Ellipse Detection for Visual Cyclists Analysis “In the Wild”

Abdelrahman Eldesokey, Michael Felsberg and Fahad Shahbaz Khan

The self-archived postprint version of this conference article is available at
Linköping University Institutional Repository (DiVA):

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-145372>

N.B.: When citing this work, cite the original publication.

The original publication is available at www.springerlink.com:

Eldesokey, A., Felsberg, M., Khan, F. S., (2017), Ellipse Detection for Visual Cyclists
Analysis “In the Wild”, *Computer Analysis of Images and Patterns*, , 319-331.

https://doi.org/10.1007/978-3-319-64689-3_26

Original publication available at:

https://doi.org/10.1007/978-3-319-64689-3_26

Copyright: Springer Verlag (Germany)

<http://www.springerlink.com/?MUD=MP>

Ellipse Detection for Visual Cyclists Analysis “In the Wild”

Abdelrahman Eldesokey, Michael Felsberg, and Fahad Shahbaz Khan

Computer Vision Laboratory, Linköping University, Sweden
{[abdelrahman.eldeskey](mailto:abdelrahman.eldeskey@liu.se),[michael.felsberg](mailto:michael.felsberg@liu.se),[fahad.khan](mailto:fahad.khan@liu.se)}@liu.se

Abstract. Autonomous driving safety is becoming a paramount issue due to the emergence of many autonomous vehicle prototypes. The safety measures ensure that autonomous vehicles are safe to operate among pedestrians, cyclists and conventional vehicles. While safety measures for pedestrians have been widely studied in literature, little attention has been paid to safety measures for cyclists. Visual cyclists analysis is a challenging problem due to the complex structure and dynamic nature of the cyclists. The dynamic model used for cyclists analysis heavily relies on the wheels. In this paper, we investigate the problem of ellipse detection for visual cyclists analysis in the wild. Our first contribution is the introduction of a new challenging annotated dataset for bicycle wheels, collected in real-world urban environment. Our second contribution is a method that combines reliable arcs selection and grouping strategies for ellipse detection. The reliable selection and grouping mechanism leads to robust ellipse detections when combined with the standard least square ellipse fitting approach. Our experiments clearly demonstrate that our method provides improved results, both in terms of accuracy and robustness in challenging urban environment settings.

1 Introduction

Visual cyclists analysis is gaining considerable attention, especially due to the growing demand for autonomous driving safety. The analysis mainly involves understanding cyclists’ behavior and their intentions. A cyclist has a complex structure composed of a bicycle and a pedestrian. Therefore, it cannot be processed as a pedestrian nor a bicycle. The work of [35] introduced a dynamic model for cyclists with nine state parameters that define the cyclist pose in the global coordinate system. Among these parameters are the wheel base, the steering angle, and the normal vector to the rear wheel, as well as the normal to the front wheel, which can be estimated from the steering angle. Such sophisticated dynamic model requires a robust and accurate ellipse estimation for the front and the back wheel to facilitate the state estimation. The assumption is that by tracking these states, the behavior of cyclists can be analyzed and their intentions can be predicted. This would have a great impact on autonomous driving safety, allowing vehicles to interact with cyclists efficiently by knowing their current state and their intentions.

Thus, ellipse detection plays a crucial rule for visual cyclists analysis. There are plethora of good ellipse detectors in literature, which have been extensively

Fig. 1. (a) An example from the Caltech 256 dataset [14], (b) Examples for challenges in the TDCB dataset [20].

evaluated [26, 2, 11]. However, they have been evaluated on either synthetic data or clean images taken in controlled environments. To our knowledge, there exist no dataset for ellipse detection in realistic imagery acquired in uncontrolled environment. Hence, we introduce a new dataset, E-TDCB, with annotated ellipses of wheels for visual cyclists analysis in urban environment. The images in the E-TDCB dataset are taken from the *Tsinghua-Daimler Cyclist Detection Benchmark Dataset (TDCB)* [20], and we provide rich annotations of the bicycle wheels. Our motivation to generate this dataset is to provide an evaluation benchmark for ellipse fitting in real-world urban imagery that is more challenging than the standard datasets and to produce a baseline for visual cyclists analysis methods that apply a dynamic cyclist model relying on ellipse estimates. Figure 1 shows a comparison between a bicycle image from the Caltech 256 dataset [14], taken in a controlled environment, and some challenging examples from the TDCB dataset. We also introduce a novel ellipse detector which combines several ellipse fitting approaches into a light-weight detector with real-time performance, high accuracy, and robustness. We perform comprehensive experiments by evaluating our method and state-of-the-art ellipse detectors on the E-TDCB dataset. The results clearly demonstrate that our method outperforms existing state-of-the-art detectors, while providing an exceptional balance between accuracy and robustness. In summary, our contributions are:

- A new dataset with wheels annotations for visual cyclists analysis in the wild.
- A robust ellipse-based wheel detection method facilitating cyclists analysis.
- Comparison to existing state-of-the-art ellipse detectors on the new dataset.

2 Related Work

One of the few existing works on visual cyclists analysis is the recent work by Zernetsch et al. [34] that predicts the trajectory of cyclists and their intentions such as “Starting”, “Stopping”, “Waiting”, or “Passing”. The approach is based on Artificial Neural Networks that are trained on annotated tracks captured using a stereo camera and a laser scanner. Another work by Ardeshiri et al. [1] estimates and tracks the cyclist’s state based on measurements from a conventional monocular camera instead of special hardware. It applies an advanced dynamic cyclist model [35] and particle filters to predict the future state of cyclists. For the ellipse extraction, a simple method fits ellipses to bicycles wheels based on the assumption that tires have reflective rings on them, thus unsuitable

for uncontrolled imagery. A robust ellipse detector that works with uncontrolled bicycle imagery would be essential to make this method applicable in practice.

Shape matching, as a generalization to ellipse extraction, has been a frequently studied computer vision approach to object recognition tasks. Several methods have been proposed based on geometric context [3, 4], shape descriptors in the spatial domain [12], and in the frequency domain [19]. These methods have been applied to numerous challenging problems such as object recognition [25, 24], character recognition [29], traffic sign recognition [19], pedestrian detection, and motion analysis [21]. *Ellipses* are one of the most frequently observed shapes in digital imagery since they are projections of circles commonly available in real-world objects. This has prompted ellipse fitting to be one of the prerequisites for several shape matching methods employed in many applications, including facial gesture analysis [30, 15], medical image analysis [27, 31], vehicle wheels detection [7], visual cyclists analysis [1], and traffic sign detection [22].

An ellipse is defined by five parameters and a common approach was to use the Hough Transform (HT) to estimate these parameters. The approach is similar to HT line detection, but in case of ellipses, the accumulator has five dimensions instead of two. This imposes high computational and memory demands to explore the space and several attempts have been made to reduce its size. McLaughlin [23] eliminated two parameters by geometrically finding the center of the ellipse and then performing Randomized HT (RHT) [33] to obtain the other three parameters. In [32], the dimensionality of the accumulator was reduced to one dimension by randomly selecting pairs of pixels that match certain geometric constraints and estimate the center, major axis, and the angle for candidate ellipses from these pairs. The minor axis is then calculated using 1D HT. Similarly, two optimized versions of [32] were proposed in [6, 2] which require less memory, fewer computations, and are more robust against noise and false detections. All these HT-based methods require a proper selection of several control parameters that define geometric constraints.

Another approach is based on the canonical representation of ellipses and formulating the problem as a least-squares minimization. Fitzgibbon [10] introduced direct least squares fitting of ellipses by minimizing the algebraic distance between some scattered points and an ellipse hypothesis that is represented in canonical form and constrained to produce only ellipses, not parabolas or hyperbolas. A more stable version [28] addresses the problem of singularities in the design matrix [10] and produces more stable solutions for the least-squares problem. RANSAC has also been used to randomly sample a subset of the data, to fit an ellipse using least-squares minimization, and to iterate until the best ellipse is found according to some convergence criteria [17, 8].

Recently, two methods were introduced for ellipse fitting which combine the two approaches above. The first one [26] is mainly based on edge curvature and convexity to determine groups which potentially form ellipses. A modified 2D HT is used to evaluate these groups and if they fulfill some constraints, an ellipse is fitted using direct least squares [10]. The ellipse candidates are evaluated using three unique saliency measures and an ellipse is selected if its saliency score exceeds the average score for all other ellipses. The second method

was proposed by Fornaciari and others [11] and is based on grouping edges that adhere to some geometric constraints. Thus the parameter space for HT is reduced, enabling real-time performance even on smart phones with limited computational power. These two methods were shown to achieve state-of-the-art performance on synthetic data and the Caltech 256 dataset [14]. Both methods will be evaluated on the E-TDCB dataset in section 5.

3 Ellipse Fitting for Visual Cyclists Analysis

The Tsinghua-Daimler Cyclists Benchmark [20] provides a comprehensive evaluation for state-of-the-art object detectors on their dataset. The Deformable Parts Model (DPM) object detector [9] was able to achieve a remarkable performance against sophisticated methods such as F-RCNN deep networks [13]. A major advantage of a DPM is its ability to construct a flexible model for the object-of-interest. This model is composed of the prominent parts that form the object, their arrangement, and relationships between them. This makes the DPM approach highly appropriate for visual cyclist analysis as it gives an insight about the structure of the cyclist and facilitates its analysis accordingly. A part of the training process for a DPM detector is to find the optimal locations for different parts and to construct a weight filter for each part. Those filters are convolved with the image during test time and are supposed to produce high output at their corresponding parts.

3.1 Finding The Wheels

We trained a DPM model on the TDCB dataset, restricted to cyclists that are seen from the side view as described in details in section 4. The model has a root filter that locates cyclists as a whole and another eight part filters which locate different parts of the cyclist as illustrated in Figure 2(a). The state model of the cyclist relies basically on the wheels as described in section 1. Hence, we sample *potential patches* from DPM parts that are located on the wheels (shown in *red* in Figure 2(a)). The Canny edge detector [5], i.e., horizontal and vertical derivatives, is applied to potential patches to get the edge map. Each patch is processed individually on the subsequent steps.

3.2 Arcs Selection

Initially, each edge pixel p in the edge map is classified according to the sign of the orientation tangent of its gradient as follows:

$$\text{sign}(\tan \varphi(p)) = \text{sign}(G_x(p)) \cdot \text{sign}(G_y(p)) \quad (1)$$

where $G_x(p)$ and $G_y(p)$ are the horizontal and the vertical derivatives, respectively. Pixels with positive orientations are stored in a set P_{pos} while pixels with negative orientations are stored in P_{neg} . Each set is processed by an edge linking algorithm to define connected arcs. In this work we use the edge linking algorithm by Kovess [18] and store each arc a_i from $P_{\text{pos}}, P_{\text{neg}}$ in sets $A_{\text{pos}}, A_{\text{neg}}$, respectively. Arcs that have a length less than Θ_{Len} are discarded.

Fig. 2. Method overview. **(a)** DPM root filter in green, parts filters in white and selected filters in red. **(b)** Candidate patches sampled at the wheels. **(c)** Edge maps for candidate patches. **(d)** Classifying edge pixels according to positive (cyan) and negative (magenta) sign of the orientation tangent. **(e)** Arcs convexity check. **(f)** Different arcs quadrants in unique colors. **(g)** Ellipse hypothesis. **(h)** Best ellipse hypothesis. **(j)** A long arc shown in green that is split using inflexion point detection.

The convexity of each arc is checked to determine if it is up-facing or down-facing. Similarly to [11] and as illustrated in Figure 2(e), the area below and above the arc determines its convexity:

$$\text{Conv}(a_i) = \begin{cases} 1, & \text{Area}(L^i) > \text{Area}(U^i) \\ -1, & \text{Area}(U^i) > \text{Area}(L^i) \end{cases}, \quad (2)$$

Based on the arc convexity and orientation, the associated quadrant can be determined:

$$Q(a_i) = \begin{cases} I & \text{if } (a_i \in A_{\text{pos}}) \wedge (\text{Conv}(a_i) = 1) \\ II & \text{if } (a_i \in A_{\text{neg}}) \wedge (\text{Conv}(a_i) = 1) \\ III & \text{if } (a_i \in A_{\text{pos}}) \wedge (\text{Conv}(a_i) = -1) \\ IV & \text{if } (a_i \in A_{\text{neg}}) \wedge (\text{Conv}(a_i) = -1) \end{cases}, \quad (3)$$

note that the signs of G_x and G_y are not sufficient to determine the quadrant as the direction of the gradient is unknown (direction vs. orientation [16]). Figure 2(f) shows different quadrants in unique colors. Occasionally, two arcs are merged due to noise or background clutter as shown in Figure 2(j). Therefore, we detect inflexion points, where the continuity of the arc is violated. We apply a 3-steps approach for detecting those inflexion points [26]: (a) fit line segments to the arc; (b) calculate their angles with the arc; and (c) check how these angles change between line segments. A sign change for these angles indicates a change in the arc curvature and the arc is split at this point.

3.3 Arcs Grouping

Arcs are grouped into pairs in a anti-clockwise order, e.g., arc $a_i \Leftrightarrow Q(a_i) = I$ is grouped with arc $a_k \Leftrightarrow Q(a_k) = II$. This grouping is constrained to prevent irrelevant arcs from being grouped with arcs from the wheels:

$$Pr(a_i, a_k) = \begin{cases} \top & \text{if } Q(a_i, a_k) = \langle I, II \rangle \wedge \text{abs}(a_i.L_y - a_k.R_y) \leq \Theta_{\text{pair}} \\ \top & \text{if } Q(a_i, a_k) = \langle II, III \rangle \wedge \text{abs}(a_i.L_x - a_k.L_x) \leq \Theta_{\text{pair}} \\ \top & \text{if } Q(a_i, a_k) = \langle III, IV \rangle \wedge \text{abs}(a_i.R_y - a_k.L_y) \leq \Theta_{\text{pair}} \\ \top & \text{if } Q(a_i, a_k) = \langle IV, I \rangle \wedge \text{abs}(a_i.R_x - a_k.R_x) \leq \Theta_{\text{pair}} \\ \perp & \text{else ,} \end{cases} \quad (4)$$

where L/R is the leftmost/rightmost point of the arc, $\text{abs}()$ is the absolute value, and Θ_{pair} is the pairing threshold. The choice of the value of Θ_{pair} depends on the thickness of the wheel and the size of the bicycle. A good selection for this parameter prevents wrong pairings of arcs that do not belong to the same ellipse. For instance, in Figure 2(g), a high value for Θ_{pair} will cause arcs belonging to the inner rim to be grouped with arcs from the outer rim. After applying the constraints (4) to all possible pairs, only $\top$ -arcs are paired together and added to the set of pairs S_{pairs} . Eventually, all pairs that have a common arc are grouped into triplets and added to a set of triplets S_{triplets} . For example, assume two pairs of arcs $\langle a_c, a_d \rangle$ and $\langle a_d, a_e \rangle$ will be grouped into a triplet $\langle a_c, a_d, a_e \rangle$ and this triplet is added to S_{triplets} .

3.4 Ellipse Fitting, Grouping and Evaluating

For each triplet in S_{triplets} , all arc points in this triplet are used for fitting an ellipse. Direct least squares ellipse fitting [10] is used and the residual error is required to be less than Θ_{LSE} for this ellipse to be considered as an *ellipse hypothesis*, see Figure 2(g). As a second step, similar ellipses, i.e., ellipses with high overlap with each other, are grouped and their parameters are averaged to get a representing ellipse. Finally, ellipse hypotheses are evaluated by checking their intersection with the edge map. Bicycle wheels are usually occluded and cluttered by the background. Consequently, some parts of the wheels will not form an arc and will be removed as noise which will lead to an imprecise ellipse estimation. Therefore, checking for intersection with all edge pixels results in a better estimation even if some arcs are removed.

4 Dataset

The starting point of this work was to investigate whether the available ellipse detectors work in real-life applications such as visual cyclists analysis, where data is taken from uncontrolled environment, i.e. a camera mounted on the dashboard of a car. We checked the recently published datasets that match the above criteria and we found that the *Tsinghua-Daimler Cyclist Detection Benchmark Dataset (TDCB)* [20] is a perfect match. The images were recorded using a stereo camera setup mounted on a car that drove in the streets of Hong Kong. Bounding box annotations were provided for different classes of objects such as pedestrians,

Fig. 3. Examples of wheels annotation.

cyclists, motorcycles, and other objects. The dataset has many challenges that only arise in uncontrolled environments. Bicycle wheels are usually cluttered by the background and sometimes occluded by the legs of the cyclist. Motion blur occur occasionally and cause bicycle wheels edges to be smeared. Also if a bicycle is occluded by another bicycle, it becomes tricky to discriminate between their wheels. Some challenging examples are shown in Figure 1.

We provide wheels annotations on the monocular images for cyclists with bounding box aspect ratio $r < 1.25$ in *training set*, $r < 1.75$ in *test set* where r is defined as $r = \text{bbox}_{\text{height}} / \text{bbox}_{\text{width}}$. This ratio indicates that the cyclist is seen from the side and the wheels are visible. Other images from the front or the back view are ignored as they do not have visible wheels. The number of cyclists that matched the above ratio is 642 images from training set and 142 image from the test set. This means that the dataset has 1568 manually annotated ellipses for the outer rims of the wheels, while the inner rims were estimated roughly with respect to the cyclist size. Some annotated examples are shown in Figure 3.

5 Experiments

As discussed previously, a reliable ellipse fitting is essential for cyclist state estimation. Hence, we evaluate our ellipse detector as well as three state-of-the-art methods on the proposed dataset. The methods are Prasad [26], Basca [2], and Yaed [11]. An overview of each method was provided in section 2, while an overview of the dataset was given in section 4. The source code for Prasad was provided by the author while the source code for Yaed and Basca was found online.

5.1 Evaluation metrics

For evaluation, following evaluation metrics are used:

$$\text{Precision} = \frac{\text{Number of True Positive Ellipse Hypotheses}}{\text{Total Number of Ellipse Hypotheses}} \quad (5)$$

$$\text{Recall} = \frac{\text{Number of True Positive Ellipse Hypotheses}}{\text{Total Number of Ground-truth Ellipse}} \quad (6)$$

$$\text{F-Score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (7)$$

where a *true positive hypothesis* is an ellipse which has a certain overlap with any of the ground-truth ellipses. The overlap in case of ellipses is defined as $\text{overlap} =$

$1 - \frac{\text{count}(\text{XOR}(\text{Ellipse}, \text{GT}))}{\text{count}(\text{OR}(\text{Ellipse}, \text{GT}))}$. High precision indicates that the detector outputs highly confident hypotheses, high recall means that the detector is reliable in finding the ellipses, and f-score incorporates both. A good ellipse detector should combine high precision and a good recall.

5.2 Parameter Selection

Our method has three parameters. The minimum length of arc Θ_{Len} is set to 9, Θ_{pair} can be any value from 5 to 15, and Θ_{LSE} is set to 0.01. In Prasad, we change *minimum_edge_length* to 10 instead of 15 and the rest of parameters are kept unchanged. Basca has many control parameters as it is HT-based. We set *minMajorAxis* to one third of the candidate patch width, *maxMajorAxis* to the largest dimension of the candidate patch, and we only consider at most the best 25 ellipse hypothesis in the evaluation to retain levels of precision. For Yaed, we tried different combinations of its many control parameters and the best f-score was achieved at minimum arc length $Th_{\text{Len}} = 5$, minimum shortest side of arc oriented bounding box $Th_{\text{OBB}} = 1.0$, and the default remaining parameters.

5.3 Quantitative Results

Evaluation metrics are calculated for each method under different overlap ratios from 0 to 1 both on the training and the test set. The test set is more challenging as the locations of DPM filters are not perfectly aligned as in the training set. Besides, cyclists in the test set have larger yaw angles due to larger value of r which makes ellipse fitting more difficult. Figure 4 summarizes the evaluation metrics for the training and the test set. For Prasad, the final policy for selecting salient ellipses is too strict for this realistic dataset which led to very low recall. Therefore, we evaluated Prasad twice, once on the most salient hypotheses (Prasad-Best) and another on all hypotheses after grouping (Prasad-All).

As shown, our method and Yaed have a comparably high precision on both sets, while Prasad and Basca have average precision. Our explanation is that their arcs grouping criteria is based on a HT accumulator for finding the potential ellipse center for each arc. In case of concentric ellipses as in bicycle wheels, irrelevant ellipses from outer and inner rims are grouped which leads to false ellipse estimation. On the contrary, the constrained arcs grouping criteria in Yaed and our method alleviates these false groupings. The recall is high for all methods but for different reasons. For Basca, a huge number of hypotheses is produced which leads to a high recall under low thresholds only, while Prasad-All has a high recall due to considering all hypothesis in the evaluation. Yaed has slightly lower recall due to the large number of control parameters that are difficult to determine for balancing precision and recall. Finally, our method was able to retain high recalls even under high thresholds due to its relaxed constraints on dealing with all available edge information. The F-score shows that our method outperforms all other methods, achieving an outstanding balance between accuracy and robustness both on training and test set.

5.4 Qualitative Analysis

Figure 5 shows how each method performs with different challenges. Basca performs well when the wheels are visible in the ideal case, but when some challenges

Fig. 4. Evaluation metrics for Basca [2], Prasad [26], Yaed [11], and our proposed method on the training set (Left column) and test set (Right Column). The test set is more difficult than the training set as the DPM filters sometimes are not perfectly aligned in case of test set, which makes ellipse detection task more challenging.

are introduced, a considerable tuning has to be done for the control parameters of HT space. However, it does not perform well with cluttered or occluded wheels as HT-based approaches are sensitive to outliers. Prasad performs similarly to Basca and cannot handle most challenges as its ellipses saliency criteria is either too strict or in favor of small ellipses which has high circumference overlap ratio [26]. YAED noticeably performs better than the latter methods and achieves higher accuracy as it has a reliable arcs selection criteria. However it fails also to detect cluttered and occluded wheels due to its strict ellipses selection criteria and the large number of control parameters that needs to be adjusted to

Fig. 5. Performance of ellipse detectors with different challenges. For clarity, only the best hypothesis is drawn. For the Prasad method, we show the top half of all hypothesis as the best hypothesis is always a small ellipse due to its circumference overlap ratio.

each case. Finally, our method achieves an outstanding performance with all challenges. It succeeds to fit ellipses to occluded and cluttered wheels as it employs a very reliable and relaxed arcs selection criteria that encounters for all available information in the edge map, which is usually treated as noise in other approaches. Also our ellipse selection policy is suitable for real-world scenarios where images are not always ideal. Thereby, our method succeeded in combining both precision and reliability, which is essential for real-world applications.

6 Conclusion and Future Work

Due to the growing emergence of autonomous vehicles and the increasing attention to their safety measures, visual cyclists analysis needs to be investigated more for the sake of cyclists safety. The literature lacks for a realistic real-world dataset for this purpose. Therefore, we introduced a new dataset with wheels

annotations for cyclists that contributes to addressing the problem of visual cyclists analysis. We also proposed a robust and reliable method for ellipse fitting on bicycles wheels that is needed for cyclists state estimation. Our method as well as the state-of-the-art methods were evaluated on the new dataset and our method was able to outperform all other methods providing robust and reliable ellipse detection. In the future, more approaches need to be investigated for ellipse fitting. For instance, the use of edge orientation information in ellipse fitting, iterative least square minimization for refined fitting, and further investigations for the most suitable ellipse selection criteria. Also integrating our ellipse fitting method in an existing visual cyclists analysis platform would give some insights on the advantages and drawbacks of our method.

7 Acknowledgments

This work has been supported by VR (EMC², ELLIIT, starting grant [2016-05543]) and Vinnova (Cykla).

References

1. T. Ardeschiri, F. Larsson, F. Gustafsson, T. B. Schön, and M. Felsberg. Bicycle tracking using ellipse extraction. In *Information Fusion (FUSION), 2011 Proceedings of the 14th International Conference on*, pages 1–8. IEEE, 2011.
2. C. Basca, M. Talos, and R. Brad. Randomized hough transform for ellipse detection with result clustering. In *Computer as a Tool, 2005. EUROCON 2005. The International Conference on*, volume 2, pages 1397–1400. IEEE, 2005.
3. S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE trans. on pattern analysis and machine intelligence*, 24(4):509–522, 2002.
4. A. C. Berg, T. L. Berg, and J. Malik. Shape matching and object recognition using low distortion correspondences. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 1, pages 26–33.
5. J. Canny. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence*, (6):679–698, 1986.
6. A. Y. S. Chia, M. K. Leung, H.-L. Eng, and S. Rahardja. Ellipse detection with hough transform in one dimensional parametric space. In *Image Processing, 2007. ICIP 2007. IEEE International Conference on*, volume 5, pages V–333.
7. T. Cooke. A fast automatic ellipse detector. In *Digital Image Computing: Techniques and Applications (DICTA), 2010 International Conference on*, pages 575–580. IEEE, 2010.
8. F. Duan, L. Wang, and P. Guo. Ransac based ellipse detection with application to catadioptric camera calibration. In *International Conference on Neural Information Processing*, pages 525–532. Springer, 2010.
9. P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *IEEE transactions on pattern analysis and machine intelligence*, 32(9):1627–1645, 2010.
10. A. Fitzgibbon, M. Pilu, and R. B. Fisher. Direct least square fitting of ellipses. *IEEE Trans. on pattern analysis and machine intelligence*, 21(5):476–480, 1999.
11. M. Fornaciari, A. Prati, and R. Cucchiara. A fast and effective ellipse detector for embedded vision applications. *Pattern Recognition*, 47(11):3693–3708, 2014.
12. R. Gal and D. Cohen-Or. Salient geometric features for partial shape matching and similarity. *ACM Transactions on Graphics (TOG)*, 25(1):130–150, 2006.

13. R. Girshick. Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1440–1448, 2015.
14. G. Griffin, A. Holub, and P. Perona. Caltech-256 object category dataset. 2007.
15. Y.-H. Huang, B.-C. Pan, S.-L. Zheng, J. Pan, and Y. Tang. Lip-reading detection and localization based on two stage ellipse fitting. In *International Conf. Wavelet Analysis and Pattern Recognition*. IEEE, 2008.
16. B. Jähne. Digital image processing. 2002.
17. W. Kaewapichai and P. Kaewtrakulpong. Robust ellipse detection by fitting randomly selected edge patches.
18. P. Kovesi. Edge linking and line segment fitting.
19. F. Larsson, M. Felsberg, and P.-E. Forssen. Correlating fourier descriptors of local patches for road sign recognition. *IET Computer Vision*, 5(4):244–254, 2011.
20. 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 *Intelligent Vehicles Symposium (IV), 2016 IEEE*, pages 1028–1033. IEEE, 2016.
21. H. Ling and D. W. Jacobs. Shape classification using the inner-distance. *IEEE transactions on pattern analysis and machine intelligence*, 29(2), 2007.
22. H. Liu and B. Ran. Vision-based stop sign detection and recognition system for intelligent vehicles. *Transportation Research Record: Journal of the Transportation Research Board*, (1748):161–166, 2001.
23. R. A. McLaughlin. Randomized hough transform: improved ellipse detection with comparison. *Pattern Recognition Letters*, 19(3):299–305, 1998.
24. G. Mori and J. Malik. Recognizing objects in adversarial clutter: Breaking a visual captcha. In *Computer Vision and Pattern Recognition, 2003. Proceedings. 2003 IEEE Computer Society Conference on*, volume 1, pages I–I. IEEE, 2003.
25. A. Opelt, A. Pinz, and A. Zisserman. A boundary-fragment-model for object detection. *Computer Vision–ECCV 2006*, pages 575–588, 2006.
26. D. K. Prasad, M. K. Leung, and S.-Y. Cho. Edge curvature and convexity based ellipse detection method. *Pattern Recognition*, 45(9):3204–3221, 2012.
27. J. Pu, B. Zheng, J. K. Leader, and D. Gur. An ellipse-fitting based method for efficient registration of breast masses on two mammographic views. *Medical physics*, 35(2):487–494, 2008.
28. J. F. Radim Halir. Numerically stable direct least squares fitting of ellipses, 1998.
29. J. Rocha and T. Pavlidis. A shape analysis model with applications to a character recognition system. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(4):393–404, 1994.
30. T. Takegami, T. Gotoh, and G. Ohyama. An algorithm for model-based stable pupil detection for eye tracking system. *Systems and Computers in Japan*, 35(13):21–31, 2004.
31. C. Teutsch, D. Berndt, E. Trostmann, and M. Weber. Real-time detection of elliptic shapes for automated object recognition and object tracking. In *Electronic Imaging 2006*, pages 60700J–60700J. International Society for Optics and Photonics, 2006.
32. Y. Xie and Q. Ji. A new efficient ellipse detection method. In *Pattern Recognition, 2002. Proceedings. 16th International Conference on*, volume 2, pages 957–960.
33. L. Xu, E. Oja, and P. Kultanen. A new curve detection method: randomized hough transform (rht). *Pattern recognition letters*, 11(5):331–338, 1990.
34. S. Zernetsch, S. Kohnen, M. Goldhammer, K. Doll, and B. Sick. Trajectory prediction of cyclists using a physical model and an artificial neural network. In *Intelligent Vehicles Symposium (IV), 2016 IEEE*, pages 833–838. IEEE, 2016.
35. K. J. strm, R. E. Klein, and A. Lennartsson. Bicycle dynamics and control. 25(4):26–47, 2005.