

Computer Vision Applications for Automated Industrial Inspection

Atsushi Yamashita

Professor, Chukyo University, Japan

Abstract

Computer vision has emerged as a significant enabling technology for automated industrial inspection by providing machines with the ability to capture, process, and interpret visual information from manufacturing environments. Conventional inspection practices frequently depend on manual observation, sampling procedures, and human judgment, which may introduce variability, fatigue-related errors, and limitations in inspection speed. Computer vision systems address these challenges by integrating industrial cameras, illumination systems, image-processing algorithms, machine learning, and artificial intelligence to identify defects and evaluate product characteristics automatically. This study examines computer vision applications for automated industrial inspection, focusing on image acquisition, preprocessing, feature extraction, defect detection, classification, dimensional inspection, surface analysis, and real-time quality control. A structured analytical methodology based on secondary technical literature is employed to evaluate the technological capabilities, operational benefits, and implementation challenges of computer vision systems.

The analysis demonstrates that automated visual inspection can improve inspection consistency, detection speed, production traceability, and quality assurance while reducing dependence on repetitive manual inspection. Traditional image-processing approaches remain effective for controlled inspection environments, whereas deep learning techniques provide greater flexibility for complex defect patterns and variable product conditions. The integration of computer vision with robotics, edge computing, Industrial Internet of Things platforms, and intelligent manufacturing systems further enables real-time inspection and automated responses. Nevertheless, challenges related to lighting variability, image quality, limited defect datasets, model generalization, computational requirements, system integration, and false-positive or false-negative classifications remain important. Future research should emphasize explainable artificial intelligence, synthetic training data, multimodal inspection, edge-based inference, adaptive vision systems, and digital-twin integration. The study concludes that computer vision can significantly strengthen industrial quality management when imaging hardware, algorithms, machine learning models, and production workflows are designed as an integrated inspection system.

Keywords: Computer Vision, Automated Inspection, Industrial Quality Control, Machine Vision, Defect Detection, Image Processing, Deep Learning, Smart Manufacturing.

1. Introduction

Industrial manufacturing increasingly requires inspection systems that can provide rapid, consistent, and accurate evaluation of products and production processes. Conventional inspection frequently depends on human operators who visually examine components for defects, dimensional deviations, surface irregularities, assembly errors, or other quality problems. Although experienced inspectors can provide

valuable judgments, manual inspection may be affected by fatigue, repetitive workloads, lighting conditions, subjective interpretation, and production speed. The growing complexity and volume of manufacturing operations have therefore increased interest in automated inspection technologies capable of operating continuously and consistently.

Computer vision provides an important technological foundation for automated industrial inspection. A typical computer vision system combines cameras, illumination, image-acquisition hardware, processing units, and analytical algorithms to transform visual information into quality-related decisions. The system can capture images of components as they move through a production line and analyze features such as shape, color, texture, dimensions, surface conditions, and assembly characteristics. Automated decisions can then be communicated to production-control systems or robotic mechanisms.

Traditional machine vision approaches commonly rely on image preprocessing, segmentation, edge detection, feature extraction, and rule-based classification. These methods can perform effectively when inspection environments are highly controlled and defect characteristics are known in advance. However, industrial environments can involve changing lighting, complex geometries, diverse product variants, and subtle defect patterns. These conditions have encouraged the adoption of machine learning and deep learning methods capable of learning visual representations directly from inspection data.

Deep convolutional neural networks have substantially expanded the capabilities of automated visual inspection. Such models can classify components, detect defects, segment damaged regions, and identify complex patterns that may be difficult to define through manually designed image-processing rules. More advanced approaches can support object detection and pixel-level segmentation, enabling inspection systems to identify both the location and characteristics of defects. However, successful deployment depends heavily on the availability and quality of representative training data.

Computer vision is increasingly being integrated with broader smart manufacturing technologies. Industrial cameras can be connected to Industrial Internet of Things systems, while edge-computing platforms can process images near production equipment. Robotic systems can use visual information to perform sorting, positioning, rejection, or corrective actions. Digital twins and production analytics can further connect inspection information with manufacturing processes. This integration can transform inspection from an isolated quality-control activity into a continuous component of intelligent production.

Despite its advantages, automated industrial vision remains subject to several technical and operational challenges. Variations in illumination, reflections, camera positioning, lens characteristics, product orientation, and environmental conditions can influence image quality. Machine-learning models may also perform poorly when exposed to defect types or production conditions not represented in their training data. This study therefore examines computer vision applications for automated industrial inspection, with emphasis on system architecture, image-processing and AI techniques, industrial applications, benefits, challenges, and future research directions.

2. Objectives

The primary objective of this study is to examine the application of computer vision technologies for automated industrial inspection. The study focuses on how imaging systems, image-processing methods, machine learning, and deep learning can improve defect detection, dimensional verification, surface inspection, and real-time quality assurance.

2.1 Examine Computer Vision Technologies

The first objective is to examine the major technical components used in industrial computer vision systems. These include cameras, lighting systems, image-acquisition devices, preprocessing techniques, feature extraction algorithms, classification methods, and deep learning architectures.

The study also evaluates how these technologies contribute to automated inspection. Particular attention is given to the relationship between image quality, algorithm selection, computational capability, and inspection accuracy. Understanding these factors is essential for designing reliable industrial vision systems.

2.2 Evaluate Industrial Inspection Applications

The second objective is to evaluate the application of computer vision across industrial inspection tasks. These include surface-defect detection, dimensional inspection, assembly verification, component classification, and quality monitoring.

The study further considers the operational benefits of automated vision, including inspection speed, consistency, traceability, reduced manual workload, and integration with automated production systems. These factors determine the practical value of computer vision in modern manufacturing environments.

3. Literature Review

3.1 Development of Machine Vision in Manufacturing

Machine vision has traditionally been used in manufacturing to automate inspection tasks that require consistent visual evaluation. Early systems depended heavily on controlled lighting, predefined geometric measurements, thresholding, edge detection, and manually engineered features. These systems were particularly effective for repetitive products where the appearance of acceptable and defective components could be clearly defined.

The development of high-resolution cameras, faster processors, and improved image-processing algorithms expanded machine vision capabilities. Modern systems can capture detailed images at high production speeds and analyze multiple characteristics simultaneously. Machine vision has consequently become relevant to quality inspection, production monitoring, robotic guidance, and process control.

3.2 Image Processing and Defect Detection

Image processing represents a fundamental stage in automated inspection. Raw images may contain noise, uneven illumination, reflections, or irrelevant background information. Preprocessing techniques such as filtering, normalization, contrast enhancement, and geometric correction can improve the quality of subsequent analysis.

Segmentation and feature extraction techniques are then used to identify relevant product regions. Shape, texture, color, edges, contours, and statistical properties can provide information about component quality. Traditional approaches remain useful in highly controlled environments because they can be computationally efficient and relatively easy to interpret. However, their performance may decline when defect appearance changes significantly.

3.3 Deep Learning for Industrial Inspection

Deep learning has significantly changed automated visual inspection by enabling systems to learn complex visual patterns directly from training data. Convolutional neural networks have been widely

applied to industrial image classification and defect recognition. Object-detection architectures can locate multiple defects within an image, while semantic and instance segmentation methods can identify defect boundaries.

Deep learning can outperform manually engineered approaches in complex inspection environments, but it requires representative training datasets. Industrial defects may be rare, making it difficult to collect sufficient examples of every possible failure condition. Data augmentation, transfer learning, synthetic image generation, and anomaly-detection methods have therefore become important research areas.

3.4 Vision-Based Quality Control and Smart Manufacturing

Computer vision increasingly functions as part of integrated manufacturing systems rather than as a standalone inspection tool. Vision systems can communicate inspection results to manufacturing execution systems, robots, programmable controllers, and IIoT platforms. This allows quality information to influence production decisions in real time.

Integration can also support traceability. Each inspected component can be associated with inspection images, defect classifications, timestamps, production conditions, and machine information. Such records can help manufacturers identify recurring problems and connect quality outcomes with specific production parameters. Consequently, computer vision can contribute to both immediate defect detection and long-term process improvement.

4. Methodology

4.1 Research Design

This study adopts a structured qualitative and analytical research design to examine computer vision applications in automated industrial inspection. Secondary technical literature was reviewed to identify established approaches in machine vision, image processing, machine learning, deep learning, defect detection, and industrial quality control. The reviewed information was organized into technological and application-oriented themes.

The research compares conventional and AI-based inspection approaches according to their ability to detect defects, process images in real time, adapt to product variation, and integrate with industrial systems. This comparative approach provides an understanding of the conditions under which different computer vision methods are appropriate for manufacturing inspection.

4.2 Computer Vision Technology Assessment

The technological assessment covers image acquisition, illumination, preprocessing, feature extraction, classification, object detection, and image segmentation. Camera resolution, acquisition speed, image consistency, and environmental conditions are considered because they directly influence inspection quality.

Machine-learning and deep-learning approaches are evaluated according to their ability to identify complex patterns and generalize across manufacturing conditions. The assessment also considers computational requirements and the importance of edge or industrial computing for real-time inference.

4.3 Analytical Evaluation Criteria

The reviewed computer vision approaches were evaluated according to the following criteria:

- Image acquisition quality
- Detection accuracy
- Processing speed
- Defect localization capability
- Adaptability to product variation
- Dataset requirements
- Computational requirements
- False-positive and false-negative risks
- Integration with production systems
- Inspection traceability

These criteria provide a structured framework for comparing automated inspection approaches and determining their practical industrial relevance.

4.4 Data Synthesis

The final stage involved synthesizing the reviewed findings into major themes covering inspection technology, defect detection, dimensional assessment, AI-based classification, production integration, and quality management. The analysis emphasized the relationship between technological capabilities and industrial inspection requirements.

Emerging approaches involving edge computing, deep learning, synthetic data, anomaly detection, and multimodal sensing were also considered. The synthesis provides an integrated perspective on the current capabilities and future development of computer vision-based industrial inspection.

5. Analysis

5.1 Computer Vision Methods for Industrial Inspection

Industrial inspection systems employ different analytical methods depending on product characteristics and defect complexity. Traditional image-processing techniques are suitable when defects have predictable visual characteristics, whereas machine learning and deep learning provide greater adaptability for complex visual patterns.

Table 1: Major Computer Vision Approaches for Automated Industrial Inspection

Approach	Primary Function	Major Strength	Typical Application
Thresholding	Region separation	Simple and fast	Surface and color inspection
Edge Detection	Boundary identification	Efficient geometric analysis	Dimensional inspection
Feature-Based Classification	Pattern recognition	Interpretable features	Component classification

Approach	Primary Function	Major Strength	Typical Application
Machine Learning	Defect classification	Flexible decision-making	Quality classification
CNN-Based Classification	Image recognition	High pattern-learning capability	Defect identification
Object Detection	Defect localization	Identifies multiple regions	Production-line inspection
Image Segmentation	Pixel-level analysis	Precise defect boundaries	Surface-defect assessment
Anomaly Detection	Unusual-pattern identification	Useful with limited defect samples	Unknown defect detection

Interpretation: Table 1 shows that automated inspection can employ a spectrum of methods ranging from simple rule-based image processing to advanced deep learning. Traditional approaches provide speed and interpretability under controlled conditions, while deep learning is more suitable for complex visual patterns and variable defect characteristics.

Anomaly detection is particularly valuable when defective samples are limited because models can learn the appearance of normal products and identify deviations. The appropriate method should therefore be selected according to defect complexity, data availability, inspection speed, and environmental variability.

5.2 Inspection Performance and Operational Benefits

Computer vision can improve several dimensions of industrial inspection. Automated systems can operate continuously at production-line speeds and apply consistent inspection criteria to every component. Digital inspection records also support traceability and quality analysis.

Table 2: Operational Contributions of Computer Vision-Based Inspection

Inspection Area	Computer Vision Capability	Operational Contribution
Surface Inspection	Detects scratches, cracks, stains, and irregularities	Improved defect identification
Dimensional Inspection	Measures geometric features	Consistent tolerance verification
Assembly Verification	Checks component presence and position	Reduced assembly errors
Product Classification	Identifies product variants	Automated sorting

Inspection Area	Computer Vision Capability	Operational Contribution
Label Inspection	Verifies printed information	Improved traceability
Weld Inspection	Detects surface abnormalities	Enhanced structural quality control
Packaging Inspection	Identifies packaging defects	Reduced defective shipments
Process Monitoring	Continuously observes production	Faster quality intervention

Interpretation: Table 2 demonstrates that computer vision contributes to industrial inspection beyond simple defect identification. Automated dimensional measurement, assembly verification, product classification, and packaging inspection can reduce repetitive manual activities while improving consistency.

Continuous inspection also enables manufacturers to identify process problems earlier, potentially reducing the quantity of defective products produced. When inspection information is stored digitally, manufacturers can analyze historical patterns and associate quality problems with production conditions, creating opportunities for systematic process improvement.

5.3 Application Areas in Industrial Manufacturing

Computer vision is widely applicable to automotive, electronics, metalworking, food processing, pharmaceutical manufacturing, packaging, and general industrial production. In automotive manufacturing, vision systems can inspect body panels, welds, assemblies, and surface finishes. Electronics manufacturers can use high-resolution vision for component placement and solder-joint inspection.

In metalworking environments, vision systems can identify scratches, cracks, corrosion, dimensional deviations, and machining abnormalities. Packaging industries can use vision systems to verify labels, seals, product placement, and packaging integrity. The flexibility of modern vision systems allows the same basic inspection architecture to be adapted to different industrial products through changes in imaging configurations and analytical models.

6. Discussion

6.1 Accuracy and Inspection Consistency

Automated computer vision can provide highly consistent inspection because the same programmed criteria or trained model can be applied repeatedly to large numbers of components. This reduces variability associated with subjective human judgments and fatigue. Continuous operation is especially valuable for high-volume production environments where manual inspection may be difficult to sustain at required speeds.

However, inspection accuracy depends on the complete system rather than solely on the algorithm. Camera quality, illumination, positioning, focus, image resolution, and preprocessing all influence the information available to the model. A sophisticated AI model cannot compensate completely for poor image acquisition. Consequently, industrial vision development should treat hardware and software as an integrated inspection architecture.

6.2 Deep Learning and Adaptive Inspection

Deep learning has increased the ability of inspection systems to identify complex defects that are difficult to describe through fixed rules. Convolutional networks can learn visual representations from training images, while object-detection and segmentation models can identify the location and extent of defects. These capabilities are especially useful for products with variable appearances.

Adaptive inspection remains challenging because industrial conditions can change over time. New materials, machine settings, lighting conditions, and product variants may produce images that differ from the original training dataset. Continuous model validation, dataset expansion, transfer learning, and anomaly-detection methods can improve adaptability. Human expertise remains valuable for validating model decisions and handling unusual inspection cases.

6.3 Integration with Automated Manufacturing

The integration of computer vision with robotics and industrial control can convert inspection results into immediate physical actions. For example, a vision system can identify a defective product and communicate its location to a robotic mechanism that removes it from the production line. Such integration reduces the time between detection and response.

Computer vision can also support broader manufacturing analytics. Inspection outcomes can be linked with machine conditions, production parameters, and material information to identify the causes of recurring defects.

7. Challenges

7.1 Lighting and Image Acquisition

Lighting variation is one of the most important challenges in industrial computer vision. Reflections, shadows, glare, surface texture, and changing ambient conditions can alter the appearance of components. A defect that is clearly visible under one lighting condition may become difficult to detect under another.

Effective inspection therefore requires careful camera and illumination design. Industrial systems may use controlled lighting, multiple cameras, polarized illumination, structured light, or specialized imaging techniques.

7.2 Data Availability and Model Generalization

Deep-learning inspection systems require representative datasets containing sufficient examples of normal and defective products. However, industrial defects are often rare, meaning that obtaining large numbers of accurately labeled defective samples can be difficult. Class imbalance can consequently affect model training and detection performance.

Models may also fail to generalize when production conditions change. A system trained on one product configuration or manufacturing environment may perform poorly in another. Transfer learning, data augmentation, synthetic data generation, and anomaly-detection approaches can help address these limitations, but robust validation remains necessary.

7.3 Integration and Computational Requirements

Industrial vision systems must often operate under strict production-cycle constraints. High-resolution images and complex deep-learning models can require substantial computational resources. Processing delays may reduce the practicality of the system for high-speed manufacturing.

Integration with programmable logic controllers, manufacturing execution systems, robotic platforms, and existing inspection infrastructure can also be technically challenging. Communication interfaces, data formats, timing requirements, and cybersecurity must be considered. Successful deployment therefore requires cooperation among vision specialists, manufacturing engineers, automation experts, and IT professionals.

8. Future Research

8.1 Explainable and Trustworthy Industrial AI

Future research should focus on developing explainable computer vision systems that can provide understandable reasons for inspection decisions. This is particularly important in manufacturing environments where engineers may need to verify why a component has been classified as defective. Explainable models could increase operator confidence and facilitate troubleshooting.

Research should also address model robustness and reliability. Inspection systems should be capable of identifying uncertainty and recognizing situations where their predictions may be unreliable. Combining confidence estimation with human review could improve the safety and reliability of AI-based inspection.

8.2 Synthetic Data and Anomaly Detection

Synthetic data generation offers an important opportunity for addressing limited industrial defect datasets. Computer-generated defects and simulated manufacturing variations can supplement real inspection images and improve model training. Research should determine how closely synthetic images must represent real defects to provide useful training information.

Anomaly detection should also receive greater attention because industrial systems frequently have abundant examples of normal products but relatively few examples of every possible defect. Models trained primarily on normal production patterns may detect unexpected deviations without requiring extensive defect labeling. This approach could improve the ability of inspection systems to identify previously unseen defect types.

8.3 Edge-Based and Multimodal Inspection

Future computer vision systems are likely to increasingly use edge computing for real-time industrial inference. Processing images close to production equipment can reduce latency, network dependence, and the need to transmit large volumes of visual data to centralized servers. Research should optimize AI models for industrial edge hardware while maintaining acceptable detection performance.

Multimodal inspection represents another promising direction. Combining RGB images with thermal imaging, depth information, hyperspectral sensing, acoustic measurements, or 3D scanning could provide information that cannot be obtained from conventional images alone. Such systems may improve detection of hidden, subtle, or multidimensional defects.

9. Limitations

This study is primarily based on secondary technical literature and analytical evaluation rather than experimental testing of a specific industrial computer vision system. Therefore, the findings represent generalized technological patterns rather than performance measurements from a particular manufacturing facility. Actual inspection accuracy may vary considerably according to camera

configuration, lighting conditions, product geometry, defect characteristics, dataset quality, and computational infrastructure.

Another limitation is the rapid development of computer vision and artificial intelligence technologies. New architectures, imaging technologies, training methods, and edge-computing solutions continue to emerge. Consequently, some approaches may develop beyond the scope of the reviewed literature. Future experimental studies should evaluate complete inspection systems under standardized industrial conditions and compare performance using consistent accuracy, latency, reliability, and economic metrics.

10. Conclusion

Computer vision has become an important technology for automated industrial inspection because it enables manufacturing systems to capture and interpret visual information with high speed and consistency. The technology can support surface inspection, dimensional verification, assembly checking, product classification, packaging inspection, and continuous quality monitoring. Traditional image-processing techniques remain effective for controlled applications, while machine learning and deep learning provide greater flexibility for complex and variable inspection conditions. The analysis demonstrates that the effectiveness of computer vision depends on the integration of imaging hardware, illumination, preprocessing, analytical models, and production-control systems.

Deep learning has expanded the potential of automated inspection by enabling systems to learn complex visual characteristics from large datasets. Nevertheless, data availability, defect rarity, environmental variation, and model generalization remain important challenges. Industrial organizations must therefore develop representative datasets, establish reliable validation procedures, and continuously monitor model performance. Hardware design is equally important because poor lighting, reflections, camera positioning, or inadequate resolution can compromise the information supplied to even advanced AI models. Computer vision should consequently be implemented as an integrated engineering system rather than simply as a software solution.

The future of automated industrial inspection is likely to involve explainable AI, synthetic data, anomaly detection, edge computing, multimodal sensing, robotics, and digital manufacturing integration. These developments can enable inspection systems to operate with greater adaptability and provide faster responses to production abnormalities. When connected with manufacturing analytics and process-control systems, inspection data can also support root-cause analysis and continuous production improvement. Overall, computer vision offers substantial potential for creating more reliable, automated, and intelligent industrial quality-management systems, provided that technological development is accompanied by robust data management, system validation, cybersecurity, and application-specific engineering design.

References

1. Bergmann, P., Batzner, K., Fauser, M., Sattlegger, D., & Steger, C. (2021). The MVTec anomaly detection dataset: A comprehensive real-world dataset for unsupervised anomaly detection. *International Journal of Computer Vision*, 129, 1038–1059. <https://doi.org/10.1007/s11263-020-01400-4>
2. Bergmann, P., Fauser, M., Sattlegger, D., & Steger, C. (2022). Beyond dents and scratches: Logical constraints in unsupervised anomaly detection and localization. *International Journal of Computer Vision*, 130, 947–969. <https://doi.org/10.1007/s11263-022-01578-9>

3. Bergmann, P., Jin, X., Sattlegger, D., & Steger, C. (2022). The MVTec 3D-AD dataset for unsupervised 3D anomaly detection and localization. In *Proceedings of the 17th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (Vol. 5, pp. 202–213)*.
4. Czimmermann, T., Ciuti, G., Oprea, S., Nechifor, G., Rus, D., Mura, M. D., & others. (2020). Visual-based defect detection and classification approaches for industrial applications—A survey. *Sensors*, 20(5), Article 1459. <https://doi.org/10.3390/s20051459>
5. Dai, W., Mujeeb, A., Erdt, M., & Sourin, A. (2018). Towards automatic optical inspection of soldering defects. In *2018 International Conference on Cyberworlds (CW)* (pp. 375–382).
6. He, Y., Song, K., Meng, Q., & Yan, Y. (2020). An end-to-end steel surface defect detection approach via fusing multiple hierarchical features. *IEEE Transactions on Instrumentation and Measurement*, 69(4), 1493–1504.
7. Huang, Y., Qiu, C., & Yuan, K. (2020). Surface defect saliency of magnetic tile. *The Visual Computer*, 36, 85–96.
8. Li, L., Ota, K., & Dong, M. (2018). Deep learning for smart industry: Efficient manufacture inspection system with fog computing. *IEEE Transactions on Industrial Informatics*, 14(10), 4665–4673.
9. Liu, L., Ouyang, W., Wang, X., Fieguth, P., Chen, J., Liu, X., & Pietikäinen, M. (2020). Deep learning for generic object detection: A survey. *International Journal of Computer Vision*, 128, 261–318. <https://doi.org/10.1007/s11263-019-01247-4>
10. Pang, G., Shen, C., Cao, L., & van den Hengel, A. (2021). Deep learning for anomaly detection: A review. *ACM Computing Surveys*, 54(2), Article 38. <https://doi.org/10.1145/3439950>
11. Qin, J., Liu, Y., & Grosvenor, R. (2016). A categorical framework of manufacturing for Industry 4.0 and beyond. *Procedia CIRP*, 52, 173–178.
12. Ren, S., He, K., Girshick, R., & Sun, J. (2017). Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6), 1137–1149. <https://doi.org/10.1109/TPAMI.2016.2577031>
13. Roth, K., Pemula, L., Zepeda, J., Schölkopf, B., Brox, T., & Gehler, P. (2022). Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 14318–14328).
14. Tao, F., Qi, Q., Liu, A., & Kusiak, A. (2018). Data-driven smart manufacturing. *Journal of Manufacturing Systems*, 48, 157–169. <https://doi.org/10.1016/j.jmsy.2018.01.006>
15. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415. <https://doi.org/10.1109/TII.2018.2873186>
16. Wuest, T., Weimer, D., Irgens, C., & Thoben, K.-D. (2016). Machine learning in manufacturing: Advantages, challenges, and applications. *Production & Manufacturing Research*, 4(1), 23–45. <https://doi.org/10.1080/21693277.2016.1192517>
17. Zou, Y., Jeong, J., Pemula, L., Zhang, D., & Dabeer, O. (2022). SPot-the-Difference self-supervised pre-training for anomaly detection and segmentation. In *Computer Vision—ECCV 2022 Workshops*.
18. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 770–778).