



IJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** X **Month of publication:** October 2026

DOI: <https://doi.org/10.22214/ijraset.2026.85326>

www.ijraset.com

Call: ☎ 08813907089

E-mail ID: ijraset@gmail.com

Advanced Digital Image and Signal Processing Techniques for Real-Time Anomaly Detection

Monika Kailas Khairnar

Lecturer, Department of Electronics & Computer Engineering, P.E.S. Polytechnic, Chhatrapati Sambhaji Nagar, Maharashtra, India

Abstract: Real-time anomaly detection is a core requirement in industrial monitoring, surveillance, biomedical diagnostics and predictive maintenance, where unusual events must be identified within strict latency limits. This paper presents a unified framework that combines digital image processing (DIP) and digital signal processing (DSP) techniques for detecting anomalies in streaming visual and one-dimensional sensor data. The framework consists of modality-specific pre-processing, spatial and time-frequency feature extraction, feature-level fusion, unsupervised anomaly scoring and an adaptive decision threshold. Techniques reviewed and integrated include adaptive background modelling, optical flow, edge and texture analysis, the short-time Fourier transform, wavelet decomposition, statistical indicators, autoencoder reconstruction error and isolation-based scoring. Real-time design considerations such as sliding-window processing, computational complexity and embedded deployment are discussed, and an evaluation protocol based on precision, recall, F1-score, ROC-AUC and detection latency is defined. The framework is demonstrated on reproducible synthetic vibration-signal and surface-image data; the signal pipeline achieved an F1-score of 0.949 with a per-window processing time of about 2.0 ms, and the image pipeline achieved a recall of 1.000 at about 1.6 ms per frame on a general-purpose processor. A qualitative comparison of methods and a set of design guidelines are provided to help practitioners select a suitable technique for a given application and hardware budget.

Keywords: anomaly detection, digital image processing, digital signal processing, real-time systems, wavelet transform, feature fusion, autoencoder, adaptive threshold.

I. INTRODUCTION

An anomaly is an observation or pattern that deviates significantly from what is regarded as normal behaviour for a system [1]. Detecting such deviations quickly is critical in many fields. A surface crack on a production line, an irregular heartbeat, a developing bearing fault in a motor, or an intruder in a restricted area all demand an immediate response, and the cost of late detection is often high in terms of safety, money and downtime. Modern systems generate continuous streams of images, video and sensor signals at rates far beyond what human operators can examine, which motivates automated detection that runs in real time.

Digital image processing and digital signal processing provide the mathematical foundation for this task. DIP techniques extract spatial and temporal information from pixel data, such as edges, texture, motion and foreground regions [2]. DSP techniques characterise the temporal and spectral behaviour of one-dimensional signals through filtering, spectral estimation and multi-resolution analysis [3]. Each field has been widely applied on its own. In practice, however, many monitored assets are observed through more than one sensor. A rotating machine, for example, may be watched by a camera and instrumented with an accelerometer and a current sensor, and a joint analysis can reveal events that neither stream shows clearly alone.

Real-time operation adds constraints that are less prominent in offline analysis. The detector must process each incoming window or frame before the next arrives, must work within the memory and power limits of the target device, must keep the false-alarm rate low enough that operators trust it, and must adapt to slow changes in operating conditions. A method with excellent offline accuracy may therefore be unusable if it is too slow, too memory-hungry or too sensitive to drift.

The objectives of this paper are: (i) to review the principal DIP and DSP techniques used for anomaly detection; (ii) to propose a unified, modular processing pipeline suitable for real-time operation; (iii) to discuss implementation constraints on embedded and edge hardware; (iv) to define an evaluation methodology that allows fair comparison; and (v) to demonstrate the pipeline on reproducible synthetic data. The remainder of the paper is organised as follows. Section II reviews related work. Section III describes the proposed framework. Section IV discusses real-time implementation. Section V defines the evaluation methodology. Section VI describes the experimental setup, and Section VII presents and discusses the results. Sections VIII to X cover the comparative analysis, applications, challenges and conclusions.

II. LITERATURE REVIEW

A. Statistical and classical approaches

Early anomaly detection relied on statistical thresholds, control charts and distance-based methods, which assume a model of normal behaviour and flag observations that are improbable under it. Chandola et al. provide a comprehensive survey that classifies techniques into classification-based, nearest-neighbour, clustering, statistical and spectral categories, and discusses the nature of the input data, the type of anomaly (point, contextual or collective) and the availability of labels [1]. Because genuine anomalies are rare and diverse, labelled examples are seldom available for all fault types, which makes semi-supervised and unsupervised techniques particularly attractive. One-class support vector machines, for instance, estimate the support of the distribution of normal data and flag samples that fall outside it [10].

B. Image-based techniques

For video surveillance and visual inspection, adaptive Gaussian mixture background models remain a standard method of separating a normally static scene from unusual foreground activity, and they were designed explicitly for real-time tracking [5]. Optical flow methods estimate the apparent motion of pixels between frames and can expose abnormal speed, direction or crowd behaviour [13]. For surface-defect inspection, edge detectors such as the Canny operator [12] and local texture descriptors highlight scratches, cracks and blemishes against a regular background. More recently, convolutional autoencoders trained only on normal footage have been used to learn temporal regularity in video; frames or regions that the network reconstructs poorly are marked as anomalous [9].

C. Signal-based techniques

For one-dimensional signals such as vibration, current, acoustic emission and biomedical signals, the Fourier transform reveals the frequency content associated with specific faults, and the short-time Fourier transform (STFT) extends this to signals whose spectrum changes with time [3]. The wavelet transform provides multi-resolution analysis that is well suited to non-stationary signals and short transients, because it offers good time resolution at high frequencies and good frequency resolution at low frequencies [4], [11]. In addition, simple time-domain indicators such as root-mean-square (RMS) value, kurtosis and crest factor are inexpensive to compute and are widely used in condition monitoring, since impulsive faults raise kurtosis and crest factor well before the overall energy changes noticeably.

D. Machine-learning and deep-learning approaches

Isolation forests detect anomalies by measuring how easily a sample can be isolated using random partitions; anomalies require fewer splits, and the method has linear time complexity with a small memory footprint, which suits real-time use [6]. Autoencoders trained on normal data only have been used for non-linear dimensionality reduction and anomaly scoring, with reconstruction error acting as the score [8]. Pang et al. review deep anomaly detection and identify open challenges, including scarcity of labels, class imbalance, robustness to contamination of the training data, and interpretability [7]. Deep models can capture complex normal behaviour, but their computational cost must be weighed against the resources of the target platform.

E. Research gap

Most published work treats image-based and signal-based detection separately, and many studies report offline accuracy without reporting latency or processing cost. There is therefore a need for a clear, modular pipeline that combines both families of techniques, that is explicit about its real-time properties, and that is evaluated with metrics that reflect operational needs such as false-alarm rate and detection delay. This paper addresses that need.

III. PROPOSED FRAMEWORK

The proposed framework, shown in Fig. 1, has two parallel front-ends, one for images and one for one-dimensional signals, that feed a common fusion and decision back-end. It is intentionally modular so that each block can be replaced or omitted according to the application and the available computing resources.




Fig. 1. Block diagram of the proposed DIP–DSP anomaly detection pipeline.

A. Image pre-processing and feature extraction

Incoming frames are converted to grayscale or a suitable colour space, denoised using median or Gaussian filtering, and contrast-normalised, for example by histogram equalisation. A region of interest (ROI) is selected to reduce the number of pixels processed. Features are then computed. The foreground ratio from a background model indicates how much of the scene is unusual; optical-flow magnitude and direction histograms describe motion; edge density and local texture statistics describe surface condition; and, where resources allow, embeddings from a lightweight convolutional network provide a learned description. For the surface-inspection case, a simple and fast residual approach is also effective: the local background is estimated by heavy Gaussian smoothing and the absolute difference between the image and its background forms a residual map, in which defects appear as bright regions.

B. Signal pre-processing and feature extraction

The sensor signal $x[n]$ is band-pass filtered to the band of interest, detrended to remove drift, and segmented into overlapping windows of length N samples with hop size H . For each window the following features are computed: RMS value, kurtosis, crest factor, spectral centroid, spectral entropy, and the relative energies of the sub-bands obtained from a multi-level discrete wavelet transform. The RMS value of a window is given by

$$RMS = \sqrt{(1/N) \sum x^2[n]} \quad (1)$$

and the relative energy of the wavelet detail band at decomposition level j is

$$E_j = \sum_k |d_j[k]|^2 / \sum_i \sum_k |d_i[k]|^2 \quad (2)$$

where $d_j[k]$ denotes the detail coefficients at level j . The spectral entropy is computed from the normalised power spectral density $P_n(f)$ as $H_s = -\sum P_n(f) \ln P_n(f)$; a spectrum dominated by a few lines has low entropy, whereas a broadband disturbance raises it. Together, these features respond to changes in energy, impulsiveness and spectral distribution, which cover the most common signal manifestations of faults.

C. Feature-level fusion

Image and signal features are time-aligned using timestamps, normalised using the mean and standard deviation estimated on normal data (z-score), and concatenated into a single feature vector $f(t)$. Fusion at feature level allows the detector to learn joint patterns; for example, a rise in vibration that coincides with visible motion at the same location may be more significant than either event alone. Because sensors can fail or drop out, the design includes a fallback in which the detector continues with the remaining stream, using a separately trained model for each available combination of inputs.

D. Anomaly scoring

Three complementary scorers are suitable for the fused feature vector. The first is a statistical distance, the Mahalanobis distance of $f(t)$ from the mean of the normal training data, which accounts for correlation between features. The second is the reconstruction error of an autoencoder trained on normal data [8]. The third is the isolation-forest score [6]. The reconstruction error for an input vector f and its reconstruction $\hat{f}$ is

$$e(t) = (1/D) ||f(t) - \hat{f}(t)||^2 \quad (3)$$

where D is the feature dimension. Scores are normalised to the range $[0, 1]$ and may be averaged or combined by a weighted rule when more than one scorer is used. A single inexpensive scorer is often sufficient and is preferable on constrained hardware.

E. Adaptive thresholding and decision

A fixed threshold degrades when operating conditions drift. An adaptive threshold is therefore defined from the running mean $\mu(t)$ and standard deviation $\sigma(t)$ of the anomaly score over a recent window of normal operation:

$$\tau(t) = \mu(t) + k \sigma(t) \quad (4)$$

To avoid contaminating the threshold, windows flagged as anomalous are not added to the running statistics. An alarm is raised only if the score exceeds $\tau(t)$ for M consecutive windows, which suppresses isolated false alarms caused by noise. The parameters k and M control the trade-off between sensitivity, false-alarm rate and detection delay: a larger k or M reduces false alarms but delays detection.

IV. REAL-TIME IMPLEMENTATION CONSIDERATIONS

- 1) Sliding-window processing. The latency of the signal pipeline is bounded below by the window length, since an event can only be recognised once enough of it has been observed, plus the processing time per window. Overlapping windows reduce the delay at the cost of extra computation. The window length and hop size should be chosen so that the window is long enough to resolve the lowest frequency of interest yet short enough to meet the application's response deadline.
- 2) Computational complexity. The FFT costs $O(N \log N)$ per window and the discrete wavelet transform costs $O(N)$, so both are suitable for low-power processors. The cost of image processing grows with resolution; region-of-interest cropping, down-sampling and frame skipping are effective ways to meet frame-rate requirements. Isolation forests have low scoring cost, whereas deep networks need quantisation, pruning or compact architectures to run on edge devices.
- 3) Embedded deployment. On microcontrollers and FPGAs, fixed-point arithmetic, lookup tables and pipelined filters reduce power consumption and latency. On edge devices, quantised neural networks keep model size small. Processing close to the sensor also reduces communication bandwidth and can preserve privacy, since raw images need not leave the device; only alarms and compact features are transmitted.
- 4) Concept drift. Normal behaviour changes with temperature, wear, load and, for cameras, illumination. The adaptive threshold handles slow drift in the score, but periodic re-estimation of the normalisation statistics and of the model, using data confirmed as normal, is needed to keep the detector reliable over months of operation.
- 5) Robustness and reliability. Real deployments must handle missing samples, sensor saturation and dropped frames. Input validity checks, watchdog timers and logging of every alarm with the associated features support maintenance and later model improvement.

V. EVALUATION METHODOLOGY

Because anomalies are rare, accuracy alone is misleading; a detector that never raises an alarm can appear highly accurate. The following metrics are therefore recommended, where TP, FP, TN and FN denote true positives, false positives, true negatives and false negatives:

$$\text{Precision} = TP / (TP + FP), \quad \text{Recall} = TP / (TP + FN) \quad (5)$$

$$F1 = 2 \cdot \text{Precision} \cdot \text{Recall} / (\text{Precision} + \text{Recall}) \quad (6)$$

In addition, the area under the ROC curve (ROC-AUC) measures threshold-independent separability of normal and anomalous samples. For pixel-level localisation in images, the intersection over union (IoU) between the detected and true defect masks is used. Operational metrics are equally important: detection latency, defined as the time from anomaly onset to the first alarm, the false-alarm rate per unit time, and the processing time per window or frame compared with the available time budget. Training should use normal data only, with a held-out test set that contains labelled anomalies and is never used for choosing parameters, and results should be reported over multiple random seeds where randomness is involved.

VI. EXPERIMENTAL SETUP

To demonstrate the framework in a fully reproducible way, synthetic data were generated with a fixed random seed. Synthetic data allow the exact location of every anomaly to be known, which makes the metrics unambiguous; they do not replace validation on real measurements, and the results below should be read as a demonstration of the pipeline rather than as a benchmark.

A. Synthetic Vibration Signal

A 60-second signal was generated at a sampling frequency of 2 kHz. The normal condition consisted of a 50 Hz component, a 120 Hz component of half the amplitude, and additive Gaussian noise. Three anomalous intervals were then inserted: a train of impulsive bursts (resembling a localised bearing defect) from 12 s to 14 s, an additional 310 Hz tone from 30 s to 33 s, and a 60% increase in amplitude from 47 s to 49 s. The signal and its spectrogram are shown in Fig. 2. Windows of $N = 1000$ samples (0.5 s) with a hop of $H = 500$ samples were used, giving 239 windows, of which 31 overlapped an anomaly by more than 30%. Eight features per window (RMS, kurtosis, crest factor, spectral centroid, spectral entropy and relative energies of the wavelet bands, using a 4-level Daubechies-4 decomposition) were computed. The first 11.5 seconds, free of anomalies, were used for training; all remaining data were used for testing.




Fig. 2. Synthetic vibration signal (top, anomalous intervals shaded) and its spectrogram (bottom).

Three scorers were compared: the Mahalanobis distance, an isolation forest with 100 trees, and a baseline that uses only the RMS value. Each used the adaptive threshold of Eq. (4) with $k = 3$, a running window of 40 normal scores and $M = 2$ consecutive exceedances.

B. Synthetic surface images

Textured surface images of 128×128 pixels were generated from smoothed Gaussian noise with additive sensor noise. Defective images contained one or two circular blemishes of random polarity and, in about half of the cases, a thin line-shaped scratch. Thirty defect-free images were used to estimate the normal distribution of the residual map, and the threshold was set to the mean plus four standard deviations. The test set contained 40 defective and 40 defect-free images. Detection used the background-residual method of Section III-A (Gaussian background with $\sigma = 6$, followed by smoothing, thresholding, morphological opening and dilation); an image was declared defective if the detected region exceeded 12 pixels.

VII. RESULTS AND DISCUSSION

A. Signal pipeline

Fig. 3 shows the anomaly score and the adaptive threshold for the Mahalanobis scorer. The score stays close to its baseline during normal operation and rises clearly above the threshold during all three anomalous intervals. Table I summarises the performance of the three scorers.




Fig. 3. Mahalanobis anomaly score, adaptive threshold $\tau(t)$ and true anomalous intervals (shaded).

Table I. Signal Pipeline Performance On The Synthetic Data

Scorer	Precision	Recall	F1-score	ROC-AUC
Mahalanobis distance	1.000	0.903	0.949	1.000
Isolation forest	1.000	0.645	0.784	0.898
RMS only (baseline)	1.000	0.871	0.931	1.000

The Mahalanobis scorer on the fused feature set gave the best result ($F1 = 0.949$). The isolation forest had equal precision but lower recall (0.645) and a lower ROC-AUC, because with only eight correlated features and a short training set its random partitions separate the tone and amplitude anomalies less cleanly. The RMS-only baseline performed well here (0.931) because every inserted anomaly raised the signal energy to some extent; it would be expected to do worse for anomalies that change the spectrum without changing the energy, which is the situation where the spectral and wavelet features contribute most. This synthetic test is not designed to prove that, so it is noted as a limitation rather than as a finding.

Detection latency, measured from the start of each anomalous interval to the first alarm, was 0.5 s, 0.5 s, 0.5 s for the three events, which equals one window length: with $M = 2$ and 50% overlap the alarm is raised as soon as the second window that contains the event has been completed. The median processing time for computing all features of one window was 2.0 ms on a general-purpose processor in the test environment, which is well below the 250 ms hop interval, so the pipeline satisfies the real-time requirement with a large margin. Timings on embedded hardware will differ and must be measured on the target device.

B. Image pipeline

Fig. 4 shows an example of the image pipeline. The residual map suppresses the slowly varying background texture and concentrates the response on the defects, and the thresholded, morphologically cleaned mask marks their location.




Fig. 4. Surface-defect detection: (a) input image, (b) residual map, (c) detected mask.

TABLE II. IMAGE PIPELINE PERFORMANCE ON THE SYNTHETIC DATA

Method	Precision	Recall	F1-score	Mean IoU
Background residual	0.541	1.000	0.702	0.310

All defective images were detected (recall = 1.000, image-level ROC-AUC = 0.938), but the precision was lower (0.541), meaning that a number of defect-free images produced small false regions that passed the 12-pixel size rule. The mean IoU of 0.310 shows that the detected regions overlap only partly with the true defects; thin scratches in particular are partly removed by the morphological opening, and the dilation enlarges small blobs. These results show the expected trade-off of a simple, fast method: it is sensitive and cheap (about 1.6 ms per frame at 128×128 pixels) but needs a tighter decision rule or a learned model to reduce false alarms. Raising the threshold multiplier or the minimum region size would improve precision at the expense of recall; this should be tuned on a validation set and not on the test set.

C. Discussion

Three observations follow. First, statistical scoring on a compact, well-chosen feature set is both accurate and cheap enough for real-time use on modest hardware. Second, the adaptive threshold together with the requirement of consecutive exceedances controls false alarms without learning a model of the anomalies themselves, which is helpful because anomalies are unknown in advance. Third, the image residual method illustrates that classical DIP operations can give high sensitivity at very low cost, but that an additional stage (a learned classifier or an autoencoder) is needed when low false-alarm rates are essential. These findings are consistent with the qualitative comparison in the next section, but they are limited to synthetic data with simple anomaly types and must be confirmed on real-world datasets.

VIII. COMPARATIVE ANALYSIS OF TECHNIQUES

Table III gives a qualitative comparison of the techniques discussed in this paper, based on the characteristics reported in the cited literature and on the observations above. Actual performance depends on the data and the hardware and should be verified experimentally for each application.

Table III. Qualitative Comparison Of Techniques

Technique	Data	Strength	Cost
Background model (GMM) [5]	Video	Adapts to slow scene change	Low–Med
Optical flow [13]	Video	Captures abnormal motion	Med–High
Edge / texture [12]	Image	Detects surface defects	Low
STFT / FFT [3]	1-D signal	Frequency-fault signatures	Low
Wavelet [4], [11]	1-D signal	Transients, non-stationary signals	Low–Med
One-class SVM [10]	Features	Boundary of normal data	Med
Isolation forest [6]	Features	Fast, few parameters	Low
Autoencoder [8], [9]	Image / signal	Learns complex normal behaviour	Med–High

In general, the choice follows the resource budget and the nature of the anomaly. For impulsive faults in vibration, wavelet and kurtosis-based features are efficient. For changes in spectral content, STFT or spectral-entropy features are preferable. For visual defects on regular surfaces, residual or texture methods are a good starting point, and for complex scenes with variable normal appearance, learned models are generally required. Combining a cheap first-stage detector with a more expensive second-stage verifier is a practical way to keep both latency and false alarms low.

IX. APPLICATIONS

- 1) Industrial inspection and maintenance. Visual defect detection on products can be combined with vibration or motor-current signature analysis for early detection of bearing, gear and rotor faults, enabling condition-based maintenance in place of fixed schedules.
- 2) Healthcare. Arrhythmia detection in ECG signals, seizure detection in EEG and the detection of abnormal regions in medical images are anomaly-detection problems with strict latency and reliability requirements, especially for wearable devices.
- 3) Security and traffic. Abnormal crowd movement, intrusion, abandoned objects and wrong-way vehicles can be detected from video using background modelling and optical flow, and acoustic sensors can add events such as glass breaking or gunshots.
- 4) Smart infrastructure and energy. Structural health monitoring of bridges and buildings uses vibration and acoustic signals along with camera inspection, and power systems use anomaly detection on current and voltage waveforms to identify faults and power-quality disturbances.

X. CHALLENGES AND FUTURE SCOPE


Key challenges include the scarcity and imbalance of labelled anomalies, variation of operating conditions, noise and illumination change, the limited explainability of deep models, and the resource limits of edge hardware. The results in this paper come from synthetic data, and a clear next step is to evaluate the framework on public image and sensor benchmarks and on measurements from a real machine or production line. Further work includes implementing the pipeline on a low-cost embedded platform such as a Raspberry Pi, a microcontroller or an FPGA board and measuring latency and power; investigating self-supervised and transfer learning to reduce the need for training data; extending the framework with learned image features and an autoencoder-based scorer; and adding explainable outputs that localise and describe the likely cause of an anomaly to help maintenance staff.

XI. CONCLUSION

This paper presented a modular framework that combines digital image processing and digital signal processing for real-time anomaly detection. By pairing spatial and time-frequency features, fusing them, and applying unsupervised scoring with an adaptive threshold, the framework can serve varied applications while meeting latency constraints. A demonstration on reproducible synthetic data showed that a Mahalanobis-distance detector on eight signal features reached an F1-score of 0.949 with a processing time of about 2.0 ms per window, and that a background-residual image detector found every defective surface image at about 1.6 ms per frame, albeit with a lower precision of 0.541. The design guidelines, implementation considerations and evaluation protocol provide a practical basis for researchers and engineers. Validation on real data and on embedded hardware will be the focus of subsequent work.

REFERENCES

- [1] V. Chandola, A. Banerjee and V. Kumar, "Anomaly detection: A survey," *ACM Computing Surveys*, vol. 41, no. 3, pp. 1–58, 2009.
- [2] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 4th ed. New York, NY, USA: Pearson, 2018.
- [3] A. V. Oppenheim and R. W. Schaffer, *Discrete-Time Signal Processing*, 3rd ed. Upper Saddle River, NJ, USA: Pearson, 2010.
- [4] S. Mallat, *A Wavelet Tour of Signal Processing*, 3rd ed. Burlington, MA, USA: Academic Press, 2009.
- [5] C. Stauffer and W. E. L. Grimson, "Adaptive background mixture models for real-time tracking," in *Proc. IEEE CVPR*, 1999, pp. 246–252.
- [6] F. T. Liu, K. M. Ting and Z.-H. Zhou, "Isolation forest," in *Proc. IEEE ICDM*, 2008, pp. 413–422.
- [7] G. Pang, C. Shen, L. Cao and A. van den Hengel, "Deep learning for anomaly detection: A review," *ACM Computing Surveys*, vol. 54, no. 2, pp. 1–38, 2021.
- [8] M. Sakurada and T. Yairi, "Anomaly detection using autoencoders with nonlinear dimensionality reduction," in *Proc. MLSA Workshop*, 2014, pp. 4–11.
- [9] M. Hasan, J. Choi, J. Neumann, A. K. Roy-Chowdhury and L. S. Davis, "Learning temporal regularity in video sequences," in *Proc. IEEE CVPR*, 2016, pp. 733–742.
- [10] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola and R. C. Williamson, "Estimating the support of a high-dimensional distribution," *Neural Computation*, vol. 13, no. 7, pp. 1443–1471, 2001.
- [11] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia, PA, USA: SIAM, 1992.
- [12] J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-8, no. 6, pp. 679–698, 1986.
- [13] B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artificial Intelligence*, vol. 17, nos. 1–3, pp. 185–203, 1981.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)