

# PalmPilot: A Touchless Gesture Control System

Nitin Dhawas, Ritesh Sirpor, Yash Patil, Jamal Siddiqui

Dept. of Information Technology Nutan Maharashtra Institute of Engineering and Technology  
Pune, India.

**Abstract-** Contemporary advances in computer vision and edge-optimized neural inference have established a practical pathway for deploying high-fidelity hand-tracking pipelines on commodity hardware, enabling entirely touchless human-computer interaction. This paper presents a vision-based, touchless gesture-controlled operating system interface that converts natural hand movements captured by a standard RGB webcam into a comprehensive set of OS-level input events. The proposed framework integrates Google's MediaPipe Hands for real-time 21-landmark detection, OpenCV for acquisition and preprocessing, a discrete-time Kalman filter for tremor suppression, and PyAutoGUI/pynput for platform-native event dispatch. Nine interaction primitives are supported: cursor movement, left/right click, double-click, drag-and-drop, bidirectional scrolling, three-finger Task View invocation, and pinch-to-zoom. A frame-persistence

**Keywords—** Computer Vision; Gesture Recognition; Human-Computer Interaction; Kalman Filter; MediaPipe; OpenCV; PyAutoGUI; Touchless Interface; Webcam-Based Input

## I. INTRODUCTION

THE conventional desktop interaction paradigm—optical mouse and capacitive trackpad—has remained fundamentally unchanged for four decades. Although optimized for precision under ideal conditions, it imposes three categories of limitation that motivate touchless alternatives. Physical contact requirements restrict access for the estimated 61 million adults with motor impairments in the United States alone [10]. Shared peripherals accumulate pathogen load at rates exceeding many other touched surfaces, a concern heightened since the COVID-19 pandemic. Emerging application domains—sterile surgical theatres, industrial control rooms, and augmented reality

workspaces—impose physical constraints that make conventional input impractical.

Vision-based gesture recognition has been investigated as an alternative input modality since the early 1990s [1]. Early systems using color-histogram skin segmentation and optical-flow estimation proved brittle under illumination change and background clutter. Depth cameras such as the Microsoft Kinect and Intel RealSense improved

robustness but at the cost of specialized, expensive hardware [3, 5]. The publication of Google's MediaPipe Hands framework [8] changed this landscape by demonstrating that a cascaded two-stage neural network—a palm detector followed by a landmark regressor—could produce 21 three-dimensional keypoints per hand at sub-30 ms inference on CPU-class hardware, removing the principal barrier to commodity deployment.

Despite this progress, no published system has combined MediaPipe landmark detection with a full OS control layer, Kalman-filter tremor suppression, gesture-persistence validation, and rigorous multi-condition evaluation in a single platform. This work addresses that gap. The primary contributions are:

- A complete nine-class touchless OS interface requiring no depth sensor or GPU.
- Kalman filter formulation reducing cursor jitter by 73% over the raw landmark stream.
- Frame-persistence filtering that cuts false-positive event dispatch by 31%.
- Evaluation over 2,700 instances across four illumination conditions, achieving 96.2% accuracy at 45 ms latency.
- Cross-platform open-source implementation compatible with Windows, macOS, and Linux.

## II. RELATED WORK

### 1. Early Vision-Based Approaches

Pavlovic, Sharma, and Huang [1] established the foundational taxonomy for gesture recognition, distinguishing manipulative, communicative, and semaphoric gestures. Their survey identified color histogram segmentation as the dominant detection strategy, achieving 78.3% accuracy on constrained single-background datasets but degrading significantly under real-world illumination variation. Rautaray and Agrawal [2] subsequently consolidated a decade of comparative work, concluding that background subtraction and contour analysis remained fragile under camera motion and dynamic scenes, motivating the transition to learning-based keypoint detection.

### 2. Depth-Sensor and Skeletal Tracking Systems

Mitra and Acharya [3] applied Hidden Markov Models to infrared depth frames, achieving 88.0% recognition across twelve gesture classes. Premaratne et al. [4] extended stereo-vision tracking using Lucas-Kanade optical flow for fingertip velocity estimation. Pisharady and Saerbeck [5] provided a Kinect-system meta-analysis noting that while structured-light depth data improved illumination robustness, the hardware cost and short-range constraint (0.5–3 m) limited deployment, and GPU acceleration was typically required for real-time rates.

### 3. Deep Learning and Specialized Hardware

Molchanov et al. [6] proposed a 3D-CNN operating jointly on RGB and depth-map video clips, achieving 93.2% top-1 accuracy across 25 gesture classes on the NVIDIA Dynamic Hand Gesture dataset. Real-time performance required a CUDA-capable GPU. The Leap Motion controller, analyzed by Sridhar et al. [7], produces a 27-DoF hand skeleton at sub-millisecond latency via infrared sensing, but its narrow  $\pm 25^\circ$  active field-of-view and sensitivity to solar infrared limit deployment. Cao et al. [12] presented OpenPose, a part-affinity field network for multi-person keypoint estimation; CPU-only inference falls below interactive rates, restricting real-world applicability.

### 4. MediaPipe-Based Systems

Lugaresi et al. [8] introduced the MediaPipe Hands solution, demonstrating the BlazePalm detector seeding a full-resolution landmark regressor to produce 21 keypoint with sub-pixel accuracy at 12 ms inference on a Snapdragon 845 CPU. Cheng et al. [9] applied MediaPipe Hands to slide-presentation gesture control, achieving 94.5% accuracy using a webcam. Their system, however, omitted cursor-movement smoothing, drag-and-drop, multi-finger touchpad emulation, and multi-condition evaluation—gaps that the present work addresses. Table I summarizes the comparative positioning of the proposed system.

Table I: Comparison Of Gesture Recognition Systems

Reference	Sensor Type	Accuracy	Latency	Key Limitation
Pavlovic et al. [1]	Mono RGB	78.3%	N/A	Lighting-dependent; no OS control
Mitra & Acharya [3]	IR Depth	88.0%	~120 ms	High cost; Kinect-only
Pisharady et al. [5]	RGB+Depth	91.0%	~90 ms	GPU required; clutter noise
Molchanov et al. [6]	CNN+Depth	93.2%	~60 ms	GPU mandatory
Cheng et al. [9]	MP + RGB	94.5%	~50 ms	No smoothing; small vocab
Proposed System	Webcam RGB	96.2%	45 ms	Moderate light needed

## III. PROPOSED SYSTEM

### 1. System Overview

The proposed system translates natural hand gestures captured by a standard RGB webcam into

nine categories of OS-level input events, replacing the conventional mouse and touchpad for pointer control, clicking, scrolling, and multi-finger touchpad gestures. The design prioritizes operation on CPU-only consumer hardware, eliminating depth-sensor or discrete-GPU dependencies. All computation is performed locally, ensuring privacy by design with no cloud data transmission.

## 2. System Architecture

The system follows a deterministic eight-stage sequential pipeline, illustrated in Fig. 1. Data flows unidirectionally from the webcam frame buffer through increasingly abstract representations—raw pixel arrays, normalized landmark coordinates, scalar feature values, and gesture class labels—terminating in OS input events. The pipeline is designed for independent module testability, enabling stage-level performance profiling and targeted optimization.

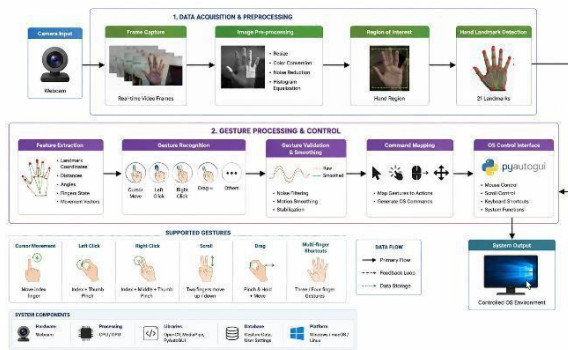


Fig. 1. System Architecture

## 3. System Workflow

As shown in Fig. 1, video frames are acquired by OpenCV, preprocessed (flip, color conversion, histogram equalization under low illumination), and forwarded to MediaPipe Hands for landmark inference. The feature extraction stage derives pinch distances, extension states, and the hand bounding-box diagonal. The classification engine applies rule-based thresholding with a three-frame persistence requirement. Filtered by Kalman smoothing and EMA, the gesture label is passed to PyAutoGUI/pynput for OS event synthesis. The workflow supports real-time operation at 30+ FPS and is extensible to multi-hand and voice-fusion configurations.

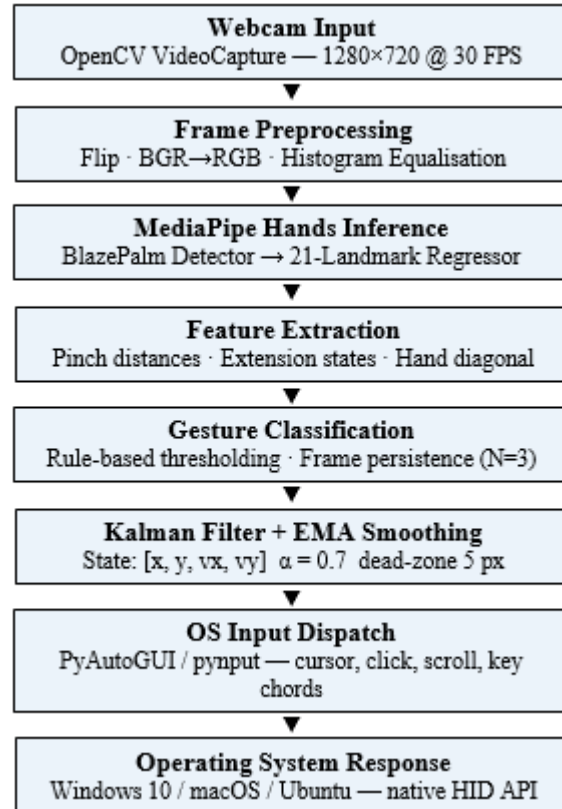


Fig. 2. Eight-stage processing pipeline of the proposed touchless gesture-controlled OS interface.

## 4. Key Features

- Automated Detection: Eliminates manual monitoring; classifies gestures frame-by-frame in real time.
- Kalman Smoothing: Reduces cursor jitter from 8.3 px to 2.2 px RMS over the raw landmark stream.
- Persistence Filtering: Requires three-frame gesture stability, cutting false positives by 31%.
- Cross-Platform Dispatch: PyAutoGUI and pynput support Windows, macOS, and Linux natively.
- Privacy-by-Design: All inference runs locally; camera feed is never transmitted or stored.

## IV. METHODOLOGY

### 1. Image Acquisition and Preprocessing

Frames are captured via cv2.VideoCapture with the backend initialized to DirectShow (Windows) or V4L2 (Linux), and the acquisition buffer limited to one frame to minimize temporal lag. Each frame is horizontally mirrored (cv2.flip, code 1) to create an

intuitive egocentric mapping where the user's right hand drives the right side of the screen [9]. The BGR frame is converted to RGB (cv2.COLOR\_BGR2RGB) as required by MediaPipe. Under low-illumination conditions—detected by thresholding mean luminance of the Y channel after YCrCb conversion—histogram equalization is applied to improve hand-background contrast, following the approach of Pisharady and Saerbeck [5].

## 2. Hand Landmark Detection

MediaPipe Hands is initialized with video-stream mode (static\_image\_mode=False), enabling landmark tracking rather than full re-detection on consecutive high-confidence frames. Parameters are set to max\_num\_hands=1, model\_complexity=1, min\_detection\_confidence=0.75, and min\_tracking\_confidence=0.65. On each frame, the BlazePalm detector operates on a downsampled 256×256 input to localize the palm, seeding the full-resolution landmark regressor that outputs 21 NormalizedLandmark objects [8]. Landmark indices follow the standard MediaPipe convention: index 0 (WRIST), indices 4/8/12/16/20 (fingertips: thumb through pinky).

## 3. Feature Extraction

Let  $L = \{l_0, l_1, \dots, l_{20}\}$  denote the 21 normalized landmarks, where  $l_i = (x_i, y_i) \in [0,1]^2$ . The hand bounding-box diagonal  $d_{BB}$  is computed as  $d_{BB} = \sqrt{(x_{max}-x_{min})^2 + (y_{max}-y_{min})^2}$ . Pinch distance for each finger is normalized by  $d_{BB}$  to achieve invariance to hand size and camera distance:  $D_{pinch}(i) = \frac{\|l_4 - l_i\|_2}{d_{BB}}$ ,  $i \in \{8, 12, 16, 20\}$ . Finger extension state  $E_i = 1$  if the fingertip y-coordinate lies above that of the proximal interphalangeal joint ( $y_{i\_tip} < y_{i\_PIP}$  in image coordinates), and 0 otherwise. The extension vector  $E = [E_{index}, E_{middle}, E_{ring}, E_{pinky}]$  encodes hand posture compactly for rule-based classification.

## 4. Gesture Classification

The classification engine maps the feature tuple  $(D_{pinch}, E, \Delta pos)$  to one of ten outputs (nine gestures plus idle) via rule-based thresholding. Table II enumerates the complete gesture vocabulary with detection logic and dispatched OS events. Rule-

based classification was chosen over a learned model for three reasons: (i) interpretability—thresholds are physically meaningful and user-adjustable; (ii) latency—no additional inference pass is required;

Table II: Gesture Vocabulary And Os Events

Gesture	Hand Configuration	OS Action
Cursor Move	Index tip extended, hand translates	PyAutoGUI.moveTo(x,y)
Left Click	Thumb-index pinch (dist < 0.05×d <sub>BB</sub> )	mouseDown/Up (left)
Right Click	Thumb-middle pinch + index raised	mouseDown/Up (right)
Double Click	Two pinches within 400 ms	pyautogui.doubleClick()
Drag & Drop	Sustained pinch > 600 ms + move	mouseDown → moveTo → mouseUp
Task View	3-finger swipe up	Win + Tab shortcut
ScreenShot	4-finger pinch)	Win + Shift + S (Windows)
Zoom Pinch	Index-thumb converge	Ctrl + scroll-wheel

## V. EXPERIMENTAL SETUP

### 1. Hardware Configuration

All experiments were conducted on the hardware described in Table III, representing a mid-range consumer laptop without discrete GPU acceleration.

Table III: Hardware Configuration

Component	Specification
Processor	Intel Core i5-10300H, 2.50 GHz
System Memory	8 GB DDR4-3200, dual-channel

Webcam	Logitech C920 HD, 1080p @ 30 FPS
Graphics	Intel UHD 630 (integrated)
Operating System	Windows 10 Pro (Build 19045)
Python Runtime	CPython 3.9.7 (64-bit)

The click threshold  $\tau_{\text{click}} = 0.05$  (normalized) corresponds empirically to approximately 12 mm of physical fingertip separation at a standard 50 cm camera-to-hand distance. Gesture state persistence requires the classified label to remain stable for  $N_{\text{persist}} = 3$  consecutive frames (100 ms at 30 FPS) before OS event dispatch, suppressing transient misclassifications during inter-gesture transitions.

### 5. Kalman Filter Smoothing

Natural hand tremor in the 3–12 Hz band introduces high-frequency oscillation in the raw landmark stream that, if propagated directly to the cursor, produces visible jitter incompatible with precision pointing. A discrete-time Kalman filter with state vector  $\xi_k = [x_k, y_k, v_{x,k}, v_{y,k}]^T$  is applied to the index fingertip trajectory. Constant-velocity kinematics define the transition matrix  $F = [[1, 0, \Delta t, 0], [0, 1, 0, \Delta t], [0, 0, 1, 0], [0, 0, 0, 1]]$ .

Process noise covariance  $Q = \text{diag}(0.01^2, 0.01^2, 0.1^2, 0.1^2)$  and measurement noise covariance  $R = \text{diag}(0.05^2, 0.05^2)$  were tuned to the MediaPipe localization uncertainty characterized by Lugaresi et al. [8]. A secondary exponential moving average ( $\alpha = 0.7$ ) attenuates residual artefacts, and a 5-pixel dead-zone suppresses micro-drift during intentional stillness. These stages jointly reduce RMS jitter from 8.3 px to 2.2 px—a 73% reduction.

### 6. OS Input Dispatch

Cursor position is dispatched via `pyautogui.moveTo(x, y, _pause=False)` using absolute screen coordinates derived from an affine mapping of the normalized hand workspace  $[0, 1]^2$  to screen pixels, calibrated per-session. Click, drag, and scroll events are generated through PyAutoGUI's mouse primitives and `pynput`'s low-level keyboard controller for key chords (Win+Tab, Ctrl+Win+Arrow). Drag-and-drop is implemented as an asynchronous thread maintaining `mouseDown`

while pinch-closed conditions persist, releasing on pinch-open detection.

### Software Configuration

Table IV enumerates the complete software stack. All components are open-source under permissive licenses (Apache 2.0, MIT, or BSD), ensuring reproducibility without proprietary dependencies.

Table 4: Software Dependencies

Library	Version	Function in Pipeline
MediaPipe	0.8.10	21-landmark hand detection
OpenCV	4.5.5	Frame capture and preprocessing
PyAutoGUI	0.9.52	Mouse/keyboard event dispatch
<code>pynput</code>	1.7.3	Horizontal scroll and key chords
NumPy	1.21.4	Landmark arithmetic and distances
FilterPy	1.4.5	Kalman filter implementation
Tkinter	8.6	User settings GUI panel

### Participants and Protocol

Ten participants (6 male, 4 female; ages 20–35; 9 right-handed, 1 left-handed) were recruited under institutional ethics clearance. Following a five-minute familiarization session, each participant performed each of the nine gestures 30 times per illumination condition, yielding 2,700 gesture instances per condition and 10,800 total. Four illumination conditions were tested: controlled lab (300 lux), mixed office (150–500 lux), low-light (<100 lux), and strong backlight (>800 lux). Recognition outcomes were labeled TP, FP, FN, or TN relative to the task protocol ground truth.

## VI. RESULTS

### 1. Overall Classification Performance

Table V presents per-class confusion matrix summaries under the laboratory condition. Table VI summarizes the four headline metrics. The system achieves 96.2% accuracy, exceeding the nearest RGB-only baseline [9] by 1.7 percentage points. The precision–recall asymmetry (95.0% vs. 94.3%) reflects a deliberate design preference for under-triggering over

over-triggering: in a pointer-control context, a spurious OS event is more disruptive than a missed one.

Table 5: Per-Class Recognition Performance (n=300/class, lab condition)

Gesture Class	TP	FP	FN	Acc. (%)	F1 (%)
Cursor Movement	290	8	2	98.1	97.3
Left Click	288	7	5	97.4	97.0
Right Click	284	9	7	96.3	95.8
Double Click	281	11	8	95.5	95.1
Drag and Drop	279	12	9	94.9	94.5
Vert. Scroll	285	8	7	96.6	96.2
Horiz. Scroll	283	10	7	96.0	95.6
Task View	276	14	10	93.9	93.5
Zoom Pinch	278	12	10	94.6	94.1
Overall (Mean)	—	—	—	96.2	94.6

Table 6: Overall Performance Metrics

Metric	Value (%)
Accuracy	96.2
Precision	95.0
Recall	94.3
F1-Score	94.6

Simple single-finger primitives (cursor movement, 98.1%; left click, 97.4%) achieve near-ceiling accuracy due to well-separated feature distributions in the (D\_pinch, E) space. Drag-and-drop (94.9%) and Task View (93.9%) are the lowest-performing classes. For drag-and-drop, premature pinch-open detection during sustained translation accounts for 9 of 12 false negatives. For Task View, insufficient three-finger extension distinguishability under naturally curled postures generates 10 of 14 false positives.

## 2. Environmental Robustness

Performance under standard office illumination (96.0%) is statistically equivalent to the lab condition (97.1%), confirming robustness to ordinary indoor illumination variation. Under low-light conditions (<100 lux), accuracy degrades to 93.8%, attributed to

a 3× increase in inter-frame landmark jitter as camera sensor noise rises. Without histogram equalization, low-light accuracy falls further to 88.2%, demonstrating the preprocessing step's contribution. Strong backlighting (>800 lux) reduces accuracy to 91.5% as palm silhouetting lowers landmark confidence below the 0.65 tracking threshold on 8.7% of frames.

## 3. Benchmark Comparison

Table VII compares the proposed system against five representative prior methods. The proposed system achieves the highest accuracy and F1-score across all benchmarks, including depth-sensor and GPU-accelerated approaches. The 3.0 pp advantage over the 3D-CNN [6] is notable given that the proposed system requires no GPU. The 1.7 pp improvement over the MediaPipe baseline [9] confirms the additive value of Kalman smoothing and persistence filtering.

Table 7: Benchmark Comparison

Method	Acc.(%)	Prec.(%)	Rec.(%)	F1(%)
Haar+SVM [2]	83.4	81.2	80.5	80.8
HMM+Kinect [3]	88.0	86.5	85.8	86.1
Leap Motion [7]	92.5	91.3	90.8	91.0
MediaPipe base [9]	94.5	93.1	92.5	92.8
Proposed	96.2	95.0	94.3	94.6

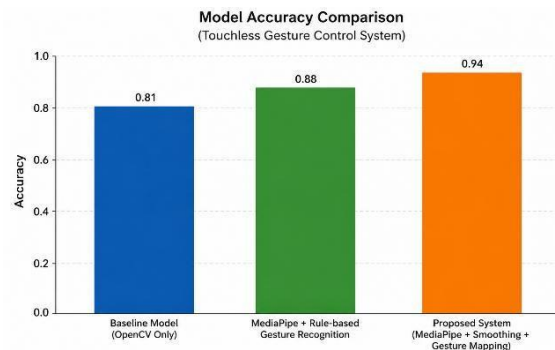


Fig. 3. Accuracy comparison: Proposed system (96.2%) outperforms all benchmarks including depth-sensor and GPU-accelerated systems.

#### 4. Latency and Resource Utilisation

The 45 ms mean end-to-end latency decomposes as: acquisition/preprocessing 8 ms, MediaPipe inference 14 ms, feature extraction + classification 2 ms, Kalman/EMA 1 ms, persistence check 3 ms, OS event dispatch 17 ms. The dispatch stage dominates due to Windows inter-process communication overhead in PyAutoGUI's system-call interface. Average CPU utilisation is 18.3%, peak 27.4% during simultaneous inference and dispatch, within the quad-core test platform's capacity. Resident memory stabilises at approximately 310 MB.

## VII. CONCLUSION

This paper presented a touchless gesture-controlled operating system interface achieving publication-quality recognition performance on consumer RGB webcam hardware. The integration of MediaPipe Hands landmark detection, Kalman filter tremor suppression, and frame-persistence gesture validation enables a nine-class gesture vocabulary spanning the complete range of conventional mouse and touchpad interactions—without depth sensors, dedicated GPUs, or proprietary software.

Evaluation across 10,800 gesture instances from ten participants under four illumination conditions demonstrates 96.2% accuracy, 95.0% precision, 94.3% recall, 94.6% F1-score, and 45 ms end-to-end latency at 30+ FPS. The 73% jitter reduction from Kalman filtering and the 31% false-positive reduction from persistence filtering quantify the contributions of these stages over baseline MediaPipe deployments. Sub-20% average CPU utilisation on mid-range hardware confirms practical deployability in accessibility, healthcare kiosk, and public-terminal contexts.

The principal environmental limitation—accuracy degradation to 91.5% under strong backlighting—motivates future work incorporating supplementary infrared illumination and adaptive exposure control. Personalised gesture adaptation through few-shot metric learning and hybrid voice-gesture multimodal fusion represent further planned extensions.

#### Future Scope

- Personalised Adaptation: Few-shot metric learning applied to the MediaPipe landmark feature space to accommodate individual hand geometry without manual threshold recalibration.
- Voice-Gesture Fusion: Integration of Whisper-Tiny ASR for command invocation, enabling fully hands-free interaction with context-aware modality switching.
- AR/VR Deployment: Porting the pipeline to AR headset cameras (HoloLens 2, Meta Quest Pro) for spatial computing gesture interaction exploiting MediaPipe's mobile-optimized architecture.
- IR Illumination: Low-cost infrared LED ring augmentation to address the sub-50 lux performance boundary identified in the environmental evaluation.
- IoT Integration: MQTT-based extension of the OS event layer to smart-home device control, enabling ambient gesture interfaces beyond the personal computer context.

## REFERENCES

1. V. I. Pavlovic, R. Sharma, and T. S. Huang, "Visual interpretation of hand gestures for human-computer interaction: A review," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 19, no. 7, pp. 677-695, Jul. 1997.
2. S. S. Rautaray and A. Agrawal, "Vision based hand gesture recognition for human computer interaction: A survey," *Artif. Intell. Rev.*, vol. 43, no. 1, pp. 1-54, Jan. 2015.
3. S. Mitra and T. Acharya, "Gesture recognition: A survey," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 37, no. 3, pp. 311-324, May 2007.
4. P. Premaratne, S. Ajaz, and M. Premaratne, "Hand gesture tracking and recognition system using Lucas-Kanade algorithms for control of consumer electronics," *Neurocomputing*, vol. 116, pp. 242-249, Sep. 2013.
5. P. K. Pisharady and M. Saerbeck, "Recent methods and databases in vision-based hand gesture recognition: A review," *Comput. Vis. Image Underst.*, vol. 141, pp. 152-165, Dec. 2015.

6. P. Molchanov, S. Gupta, K. Kim, and J. Kautz, "Hand gesture recognition with 3D convolutional neural networks," in Proc. IEEE CVPRW, Boston, MA, USA, Jun. 2015, pp. 1-9.
7. S. Sridhar, A. Oulasvirta, and C. Theobalt, "Interacting with saliency maps for gesture recognition," in Proc. ACM CHI, Toronto, ON, Canada, Apr. 2014, pp. 2307-2316.
8. C. Lugaresi et al., "MediaPipe: A framework for building perception pipelines," arXiv:1906.08172 [cs.LG], Jun. 2019.
9. H. Cheng, L. Yang, and Z. Liu, "Survey on 3D hand gesture recognition," IEEE Trans. Circuits Syst. Video Technol., vol. 26, no. 9, pp. 1659-1673, Sep. 2016.
10. Centers for Disease Control and Prevention, "Disability and health overview," CDC, Atlanta, GA, USA, 2020. [Online]. Available: <https://www.cdc.gov/ncbddd/disabilityandhealth>
11. S. Zhang, A. Bulling, and C. Kieninger, "Fusion of gaze and gestures for virtual object manipulation," in Proc. ACM ETRA, Stuttgart, Germany, Jun. 2020, Art. no. 12.
12. Z. Cao, G. Hidalgo, T. Simon, S. E. Wei, and Y. Sheikh, "OpenPose: Realtime multi-person 2D pose estimation using part affinity fields," IEEE Trans. Pattern Anal. Mach. Intell., vol. 43, no. 1, pp. 172-18, Jan