

Book series



*2<sup>nd</sup> International Scientific Conference*

*“ALFATECH – Smart Cities and modern technologies – 2025”*

**ALFATECH 2025 Conference program**

Friday 28<sup>th</sup> of February 2025

PAPER:

**Person Detection and Re-Identification in Open-World  
Settings of Retail Stores and Public Spaces**

AUTHORS:

**Branko Brklijač** (Faculty of Technical Sciences, University of Novi Sad, Novi Sad)\*

**Milan Brklijač** (Faculty of Finance, Banking and Auditing, Alfa BK University, Belgrade)

# Person detection and re-identification in open-world settings of retail stores and public spaces

Branko Brkljač

*Dept. of Power, Electronic and Telecommunication Engineering  
Faculty of Technical Sciences, University of Novi Sad  
Novi Sad, Republic of Serbia  
brkljacb@uns.ac.rs*

Milan Brkljač

*Faculty of Finance, Banking and Auditing  
Alfa BK University  
Novi Beograd, Republic of Serbia  
milan.brkljac@alfa.edu.rs*

**Abstract**—Practical applications of computer vision in smart cities usually assume system integration and operation in challenging open-world environments. In the case of person re-identification task the main goal is to retrieve information whether the specific person has appeared in another place at a different time instance of the same video, or over multiple camera feeds. This typically assumes collecting raw data from video surveillance cameras in different places and under varying illumination conditions. In the considered open-world setting it also requires detection and localization of the person inside the analyzed video frame before the main re-identification step. With multi-person and multi-camera setups the system complexity becomes higher, requiring sophisticated tracking solutions and re-identification models. In this work we will discuss existing challenges in system design architectures, consider possible solutions based on different computer vision techniques, and describe applications of such systems in retail stores and public spaces for improved marketing analytics. In order to analyse sensitivity of person re-identification task under different open-world environments, a performance of one close to real-time solution will be demonstrated over several video captures and live camera feeds. Finally, based on conducted experiments we will indicate further research directions and possible system improvements.

**Index Terms**—person re-identification (ReID), computer vision, retail stores, public spaces, marketing analytics

## I. INTRODUCTION

Modern computer vision techniques are solving complex technical challenges posed by applications in open-world settings. This refers to scenarios in which the system must overcome uncontrolled and difficult to predict characteristics of the operating environment. This is also the case with person detection [1] and person re-identification [2] tasks, which have traditionally been researched in the context of autonomous vehicles [3], robotics [4], human-machine interfaces [5], crowd analysis [6], and most commonly security applications [7]. The concept of smart cities [8] and the integration of such vision systems [9]–[11] into complex decision-making frameworks [12]–[14] have introduced the need for efficient and robust solutions, with the ability to make necessary trade-offs between operational costs and system performance on a case-by-case basis. The common goal is providing necessary insights into persons’ behaviour based on the ability to detect and possibly re-identify the same person at different points in time and space [15]. Such applications also raise the questions of

ethical and privacy nature, which is a significant line of current research [16], [17]. Nevertheless, integration of behaviour analysis systems in retail stores and public spaces is gaining momentum [18] and is expected to have significant impact on insightful business operations and space management.

Approaching potential customers based on research analysis of their behaviour can be considered a key ingredient for designing successful marketing strategies and campaigns. It is also manifestation of the general trend of tailored campaigns and data-driven decisions emerging across different modalities of interaction with retail shoppers and public space users. In other words, persons’ movement patterns and time spent in different parts of a store or elsewhere directly answer the questions: “How do consumers shop?”, or “How people engage in public space?”. Data driven insights into these questions are foundations for comprehensive understanding of their behaviour. Thus, behavioural constructs in different conceptual frameworks, like e.g. emotional regulation consumption or compensatory consumption [19], can be better related with the motivational and emotional aspects of the shopping process and public infrastructure engagement. It is also the kind of data that cannot be obtained from point-of-sale [20] or simple visitor statistics [21]. Therefore, spatial analysis of human behavior can be considered one of the complementary technologies for achieving optimal decisions and reducing the cost of adapting to individual needs of customers or users. In this paper, we provide a brief overview of methods and approaches for addressing design challenges for person detection and person re-identification in open-world environments and describe their use in retail sector and public spaces. We also present a demonstration of one cost efficient, close to real-time solution based on OAK-D lite [22] embedded vision platform [23] and discuss advantages and limitations of such hardware implementation and baseline algorithmic approach.

The rest of the paper is organized as follows. In Section II we break down the landscape of existing methods for person detection and person re-identification tasks and provide a structured overview of research directions that have been established by pursuing different solution architectures under the assumptions of specific imaging modalities, identification scenarios, robustness requirements and model learning

strategies. In Section III we go into implementation details of developed low-cost solution and indicate advantages and limitations of the adopted fast prototyping system design approach. Section IV presents conducted experiments and discusses their implications for possible applications in retail marketing analytics and public space management. Finally, in Section V we indicate directions for future research and possible system improvements. Code implementation is available at: <https://github.com/brkljac/personReID>.

## II. RELATED WORK

Person detection [1], [10], [11], alongside face detection [24], [25], is probably one of the most studied object detection tasks in the literature [26]. Similarly, person re-identification (ReID) [2], alongside face recognition [27], is one of the long-standing recognition problems in vision [28], [29]. Its goal is to retrieve information whether the specific person has already appeared in video or over multiple camera feeds in the past. In such case, re-identification is considered as successful if we can correctly assign person's identity with some of the identities already present in the gallery, or create a new identity entry in the case of novel person in the scene. Described working scenario corresponds to the open-set ReID task, which is different from the closed-set setting in which it is assumed that the query object already exists in the gallery. These two problems, detection and re-identification, can be approached either traditionally, i.e. separately [30], [31], in a sequential manner, or jointly, i.e. in an end-to-end manner [32], [33], which is a characteristic of more recent ReID approaches. In the following we provide a brief overview of methods utilized in these two specific tasks of interest.

### A. Detection task

In general, besides the problem of defining object categories [34], object detection approaches solve two fundamental problems in vision: localization and classification of object instances inside the scene. Consequently, binary classification of persons in the image can be formulated as a two stage: localization and detection problem, or as a single stage, i.e. one step regression task [26]. As there are numerous approaches for person detection, their applicability depends on the cost-performance trade-off put in front of system designer. The role of the detector in the broader context of the system in which it operates is also important. In this sense, the choice of an adequate working point on the detector's receiver operating characteristic (ROC) is also a matter of compromise and the requirements set regarding the subsequent processing of generated information. Thus, precision and recall, or the relationship between detector's sensitivity (hit rate) and specificity (false positives rate) play a crucial role in the design choice of accompanying or downstream processing elements that are relying on detector's decisions. However, in some cases, greedy approaches in which the possibility of a miss is avoided by increasing the sensitivity of the detector, despite a higher number of false alarms, are also justified if it is possible to adequately use such results. Such techniques may include

the use of detection or confidence scores in the case of still images, spatial filtering based on the temporal history of the object's locations in the case of video signal, or additional side information obtained from other sources. In the given context of public spaces monitoring and retail analytics, it is interesting to note that usual detection paradigm can be modified in order to overcome some of the difficulties posed by the application scenario. Thus, a new object detection paradigm called Locount is proposed in [35], with the goal of localizing groups of objects of interest and estimating number of instances within the group. It is similar to the type of problems considered in [6], which indicates that the task of person detection is still of active research interest. This is especially true in the case of open-world environments, where adaptation and resilience to various aggravating circumstances such as illumination variation, image noise, presence of occlusions, and variations in scale or resolution of persons in the video are often needed. It is achieved through advanced learning techniques, data augmentation, or combining of different signal modalities, [36], [37].

### B. Re-identification task

According to [38], image-based person re-identification methods can be broadly grouped by: 1) signal modality (e.g. visible, infrared [39], thermal [40], or cross-modal [41]); 2) learning approaches (e.g. whether the recognition is based on metric learning [42], [43] or more complex deep ReID architectures and learning frameworks [31], [44]); and 3) how generalization across multiple domains, i.e. cross-domain ReID is addressed [45]. Video-based person ReID methods [46]–[48] that exploit temporal information in the learning process are also seen as a separate line of research, with a similar division of methods with respect to the above-mentioned properties. Other aspects in relation to which ReID solutions are considered include different types of challenges where standard methods exhibit certain weaknesses, leading to classes of special, specific methods that effectively overcome the limitations of standard solutions. Thus, the main characteristics of effective ReID solutions are invariance to different variations of the input signal and robust hierarchical feature representations that lead to accurate identity matching.

Resolution mismatch [49], when a low resolution query image is matched against the high resolution gallery samples, is a common example of ReID problem that is affected by the scale of the object and image acquisition setup. Opposite situations of high resolution query images are also common, giving rise to resolution independent features [50] and multiscale pyramid representations [47].

The most common person ReID goal is a short-term identification, where it is assumed that the pedestrian's clothing remains unchanged over time. While such assumptions are fully justified in the case of described retail application scenarios, a more general solution would be the invariance of ReID to changes in person's clothing [51] and appearance [52], when captured during different activities. Thus, a step forward ideal long-term or lifelong person ReID would be

invariance to cloth-changing, which could be regarded as more applicable in real-world scenarios. In [53] authors claim that the most reliable discriminative characteristic for above-mentioned invariance is unclothed body shape, i.e. effective representation of the 3D shape and texture of the human body. It is closely related to other biometric identification methods based on human gait recognition [4], [54], which are complementary to classical image- or video-based ReID [55].

Significant class of adaptation methods also includes occluded person ReID [56]–[59]. A common strategy is to steer ReID toward decisions based on features that comprehend or do not take into account the hidden part of the person. Such examples include learning the dynamic prototype mask [56], feature attention mask [60], saliency-guided patch transfer [58], using multiple views of the same person [57], or relying on part based representations [61] and fine-grained image analysis [62], see e.g. [63], [64]. There have also been attempts to partition feature space [65], while the standard learning strategy for person ReID has been the triplet loss framework [66], [67] and its extensions [68]. Alternatives include various deep metric learning approaches [42] and ranking optimization at inference stage [44].

### III. PROPOSED PERSON REID IMPLEMENTATION

In order to analyze the sensitivity of person re-identification task under open-world setting of retail stores and public spaces we have implemented one cost efficient, close to real-time solution. It is based on OAK-D lite [23] device and pre-trained person detection and person ReID models made using OpenVINO [69] hardware acceleration deep learning framework. The corresponding camera device is shown in Fig. 1a and its main characteristics are summarized in Table I.

TABLE I: Main characteristics of OAK-D lite platform

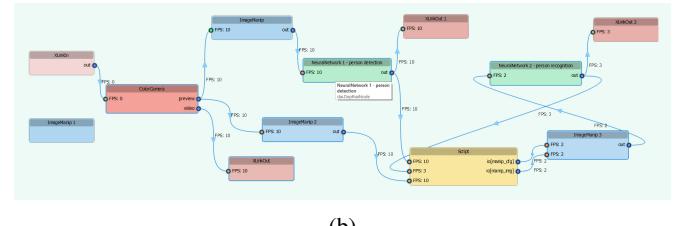
Feature	Specification
Processor	Intel Movidius Myriad X VPU
Cameras	1x 13 MP RGB (IMX214, rolling shutter) 2x 0.31 MP mono (OV7251, global shutter)
Stereo baseline	75 mm
AI performance	4 TOPS (1.4 TOPS for AI)
Video output	Up to 4K@30fps (H.264/H.265/MJPEG)
Depth perception	Stereo, 300,000-point, 200+ FPS
Connectivity	USB 3.1 Gen1 Type-C
Dimensions (WxHxD)	91 mm x 28 mm x 17.5 mm
Weight	61 g
Variants	Fixed-focus / Auto-focus (RGB cam)
Mounting	1/4"-20 tripod, VESA (7.5 cm, M4)

As can be seen from Table I, it is a versatile and heterogeneous computing platform. It is suitable for applications in different vision tasks that require real time stream processing and advantages of using embedded hardware implementation, such as high level of integration, small form factor, modular design, low power consumption and relatively low cost, taking into account that the device integrates both imaging devices and signal processor. Besides providing the options of choosing between the global and rolling shutter, on board cameras are

also designed to be geometrically precisely aligned in a linear arrangement to facilitate calibration and accurate distance measurements. Nevertheless, in the given experiments we have relied only on the main high resolution RGB camera and multiple stream processing by frame grabbing from created camera source in the constructed processing graph, Fig. 1b.



(a)



(b)

Fig. 1: OAK-D lite [23] vision platform: (a) three cameras, ISP and AI hardware accelerator on a single board, (b) implemented processing graph.

OAK-D lite platform integrates multiple specialized processing units and sensors that collaboratively handle different computational tasks, optimizing performance and efficiency. Thus, Movidius Myriad X Vision Processing Unit (VPU) integrates several system components two of which are of main interest for implemented person ReID. First one is Image Signal Processor (ISP), on-chip hardware accelerator responsible for different camera pipeline functions and high-throughput image pre-processing tasks such as denoising and color correction. The second one is Neural Compute Engine (NCE), neural network hardware accelerator, which enables efficient deployment of trained deep learning models or any other computational graphs that are possible to define in supported ONNX [70] or OpenVINO [69] network exchange formats. Its advantage is small power consumption, which is especially suitable for edge processing and deployment in environments that require continuous system operation. Note that although the main processing (person detection and person ReID) was done on embedded platform, experimental setup also required a host device (CPU) for higher-level program control and application logic (visualization, communication and I/O interface).

For software implementation, besides the device drivers and

other utilities provided by Luxonis corporation under the MIT license in official DepthAI code repository [71], we have used OpenCV [72] and Python programming language. We note that there are also C++ device interfaces and code implementations, available in [73]. The main algorithmic solution was taken from OpenVINO repository and consists of the pre-trained person detection model, available in [74], and person ReID model available from [75]. As a starting point for developing additional functionalities we have utilized demonstration script in [76], from which the provided code implementation was made, Section I. Possible alternative for baseline implementation could also be some general purpose ReID toolbox, like the one proposed in [77].

#### IV. EXPERIMENTAL RESULTS AND DISCUSSION

The experiments included visual analysis of system performance in various operating environments. We have analyzed ability of system to cope with person ReID task in outdoor and indoor space, under different illumination and in the presence of varying number and dynamics of people. General conclusions is that in all considered scenarios system managed to achieve close to real time operation, but never more than  $\approx 12$  frames per second (fps). On the other hand, under certain difficult conditions, the fps dropped even to  $\approx 4$  fps. This included low-light conditions, but mainly the presence of very dynamic background noise in the scene. For example, in an outdoor scene where a significant portion of the background was covered by leaves, in the presence of wind, the background dynamics were very pronounced and caused additional slowdown, Fig. 2a. However, even in such cases the final ReID result was good and without any false detections in the described background area.

The most challenging ReID were cases where the same person changed orientation towards the camera while walking or when entering and exiting the scene. This led to the creation of multiple identities for the same person, which is a common problem with person ReID. During experiments in an indoor

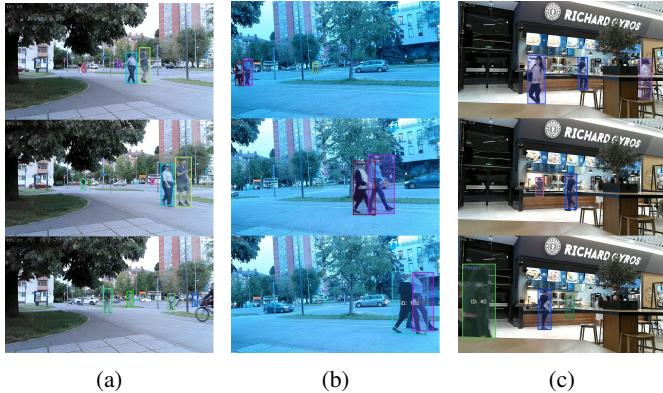


Fig. 2: (a) successful ReID, but with low fps rate due to dynamic background noise; (b) successful ReID under low light conditions, but with identity loss after change of person's orientation at the end of sequence; (c) retail store application.

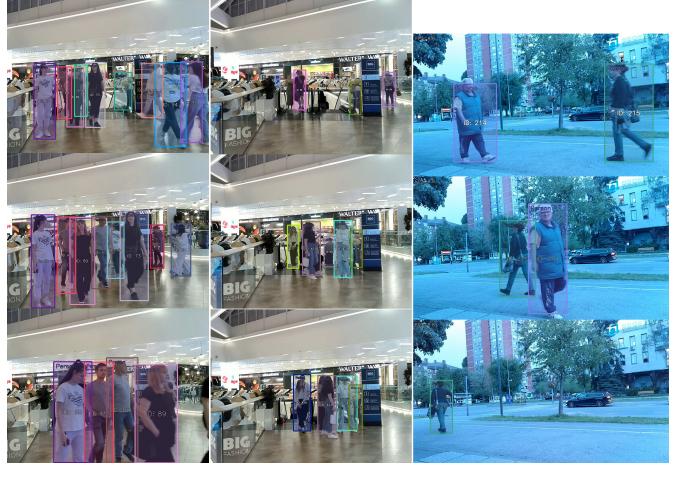


Fig. 3: Experimental results visualization: (a) crowded indoor scene: successful ReID, but with low fps; (b) false person detection, but with correct ReID; (c) low light operation.

retail space, it was found that certain object arrangements cause false person detection and initialize a correct but false person ReID in subsequent frames. This indicates that person detection module could be subject to further improvements in terms of robustness.

Most of the situations described above are illustrated in Fig. 2 and Fig. 3, while the corresponding video files are available in the given code repository, Section I.

The observed variability in people's appearance could be overcome by model training on datasets like [78], [79], which contain a sufficiently diverse range of appearances (positions of the person relative to the camera, changes in illumination and different human activities). Another research direction are more frequent ReID model updates through better use of historical appearances and without re-indexing of images in the gallery [80].

As already mentioned, crowded scenes and cluttered object detection sometimes require alternative formulations of standard vision tasks [6], [35]. On the side of possible applications, further analysis of the data collected by ReID systems could reveal a more complex behavior patterns [81], [82], i.e. go beyond simple statistics like time spent in different parts of a retail store. On the other hand, unsupervised representation learning [83] seems to be key to addressing the described challenges and increasing the robustness of ReID.

#### V. CONCLUSIONS

In this paper, we analyzed person re-identification (ReID) task from both technical and application perspectives. Conducted experiments highlighted the need to improve existing methods in order to make them fully applicable as analytical tools in retail stores and public spaces. Consequently, incorporating perceptual features such as image depth, audio signals, or other sensing modalities should be expected to enhance

ReID performance and reduce false positives in indoor environments. Nevertheless, the general conclusion is that the cost efficient person ReID solutions, like the one analyzed in the paper, are already potentially applicable in the given context.

#### ACKNOWLEDGMENT

The authors would like to thank conference organizers for their support throughout the manuscript preparation process. This research was partially supported by the Science Fund of the Republic of Serbia, project "Multimodal multilingual human-machine speech communication - AI SPEAK", grant no. 7449.

#### REFERENCES

- [1] P. Dollar, 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, 2011.
- [2] S. Gong, M. Cristani, Y. Shuicheng, and C. Change Loy, "Person re-identification," 2014, ISBN 978-1-4471-6295-7.
- [3] E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda, "A survey of autonomous driving: Common practices and emerging technologies," *IEEE access*, vol. 8, pp. 58443–58469, 2020.
- [4] J. Kwon, Y. Lee, and J. Lee, "Comparative study of markerless vision-based gait analyses for person Re-identification," *Sensors*, vol. 21, no. 24, p. 8208, 2021.
- [5] N. K. S. Behera, T. K. Behera, M. Nappi, S. Bakshi, and P. K. Sa, "Futuristic person re-identification over internet of biometrics things (iobt): Technical potential versus practical reality," *Pattern Recognition Letters*, vol. 151, pp. 163–171, 2021.
- [6] S. Chen, A. Fern, and S. Todorović, "Person count localization in videos from noisy foreground and detections," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 1364–1372.
- [7] S. Gong, C. C. Loy, and T. Xiang, "Security and surveillance," in *Visual analysis of humans: Looking at people*. Springer, 2011, pp. 455–472.
- [8] H. F. Yang, J. Cai, C. Liu, R. Ke, and Y. Wang, "Cooperative multi-camera vehicle tracking and traffic surveillance with edge artificial intelligence and representation learning," *Transportation Research part C: Emerging Technologies*, vol. 148, p. 103982, 2023.
- [9] J. Almazan, B. Gajic, N. Murray, and D. Larlus, "Re-ID done right: Towards good practices for person re-identification," *arXiv preprint arXiv:1801.05339*, 2018.
- [10] M. Shili, O. Sohaib, and S. Hammedi, "Yolo v5 and deep simple online and real-time tracking algorithms for real-time customer behavior tracking and retail optimization," *Algorithms*, vol. 17, no. 11, p. 525, 2024.
- [11] A. F. Del Carpio, "Analyzing computer vision models for detecting customers: A practical experience in a Mexican retail," *International Journal of Advances in Intelligent Informatics*, vol. 10, no. 1, pp. 131–147, 2024.
- [12] M. Quintana, J. M. Menéndez, F. Alvarez, and J. P. López, "Improving retail efficiency through sensing technologies: A survey," *Pattern Recognition Letters*, vol. 81, pp. 3–10, 2016.
- [13] X. Liu, Y. Jiang, P. Jain, and K.-H. Kim, "TAR: Enabling fine-grained targeted advertising in retail stores," in *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*, 2018, pp. 323–336.
- [14] A. Ghose, B. Li, R. Li, and K. Xu, "Real-time purchase prediction using retail video analytics," in *ICIS 2022 Proceedings*, vol. 6, 2022, pp. 1–17. [Online]. Available: [https://aisel.aisnet.org/icis2022/data\\_analytics/data\\_analytics/6](https://aisel.aisnet.org/icis2022/data_analytics/data_analytics/6)
- [15] M. Martini, M. Paolanti, and E. Frontoni, "Open-world person re-identification with RGBD camera in top-view configuration for retail applications," *IEEE Access*, vol. 8, pp. 67756–67765, 2020.
- [16] A. Van Rooijen, H. Bouma, R. Pruijm, J. Baan *et al.*, "Anonymized person re-identification in surveillance cameras," in *Counterterrorism, Crime Fighting, Forensics, and Surveillance Technologies IV*, vol. 11542. SPIE, 2020, pp. 63–67.
- [17] K. Hashemifar, F. Florez-Revuelta, G. Lacey *et al.*, "A fallen person detector with a privacy-preserving edge-AI camera," pp. 262–269, 2023.
- [18] R. Pietrini, "Deep understanding of shopper behaviours and interactions in intelligent retail environment," 2019, doctoral thesis, Università Politecnica delle Marche, Ancona, Italy. [Online]. Available: <http://hdl.handle.net/11566/274602>
- [19] L. Lee *et al.*, "The emotional shopper: Assessing the effectiveness of retail therapy," *Foundations and Trends® in Marketing*, vol. 8, no. 2, pp. 69–145, 2015.
- [20] T. Yamashiro, Y. Honma, R. Togo, T. Ogawa, and M. Haseyama, "Customer interest estimation method in real store using re-identification and 3D posture estimation models," in *International Workshop on Advanced Imaging Technology (IWAIT) 2023*, vol. 12592. SPIE, 2023, pp. 42–47.
- [21] F. Becattini, G. Becchi, A. Ferracani, A. D. Bimbo, L. L. Presti, G. Mazzola *et al.*, "I-mall an effective framework for personalized visits: Improving the customer experience in stores," in *Proceedings of the 1st Workshop on Multimedia Computing towards Fashion Recommendation*, 2022, pp. 11–19.
- [22] (2025) Luxonis. OAK-D lite. [Online]. Available: <https://docs.luxonis.com/hardware/products/OAK-D%20Lite> (Accessed 2025-01-11).
- [23] (2022) Luxonis. OAK-D lite datasheet. [Online]. Available: [https://github.com/luxonis/depthai-hardware/tree/master/DM9095\\_OAK-D-LITE\\_DepthAI\\_USB3C/Datasheet](https://github.com/luxonis/depthai-hardware/tree/master/DM9095_OAK-D-LITE_DepthAI_USB3C/Datasheet) (Accessed 2025-01-11).
- [24] S. Zafeiriou, C. Zhang, and Z. Zhang, "A survey on face detection in the wild: Past, present and future," *Computer Vision and Image Understanding*, vol. 138, pp. 1–24, 2015.
- [25] Y. Feng, S. Yu, H. Peng, Y.-R. Li, and J. Zhang, "Detect faces efficiently: A survey and evaluations," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 4, no. 1, pp. 1–18, 2022.
- [26] Z. Zou, K. Chen, Z. Shi, Y. Guo, and J. Ye, "Object detection in 20 years: A survey," *Proceedings of the IEEE*, vol. 111, no. 3, pp. 257–276, 2023.
- [27] X. Wang, J. Peng, S. Zhang, B. Chen, Y. Wang, and Y. Guo, "A survey of face recognition," *arXiv preprint arXiv:2212.13038*, 2022.
- [28] Y. Qian, J. Barthelemy, E. Karupiah, and P. Perez, "Identifying re-identification challenges: Past, current and future trends," *SN Computer Science*, vol. 5, no. 7, p. 937, 2024.
- [29] A. Zahra, N. Perwaiz, M. Shahzad, and M. M. Fraz, "Person re-identification: A retrospective on domain specific open challenges and future trends," *Pattern Recognition*, vol. 142, p. 109669, 2023.
- [30] E. Yaghoubi, A. Kumar, and H. Proença, "SSS-PR: A short survey of surveys in person re-identification," *Pattern Recognition Letters*, vol. 143, pp. 50–57, 2021.
- [31] Z. Ming, M. Zhu, X. Wang, J. Zhu, J. Cheng, C. Gao, Y. Yang, and X. Wei, "Deep learning-based person re-identification methods: A survey and outlook of recent works," *Image and Vision Computing*, vol. 119, p. 104394, 2022.
- [32] L. Zheng, H. Zhang, S. Sun, M. Chandraker, Y. Yang, and Q. Tian, "Person re-identification in the wild," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1367–1376.
- [33] T. Xiao, S. Li, B. Wang, L. Lin, and X. Wang, "Joint detection and identification feature learning for person search," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 3415–3424.
- [34] B. Alexe, "Objectness and its use for visual learning," 2013, doctoral thesis, ETH Zurich, Switzerland. [Online]. Available: <https://doi.org/10.3929/ethz-a-009906832>
- [35] Y. Cai, L. Wen, L. Zhang, D. Du, and W. Wang, "Rethinking object detection in retail stores," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 2, 2021, pp. 947–954.
- [36] S. Li and R. Hishiyama, "An indoor people counting and tracking system using mmWave sensor and sub-sensors," *IFAC-PapersOnLine*, vol. 56, no. 2, pp. 7096–7101, 2023.
- [37] C. A. Luna, C. Losada-Gutiérrez, D. Fuentes-Jimenez, and M. Mazo, "People re-identification using depth and intensity information from an overhead camera," *Expert Systems with Applications*, vol. 182, p. 115287, 2021.
- [38] A. Yadav and D. K. Vishwakarma, "Deep learning algorithms for person re-identification: State-of-the-art and research challenges," *Multimedia Tools and Applications*, vol. 83, no. 8, pp. 22005–22054, 2024.
- [39] X. Lin, J. Li, Z. Ma, H. Li, S. Li, K. Xu, G. Lu, and D. Zhang, "Learning modal-invariant and temporal-memory for video-based visible-infrared person re-identification," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 20973–20982.

[40] Q. Li, C. Zhang, Q. Hu, H. Fu, and P. Zhu, “Confidence-aware fusion using Dempster-Shafer theory for multispectral pedestrian detection,” *IEEE Transactions on Multimedia*, vol. 25, pp. 3420–3431, 2022.

[41] N. Huang, J. Liu, Y. Miao, Q. Zhang, and J. Han, “Deep learning for visible-infrared cross-modality person re-identification: A comprehensive review,” *Information Fusion*, vol. 91, pp. 396–411, 2023.

[42] G. Zou, G. Fu, X. Peng, Y. Liu, M. Gao, and Z. Liu, “Person re-identification based on metric learning: A survey,” *multimedia tools and applications*, vol. 80, no. 17, pp. 26 855–26 888, 2021.

[43] N. Wojke and A. Bewley, “Deep cosine metric learning for person re-identification,” in *IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2018, pp. 748–756.

[44] M. Ye, J. Shen, G. Lin, T. Xiang, L. Shao, and S. C. Hoi, “Deep learning for person re-identification: A survey and outlook,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 6, pp. 2872–2893, 2021.

[45] J. Deng, D. Xu, W. Li, and L. Duan, “Harmonious teacher for cross-domain object detection,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 23 829–23 838.

[46] C.-T. Liu, J.-C. Chen, C.-S. Chen, and S.-Y. Chien, “Video-based person re-identification without bells and whistles,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 1491–1500.

[47] X. Zang, G. Li, and W. Gao, “Multidirection and multiscale pyramid in transformer for video-based pedestrian retrieval,” *IEEE Transactions on Industrial Informatics*, vol. 18, no. 12, pp. 8776–8785, 2022.

[48] L. Wu, Y. Wang, J. Gao, and X. Li, “Where-and-when to look: Deep siamese attention networks for video-based person re-identification,” *IEEE Transactions on Multimedia*, vol. 21, no. 6, pp. 1412–1424, 2018.

[49] J. Jiao, W.-S. Zheng, A. Wu, X. Zhu, and S. Gong, “Deep low-resolution person re-identification,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.

[50] L. Zhang, Y. Xu, L. Zhao, and F. Qin, “Resolution independent person re-identification network,” *IET Computer Vision*, 2022.

[51] P. Xu and X. Zhu, “Deepchange: A long-term person re-identification benchmark with clothes change,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 11 196–11 205.

[52] G. Zhang, Z. Luo, Y. Chen, Y. Zheng, and W. Lin, “Illumination unification for person re-identification,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 10, pp. 6766–6777, 2022.

[53] F. Liu, M. Kim, Z. Gu, A. Jain, and X. Liu, “Learning clothing and pose invariant 3d shape representation for long-term person re-identification,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 19 617–19 626.

[54] A. Sepas-Moghaddam and A. Etemad, “Deep gait recognition: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 1, pp. 264–284, 2022.

[55] W. Li, S. Hou, C. Zhang, C. Cao, X. Liu, Y. Huang, and Y. Zhao, “An in-depth exploration of person re-identification and gait recognition in cloth-changing conditions,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 13 824–13 833.

[56] L. Tan, P. Dai, R. Ji, and Y. Wu, “Dynamic prototype mask for occluded person re-identification,” in *Proceedings of the 30th ACM international conference on multimedia*, 2022, pp. 531–540.

[57] N. Dong, S. Yan, H. Tang, J. Tang, and L. Zhang, “Multi-view information integration and propagation for occluded person re-identification,” *Information Fusion*, vol. 104, p. 102201, 2024.

[58] L. Tan, J. Xia, W. Liu, P. Dai, Y. Wu, and L. Cao, “Occluded person re-identification via saliency-guided patch transfer,” in *Proceedings of the AAAI conference on artificial intelligence*, vol. 38, no. 5, 2024, pp. 5070–5078.

[59] E. Ning, Y. Wang, C. Wang, H. Zhang, and X. Ning, “Enhancement, integration, expansion: Activating representation of detailed features for occluded person re-identification,” *Neural Networks*, vol. 169, pp. 532–541, 2024.

[60] G. Ding, S. Khan, Z. Tang, and F. Porikli, “Feature mask network for person re-identification,” *Pattern Recognition Letters*, vol. 137, pp. 91–98, 2020, learning and Recognition for Assistive Computer Vision.

[61] S. Agarwal, A. Awan, and D. Roth, “Learning to detect objects in images via a sparse, part-based representation,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 26, no. 11, pp. 1475–1490, 2004.

[62] X.-S. Wei, Y.-Z. Song, O. Mac Aodha, J. Wu, Y. Peng, J. Tang, J. Yang, and S. Belongie, “Fine-grained image analysis with deep learning: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 12, pp. 8927–8948, 2021.

[63] V. Somers, C. De Vleeschouwer, and A. Alahi, “Body part-based representation learning for occluded person re-identification,” in *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 2023, pp. 1613–1623.

[64] B. X. Nguyen, B. D. Nguyen, T. Do, E. Tjiputra, Q. D. Tran, and A. Nguyen, “Graph-based person signature for person re-identifications,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 3492–3501.

[65] Y. Zhai, X. Guo, Y. Lu, and H. Li, “In defense of the classification loss for person re-identification,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2019, pp. 1526–1535.

[66] S. Ding, L. Lin, G. Wang, and H. Chao, “Deep feature learning with relative distance comparison for person re-identification,” *Pattern Recognition*, vol. 48, no. 10, pp. 2993–3003, 2015.

[67] A. Hermans, L. Beyer, and B. Leibe, “In defense of the triplet loss for person re-identification,” *arXiv preprint arXiv:1703.07737*, 2017.

[68] E. P. Ijjina, R. Medipelly, S. K. Beerukuri, S. Vinnakota, and V. C. Nelakurthi, “Person re-identification using vision transformer and centroid triplet loss,” *Multimedia Tools and Applications*, vol. 83, no. 29, pp. 73 777–73 788, 2024.

[69] Intel. (2025) OpenVINO™ Toolkit (Open Visual Inference and Neural Network Optimization). [Online]. Available: <https://docs.openvino.ai/2025/index.html> (Accessed 2025-01-11).

[70] (2025) Open Neural Network Exchange - ONNX. [Online]. Available: <https://onnx.ai> (Accessed 2025-01-11).

[71] (2025) Luxonis. DepthAI Python API utilities, examples, and tutorials. [Online]. Available: <https://github.com/luxonis/depthai> (Accessed 2025-01-11).

[72] G. Bradski *et al.* (2025) OpenCV: Open Computer Vision library. [Online]. Available: <https://opencv.org/> (Accessed 2025-01-11).

[73] (2025) Luxonis. DepthAI C++ library. [Online]. Available: <https://github.com/luxonis/depthai-core> (Accessed 2025-01-11).

[74] (2022) OpenVINO person detection model. [Online]. Available: [https://github.com/openvino toolkit/open\\_model\\_zoo/tree/master/models/intel/person-detection-retail-0013](https://github.com/openvino toolkit/open_model_zoo/tree/master/models/intel/person-detection-retail-0013) (Accessed 2025-01-11).

[75] (2022) OpenVINO person ReID model. [Online]. Available: [https://github.com/openvino toolkit/open\\_model\\_zoo/tree/master/models/intel/person-reidentification-retail-0288](https://github.com/openvino toolkit/open_model_zoo/tree/master/models/intel/person-reidentification-retail-0288) (Accessed 2025-01-11).

[76] (2023) Depthai experiment: Pedestrian reidentification. [Online]. Available: <https://github.com/luxonis/depthai-experiments/tree/master/gen2-pedestrian-reidentification#pedestrian-reidentification> (Accessed 2025-01-11).

[77] L. He, X. Liao, W. Liu, X. Liu, P. Cheng, and T. Mei, “Fastreid: A Pytorch toolbox for general instance re-identification,” in *Proceedings of the 31st ACM International Conference on Multimedia*, 2023, pp. 9664–9667.

[78] S. Yıldız and A. N. Kasim, “ENTIRE-ID: An extensive and diverse dataset for person re-identification,” in *18th International Conference on Automatic Face and Gesture Recognition*. IEEE, 2024, pp. 1–5.

[79] L. Zheng, Z. Bie, Y. Sun, J. Wang, C. Su, S. Wang, and Q. Tian, “MARS: A video benchmark for large-scale person re-identification,” in *14th European Conference on Computer Vision, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VI 14*. Springer, 2016, pp. 868–884.

[80] Z. Cui, J. Zhou, X. Wang, M. Zhu, and Y. Peng, “Learning continual compatible representation for re-indexing free lifelong person re-identification,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 16 614–16 623.

[81] P. Chophuk, P. Boonmee, S. Jiaraswan, S. Jearaswan, and P. Bookprakong, “Theft detection by patterns of walking behavior using motion-based artificial intelligence,” in *International Workshop on Advanced Imaging Technology*, vol. 12592. SPIE, 2023, pp. 140–145.

[82] G. Abbattista, M. Chimenti, V. Dentamaro, P. Giglio, D. Impedovo, G. Pirlo, and G. Rosato, “A biometric-based system for unsupervised anomalous behaviour detection at the pawn shop,” *Cyber-Physical Systems*, vol. 9, no. 4, pp. 338–356, 2023.

[83] H. Chen, Y. Wang, B. Lagadec, A. Dantcheva, and F. Bremond, “Learning invariance from generated variance for unsupervised person re-identification,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 6, pp. 7494–7508, 2022.