


Overview


- Where is a person **heading to** and **will she cross** the road?
- Multi-object detection & tracking
Compressed SSD/AlexNet with AVX2
Geometric cues for bipartite matching
- Intent prediction
Combination of **motion extrapolation** and statistical **global motion model**
- Long-term proof-of-concept
1 month on a busy road with 3,000 pedestrians/day
- Paved the road for **traffic flow optimizations**





Evaluation

- Annotated over **2,000 images** (60/10/30 split) for detector evaluation
- Validated **80,000 trajectories** for intent prediction
- Manually annotated periodic snapshots (**every 30 minutes**)
- High prediction accuracy:

Motion Model	Acc _{0.2}	Acc _{0.5}	Acc _{0.65}	Acc _{0.8}
EX Model	0.645	0.488	0.232	0.036
GM Model	0.977	0.969	0.924	0.492
Combined Model	0.978	0.970	0.932	0.712



- Comparison to traditional push-button system:



- Quality measure & manual verification:

Motion Model	mLPWD [%]	mLSWD [sec]
EX Model	0.851	14.05
GM Model	0.209	2.06
Combined Model	0.185	1.93

	#	Recall ↑	Precision ↑
Traffic Light 1	1,578	0.992	0.972
Traffic Light 2	1,569	0.995	0.962
Both	3,147	0.994	0.967