Research Paper Reading Group

Pilot Session 3

Previous sessions

Session 1

Types

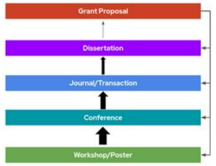

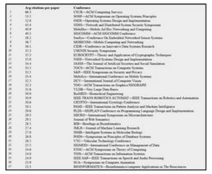

Tiers

Format

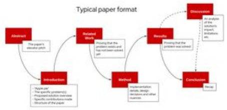

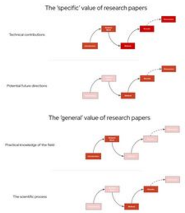

Value

Session 2

Common indicators

Citations	Higher citation count typically indicates higher quality work.
Year	Latest works tend to advance/supercede earlier efforts and/or leverage newer/state-of-the-art/defacto technologies.
Authors/Organization	Certain authors/organizations/groups in a field can have a reputation for actively producing high quality work.
Venue	Certain venues can be much more competitive/prestigious than others, and thus only accept state-of-the-art research.
Abstract	The abstract helps estimate the relevance of the presented work.
Specific contributions	Overall quality of contributions made - quantity is less important since a large contribution can be broken down into smaller ones.
“Beyond the scope”	In order to validate their contribution, some papers might be making impractical assumptions about the state of the technology.
Introduction length	If less work has been done, it typically takes longer to frame a specific enough problem for which the contributions are significant.
Scientific process	There must be some scientific process of reasonable complexity involved - basic “hit & trial” based approaches tend to be fragile.
System diagrams	Most good papers have one or more block diagrams / flow charts that illustrate their work and its different pieces.
Total figures	More illustrations can correspond to more work done (unless it is a survey paper).
Experiments done	Enough experiments should be done to demonstrate that the proposed contribution has been made.
Test environment	There can be a massive difference between testing something in theory, in simulation and in implementation on real hardware.
Formatting	Good papers typically go through many rounds of internal review before final submission and thus tend to have better formatting.
Open source	If the code is open source and reproducible, the credibility of the work increases substantially

Pilot overview

Session 1	4/3/2021	A perspective on research papers
Session 2	5/4/2021	Identifying worthwhile papers
Session 3	6/13/2021	Discussing research papers

Topic: Deep learning for self-driving/autonomous cars

Part 0	5 min	Papers under discussion
Part 1	5 min	Challenges highlighted in existing literature
Part 2	10 min	Proposed solutions
Part 3	5 min	Effectiveness/limitations of the proposed solutions

Papers

ID	Name	Link
1	Detecting Unexpected Obstacles for Self-Driving Cars: Fusing Deep Learning and Geometric Modeling	https://arxiv.org/pdf/1612.06573.pdf
2	Dynamic Occupancy Grid Prediction for Urban Autonomous Driving: A Deep Learning Approach with Fully Automatic Labeling	https://arxiv.org/pdf/1705.08781.pdf
3	DeepPicar: A Low-cost Deep Neural Network-based Autonomous Car	https://arxiv.org/pdf/1712.08644.pdf
4		
5		
6		
7		
8		
9		
10		

Challenges highlighted in existing literature

ID	Paper	Challenge
a	1	Detection of small sized and unexpected road hazards is important. Use of stereo cameras and geometric cues is insufficient.
b	2	Long Term citation prediction in autonomous driving is a major challenge since the strategic and forward-looking behavior it enables is hard to compensate with the fast reaction and precise precision of machines.
c	3	Real car-based testbeds are not only highly expensive, but also poses serious safety concerns that hinder development and exploration.
d		
e		
f		
g		
h		
i		
j		

Proposed solutions

ID	Paper	Chal.	Solution
i	1	a	- Combine stereo imaging (which leverages geometric cues) with CNNs (which leverages contextual and appearance cues). - Figure 2 in the paper gives an overview of the pipeline. - Uses stixels for object representation - Stixels approximately represent an object as a bunch of vertical bars of varying height - reduced the amount of data needed to represent the object. - 50% performance gain, reducing false positives by 13%, detection rate of over 90% for distances up to 50m for the Lost and Found dataset.
ii	2	b	- Use supervised learning as a long term predictor - but using unlabeled data (which is what you'd use in unsupervised learning) - Input to the model is a segmented Dynamic Occupancy Grid Map (or DOGMa) generated by a bayesian filtering technique.. - Output of the model is T channels, where 1 channel is the static occupancy probability, and the remaining T-1 channels are the dynamic occupancy probabilities at each time step. Segmentation is needed since 99.75% of environment data is static, so dynamic objects may be ignored during training. Cannot use simple velocity data since misclassification can occur e.g. between growing static parts and actual moving objects. - Network can predict up to 3s for complex scenarios.
iii	3	c	- Hardware used : Raspberry Pi 3, webcam (200x66 RGB pixel image), motor driver and a 1:24 scale RC car. Car shown on Figure 2. - Training is done by a human driver controlling the RC car. The webcam video is recorded and timestamped. Control commands issued by the driver are also recorded. These are used to train the model offline on a GPU. Track is shown in Figure 4. - Algorithm: read frame -> pre process -> inference -> steering motor control -> wait - Only three discrete steering angles supported even through the network can generate continuous angles. - Single control loop takes upto 23.38ms (40 Hz control frequency). By comparison, DAVE-2 could only do 30.
iv	2	a	- Leveraging multiple types of sensors and using sensor fusion instead of probability graph fusion.

Effectiveness/limitations of the proposed solutions

ID	Paper	Chal.	Sol.	Effectiveness / Limitation
α				
β				
γ				
Δ				
ϵ				
Z				
H				
θ				

Sign-up / Comments / Suggestions / Feedback

https://forms.gle/6Y2ZBH2Bq2y5Qmie7

Thank you!