

# Augmented Vision - Umwelterfassung und Selbstlokalisierung



Didier Stricker

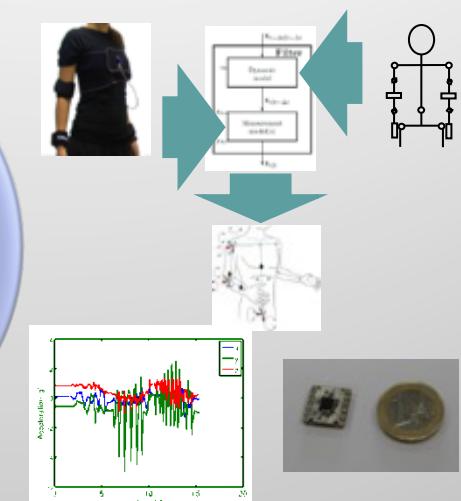
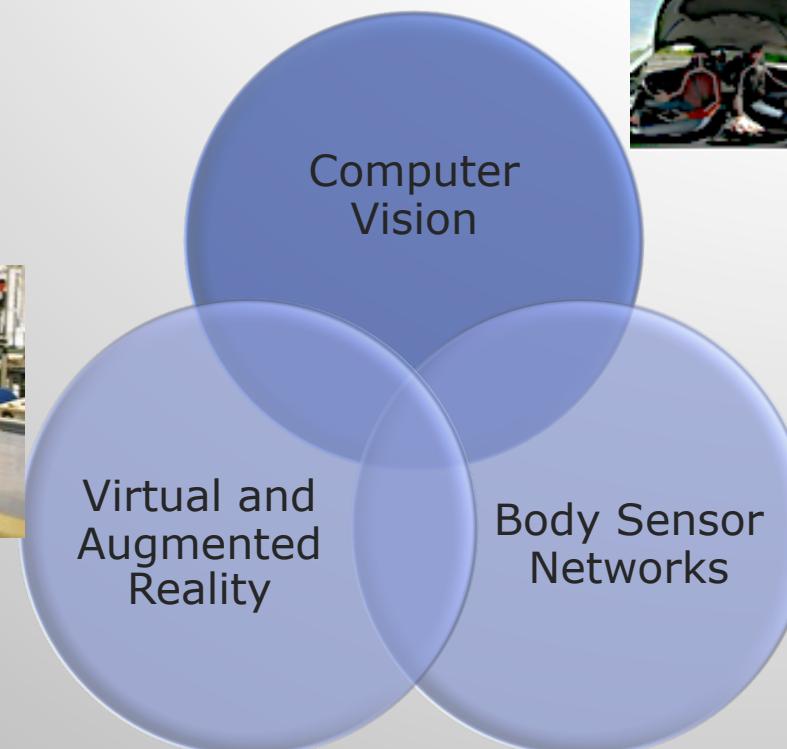
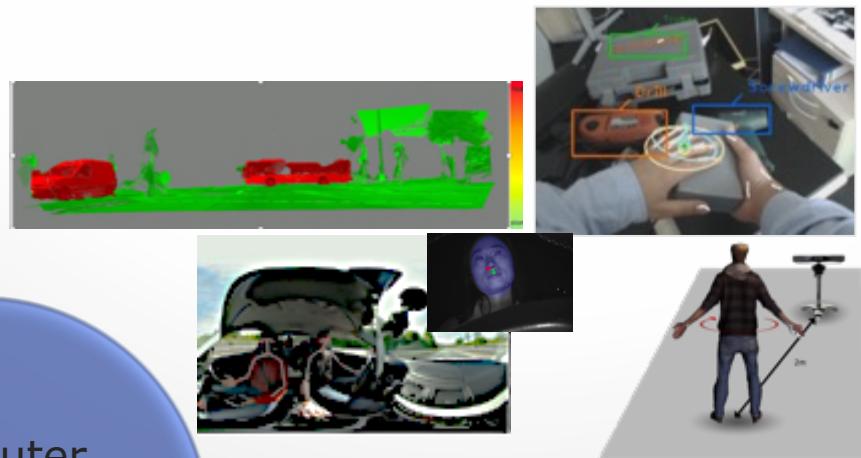


Didier.Stricker@dfki.de



# Forschungsbereich Augmented Vision @ DFKI Kaiserslautern

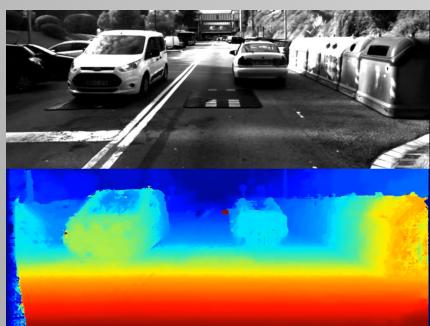
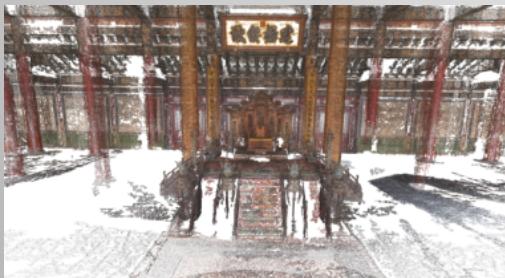
- Leitung: Prof. Didier Stricker
- 50 Doktoranden und wiss. Mitarbeiter
- Drei zusammenhängende Fachgebiete



# Visuelle Wahrnehmung

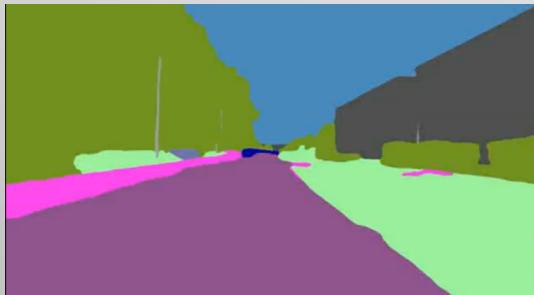
## Geometry

- Multiple view geometry
- Real-time stereo



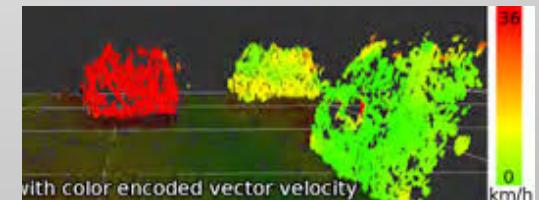
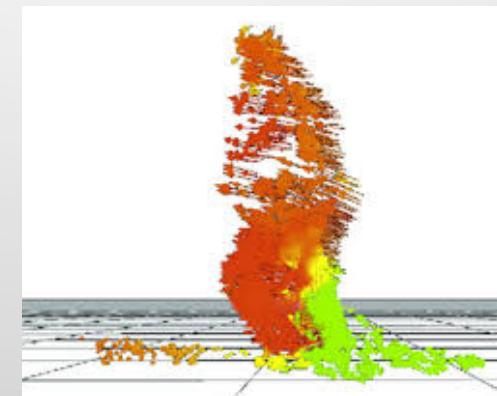
## Semantic

- Object recognition
- Semantic Segmentation



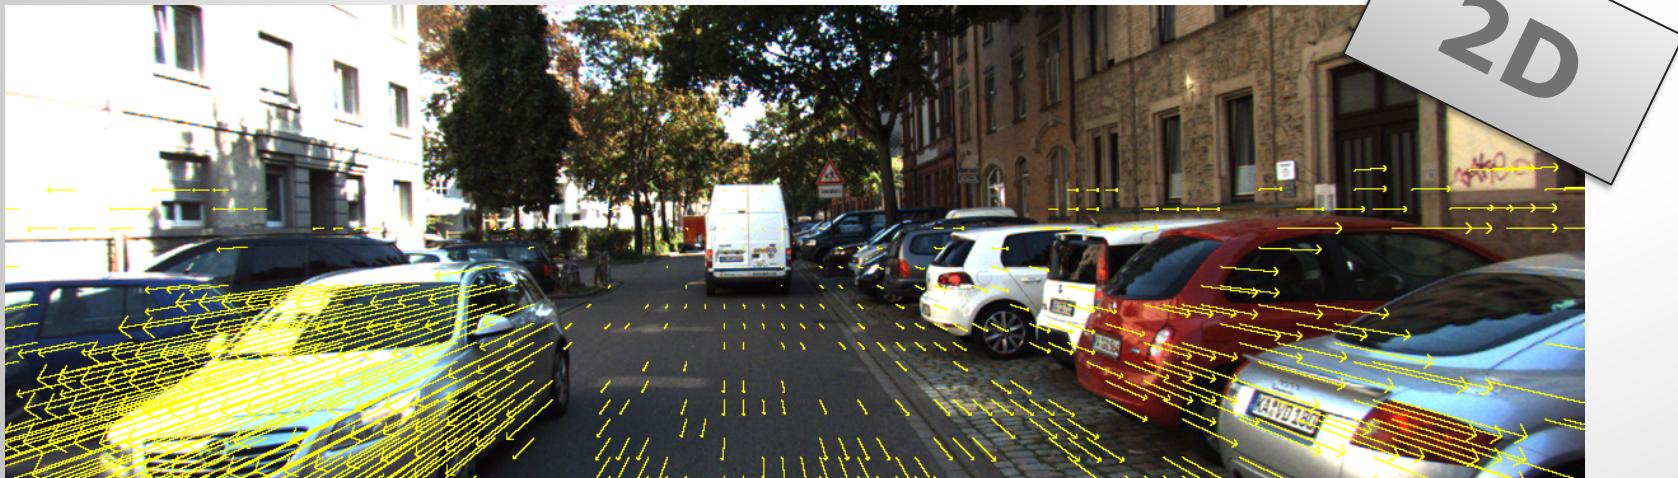
## Dynamics

- **Dense 4D reconstruction**
- **Temporal integration**  
(in order to understand events!)



# Optical Flow vs. Scene Flow

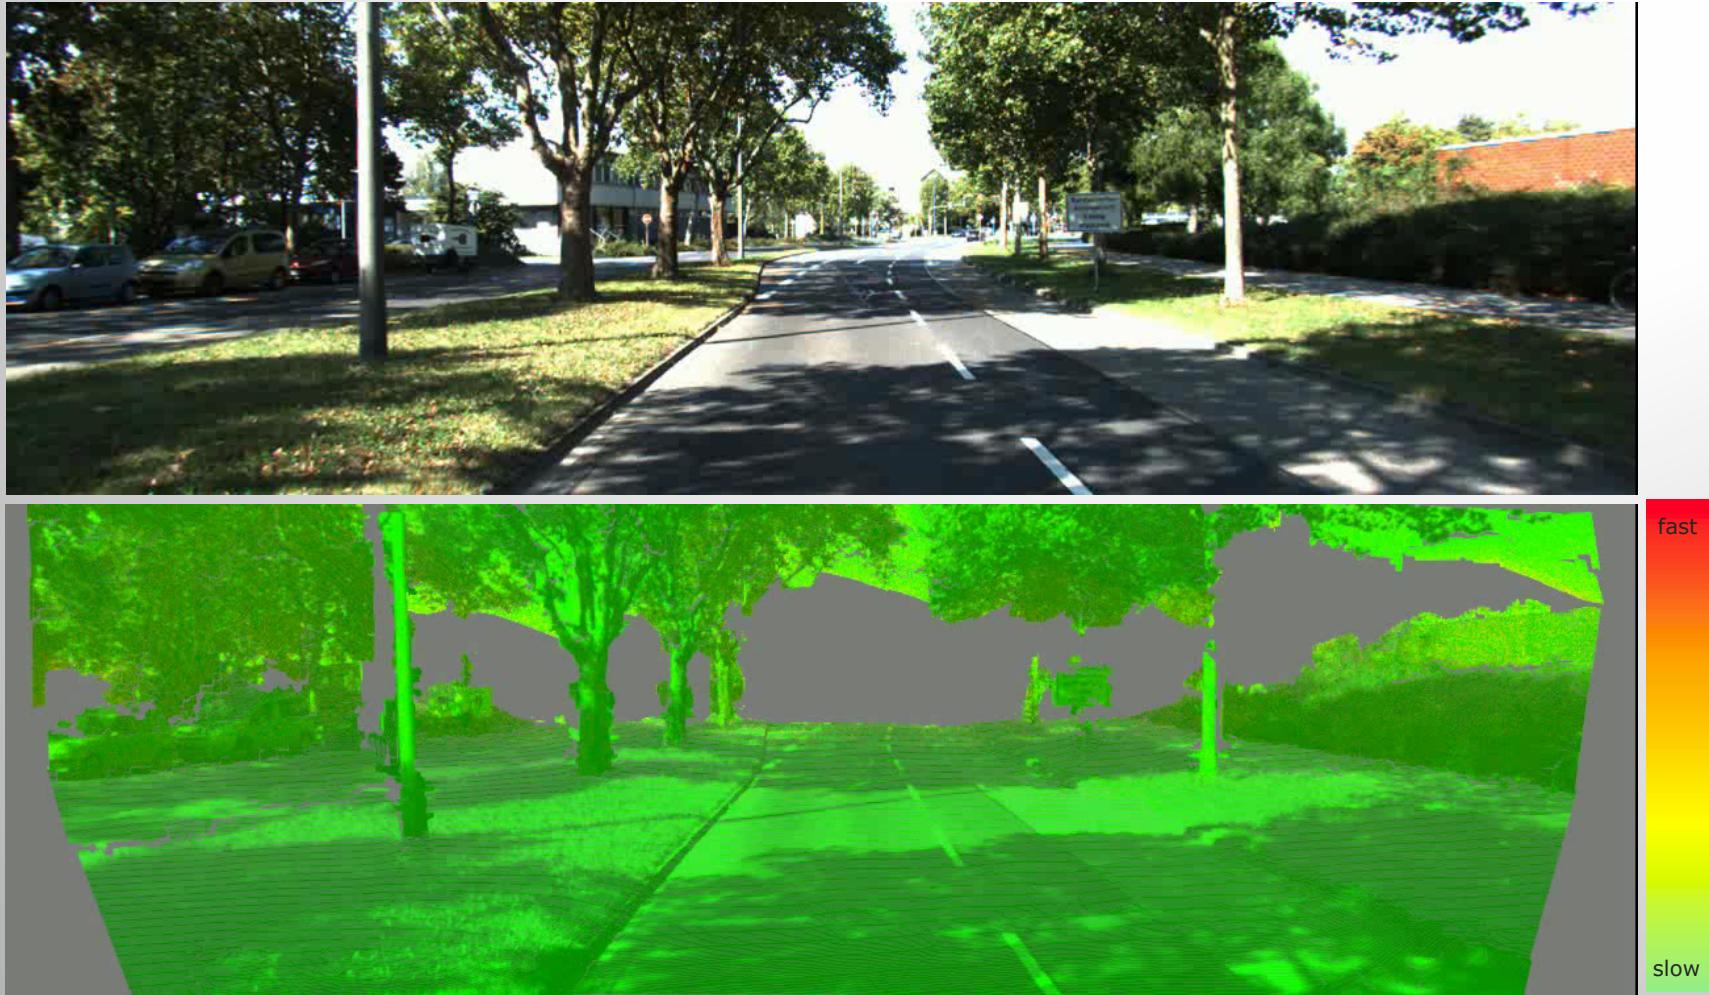
## Optical Flow



## Scene Flow



## 3D Erfassung der Geometrie und Dynamik



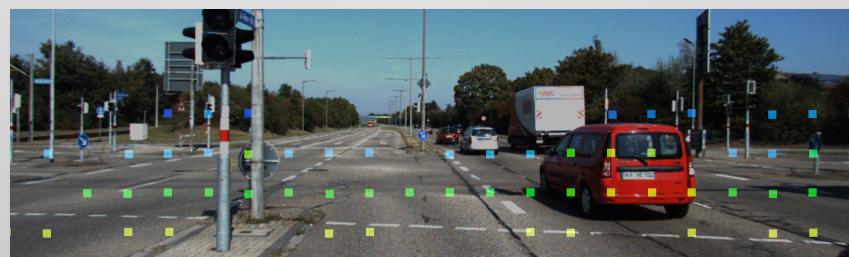
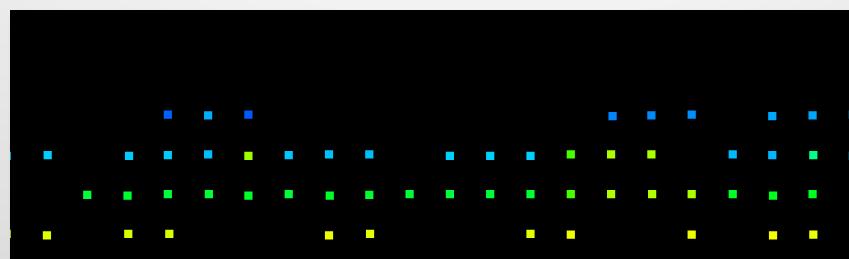
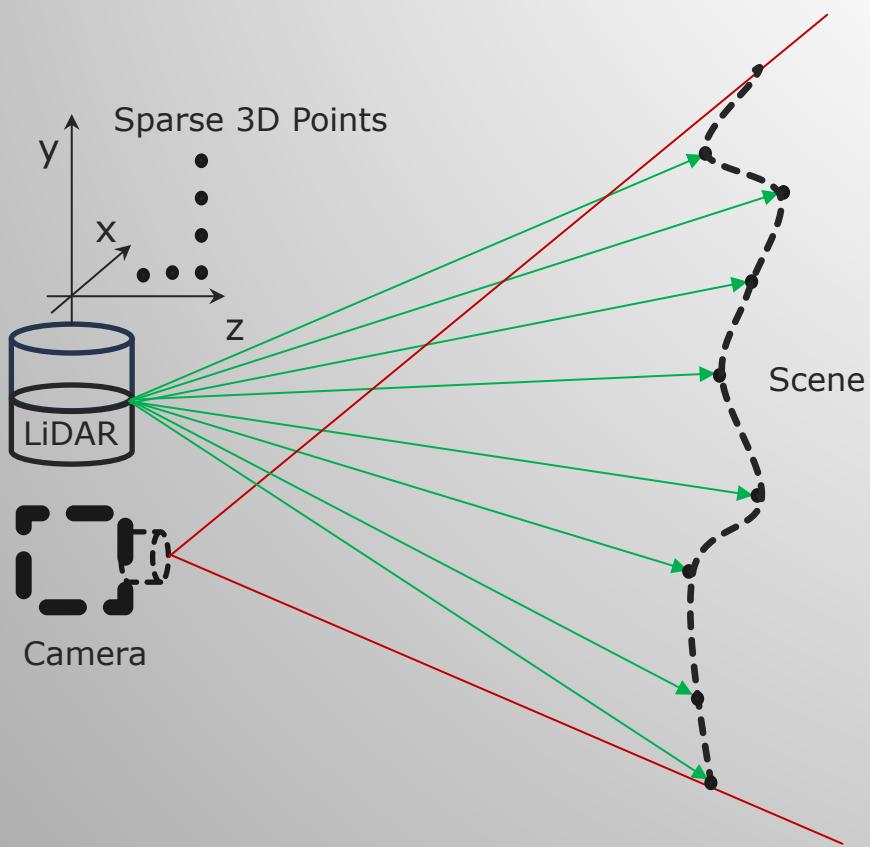
René Schuster, Oliver Wasenmüller, Georg Kuschk, Christian Bailer and Didier Stricker. „Towards Flow Estimation in Automotive Scenarios“, ACM Chapters Computer Science in Cars Symposium (CSCS), 2017.

René Schuster, Christian Bailer, Oliver Wasenmüller and Didier Stricker. "Combining Stereo Disparity and Optical Flow for Basic Scene Flow". Commercial Vehicle Technology Symposium (CVTS), Springer, 2018.

René Schuster, Oliver Wasenmüller, Georg Kuschk, Christian Bailer and Didier Stricker. „SceneFlowField: Dense Interpolation of sparse scene flow correspondences“, IEEE Winter Conference on Computer Vision (WACV), 2018.

## Fusion: Kamera-LiDAR

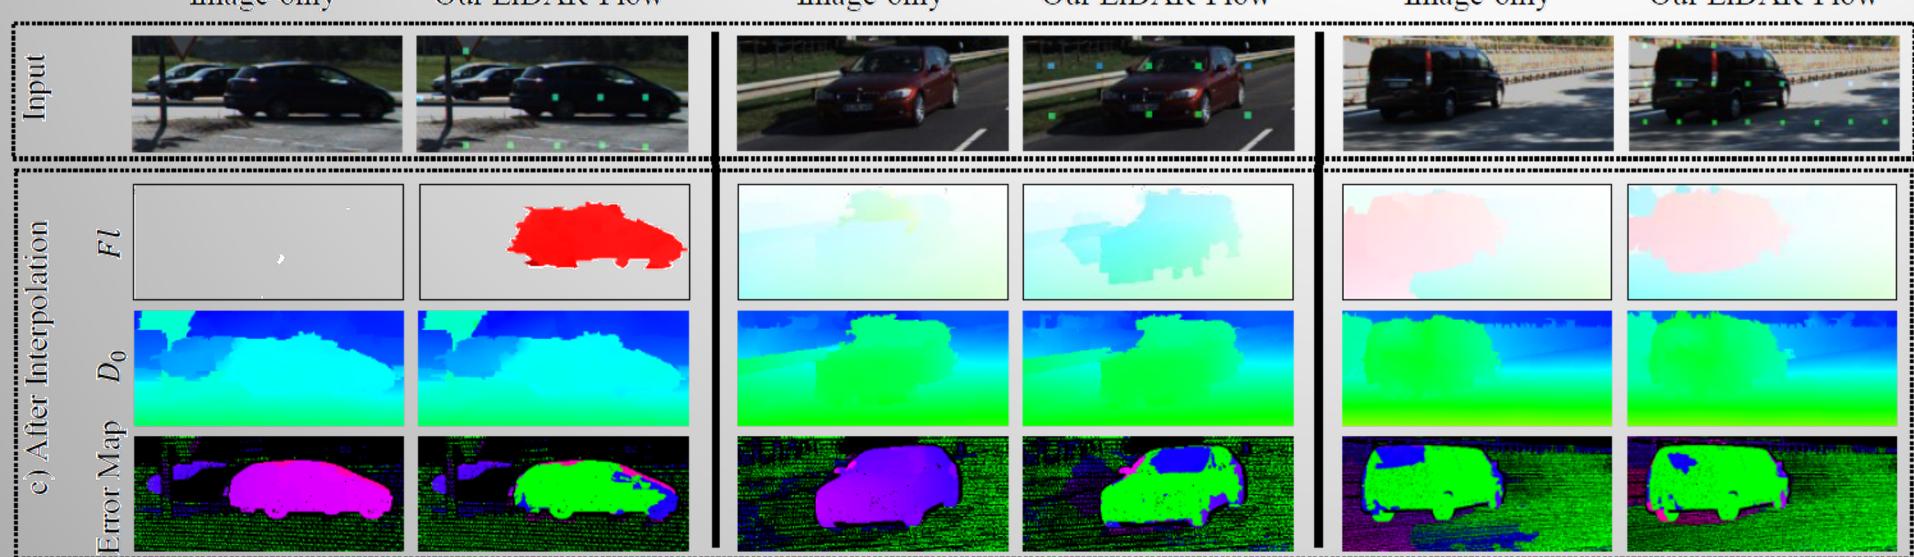
### **Sparse** MEMS-LiDAR



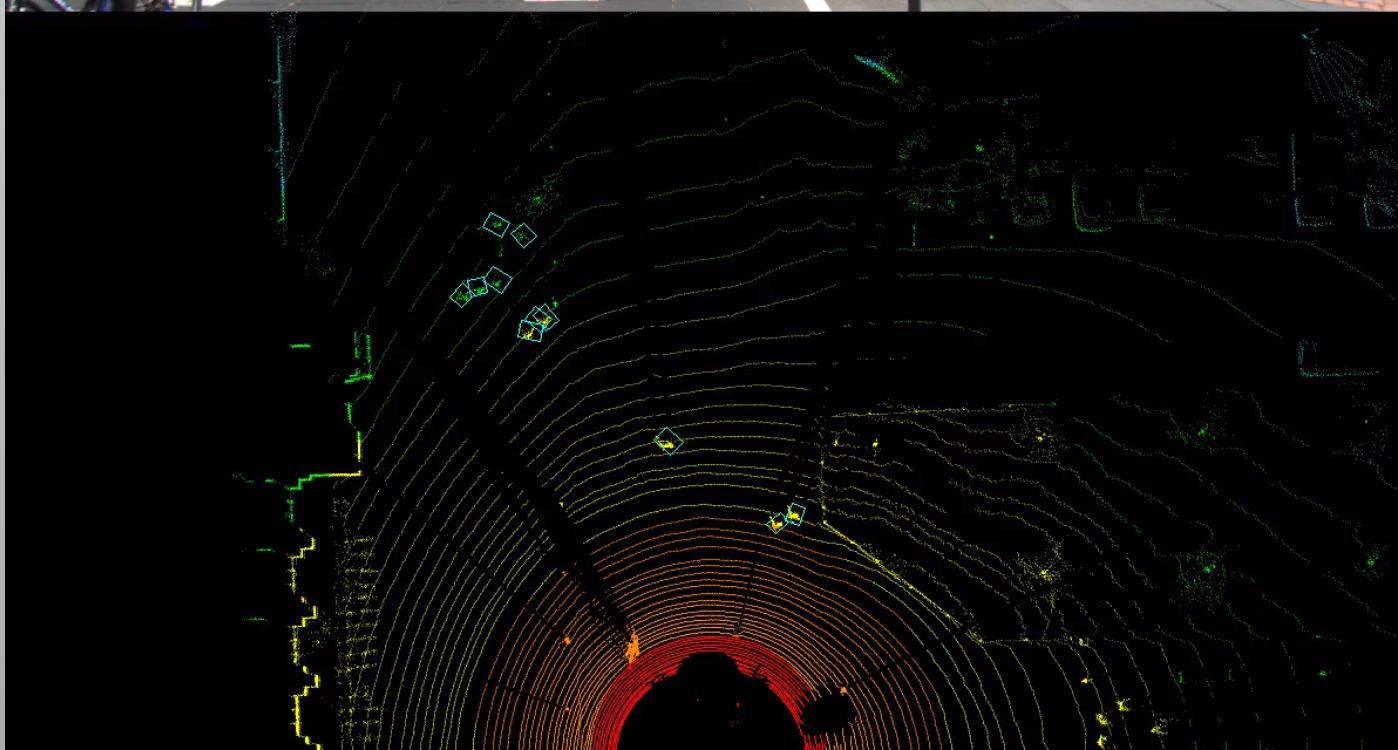
Ramy Battrawy, René Schuster, Oliver Wasenmüller, Qing Rao, and Didier Stricker, "LiDAR-Flow: Dense Scene Flow Estimation from Sparse LiDAR and Stereo Images," in IEEE International Conference on Intelligent Robots and Systems (IROS), 2019

## Fusion: Kamera-LiDAR

	SF Outliers [%]		
	Seeds	Window	Dense
SFF++ [4]	16.05	13.89	14.51
<b>LiDAR-Flow</b>	<b>7.88</b>	<b>8.70</b>	<b>11.52</b>
Average no. of evaluated points	88	6820	91875



## Problem: 3D Fußgänger Detektion



## Ergebnis: Regen



## Ergebnis: Nachts



# Odometrie: Lokalisierung mit LIDAR (ausführliche Evaluierung, KITTI Daten)

— Ground Truth Path

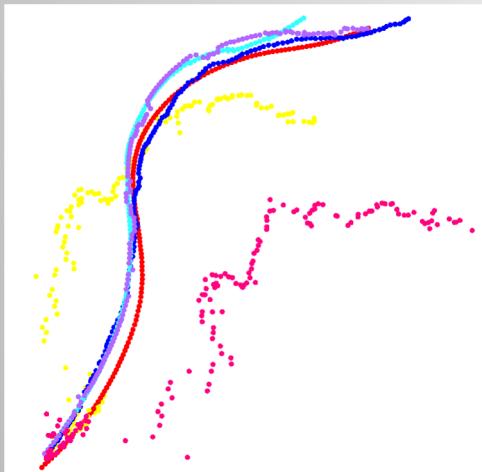
— FGA  
(ours)

— ICP

— GMMReg

— FilterReg

— CPD



Avg. runtime/registration (CPU) - 75.8 sec.

Avg. runtime/registration (GPU) - 0.81 sec.

Avg. iterations/registration - 36

#Points(M+N) per frame - 234.2k

#registration pairs: 153

Avg. runtime/registration (CPU) - 70.3 sec

Avg. runtime/registration (GPU) - 0.79 sec.

Avg. iterations/registration - 32

#Points(M+N) per frame - 232.1k

#registration pairs - 446

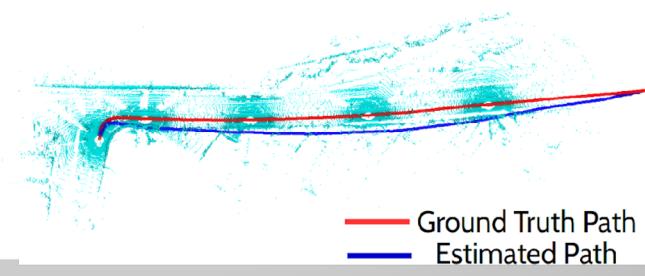
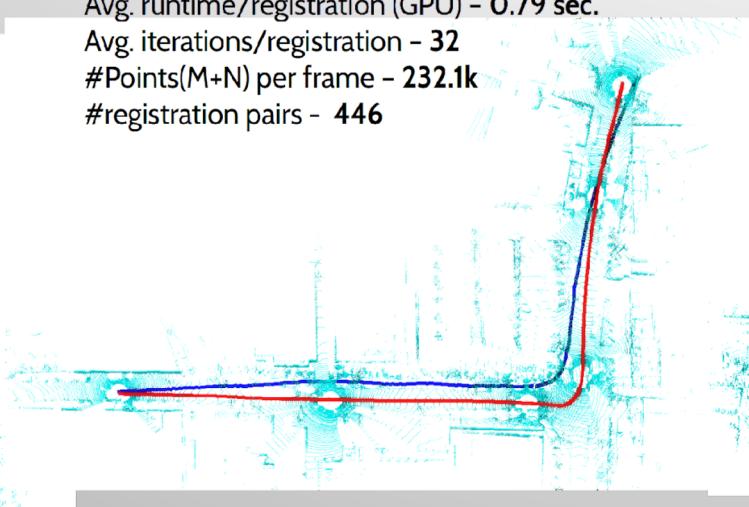
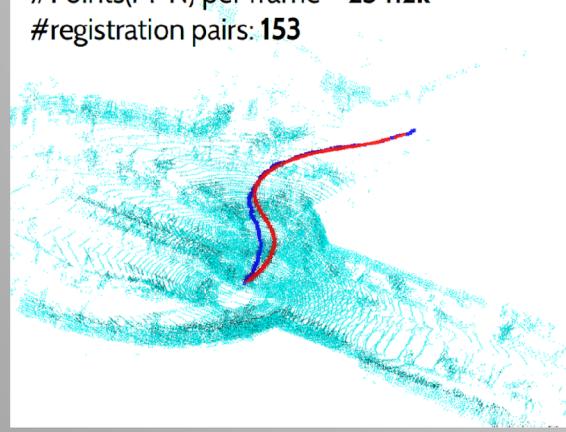
Avg. runtime/registration (CPU) - 67.28 sec

Avg. runtime/registration (GPU) - 0.65 sec.

Avg. iterations/registration - 34

#Points(M+N) per frame: 224.4k

#registration pairs: 313



(a). 2011\_09\_26\_drive\_0005 Sequence

(b). 2011\_09\_26\_drive\_0009 Sequence

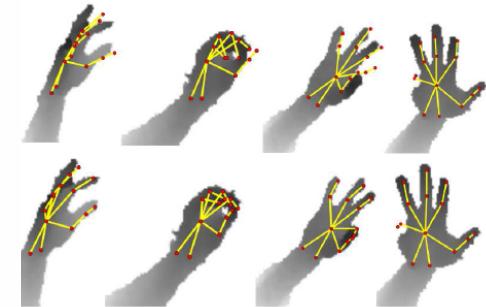
(c). 2011\_09\_26\_drive\_0014 Sequence

— Ground Truth Path  
— Estimated Path

## Human Pose Estimation



# Eingebettete KI: FPGA



Full-Precision Model



Quantized Model



Design HW  
Architecture



Implement on FPGA

	Image Time [ms]	Effective Power [W]	Energy per Image [mJ]
FPGA	1.669	0.4	0.6676
GPU	7.01	54.98	385.41

- **4.2 Mal schneller**
- **Faktor 577 bei Energieeffizienz!**
- **Genauigkeit = ca. -2%**

# DeepGUI : DragnDrop Deep Learning

The screenshot shows the DeepGUI interface. On the left, there is a vertical palette of components: Conv2D (blue), Conv3D (red), LSTM (purple), Pool2D (yellow-green), Pad2D (dark blue), Linear (black), ReLU (pink), Flatten (orange), split (cyan), + (light blue), and End (red). In the center, a neural network graph is displayed as a sequence of nodes: Start → Pad2D → Conv2D → ReLU → Pool2D → Conv2D → ReLU → Pool2D → Flatten → Linear → ReLU → Linear → End. Below the graph, there is a text panel showing configuration for a Conv2D layer:

```
text
Conv2D
in_channels 1
out_channels 8
in_height 128
in_width 128
kernel_height 5
kernel_width 5
```

At the bottom, there are two buttons: "Generate PyTorch Code" and "Generate Vivado HLS Code".

The screenshot shows the generated code side-by-side. On the left is the PyTorch code, and on the right is the Vivado HLS code. Both codes are identical, representing the same neural network structure.

```
PyTorch (GPU)
HLS C++ (FPGA)
```

```
Generate PyTorch Code | Generate Vivado HLS Code
Save Load
```

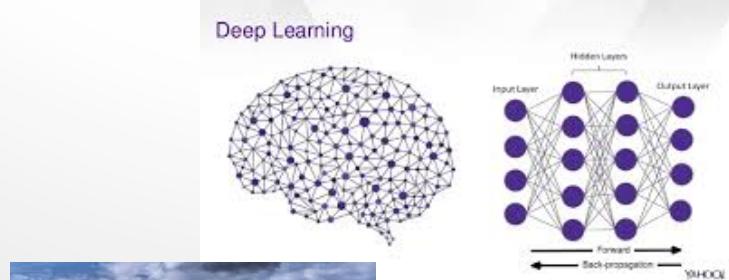
```
PyTorch HLS dataflow
import torch
import torch.nn as nn
class GENERATED_MODEL(nn.Module):
    def __init__(self):
        super(GENERATED_MODEL, self).__init__()
        pad2d_1 = nn.ZeroPad2d((2, 2, 2, 2))
        conv2d_1 = nn.Conv2d(1, 8, (5, 5))
        relu_1 = nn.ReLU()
        maxpool2d_1 = nn.MaxPool2d((2, 2))
        conv2d_2 = nn.Conv2d(8, 8, (5, 5))
        relu_2 = nn.ReLU()
        maxpool2d_2 = nn.MaxPool2d((2, 2))
        linear_1 = nn.Linear(6728, 1024)
        relu_3 = nn.ReLU()
        linear_2 = nn.Linear(1024, 93)
    def forward(self, x):
        out = self.pad2d_1(x)
        out = self.conv2d_1(out)
        out = self.relu_1(out)
        out = self.maxpool2d_1(out)
        out = self.conv2d_2(out)
        out = self.relu_2(out)
        out = self.maxpool2d_2(out)
        out = out.view(out.size(0), -1)
        out = self.linear_1(out)
        out = self.relu_3(out)
        out = self.linear_2(out)

static M5 stream<float> input_dpy> INPUT STREAM("input stream");
static M5 stream<float> zeroPad2d_1 dpy> pad2d_1 STREAM("zeroPad2d_1 stream");
static M5 stream<float> conv2d_1 dpy> conv2d_1 STREAM("conv2d_1 stream");
static M5 stream<float> maxpool2d_1 dpy> maxpool2d_1 STREAM("maxpool2d_1 stream");
static M5 stream<float> conv2d_2 dpy> conv2d_2 STREAM("conv2d_2 stream");
static M5 stream<float> maxpool2d_2 dpy> maxpool2d_2 STREAM("maxpool2d_2 stream");
static M5 stream<float> linear_1 dpy> linear_1 STREAM("linear_1 stream");
static M5 stream<float> maxpool2d_3 dpy> maxpool2d_3 STREAM("maxpool2d_3 stream");
static M5 stream<float> linear_2 dpy> linear_2 STREAM("linear_2 stream");

#PyTorch HLS dataflow
static M5 stream depth=1 variable IMP1 STREAM;
#PyTorch HLS stream depth=1 variable pad2d_1 STREAM
#PyTorch HLS stream depth=1 variable conv2d_1 STREAM
#PyTorch HLS stream depth=1 variable conv2d_2 STREAM
#PyTorch HLS stream depth=1 variable maxpool2d_1 STREAM
#PyTorch HLS stream depth=1 variable conv2d_3 STREAM
#PyTorch HLS stream depth=1 variable maxpool2d_2 STREAM
#PyTorch HLS stream depth=1 variable linear_1 STREAM
#PyTorch HLS stream depth=1 variable maxpool2d_3 STREAM
#PyTorch HLS stream depth=1 variable linear_2 STREAM
#PyTorch HLS stream depth=1 variable linear_3 STREAM
#PyTorch HLS stream depth=1 variable linear_4 STREAM
```

## Ausblick

- Fusion von **Modalitäten** und **Methoden**
- Eingebettete Lösung oder „**Edge AI**“ wird eine zentrale Rolle spielen
- „**Continual Learning**“ - Lernen mit wenigen Beispielen und mit reduzierter Hardware
- Von PKW zu Arbeitsmaschinen: **Autonomes Fahren und autonomes Arbeiten (Rauhe Umgebungen)**



# Thank you for your attention!



DFKI GmbH  
Department Augmented Vision  
Trippstadterstr. 122  
D-67663 Kaiserslautern



Didier Stricker



Didier.stricker@dfki.de



<http://av.dfki.de/>