



Redefining the Driver's Attention Gauge in Semi-Autonomous Vehicles

Raja Hasnain Anwar
University of Massachusetts Amherst

Fatima Muhammad Anwar
University of Massachusetts Amherst

Muhammad Kumail Haider
Meta Platforms

Alon Efrat
University of Arizona

Muhammad Taqi Raza
University of Massachusetts Amherst

ABSTRACT

Driver distraction caused by over-reliance on automotive technology is one of the leading causes of accidents in semi-autonomous vehicles. Existing driver's attention-gauging approaches are intrusive and as such emphasize constant driver engagement. In case of an urgent traffic event, they fail to measure the event's criticality and subsequently generate timely alerts. In this paper, we re-position the driver's attention-gauging approach as a way to improve the driver's situational awareness during critical situations. We exploit how a vehicle captures its surroundings information to convert an automotive decision into defining the criticality and timeliness of an alert. For this, we identify the relationship between the traffic event, the type of automotive sensing technologies, and its processing resources to capture that event to design the driver's attention gauge. We evaluate the timeliness of alerts for different traffic scenarios over a prototype built using NVIDIA Jetson Xavier AGX and Carla. Our results show that we can improve the timeliness of an alert by up to 75x as compared to existing state-of-the-art approaches, while also providing feedback on its criticality.

CCS CONCEPTS

• Human-centered computing; • Applied computing;

KEYWORDS

semi-autonomous vehicles, human-computer interaction.

ACM Reference Format:

Raja Hasnain Anwar, Fatima Muhammad Anwar, Muhammad Kumail Haider, Alon Efrat, and Muhammad Taqi Raza. 2023. Redefining the Driver's Attention Gauge in Semi-Autonomous Vehicles. In *Proceedings of the Int'l ACM Conference on Modeling Analysis and Simulation of Wireless and Mobile Systems (MSWiM '23)*, October 30-November 3, 2023, Montreal, QC, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3616388.3617544>

1 INTRODUCTION

Semi-autonomous vehicles are a step forward toward self-driving cars. They can navigate on the road, including highways, streets,

Table 1: Inattentive driving remains one of the major causes of both fatal and non-fatal crashes [9]

Causes	% Fatal Crashes	% Non-Fatal Crashes
Intoxication	37	6.9
Over speeding	23.1	17.5
Inattention	14.4	22.1
Inexperience	2	35.8
Others	23.5	17.7

and school zones, on their own without the driver's direct input. Their lack of full autonomy requires that the driver remains attentive at all times and takes control of the vehicle if it fails to handle an unexpected traffic situation. However, most drivers do not realize that the automation is not correctly responding to challenging situations. Eventually, the vehicle suddenly hands its control over to the driver who is not prepared to take action. A recent study [8] reports that the driver usually takes up to 1.5 seconds longer to apply the brakes when driving in a semi-autonomous vehicle compared to manually operating the vehicle. Table 1 shows that the lack of driver's attention and the resultant slow response is one of the highest root causes of accidents only behind intoxicated driving and overspeeding.

There are several factors that influence the driver's response; awareness of the surroundings, concentration on driving, time available to take action, and road conditions. Previous studies suggest various approaches (e.g., facial recognition, gaze and head mapping, and wearable devices) to alert the driver when found to be distracted from the road. These approaches are ineffective as they do not provide scenario-specific feedback and enough time for the driver to react. We argue that situational awareness is more important than keeping the driver always attentive. The driver should be alerted only when human input is required. The level of attentiveness also varies from one situation to the other. A scenario involving a jaywalker is more critical than one involving a bumpy road. In this paper, we develop a scenario-specific driver's attention gauge that provides *objective* and *timely* feedback to the driver. Our key intuition is to attain the driver's attention (i) prior to the occurrence of a critical traffic event, and (ii) corresponding to the required level of attention for that event.

Developing such an alert system is challenging because automotive systems installed in vehicles do not expose scenario-specific information to the driver. We exploit the vehicular mechanisms by which the traffic situations are evaluated and handled by the Electronic Control Unit (ECU), the main controller of the vehicle.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MSWiM '23, October 30-November 3, 2023, Montreal, QC, Canada
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0366-9/23/10...\$15.00
<https://doi.org/10.1145/3616388.3617544>




Figure 1: Integration of our proposed driver attention-gauging approach. The driver gets real-time feedback on the required level of alertness through a GUI.

The ECU may simply be using an RGB camera to detect and track objects on the road. However, if the traffic scenario is abnormal and critical (e.g., pedestrian jaywalking or a road bump), the ECU activates other sensors (Lidar, Radar, and/or Depth camera) to extract more fine-grained features and patterns of the moving objects. The involvement of additional sensors means the ECU systems usage ramps up as it engages additional environmental information and algorithms to execute more complex safety maneuvers. Our idea is to use these usage patterns as a pre-cursor to classify a critical traffic scenario and engage the required level of attention from the driver. If we can estimate the likelihood of a traffic scenario before it is evaluated by the ECU, we can alert the driver promptly; hence providing extra valuable time for situational awareness.

Another key challenge is identifying *criticality* of the alert to provide explicit feedback to the driver. Our driver's attention gauging metric is based on the observation that the more urgent or critical the scenario is, the more sensors and algorithms are involved to evaluate and handle it. For instance, in order to avoid collision with a jaywalker, the vehicle needs to scan the whole environment, track the movement of objects (i.e., other vehicles and the jaywalker), and decide a course of action, e.g., either stop or change the lane (see Figure 1). Since the processing of every sensor is unique, it generates a unique ECU system usage pattern.

We implement our proposed approach on an ECU, i.e., NVIDIA Jetson Xavier AGX, and Carla traffic simulator [7]. Our method performs better than the conventional methods while providing up to 75x faster feedback to the driver. Our approach is also easier to implement as it does not require any external sensory devices.

2 EVENT-BASED ATTENTION GAUGING APPROACH

We propose that the driver's attention should be a function of critical external events experienced by semi-autonomous vehicles. The criticality of an event should determine the required level of the driver's attention. We observed that such events are captured by a range of sensors on a vehicle. The kind and number of sensors used to process an event depend on the nature of the event. For instance, pedestrian jaywalking requires a greater number of sensors operating at their maximum capacity as compared to the road bump event where only a few sensors operating at medium capacity are

needed to detect and react. These sensing operations are processed at a central Electronic Control Unit (ECU) in a vehicle. Our key observation is that ECU resource utilization varies with the kind and amount of sensor data processing, which in turn is dependent on the kind of event. For example, high CPU, memory, and cache utilization of ECU indicate increased processing and hence the presence of a critical event. Our work sets out to establish three important correlations: 1) ECU Utilization is dependent on sensing technologies, 2) the kind of sensing technologies used to indicate the criticality of an event, and 3) an event's criticality decides the level of user attention, i.e., how early and how urgently the driver should be made attentive.

2.1 Key Insights

Our approach to gauge the required level of driver's attention is based on two key insights.

2.1.1 Varying Criticality of Traffic Scenarios. The unpredictability of critical traffic scenarios makes semi-autonomous vehicles susceptible to road accidents – requiring intervention from the driver in challenging situations. The nature of these traffic scenarios and the required response make them unique. For example, when driving in poor weather, the semi-autonomous vehicle's priority is to remain safely within lanes; whereas, with the sudden appearance of a jaywalker, the vehicle applies the brakes and/or changes its course safely. These are critical scenarios that mandate that the driver should be attentive in a *timely* manner to take over. However, two critical events cannot require the same amount of time and course of safety action. Therefore, it is important to alert the driver according to the *criticality* of the event. We observe that if the vehicle's actions are mapped to a particular combination of sensor devices, we can predict the arrival of the key traffic event. This will lead the vehicle to engage the driver earlier and make them attentive before the occurrence of the traffic scenario.

2.1.2 Traffic Scenarios Are Processed Differently. Typically, a semi-autonomous vehicle uses combinations of four types of sensors: Radar, Lidar, RGB, and Depth, to make driving decisions. Table 2 lists critical traffic scenarios and the corresponding sensors used. These sensors operate at varying capacities and consume different ECU system resources, i.e., CPU, memory, network, and disk. The difference in resource consumption helps us identify traffic events at run-time. As shown in Figure 2(a), the CPU load in pedestrian detection is higher compared to detecting the lane departure scenario. This is mainly because pedestrian tracking activates all types of sensors as this is a more complex and critical scenario. Adjusting to unexpected lane departure does not require similar maneuvering.

We use these insights to propose a driver attention-gauging technique. Our key intuition is to map the systems logs to the traffic scenarios. Such a reverse engineering approach provides pre-event feedback to the driver, which is faster than the existing post-event approach that waits for the event to be detected and evaluated before alerting the driver.

Our empirical analysis, confirms that different traffic scenarios indeed exhibit different CPU usage patterns. Figure 2 shows (a) the time series of the change in CPU load as different events happen during the experiment drive and (b) the boxplot of CPU load spread

Table 2: Sensing technologies are used according to the traffic scenario [9]

Traffic Scenarios	Vehicle Actions	Sensors
Jaywalking	Emergency braking; scan and track objects for safe maneuvering	RGB, Depth, Radar, Lidar
Reckless driving	Blind spot monitoring; lane keeping; safe distancing	RGB, Radar
Sudden lane change	Scan environment for previously occluded objects; object tracking	RGB, Depth, Radar
Sudden stop	Emergency braking; scan environment for safe maneuvering	RGB, Lidar
Poor weather	Use alternate sensors as primary because RGB may fail	Lidar, Radar
Road bump	Scan for inconsistencies in road and surroundings	RGB, Depth

corresponding to each scenario independently. The time-series results manifest the CPU load spikes as soon as the ECU starts evaluating the critical traffic scenario; whereas the boxplot results show the min, max, and median, i.e., the spread of each spike in the CPU load for complete processing of each scenario.



**(a) Joint experiment run****(b) Separate experiment runs**

Figure 2: Variation in CPU load: (a) shows the resource usage signature of six major critical traffic events occurring in sequential order. (b) shows resource usage behavior of six critical traffic events averaged over multiple simulations.

2.2 Modeling Traffic Events to Gauge Attention

We put forward a Markov-Poisson model that illustrates how to characterize anomalous events (high resource usage) from normal scenarios. The ECU generates time-series data where time is discrete and $X(t)$ is a measurement of the event over the time interval $(t - 1, t)$. Our theoretical model consists of a time-varying Poisson process using ECU systems usage $X(t)$ (i.e., CPU, memory, network I/O, and disk I/O) [1]. It is a Markovian process where the conditional probability of a future traffic scenario is only dependent on the present state. For example, when a pedestrian suddenly starts jaywalking, it will result in the triggering of emergency braking and the required action from the driver to strategically change the course of the vehicle. In this example, the action of the driver is only dependent on the present traffic state, i.e., the pedestrian stepping on the road. We assume that the traffic scenarios occur randomly in time and obey the following conditions:

- The probability of occurrence of the event in a given time frame is independent of the past events.
- The probability of two or more occurrences of different events in a very small time interval is negligible.
- The amount of usage of systems resources in a particular time interval manifests the urgency of the event.

Now the probability distribution of the systems resources of the given event in a fixed time interval is Poisson distributed with mean $\mu = \lambda t$, where λ is the rate of the system usage per unit of time, and t is the length of the time interval. We frame our problem by using the Poisson distribution parameters and the Maximum Likelihood Estimation method to estimate the likelihood of an urgent event from resource usage. In this way, our approach takes the current system usage values as a precursor of the upcoming traffic event, and timely alerts the driver corresponding to the required level of alertness.

3 IMPLEMENTATION AND EVALUATION

3.1 Systems Implementation

Our testbed consists of an ECU equipped with NVIDIA Jetson Xavier AGX to execute semi-autonomous vehicle control models. The functionalities of RGB, Depth, Lidar, and Radar are implemented using the models: NVIDIA Autopilot [2], NVIDIA Autopilot [2], Cylinder3D [24], RODNet [23], respectively. The RGB and Depth models are trained to control the car acceleration, steering, and braking while Lidar and Radar models are used to detect and track objects in the surroundings. The Jetson module is connected to a

general-purpose computer (with a 2.8 GHz processor and 32 GB memory) over a WiFi connection.

We emulate the driving of a semi-autonomous vehicle in Carla Simulator [7] version 0.9.11 installed on a general-purpose computer. We use the autopilot system in the Carla simulator to systematically collect the vehicular data. This feature employs a control loop on vehicular sensors' feedback imitating the human driver's/semi-autonomous vehicular actions: such as stopping at a stop sign/red light, lane change, and maintaining a specific distance from the front vehicle. All the map, vehicle, and sensor data is sent over WiFi to the Jetson for processing. While the regular traffic scenario was running, we launched specific traffic scenarios (as shown in Table 2) one by one and measured the corresponding Jetson systems logs.

3.2 Evaluation

We evaluate our proposed approach on how timely the alert decision is made, and what level of alertness is conveyed to the driver. We evaluate the critical traffic scenarios as described in Table 2.

How quickly our approach can trigger an alert to the driver?

We compare the timeliness of the alert in our approach with the baseline. The baseline approach is the output of state-of-the-art models currently implemented in semi-autonomous vehicles. It observes the traffic scenario, evaluates it for being critical, and then generates an alert to the driver. However, in our proposed approach, we can generate the alert as soon as the environmental conditions reflect the criticality of the traffic scenario, without waiting for the event to be evaluated.

Table 3: The timeliness of different traffic scenarios are compared: proposed approach vs state-of-the-art

Traffic Scenarios	Average Processing Time (ms)	
	State-of-the-art	Our approach
Jaywalking	1582.47	20.88
Reckless driving	720.27	17.43
Sudden lane change	1465.92	19.51
Sudden stop	64.16	16.93
Poor weather	749.56	18.5
Road bump	26.57	16.06

Table 3 shows the time to generate alerts for different critical events/traffic scenarios. The results show that our approach is up to 75x better when compared to the baseline. The more critical the traffic event is, the more timely alert is generated in our approach as compared to the state-of-the-art. There are several factors that influence the time in the contemporary approach. To evaluate the critical event, the ECU implements several deep learning AI models where each model consumes a certain amount of system resources to generate an output and make a decision for handling the situation. For example, the Lidar model processes the sensor feeds as a 3D model of the environment and runs object detection and tracking tasks. This results in more CPU processing and memory as compared to the RGB model which deals with only a small portion of the environment. Further, in order to execute every model, the ECU needs to perform pre-processing of sensor data, model

forward pass, and then land on a decision. Such processing chains take 100s of milliseconds that can be avoided using our approach. Our approach is based on the Markovian-Poisson process that calculates the probability of systems resources (i.e., the spike) at the next time instance from the current time. It only takes just over 16 milliseconds (i.e., the delay between two consecutive clock ticks) to evaluate such probability distribution on a given time instance.

How accurately our approach can detect the critical traffic situation? We test our proposed theoretical model (i.e., Markovian-Poisson model) in the wild by creating an arbitrary number of critical traffic scenarios on a Carla simulator map [7]. The order and frequency of the scenarios are randomly selected where one scenario can occur more often than the others. On every clock tick, we use the Markov model to select the most probable value of λ for the present state. We then use this λ to estimate the probability of occurrence of the corresponding critical scenario.

Figure 3 shows the comparison of the observed probability of occurrence for every scenario in an experiment and the average probability of the same scenario estimated by our theoretical model using the system resource data of the same experiment. The theoretical probability of an event is an average of all estimated probabilities. The high similarity between observed and estimated probabilities indicates that (a) our model correctly selects the appropriate λ using the present state as prior, and (b) it estimates the likelihood of occurrence of a critical traffic scenario with high accuracy and confidence. Once we have the likelihood of a critical scenario, we use a linear function to calculate the required level of alertness to engage the driver in handling the traffic scenario and display it to the driver (Figure 1).




Figure 3: Detecting the future occurrence of a particular traffic scenario: comparing empirically observed and theoretically calculated. Our proposed probabilistic model accurately estimates the likelihood of each traffic scenario by using the system state from only the prior clock tick.

4 RELATED WORK

Our work is in contrast with recent efforts on drivers' distraction causes, driver performance measurements, vehicular safety, and driver assisting techniques.

Numerous studies have investigated the causes of driver distraction. [11] uses Thematic and Correspondence Analysis to understand what information is needed to engage the driver. It declares communication errors as the major cause of drivers' ineffective decisions. Similarly, [16] blame poor visibility and safety-related behaviors of drivers to be the major causes of accidents. This paper indeed strengthens these claims and states that we cannot assume that drivers always remain attentive.

Various studies show how challenging it is to ensure constant driver engagement while driving. [4, 6, 14] show that the drivers are not ready to take over under unexpected scenarios leading to accidents and call for traffic scenario prediction-based approaches.

While other works propose the use of technologies to address drivers' distraction issues, [3, 13, 19] propose the use of virtual reality (VR) to generate gaze-aware warning cues to drivers. [15, 18, 22] uses visual, auditory, and 3D environmental cues to predict the driver's attention, respectively. They make the driver attentive if the recorded observations deviate from a pre-defined threshold. [5, 10, 20] blame the use of mobile devices to be the major cause of drivers' distractions. These works call for analyzing mobile data during mobility to engage the drivers. In contrast to the above works, this paper does not require analysis of any external devices, rather it uses the ECU processing information to decide the level of attention required from the driver. Instead of taking a binary decision of whether to make the driver attentive or not, this paper proposes a traffic scenario-aware driver's attention gauging approach.

Some other works take more cautious approaches and call for the limited role of semi-autonomous vehicles. [12, 21] call for manual switches between the semi-autonomous vehicles and the driver. [17] call for suggestion and rule-based approach under safety-critical situations. However, these works do not address the major issues of a challenging traffic scenario when the vehicle is in autonomous mode, and the driver is required to intervene.

5 CONCLUSION

In this paper, we disclose that existing driver's attention-gauging approaches are ineffective and unsuitable in the wild. We propose a situational awareness-based driver's attention approach. We leverage the vehicular ECU systems data in designing timely and criticality-aware driver's attention gauge. This approach relaxes the requirement of being constantly attentive while driving, and only engages the driver when a safety-critical driving situation arises.

ACKNOWLEDGEMENT

This work is partially funded by the NSF award # 2051621.

REFERENCES

- [1] J Bernardo, M Bayarri, J Berger, A Dawid, D Heckerman, AFM Smith, and M West. 2003. The markov modulated poisson process and markov poisson cascade with applications to web traffic modeling. *Bayesian Statistics* (2003).
- [2] Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao, and Karol Zieba. 2016. End to End Learning for Self-Driving Cars. arXiv:1604.07316 [cs.CV]
- [3] Efe Bozkir, David Geisler, and Enkelejda Kasneci. 2019. Assessment of driver attention during a safety critical situation in VR to generate VR-based training. In *ACM Symposium on Applied Perception* 2019. 1–5.
- [4] Christian Braunagel, Wolfgang Rosenstiel, and Enkelejda Kasneci. 2017. Ready for take-over? A new driver assistance system for an automated classification of driver take-over readiness. *IEEE Intelligent Transportation Systems Magazine* 9, 4 (2017), 10–22.
- [5] Mitchell L Cunningham and Michael A Regan. 2018. Driver distraction and inattention. In *Safe mobility: Challenges, methodology and solutions*. Emerald Publishing Limited.
- [6] Nachiket Deo and Mohan M Trivedi. 2019. Looking at the driver/rider in autonomous vehicles to predict take-over readiness. *IEEE Transactions on Intelligent Vehicles* 5, 1 (2019), 41–52.
- [7] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. CARLA: An Open Urban Driving Simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*. 1–16.
- [8] Alexander Eriksson and Neville A Stanton. 2017. Takeover time in highly automated vehicles: noncritical transitions to and from manual control. *Human factors* 59, 4 (2017), 689–705.
- [9] Center for Transportation Research: The University of Texas at Austin. 2017. An Assessment of Autonomous Vehicles: Traffic Impacts and Infrastructure Needs—Final Report. Retrieved September 4, 2021 from <https://library.ctr.utexas.edu/ctr-publications/0-6847-1.pdf>
- [10] Marko Gjorgievski, Sheila Sprague, Harman Chaudhry, Lydia Ginsberg, Alick Wang, Mohit Bhandari, and Bill Risteovski. 2020. Distracted driving on YouTube: Categorical and quantitative analyses of messages portrayed. *JMIR public health and surveillance* 6, 1 (2020), e14995.
- [11] Lea Theresa Gröber, Matthias Fassl, Abhilash Gupta, and Katharina Krombholz. 2021. Investigating Car Drivers' Information Demand after Safety and Security Critical Incidents. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [12] Sylwia I Kaduk, Aaron PJ Roberts, and Neville A Stanton. 2021. The manual shift in phase: the impact of circadian phase on semi-autonomous driving. What can we learn from current understanding in manual driving? *Theoretical Issues in Ergonomics Science* 22, 1 (2021), 103–123.
- [13] Hyungil Kim, Xuefang Wu, Joseph L Gabbard, and Nicholas F Polys. 2013. Exploring head-up augmented reality interfaces for crash warning systems. In *Proceedings of the 5th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. 224–227.
- [14] Alexander Lotz and Sarah Weissenberger. 2018. Predicting take-over times of truck drivers in conditional autonomous driving. In *International Conference on Applied Human Factors and Ergonomics*. Springer, 329–338.
- [15] João Henrique Catarino Cardoso Loureiro. 2019. Testing driver attention in virtual environments through audio cues. (2019).
- [16] Peter Mikoski, Gian Zlupko, and D Alfred Owens. 2019. Drivers' assessments of the risks of distraction, poor visibility at night, and safety-related behaviors of themselves and other drivers. *Transportation research part F: traffic psychology and behaviour* 62 (2019), 416–434.
- [17] Richard Moon. 2020. *A Suggestion-based Advanced Driving Assistance System for Semi-autonomous Vehicles in Safety Critical Situations*. Ph.D. Dissertation. University of Colorado at Boulder.
- [18] Alberto Morando, Trent Victor, and Marco Dozza. 2018. A reference model for driver attention in automation: Glance behavior changes during lateral and longitudinal assistance. *IEEE Transactions on Intelligent Transportation Systems* 20, 8 (2018), 2999–3009.
- [19] Minh Tien Phan, Indira Thouvenin, and Vincent Frémont. 2016. Enhancing the driver awareness of pedestrian using augmented reality cues. In *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 1298–1304.
- [20] Syeda Rizvi, Francesca Favaro, and Nazanin Nader. 2020. Driver's Situational Awareness and Impact of Phone Interface Modality in Conventional and Semi-autonomous Vehicles. In *International Conference on Intelligent Human Systems Integration*. Springer, 945–951.
- [21] Franco van Wyk, Anahita Khojandi, and Neda Masoud. 2020. Optimal switching policy between driving entities in semi-autonomous vehicles. *Transportation Research Part C: Emerging Technologies* 114 (2020), 517–531.
- [22] MinJuan Wang, Sus Lundgren Lyckvi, Chenhui Chen, Palle Dahlstedt, and Fang Chen. 2017. Using Advisory 3D Sound Cues to Improve Drivers' Performance and Situation Awareness. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 2814–2825.
- [23] Yizhou Wang, Zhongyu Jiang, Xiangyu Gao, Jenq-Neng Hwang, Guanbin Xing, and Hui Liu. 2021. Rodnet: Radar object detection using cross-modal supervision. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 504–513.
- [24] Xinge Zhu, Hui Zhou, Tai Wang, Fangzhou Hong, Yuxin Ma, Wei Li, Hongsheng Li, and Dahua Lin. 2021. Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9939–9948.