effectiveness on another dataset.
- Further improve the methods by understanding and overcoming the drawbacks, consider using images too.
- Finetuning the network trained with the SSL method on a downstream task (semantic segmentation)

Requirements:

- Basic understanding of classical computer vision.
- Good understanding of neural networks.
- Knowledge of the Python programming language.

Tutor: Fabjan Kallasi, Mario Cerino

Duration: 6 months

References:

- Wu, Xiaoyang, et al. "Sonata: Self-supervised learning of reliable point representations." *Proceedings of the Computer Vision and Pattern Recognition Conference*. 2025.

Internship 16

Title: Mesh-Based 3D Gaussian Splatting

Description: 3D Gaussian Splatting (GS) has recently emerged as a promising technique in autonomous driving and dynamic scene representation. However, most existing approaches rely primarily on sequence-based data and do not effectively leverage prior knowledge.
Incorporating pre-trained models offers a compelling direction to address this limitation. In particular, image-to-3D reconstruction models capable of generating object-level meshes can be used to provide additional geometric supervision, enhancing the accuracy and robustness of Gaussian Splatting methods.
This thesis explores the integration of mesh-based priors into 3D Gaussian Splatting pipelines to improve scene understanding and reconstruction quality.

Requirements:
- Good understanding of classical computer vision, especially multi-view geometry and 3D reconstruction.
- Basic understanding of neural networks.
- Knowledge of the Python programming language. Some previous CUDA exposure would be appreciated.

Tutor: Fabjan Kallasi, Emanuele Ghelfi

Duration: 6 months

References:

- Hess, Georg, et al. "Splatad: Real-time lidar and camera rendering with 3d gaussian splatting for autonomous driving. *Proceedings of the Computer Vision and Pattern Recognition Conference*. 2025."

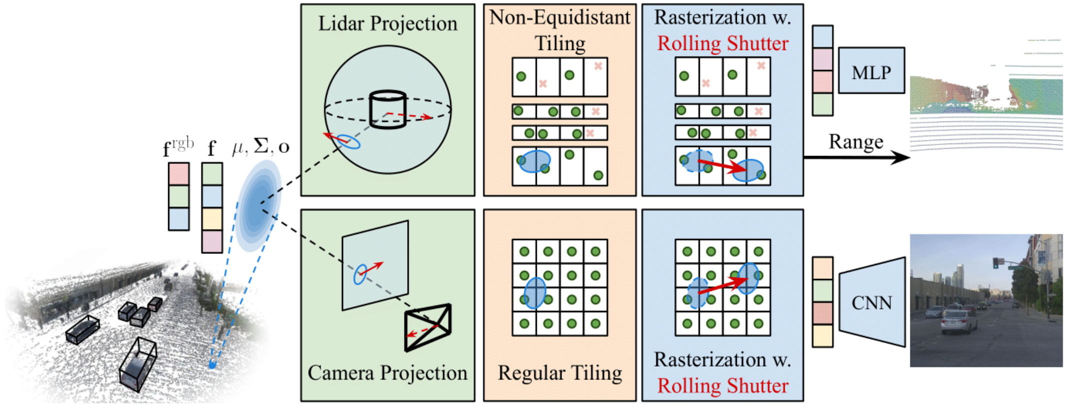

Internship 17

Title: Unknown Dynamic Objects

Description: Unknown dynamic objects refer to scene elements that are not covered by standard object detectors and therefore remain unmodeled during perception. Detecting such objects is crucial for robust scene understanding, particularly in real-world and safety-critical scenarios where unexpected motion may occur.
Self-supervised learning techniques offer a promising approach to identifying these objects by exploiting consistency and motion cues directly from the data, without relying on explicit annotations. This thesis focuses on developing self-supervised methods to detect and reason about unknown dynamic objects in complex environments.

Requirements:

- Basic understanding of classical computer vision.
- Good understanding of Neural Networks.
- Knowledge of the Python programming language.

Tutor: Fabjan Kallasi, Emanuele Ghelfi

Duration: 6 months

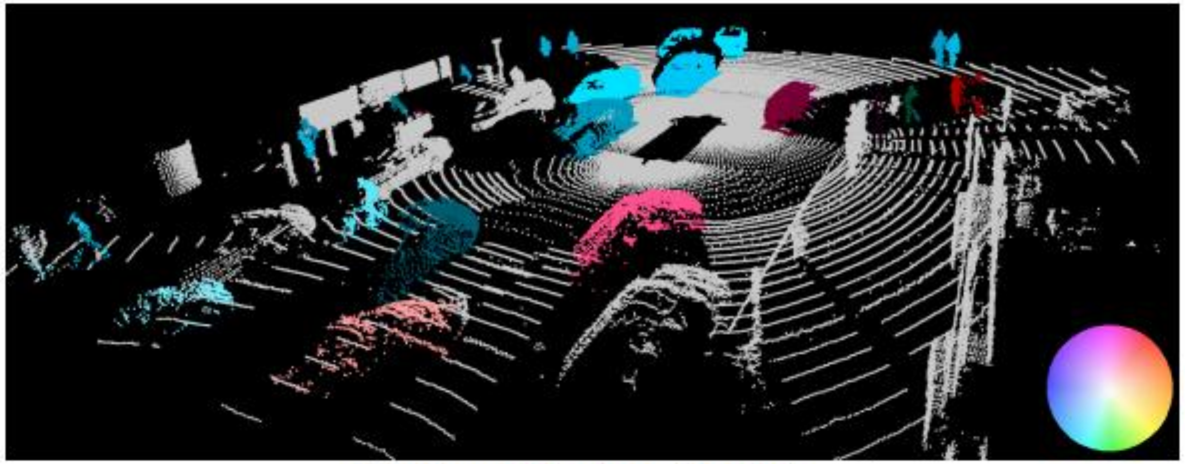

(a) Ground truth motion data,
(moving points are colored by their motion, static points are Gray)

References:

- Li, R., Shi, H., Wang, Z., & Lin, G. (2025). Weakly and self-supervised class-agnostic motion prediction for autonomous driving. IEEE Transactions on Pattern Analysis and Machine Intelligence.

Internship 18

Title: Iterative LiDAR-Camera Calibration via Surrogate Diffusion Process

Description: The LiDAR-Camera calibration task aims to find the correct extrinsic transformation matrix that aligns a LiDAR point cloud with a reference camera image. Many state-of-the-art approaches address this task through an iterative refinement strategy to progressively reduce the calibration error. This is achieved by training multiple instances of the same network, each specialized for a different range of initial mis-calibration.
Recent advances in LiDAR-Camera calibration have explored the use of surrogate diffusion processes to provide a versatile iterative framework, thereby eliminating the need to train multiple network instances. While this approach represents a promising alternative, the diffusion-based strategy has shown inconsistent performance and it may lead to degraded calibration performance.
Thus, the proposed internship will begin with a comprehensive review of the current state-of-the-art, with a particular focus on the adaptation of diffusion processes originally designed for image generation to the LiDAR-Camera calibration task. Building upon an existing baseline, the intern will then design his/her own solution, possibly by exploring latent diffusion mechanisms.

Requirements:
- Basic understanding of classical computer vision.
- Basic understanding of neural networks.
- Knowledge of the Python programming language.
- Basic knowledge of PyTorch library.

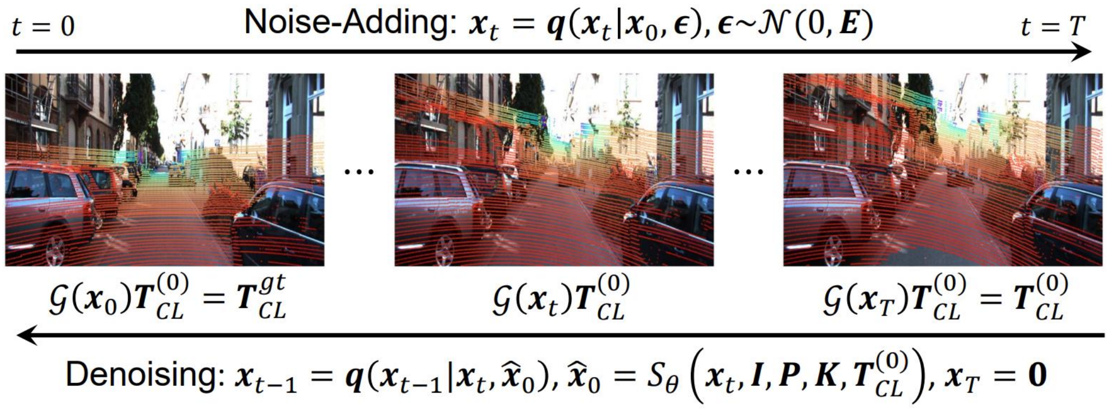

Tutor: Fabjan Kallasi, Filippo D'Addeo

Duration: 6 months

References:

- Ou, Ni, et al. "Iterative Camera-LiDAR Extrinsic Optimization via Surrogate Diffusion." *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2025.
- Lv, Xudong, et al. "LCCNet: LiDAR and camera self-calibration using cost volume network." Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021.
- Rombach, Robin, et al. "High-resolution image synthesis with latent diffusion models." Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. 2022.

Internship 19

Title: World Model for autonomous driving

Description: World Models for autonomous driving aim to learn a latent representation of the environment and its dynamics, enabling prediction of future states from sensor data. These models leverage self-supervised learning to jointly model perception, prediction, and planning, often using transformers, diffusion models, or 3D occupancy representations. They are key for simulation, decision-making, and safer trajectory planning in complex driving scenarios.

The aim of this thesis is to investigate and implement **World Models** for autonomous driving, focusing on learning compact latent representations of dynamic driving environments using self-supervised techniques. The work will explore state-of-the-art methods for predicting future scene evolution from sensory data and assess their effectiveness in capturing both spatial and temporal dynamics.
The thesis will also aim to implement and evaluate one or more modern World Model approaches on real-world driving data, analyzing their performance in terms of prediction accuracy, generalization, and suitability for downstream tasks such as motion forecasting and planning.

Requirements:
- Basic understanding of classical computer vision.
- Basic understanding of neural networks.
- Knowledge of the Python programming language.
- Basic knowledge of PyTorch library.

Tutor: Fabjan Kallasi, Mario Cerino

Duration: 6 months

References:

- Y. Guan, H. Liao, Z. Li, J. Hu, R. Yuan, Y. Li, G. Zhang, and C. Xu, **"World Models for Autonomous Driving: An Initial Survey,"** *arXiv preprint arXiv:2403.02622*, 2024
- T. Feng, W. Wang, and Y. Yang, **"A Survey of World Models for Autonomous Driving,"** *arXiv preprint arXiv:2501.11260*, 2025.

Internship 20

Title: Open Vocabulary 3D Semantic Segmentation for Autonomous Driving

Description: 3D semantic segmentation is a key task in autonomous driving perception systems, enabling scene understanding through labeling of LiDAR point clouds and multi-modal sensor data. Traditional segmentation approaches rely on supervised learning with predefined semantic categories, limiting their ability to generalize to unseen objects and "long-tail" scenarios encountered in real-world driving environments.
Advances in vision-language models (VLM) have introduced open-vocabulary perception, that is, models capable of recognizing objects beyond fixed label sets by leveraging semantic knowledge learned from large scale image-text datasets, enabling zero-shot segmentation based on textual queries rather than closed-set training labels.
After an initial review of state-of-the-art methods on 3D Open Vocabulary Semantic Segmentation, the goal of the thesis will be adapting open-vocabulary learning to outdoor 3D autonomous driving scenarios, with particular emphasis on the integration of LiDAR/RADAR point clouds into the segmentation pipeline.

Requirements:
- Basic understanding of classical computer vision.
- Basic understanding of neural networks.
- Basic experience with PyTorch or similar deep learning frameworks.

Tutor: Fabjan Kallasi, Lorenzo Cipelli

Duration: 6 months (From September '26)

References:

- OpenScene: https://arxiv.org/pdf/2211.15654
- OpenMask3D: https://arxiv.org/pdf/2306.13631
- OpenAD: https://arxiv.org/pdf/2411.17761

Internship 21

Title: Real-Time Deep Learning-based Multi-Radar Odometry

Description: Accurate ego-motion estimation is essential for autonomous driving. While camera- and LiDAR-based odometry may degrade under adverse weather, poor lighting, or low-visibility conditions, modern 4D automotive radars offer a robust sensing alternative by providing long-range measurements, Doppler velocity, RCS, and sparse 3D point clouds.

The objective of this thesis is to develop a deep learning-based odometry pipeline for a multi-radar setup using real-world 4D automotive radar data, with a focus on real-time inference. The student will review state-of-the-art radar odometry methods, implement and adapt a learning-based baseline to a multi-radar configuration, investigate lightweight neural architectures, and benchmark the proposed approach against classical geometric odometry methods.

Requirements:
- Good understanding of deep learning and neural networks.
- Good programming skills in Python.
- Basic understanding of 3D geometry and rigid transformations.
- Experience with PyTorch is appreciated.

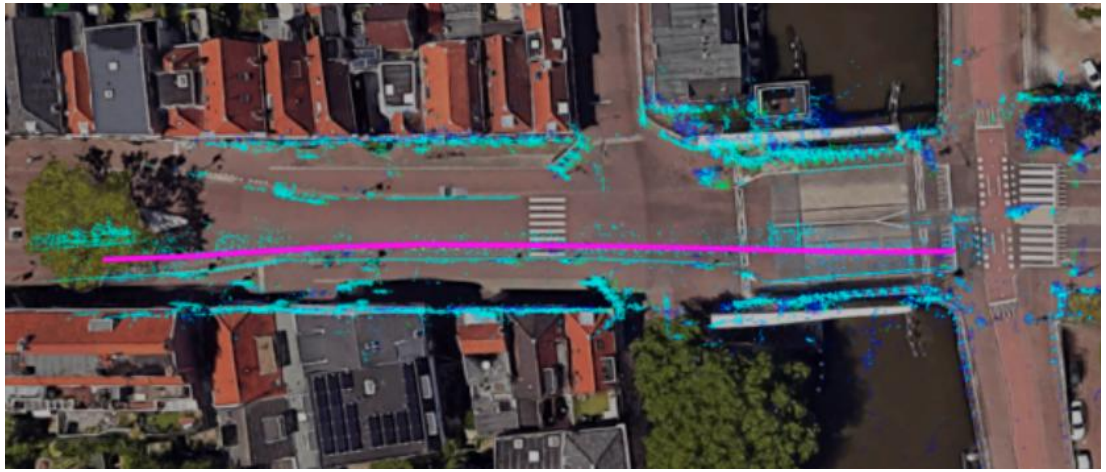

Tutor: Fabjan Kallasi, Riccardo Maiolani

Duration: 6 months

References:
1. S. Lu, H. Zhou, G. Zhuo, X. Tang, “DNOI-4DRO: Deep 4D Radar Odometry with Differentiable Neural-Optimization Iterations,” Proceedings of the AAAI Conference on Artificial Intelligence, 2026.
2. Z. Li, W. Wang, Q. Shen, Y. Zhao, Z. Fang, “Super4DR: 4D Radar-centric Self-supervised Odometry and Gaussian-based Map Optimization,” arXiv:2512.09608, 2025.
3. Z. Li, Y. Cui, N. Huang, C. Pang, Z. Fang, “CAO-RONet: A Robust 4D Radar Odometry with Exploring More Information from Low-Quality Points,” 2025 IEEE International Conference on Robotics and Automation (ICRA), 2025.

Ambarella
AI envisioned™

AMBARELLA.COM

VisLab internship opportunities

June 2026

VisLab Introduction

A **high-tech company** based in Parma, VisLab is focused on **all aspects of ADAS and autonomous driving i**ncluding machine learning, deep neural networks (DNNs), stereovision, robotics, sensor fusion, planning, system calibration, and more.

Founded in the late ‘90s as a research laboratory at the University of Parma, VisLab quickly established itself as a **pioneer in vehicle autonomy** before spinning off from the University in 2009.

In 2015, VisLab was **acquired by Ambarella**, a publicly traded semiconductor Company based in **Silicon Valley** headquartered in Santa Clara (California, USA) and with offices in Detroit (Michigan, USA), Munich (Germany), Hong Kong, Hsinchu (Taiwan), SeongNam City (South Korea), Shanghai (China), Shenzhen (China), and Yokohama (Japan).

Specializing in the development of advanced computer vision hardware processors, **Ambarella relies on VisLab for core research and development** across all the markets with a **particular focus on automotive**.

VisLab provides top level research in the most contemporary high-tech fields such as autonomous driving, robotics, and deep learning. VisLab equips and operates a **fleet of autonomous vehicles** and received the **first authorization from the Italian Ministry of Transportation** to run autonomous driving tests on open roads.

Check www.vislab.it for further information.

VisLab: 25 Years of CV and Autonomous Driving

VisLab Today

- ***Algorithm first* approach** applied to Computer Vision
- Continuous research benefits **future Ambarella chips**

1998: Mille Miglia 2000+km on Italian highways 94% autonomous steering

2005, 2007: DARPA Grand and Urban Challenge. 100% AD

2010: VIAC 15K km cross-continent drive AD following

2013: PROUD: Braive drove 13km in Parma, Level 4 demo 100% driverless AD

2014: DEEVA, 13 stereo camera systems; CV integration; 100% AD

2015: Ambarella acquires VisLab

CES 2018: camera perception stack demonstrated on CV1

2018: AD demonstration in Santa Clara office

CES 2019: AD driving on the roof track of Lingotto building in Turin

CES 2020: L4 autonomous driving demo in Las Vegas with automatic parking. CV2-based

CES 2022: L4 autonomous driving demo and new Oculii-based sensing suite.

2023: Full L4 autonomous driving stack demo on CV3

Ambarella's History / Products

- Founded in 2004
- Started: Camcorders and Broadcast Television
- Today: Action/Security Cameras and Autonomous Driving

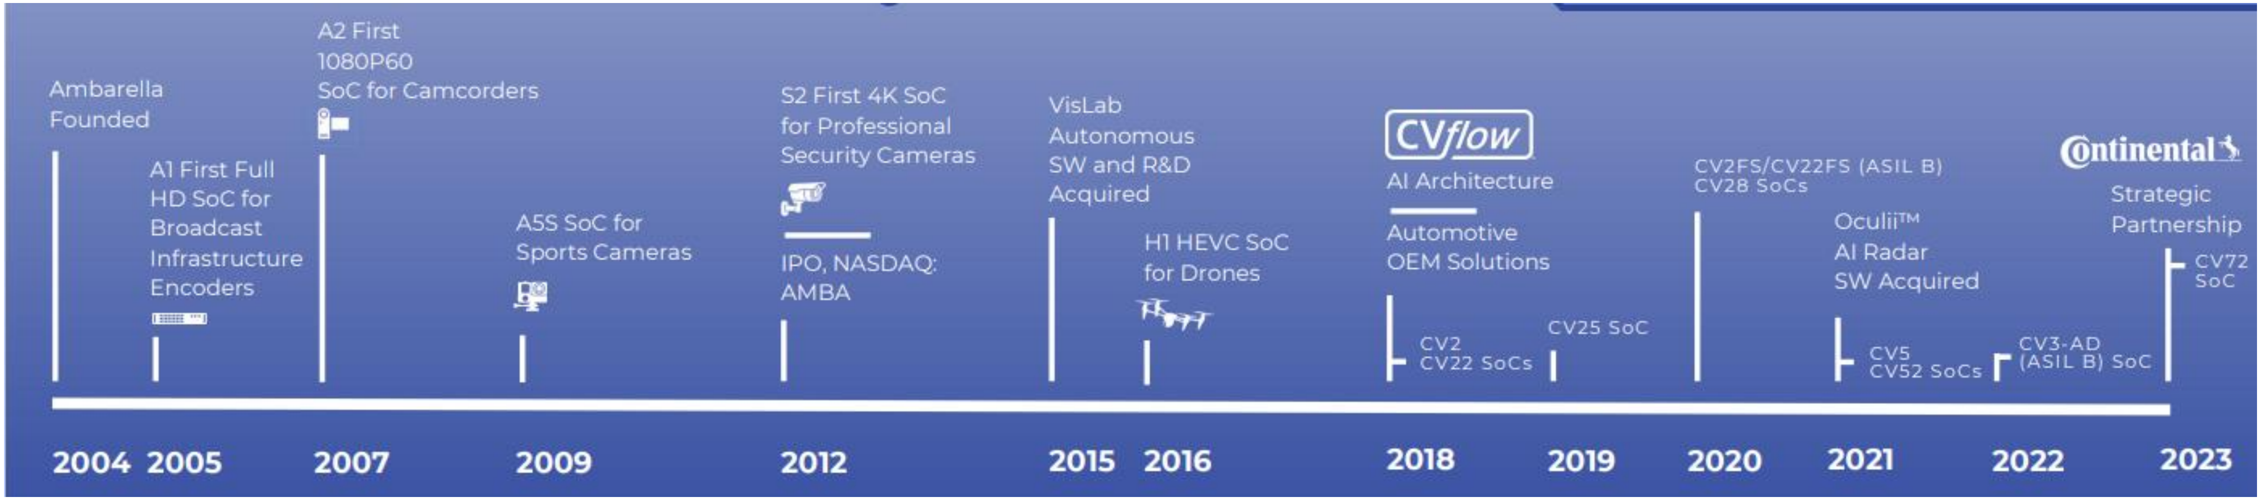

https://www.ambarella.com/wp-content/uploads/Ambarella-Fact-Sheet.pdf

You may know these guys...

GoPro Hero 3/4/5

Insta360 One RS Action

DJI Phantom 4 Pro

They run on Ambarella's silicon!

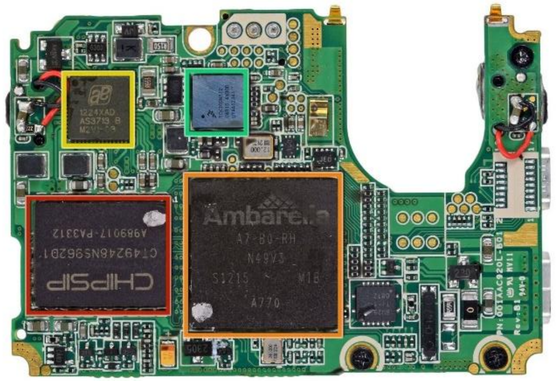

GoPro Hero 3+ Black teardown (ifixit)

Ambarella A770 camera system-on-a-chip

Ambarella's Scalable Solution Line-up

Covering the Entire Range of Sensing, Viewing and Combo Applications

Viewing / Recording		Viewing / Recording / Sensing (Powerful Computer Vision/DNN engine)						
A12AQ	H22AQ	H32AQ	CV25AQ	CV22AQ	CV2AQ	CV22FS	CV2FS	CV3 Family
28nm	14nm	10nm (Gen 1)				10nm (Gen 2)		5nm (Gen 2)

Driver and In-Cabin Monitoring

Drive Recorders

Single / Dual Channel Recorder

Multi-channel Channel Recorder and Data Logger

Electronic Mirror

Rear-View E-Mirror

Side-View E-mirror

Combined Rear-View / Side-View

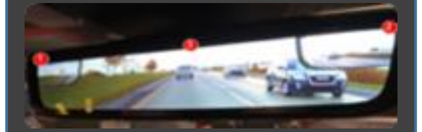

Front-facing ADAS Camera

Single Camera System

CV22FS & CV2FS are **Pin-2-Pin** compatible

Stereo NCAP+

CV3 Family

Powerful ECUs for advanced autonomy applications (L2+ to L4/L5)

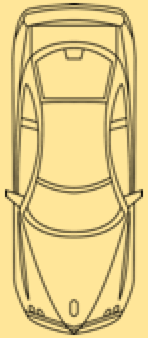

CV3-AD685 our latest Automotive AI Domain Controller

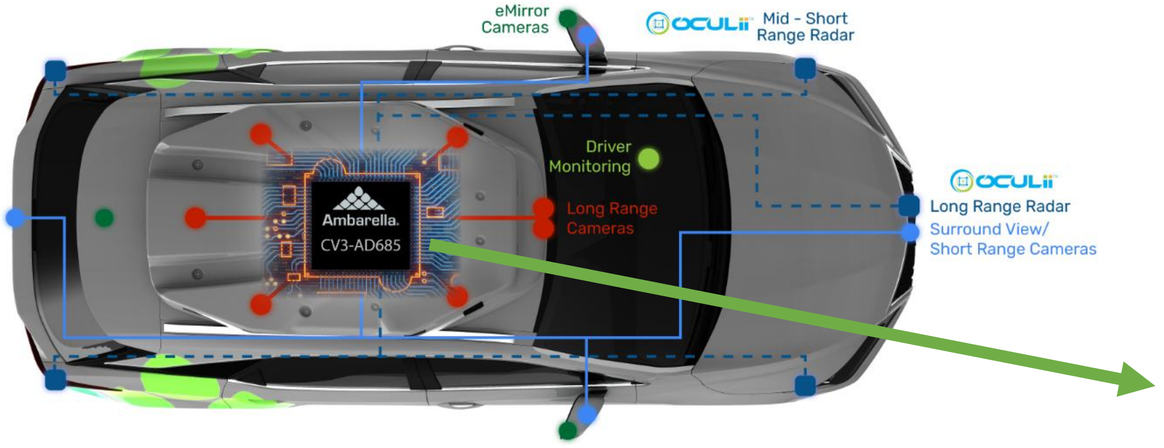

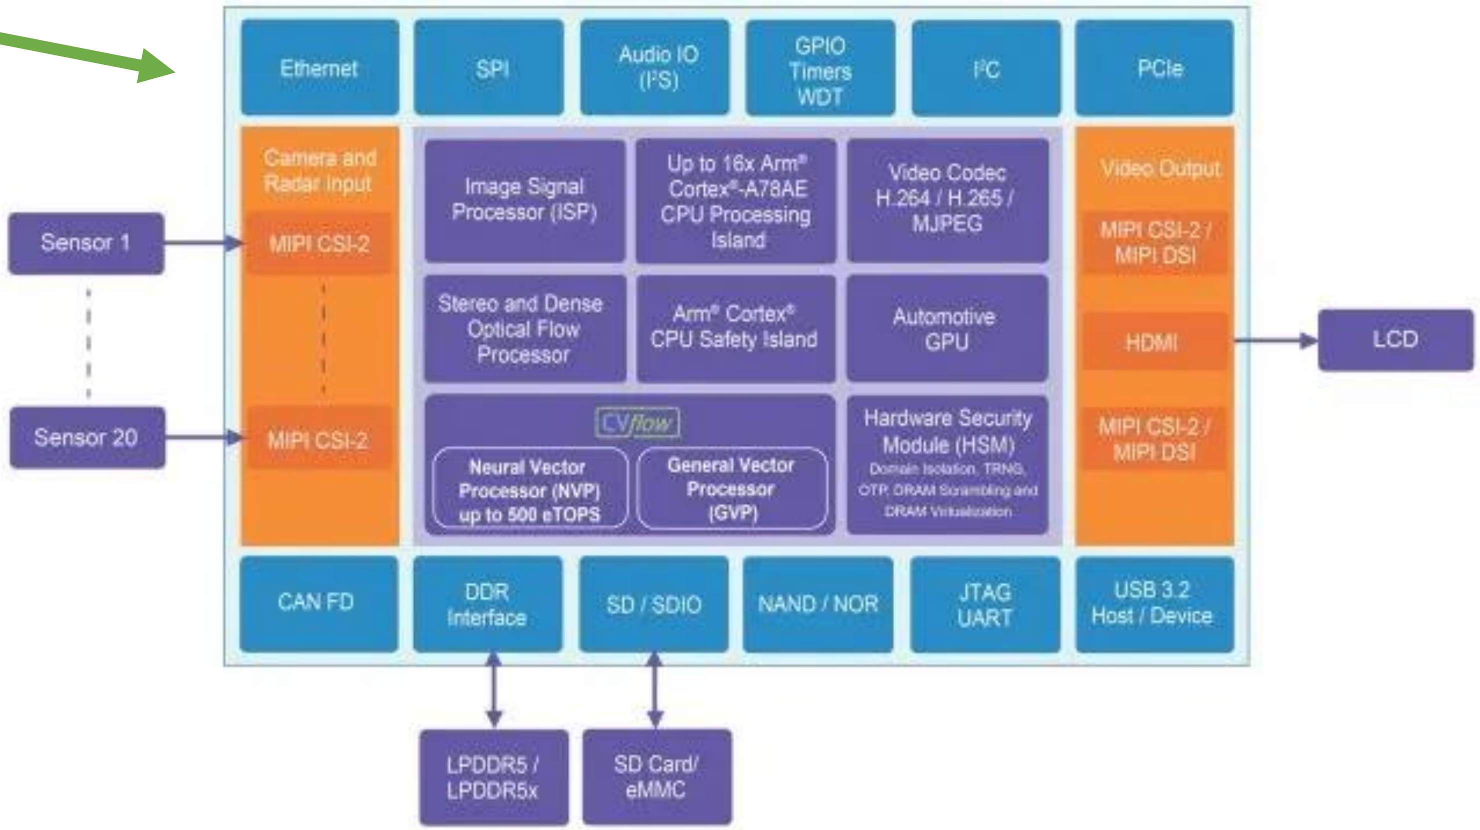

- **Neural vector processor** (NVP) with network (NN)-based processing
- **General vector processor** (GVP) for offloading classical computer vision and radar processing
- **Processing island** (ASIL B): up to 16x Cortex A78AE cores
- **Safety island** targeted to meet ASIL D requirements

VLSI Team Location

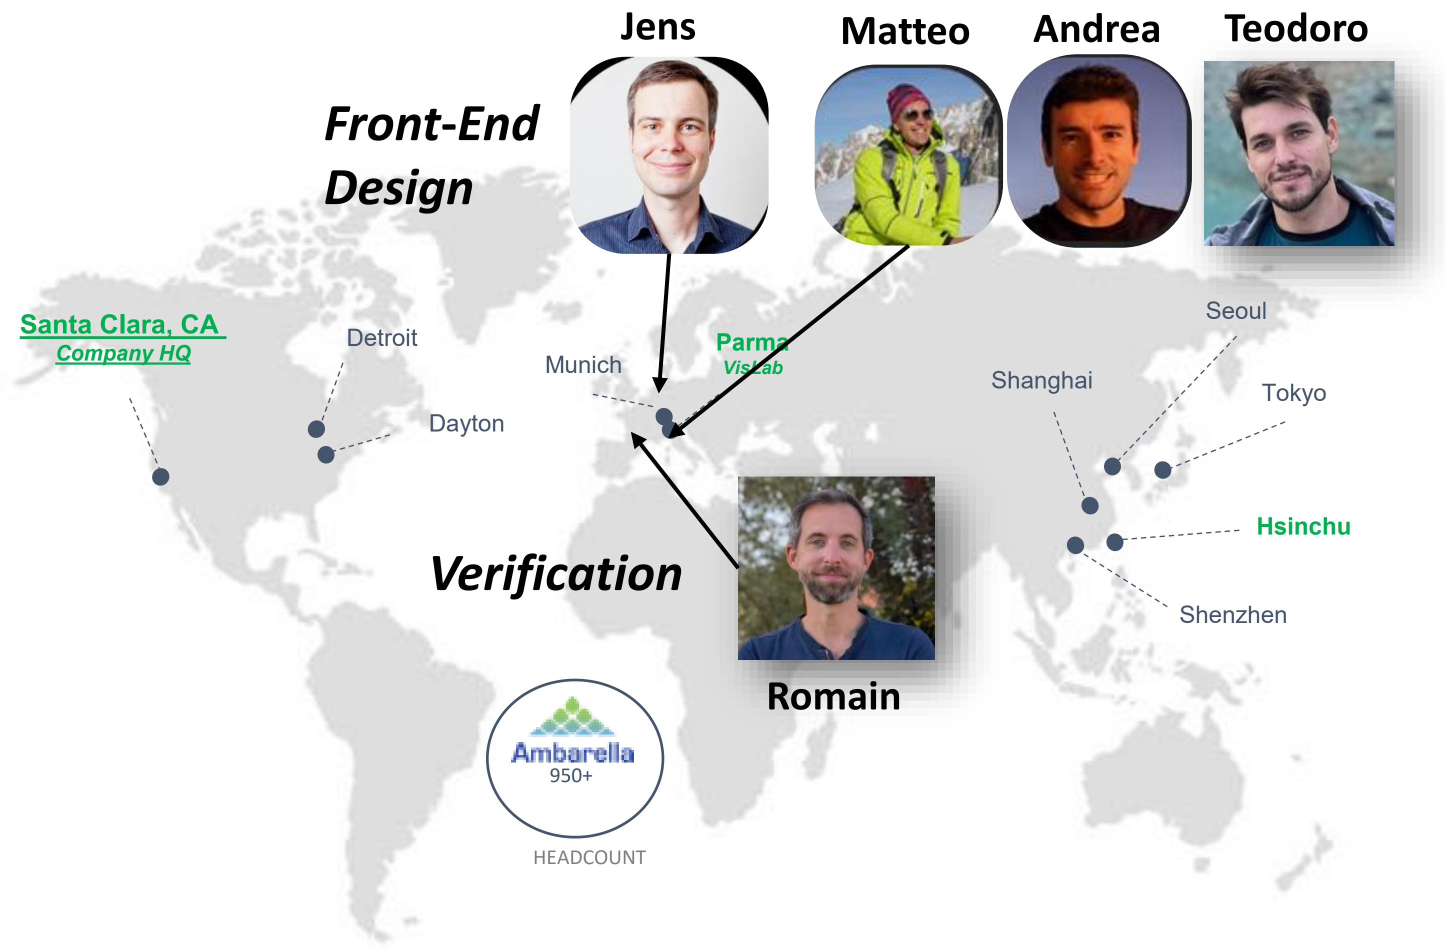

VLSI Design Sites

Internship 1

Title: Port the Open RISC Cores to RISC-V

Description: Complex modules inside Ambarella's SoCs often are equipped with one or more dedicated cores (Open RISC Cores) that are responsible for operating the module. This is done through high-level C code that's specifically compiled to target the ORC Instruction Set, alongside some specific custom instructions. The first ORC we want to transition to RISC-V is the DRAM Custom Training Engine, a simple single thread core that handles DRAM-specific commands with few custom instructions

Some of the core characteristics we look for:

- Complete RV32I instruction set.
- Custom instructions.
- Integration in AMBA simulation environment
- Simulation environment (e.g. based on Verilator).
- Synthesizable in 4nm LP technology

Field: Micro-Architecture / RTL Design

Requirements:

- Basic knowledge of computer architecture.
- Proficiency in C, assembly and Verilog/SystemVerilog.

Tutor: Matteo Mannis, Andrea Bettati

Duration: 6 months

Internship 2

Title: Wavefront Parallel Processing (WPP) architecture study for a next-gen multi-standard video encoder
Description: Modern video encoders need to support multiple standards while still optimize for PPA. The goal of the internship is to study how to apply WPP in a multi-standard (AVC, HEVC, AV1) video video codec and start the RTL development.
Some of the tasks include:

- Study parallelization of block based video codecs (AVC, HEVC, AV1)
- Evaluate 2D memory cache architectures (high level code, C++)
- Optimize prefetch, banking, eviction schemes for 2D memory cache
- Propose architecture of such 2D memory cache
- Implement Verilog RTL and experiment with synthesis in 5nm to 2nm LP technology.

Field: Micro-Architecture / RTL Design

Requirements:

- The intern should have a basic knowledge of the ASIC front-end workflow.
- Proficiency in the C++ programming language and in Verilog/SystemVerilog is also required.
- Experience in using synthesis tools is a plus.

Tutor: Jens Weber, Andrea Bettati

Duration: 6 months

Note: More video coding related topics are available if interested.

Internship 3

Title: Improve Mayo: Structural Netlist Analysis tool

Description: Mayo is an internally developed Perl script, routinely used to parse netlists and categorize the implemented logic. The goal of the internship is to revamp Mayo by:

- Re-writing parts of the tool in a modern programming language like Python.
- Implementing heavy workload routines in C++/Rust.
- Implementing intermediate save points (i.e. using a DB).
- Offloading high memory load (>400GB) to hard disk.
- Implementing unit tests.

Field: Amba Flow Improvement / Software Development

Requirements:

- The intern should have a basic knowledge of the ASIC front-end workflow.
- Proficiency in the Python and C++/Rust programming languages.

Tutor: Jens Weber, Matteo Mannis, Andrea Bettati

Duration: 3 or 6 months

Internship 4

Title: Die-to-Die Interconnect Study

Description: Die-to-die interconnects are at heart of the current industry trend of shifting away from monolithic SoCs towards chiplets architectures.
By enabling high-speed data transfers between different ICs in the same package, the interconnects effectively make chiplets possible.
The goal of the internship is to study the current state of die-to-die interconnects and evaluate their usage in Amba SoCs.
Some of the tasks include:

- Study different protocols for die-to-die interconnects (e.g. UCIE).
- Work out the benefits of standardized protocol vs. proprietary protocol
- How would a controller for die-die communication integrated into the CV HW architecture.
- Implement prototype Verilog RTL controller for die-die interconnects.

Field: Micro-Architecture / RTL Design

Requirements:

- A solid understanding of good software development practices.
- Proficiency in Verilog/SystemVerilog is also required.

Tutor: Jens Weber, Matteo Mannis, Andrea Bettati

Duration: 6 months