

# Vulnerable road user detection: state-of-the-art and open challenges

Patrick Mannion

**Abstract**—Correctly identifying vulnerable road users (VRUs), e.g. cyclists and pedestrians, remains one of the most challenging environment perception tasks for autonomous vehicles (AVs). This work surveys the current state-of-the-art in VRU detection, covering topics such as benchmarks and datasets, object detection techniques and relevant machine learning algorithms. The article concludes with a discussion of remaining open challenges and promising future research directions for this domain.

**Index Terms**—Autonomous vehicles, environment perception, object detection, object classification, vulnerable road users, pedestrian detection, cyclist detection.

## I. INTRODUCTION

RESEARCH and development work on autonomous vehicles (AVs) has accelerated in recent years, driven by technological advances in both hardware and software. However, a number of significant challenges must be addressed before widespread adoption of AVs will be possible. These include legal and regulatory issues, public perceptions of AV technologies, cost, accuracy and reliability of on-board sensing equipment, computational constraints and the limitations of current algorithms and systems for tasks such as environment perception, path planning and control.

An overview of the main tasks which must be performed by an autonomous driving system is presented in Fig. 1. Environment perception is arguably one of the most important fields for future AV development; accurate perception of the environment allows automated driver assistance systems (ADASs) to make informed decisions during functions such as adaptive cruise control, lane changing, parking and obstacle and collision avoidance. Components of the environment which must be detected by an ego vehicle include: traffic signals and signs, road markings, lane and junction topology and other road users including vehicles, cyclists and pedestrians.

Vulnerable road users (VRUs), such as pedestrians and especially cyclists, remain among the most challenging objects for AV perception systems to detect accurately [1]. These road users have little or protection during a collision when compared to vehicle occupants; this is reflected in the high proportion of VRUs among traffic fatality statistics. In recent years, traffic accidents have been the leading cause of death worldwide for people aged 18-29 [2]. Approximately 1.25 million people died on roads worldwide in 2013; almost half of these fatalities were VRUs [2]. A similar proportion of traffic fatalities during 2017 in the European Union were VRUs [3]. Therefore, improving VRU detection capabilities to human-level performance or above will be necessary over the coming

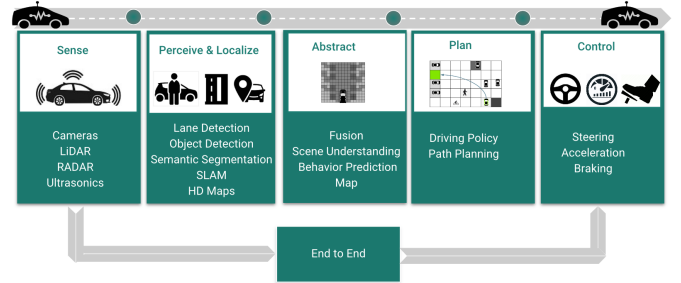


Fig. 1. Overview of autonomous driving tasks. Figure reproduced from [11].

years, if AVs are to minimise the risk that they pose to VRUs and be accepted by legislators and the public.

This article aims to provide an overview of the current state-of-the-art in VRU detection, along with a discussion of the main open challenges and promising directions for future research. Other related surveys on environment perception for AVs [4], [5], [6], pedestrian detection [7], [8] and perception benchmarks [9], [10] may also be of interest to the reader.

The next section of this paper provides a brief overview of the sensors which are used in AV perception, followed by Section III which discusses the main datasets which may be used to develop and test detection algorithms. Section IV introduces related work on generating proposals for object detection, while Section V presents an overview of the current best-performing algorithms for VRU detection. Section VI discusses the main open challenges and some promising future research directions for VRU detection. Finally, Section VII concludes this paper with some closing remarks.

## II. SENSORS FOR ENVIRONMENT PERCEPTION

Sensor types which are currently used for AV environment perception include passive sensors such as visible spectrum (VS) and infrared (IR) cameras and active sensors such as Radio Detection and Ranging (RADAR), Light Detection and Ranging (LIDAR) and ultrasonic.

Cameras provide a wide field of view, although they have limited depth perception capabilities when used individually. VS cameras are among the cheapest and most widely used sensors, however their performance degrades significantly in low-light conditions as they work by capturing the light intensity values at each pixel [6]. IR cameras capture heat signatures using relative temperatures, so they are suitable for low-light conditions. Cameras may be utilised in stereo arrangements with a lateral offset to improve depth perception capabilities [4].

Active sensors that emit signals and observe their reflections are better at providing range information than cameras, although they are more expensive. Accurate range information is important when determining the position in 3D space of a detected object in relation to the ego vehicle (i.e. when performing localisation). LIDAR may be used to capture highly accurate spatial data in the form of a 3D point cloud, but it is currently too expensive for widespread adoption on consumer vehicles.

Relying on one type of sensor alone may be problematic in certain driving conditions, therefore multiple types of sensors are typically used on AVs, combining the strengths and mitigating the weaknesses of each individual sensor type. The process of integrating information from multiple sensors is known as sensor fusion (see e.g. [12], [13], [14], [15], [16]).

### III. DATASETS AND BENCHMARKS FOR VRU DETECTION

This section gives an overview of recent datasets and benchmarks which are most relevant for VRU detection. The public availability of datasets is one of the main factors which enabled the improvements in AV perception seen in recent years. As well as removing the requirement for researchers to have access to physical AV hardware, many of these datasets have dedicated benchmarking websites where new algorithms may be independently evaluated using a standard methodology. This allows prior results to be easily replicated and compared using online leaderboards. Most of these datasets focus on detection using VS cameras; some also provide synchronised LIDAR and camera frames. The highly accurate LIDAR data serves as a ground truth source, allowing the comparison of image-based detection algorithms against 3D point cloud data.

The key features of the most important publicly available datasets are summarised in Table I. The datasets listed are EuroCity Persons (ECP) [10], CityPersons (CP) [17], Caltech [9], KITTI [18] and the Tsinghua-Daimler Cyclist benchmark (TDC) [19]. EuroCity Persons is currently the most comprehensive publicly available dataset, as it has data from all seasons in both dry and weather conditions, from 31 cities across 12 countries. As a result ECP has much greater diversity than the other datasets surveyed, including samples of a wide variety of clothing types, weather types and lighting conditions. ECP is therefore currently the best overall choice for researchers wishing to develop and test new algorithms for VRU detection; this could be supplemented by other datasets such as Caltech to obtain a larger training set.

TABLE I  
COMPARISON OF PUBLICLY AVAILABLE VRU DETECTION DATASETS

Dataset	ECP [10]	CP [17]	Caltech [9]	KITTI [18]	TDC [19]
# seasons	4	3	1	1	1
weather	dry & wet	dry	dry	dry	dry
lighting	day & night	day	day	day	day
# images	47337	5000	249884	14999	14674

### IV. OBJECT PROPOSAL GENERATION

Objects detection from images is an important task in fields such as robotics and surveillance, as well as in environment

perception for AVs. The object detection task consists of localisation (determining the position of an object in the image) and classification (determining which category an object belongs to) [20]. Traditional objection detection techniques try to identify regions of interest (ROIs) [21], i.e. areas of an image which may contain the types of objects which must be detected. ROIs may be generated using a sliding window approach, by shifting a detector over an image at different scales, although exhaustive searches of this manner can be computationally expensive. Algorithms such as selective search [22] may be used to generate ROI proposals with less computational overhead.




Fig. 2. Sample stixel-based bounding box proposals for a cyclist. Figure reproduced from [21].

As an alternative to rectangular ROIs, stixels [23] may also be used for generating object proposals from images. A sample stixel-based proposal is shown in Fig. 2. Stixels are a medium-level representation, intended to bridge the gap between individual pixels and actual objects. Individual stixels have a specific fixed width and variable height in pixels, are defined by their 3D position relative to the camera and stand vertically on the ground plane. Stixels are therefore suitable for generating proposals for objects which are mostly vertical (e.g. VRUs and vehicles).



Fig. 3. Sample LIDAR-based bounding box proposals using point clustering. Figure reproduced from [24].

Object proposals may also be generated from LIDAR point cloud data, as illustrated in Fig. 3. One commonly used

approach is to first remove all the points corresponding to the ground surface and then to group the remaining points using a nearest-neighbour clustering algorithm (e.g. a kd-tree search structure) [24]. Finally a 3D bounding box representing each LIDAR point cluster is projected onto the image plane to generate 2D object proposals.

## V. VRU DETECTION ALGORITHMS

Once good quality object proposals have been generated, a classification algorithm must be run on each of the proposals. Many of the most prominent classification methods use machine learning (ML), a process which allows a computer program to learn to improve its performance at a task with increased experience [25]. ML algorithms for AV perception are generally trained on a dataset with labelled examples like those in Section III. The goal for a designer is to develop an algorithm which can successfully generalise what it has learned on the training dataset to classify new, unseen examples (either proposals from a test set or proposals generated during a real-world deployment).

Methods such as deformable part-based models and decision forests were the main approaches used for VRU detection until recent years [7]. Artificial neural networks (ANNs) form the basis for many of the most successful detection algorithms which have been developed since. ANNs are loosely based on biological neural systems and are comprised of individual perceptrons which are interconnected. Deep learning (DL) [26] is a ML paradigm which emerged relatively recently, encompassing a broad range of techniques based on “deep” ANNs (i.e. ANNs with multiple hidden layers of perceptrons). DL architectures allow complex concepts to be built using simpler concepts [26], e.g. a person is made out of body parts, body parts are made out of contours and edges, contours and edges are detected in arrays of raw pixel input. Examples of DL-based detection algorithms which have achieved excellent performance at VRU detection from images in recent years include Fast-RCNN [27], Faster R-CNN [28], R-FCN [29], YOLO [30] and SSD [31]. Algorithms such as Pose-RCNN [24] have also been developed to leverage both image and lidar data concurrently.

## VI. OPEN CHALLENGES AND PROMISING FUTURE RESEARCH DIRECTIONS

### A. *Small & partially occluded objects*

Small or partially occluded VRUs remain among the most challenging objects for AV perception systems to detect accurately. Stixel-based proposals are efficient, but need to be improved for these cases; future work should include careful tuning of stixel and proposal parameters (e.g. stixel width and segmentation costs) to address these challenges [21].

### B. *Detection of VRU gestures*

Future AVs will need to understand at least a minimal set of gestures made by humans, e.g. in cases where a cyclist is signalling a change of direction, or when a construction worker or police officer is manually directing traffic during

construction works or at the scene of an accident. A first step in this direction would be to train AV perception systems to recognise individuals who could potentially give hand gestures, and then further classify these as VRUs giving advisory gestures (for cyclists) or directions which must be obeyed (for manual traffic control).

This presents difficulties due to differences in uniform styles across countries, as well as similarities in appearance when compared to other VRUs. Furthermore, high-visibility clothing is not exclusively worn by cyclists, police officers or construction workers, so extreme caution must be taken not to misinterpret hand movements of passers-by as genuine advisory or traffic control gestures. One possible solution in the medium term is for police and construction workers to have mobile transmitters which could communicate the need to hand back control to a human driver when manual gestures are in use, although this would not be an acceptable solution for a level 5 AV.

### C. *False positives*

Further reducing the rate of false positives (i.e. detection of an object when one is in fact not present) remains an open challenge for computer vision researchers. VRU detection presents many opportunities for false detections to happen, e.g. due to reflections, clothes displayed in shop windows and large advertisements in the scene background which feature people. Some false positives may be eliminated by taking advantage of known scene geometry constraints (e.g. pedestrians or cyclists should be on the ground plane) [10]. Object tracking between frames can also reduce the rate of false positives [10] as reflections present in one frame may not be present in the next.

### D. *Modelling and predicting VRU behaviour*

Humans drivers naturally predict the future movements of VRUs, and are able to anticipate events such as a pedestrian crossing the road suddenly. Developing models of VRU behaviour using ML techniques is an important direction for future work. Once a VRU has been detected and localised with the correct orientation, predictions of the future movements of VRUs could be integrated into the ego vehicle’s path planning algorithm. Recent research [32] recorded the movement of VRUs at an intersection and learned models for predicting VRU trajectories based on their current trajectory. Future work could adopt such models for on-board VRU behaviour prediction in AV systems.

### E. *Bayesian neural networks*

Current DL algorithms can effectively learn complex mappings between high-dimensional input data and a given set of outputs. However, a major shortcoming of these approaches is that they do not have any information about the uncertainty associated with a particular mapping. Bayesian deep learning (BDL) is a promising field which could address this shortcoming, while potentially achieving state-of-the-art-results [33]. BDL combines ideas from DL and Bayesian probability theory, so that an uncertainty value is associated with each object




Fig. 4. Sample uncertainty map generated by a BDL approach. Note how higher uncertainty values are shown on difficult surfaces, such as vehicle windows. Figure reproduced from [33].

detection in each region of an image. A sample uncertainty map generated by a BDL algorithm is shown in Fig. 4. Having an uncertainty value available may be useful to help to avoid false positives, an issue which was identified above as a major challenge for VRU detection systems. However, BDL is currently too computationally expensive to allow real-time inference on AV hardware (experiments in [33] ran slowly despite using an expensive and powerful NVIDIA Titan X GPU). Therefore, future work should investigate methods to speed up inference using this promising technique, so that it will become computationally viable for deployments on AVs.

#### F. Neuroevolution

Typical DL algorithms make use of backpropagation to learn the network weights. One alternative is to use a population-based genetic algorithm (GA), where the weights of the network are encoded in genes. The GA then uses operators inspired by biological evolution such as selection, crossover and mutation to create new genes, and thus new sets of network weights. This alternative to backpropagation has gained popularity again recently, following successes such as Uber AI Labs' demonstration that GAs are a competitive alternative to backpropagation when training deep ANNs for reinforcement learning settings [34].

One advantage of GA-based methods is that they scale extremely well across large numbers of cores (tests were conducted across up to 720 CPUs in [34]), as each solution in a new generation may be tested on its own core, allowing many candidate solutions to be evaluated at once. GAs are also effective at escaping local optima, especially quality-diversity algorithms [35] such as novelty search. Neuroevolution could feasibly be used when training deep ANNs for object detection, potentially improving training speed and reducing the time needed to develop and test new algorithms for VRU detection.

#### G. V2X communication

As with human drivers, even the best AV perception systems which are based on line-of-sight will not be able to detect VRUs that are fully occluded (e.g. behind another vehicle, a building or another VRU). Future AVs are likely to be connected to other vehicles (V2V), to infrastructure (V2I) and possibly even to VRUs (V2P). Prototype V2V systems

such as Valeo XtraVue<sup>1</sup> allow sensor data to be transmitted between adjacent vehicles, which could reduce the number of objects in a scene which are occluded by other vehicles. The fusion of data from on-board sensors with data received from V2P communications for improved pedestrian detection is also currently being investigated (see e.g. [36]). The increased availability of information resulting from these methods is likely to significantly improve VRU detection capabilities.

## VII. CONCLUSION

This paper presented an overview of recent research and the current state-of-the-art in VRU detection, followed by a discussion of the main open challenges and promising future research directions. AV systems are still very much at the prototype stage and are unlikely to secure regulatory approval for widespread deployment unless they are at least as safe as human drivers. The most fundamental requirement for safe AVs is a fast and accurate environment perception system. Computer vision engineers must pay special attention to VRUs as they are among the most difficult objects to detect, as well as having a high risk of death or serious injury in the event of a collision. Focusing research efforts on promising directions such as those listed in Section VI is essential if the performance of AV perception systems for VRU detection is to meet or exceed that of human drivers in the coming years.

## REFERENCES

- [1] P. Fairley, "Self-driving cars have a bicycle problem," *IEEE Spectrum*, vol. 54, no. 3, pp. 12–13, March 2017.
- [2] "Global status report on road safety 2015," World Health Organization (WHO), Tech. Rep., 2015. [Online]. Available: [http://www.who.int/violence\\_injury\\_prevention/road\\_safety\\_status/2015/en/](http://www.who.int/violence_injury_prevention/road_safety_status/2015/en/)
- [3] N. Sajj, "General safety of vehicles and protection of vulnerable road users," European Parliamentary Research Service, Tech. Rep., 2018.
- [4] N. Bernini, M. Bertozzi, L. Castangia, M. Patander, and M. Sabbatelli, "Real-time obstacle detection using stereo vision for autonomous ground vehicles: A survey," in *17th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, Oct 2014, pp. 873–878.
- [5] H. Zhu, K. Yuen, L. Mihaylova, and H. Leung, "Overview of environment perception for intelligent vehicles," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 10, pp. 2584–2601, Oct 2017.
- [6] J. Janai, F. Güneş, A. Behl, and A. Geiger, "Computer Vision for Autonomous Vehicles: Problems, Datasets and State-of-the-Art," *ArXiv e-prints*, April 2017. [Online]. Available: <https://arxiv.org/abs/1704.05519>
- [7] R. Benenson, M. Omran, J. Hosang, and B. Schiele, "Ten years of pedestrian detection, what have we learned?" in *European Conference on Computer Vision*. Springer, 2014, pp. 613–627.
- [8] S. Zhang, R. Benenson, M. Omran, J. Hosang, and B. Schiele, "How far are we from solving pedestrian detection?" in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 1259–1267.
- [9] P. Dollár, C. Wojek, B. Schiele, and P. Perona, "Pedestrian detection: An evaluation of the state of the art," *IEEE transactions on pattern analysis and machine intelligence*, vol. 34, no. 4, pp. 743–761, 2012.
- [10] M. Braun, S. Krebs, F. Flohr, and D. M. Gavrila, "The EuroCity Persons Dataset: A Novel Benchmark for Object Detection," *ArXiv e-prints*, May 2018. [Online]. Available: <https://arxiv.org/abs/1805.07193>
- [11] V. Talpert, I. Sobh, B. R. Kiran, P. Mannion, S. Yogamani, A. El-Sallab, and P. Perez, "Exploring applications of deep reinforcement learning for real-world autonomous driving systems," in *International Conference on Computer Vision Theory and Applications (VISAPP)*, February 2019. [Online]. Available: <http://insticc.org/node/TechnicalProgram/visigrapp/presentationDetails/75203>

<sup>1</sup><https://youtu.be/F4-wWfCcyK4>

- [12] C. Premebida, G. Monteiro, U. Nunes, and P. Peixoto, "A lidar and vision-based approach for pedestrian and vehicle detection and tracking," in *2007 IEEE Intelligent Transportation Systems Conference*, Sept 2007, pp. 1044–1049.
- [13] T. Ogawa, H. Sakai, Y. Suzuki, K. Takagi, and K. Morikawa, "Pedestrian detection and tracking using in-vehicle lidar for automotive application," in *2011 IEEE Intelligent Vehicles Symposium (IV)*, June 2011, pp. 734–739.
- [14] X. Han, J. Lu, Y. Tai, and C. Zhao, "A real-time lidar and vision based pedestrian detection system for unmanned ground vehicles," in *2015 3rd IAPR Asian Conference on Pattern Recognition (ACPR)*, Nov 2015, pp. 635–639.
- [15] A. Gonzalez, D. Vzquez, A. M. Lopez, and J. Amores, "On-board object detection: Multicue, multimodal, and multiview random forest of local experts," *IEEE Transactions on Cybernetics*, vol. 47, no. 11, pp. 3980–3990, Nov 2017.
- [16] Y. Zhang, S. Gu, J. Yang, M. J. Alvarez, and H. Kong, "Fusion of lidar and camera by scanning in lidar imagery and image-guided diffusion for urban road detection," in *2018 IEEE Intelligent Vehicles Symposium (IV)*, June 2018, pp. 579–584.
- [17] S. Zhang, R. Benenson, and B. Schiele, "Citypersons: A diverse dataset for pedestrian detection," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017, pp. 4457–4465.
- [18] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. IEEE, 2012, pp. 3354–3361.
- [19] X. Li, F. Flohr, Y. Yang, H. Xiong, M. Braun, S. Pan, K. Li, and D. M. Gavrila, "A new benchmark for vision-based cyclist detection," in *2016 IEEE Intelligent Vehicles Symposium (IV)*, June 2016, pp. 1028–1033.
- [20] Z.-Q. Zhao, P. Zheng, S.-t. Xu, and X. Wu, "Object detection with deep learning: A review," *arXiv preprint arXiv:1807.05511*, 2018.
- [21] F. B. Flohr, "Vulnerable road user detection and orientation estimation for context-aware automated driving," Ph.D. dissertation, Universiteit van Amsterdam, 2018.
- [22] J. R. Uijlings, K. E. Van De Sande, T. Gevers, and A. W. Smeulders, "Selective search for object recognition," *International journal of computer vision*, vol. 104, no. 2, pp. 154–171, 2013.
- [23] H. Badino, U. Franke, and D. Pfeiffer, "The stixel world—a compact medium level representation of the 3d-world," in *Joint Pattern Recognition Symposium*. Springer, 2009, pp. 51–60.
- [24] M. Braun, Q. Rao, Y. Wang, and F. Flohr, "Pose-rcnn: Joint object detection and pose estimation using 3d object proposals," in *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, Nov 2016, pp. 1546–1551.
- [25] T. M. Mitchell, *Machine learning*, ser. McGraw-Hill series in computer science. Boston (Mass.), Burr Ridge (Ill.), Dubuque (Iowa): McGraw-Hill, 1997.
- [26] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>.
- [27] R. Girshick, "Fast r-cnn," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1440–1448.
- [28] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Advances in neural information processing systems*, 2015, pp. 91–99.
- [29] J. Dai, Y. Li, K. He, and J. Sun, "R-fcn: Object detection via region-based fully convolutional networks," in *Advances in neural information processing systems*, 2016, pp. 379–387.
- [30] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 779–788.
- [31] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European conference on computer vision*. Springer, 2016, pp. 21–37.
- [32] M. Goldhammer, S. Köhler, S. Zernetsch, K. Doll, B. Sick, and K. Dietmayer, "Intentions of vulnerable road users—detection and forecasting by means of machine learning," *arXiv preprint arXiv:1803.03577*, 2018.
- [33] A. Kendall and Y. Gal, "What uncertainties do we need in bayesian deep learning for computer vision?" in *Advances in neural information processing systems*, 2017, pp. 5574–5584.
- [34] F. P. Such, V. Madhavan, E. Conti, J. Lehman, K. O. Stanley, and J. Clune, "Deep neuroevolution: genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning," *arXiv preprint arXiv:1712.06567*, 2017.
- [35] J. K. Pugh, L. B. Soros, and K. O. Stanley, "Quality diversity: A new frontier for evolutionary computation," *Frontiers in Robotics and AI*, vol. 3, 2016.
- [36] P. Merdrignac, O. Shagdar, and F. Nashashibi, "Fusion of perception and v2p communication systems for the safety of vulnerable road users," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 7, pp. 1740–1751, July 2017.