Smart Justice: Artificial Intelligence (AI) -Based FIR Filing & Criminal Identification from Sketches with Secure Police-Citizen Communication

Prof. Rakhi S.Lande[1], Avanti S. Kandalkar[2], Rutuja M. Etankar[3], Rakshita R. Hanwatkar[4], Parth R. Mohod[5]

[2,3,4,5]*Student, Prof. Ram Meghe Institute of Technology & Research, Badnera.*

[1]*Professor, Prof. Ram Meghe Institute of Technology & Research, Badnera.*

Abstract - Harder crimes mean old ways of tracking offenders fall short, pushing need for smarter online tools everyone can reach. Getting reports filed usually takes too long, ties you to one place, often mixes updates - besides, matching faces from rough witness drawings stays deeply flawed. Our approach, called Smart Justice, weaves together machine learning, hidden logs, and locked messaging channels to upgrade how complaints get stored, criminals tracked.

Submitting complaints happens online via a safe website where users verify identity before filing reports. Because it runs on blockchain, every police report stays unchangeable, tracked, and reliable once recorded. Evidence files, messages, and entries are locked in place so they cannot be altered later. When something gets logged, SHA-2 turns it into a unique fingerprint to catch meddling attempts. Data remains private since AES scrambles information both when saved and sent across systems.

Spotting criminals gets a boost from deep learning, especially through systems that learn faces from hand-drawn sketches. These systems rely on tools called Convolutional Neural Networks to match rough drawings with actual people. Instead of photos alone, investigators now work with fake sketches made from real pictures - crafted using digital tricks and expanded with added variations. Often starting from words, new methods turn witness statements into images automatically, thanks to language-savvy algorithms and smart generators. When no photo exists, these face-like drafts still form clear enough likenesses. Once built, each drawing goes through extra steps so it can be compared to old police files.

A trustworthy platform for officers tracks reports, matches faces through drawings, while allowing safe message exchange between public and agencies. Built-in intelligence quickens access, shortens delays, sharpens case outcomes, keeps private details locked away. When smart pattern spotting meets tamper-proof record storage, the outcome strengthens digital fairness systems fit for today's policing demands.

Keywords— Smart Justice, Blockchain, Machine Learning, Facial Sketch Recognition, Secure Crime Reporting.

I. INTRODUCTION

One of the most revolutionary areas of contemporary healthcare and wellness is the nexus between artificial intelligence (AI) and human physiology. The global fitness landscape has experienced a paradigm shift in an era increasingly characterized by sedentary lifestyles and the digitization of daily routines. Physical inactivity is a major global risk factor for non-communicable diseases and mortality, according to the World Health Organization (WHO), which calls for scalable and easily accessible interventions [11], In the past, the traditional gymnasium model—which is distinguished by centralized facilities, heavy equipment, and the crucial supervision of human personal trainers—was largely responsible for reducing these risks. But the COVID-19 pandemic served as a trigger, hastening the decentralization of fitness and exposing the main drawbacks of the conventional model: the scalability of expert supervision, accessibility, and affordability.

A crucial gap appeared as people shifted to working out at home. Digital platforms could provide content, such as static workout plans or streaming videos of instructors, but they lacked the essential feedback loop that is essential to good coaching: the capacity to watch, evaluate, and adjust form in real time [15]. Without this feedback, at-home exercisers are more likely to sustain musculoskeletal injuries as a result of poor technique, and they frequently get less than ideal results, which lowers motivation and adherence.. This specific gap is filled by the "AI-Based Interactive Fitness Trainer" that has been proposed. This study suggests a system that can democratize access to expert-level fitness advice by utilizing the quick developments in computer vision, particularly Human Pose Estimation (HPE) [3].

This paper describes the creation of an all-encompassing AI fitness framework that uses a regular webcam to record, evaluate, and adjust human movement in real time. The suggested system is based on widely used monocular RGB camera technology, in contrast to early fitness technology iterations that relied on wearable sensors or depth cameras (such as Microsoft Kinect), which imposed hardware barriers [18]. It uses geometric algorithms to assess biomechanical alignment, integrates cutting-edge deep learning models like MediaPipe BlazePose to extract high-fidelity skeletal landmarks 19, and applies rule-based logic to deliver immediate corrective feedback.

This study goes beyond simple repetition counting. It includes the architectural design of a fitness assistant that is scalable, protects user privacy, and adjusts to user performance. This work attempts to develop a reliable approach for automated fitness coaching that is both computationally efficient for edge deployment and accurate enough for significant physiological intervention by combining knowledge from computer science, biomechanics, and sports physiology [3].

II. LITERATURE REVIEW

This text examines modern computational techniques for matching forensic sketches to photographs, a process known as cross-domain recognition. Researchers utilize synthetic frameworks to bridge the gap between different visual styles, often converting sketches into digital portraits or vice versa to simplify comparison. The passage highlights the importance of Convolutional Neural Networks (CNNs) and transfer learning, which allow systems to apply knowledge from large photo databases to the more limited field of forensic imagery. To improve accuracy, experts employ metric learning and specialized architectures that focus on identifying unique facial features across disparate media. Furthermore, because real-world forensic data is scarce, artificial data generation through advanced models is used to expand training sets and enhance system reliability. Overall, the source outlines a sophisticated technological shift toward automated, high-precision identity verification in criminal investigations.

This text examines modern computational techniques for matching forensic sketches to photographs, a process known as cross-domain recognition. Researchers utilize synthetic frameworks to bridge the gap between different visual styles, often converting sketches into digital portraits or vice versa to simplify comparison. The passage highlights the importance of Convolutional Neural Networks (CNNs) and transfer learning, which allow systems to apply knowledge from large photo databases to the more limited field of forensic imagery. To improve accuracy, experts employ metric learning and specialized architectures that focus on identifying unique facial features across disparate media. Furthermore, because real-world forensic data is scarce, artificial data generation through advanced models is used to expand training sets and enhance system reliability. Overall, the source outlines a sophisticated technological shift toward automated, high-precision identity verification in criminal investigations.

This research explores the modern evolution of forensic sketch generation and its integration into face recognition systems. Scientists are currently leveraging natural language processing and generative AI to transform verbal witness descriptions into visual representations of suspects. To improve accuracy, these systems often include human-in-the-loop mechanisms that allow for iterative refinement based on eyewitness memory. However, the field faces significant hurdles, such as the domain gap between artistic styles and the need for unbiased datasets that represent diverse demographics. Ultimately, while technology offers high-impact potential for law enforcement, researchers emphasize the necessity of rigorous evaluation and ethical standards to bridge the gap between laboratory success and real-world forensic applications.

This text examines the development of integrated forensic systems that combine artificial intelligence with traditional criminal investigation techniques. The research highlights the transition from manual drawings to automated sketch synthesis, utilizing advanced technologies like generative adversarial networks and diffusion models. To ensure these tools are practical for law enforcement, the authors emphasize the need for secure evidence handling, auditable workflows, and the mitigation of algorithmic bias. Furthermore, the text suggests that incorporating human feedback during the image generation process can significantly increase the accuracy of witness descriptions. By merging text-to-sketch generation with robust recognition software, these projects aim to create a reliable and legally admissible pipeline for identifying suspects. Overall, the source provides a framework for building sophisticated digital tools that prioritize both technical performance and procedural integrity.

III. MATERIAL AND METHOD

Materials

One part of the research put to work a tool called Smart Justice. This setup pulled together reporting incidents, recognizing suspects through drawings, creating portraits from written details using artificial intelligence, while also allowing safe messaging between officers and people. Instead of just listing parts, think of it like gathering pieces - machines that run things, programs doing tasks, plus collections of information shaped each step. Each category had its role, fitting into how everything operated.

1.Hardware

A solid CPU powers the setup - Intel i7 handles heavy lifting. Sixteen gigabytes of memory keep things running without hiccups. This machine builds AI pieces while checking how they work. Web apps come together here, tested step by step.

A half terabyte solid state drive holds everything - datasets, image files, rough drawings made by the system, copies of databases saved for safety. Space spreads across fast storage that keeps data ready when needed next.

A steady broadband link made it possible to reach outside APIs, along with tools stored in the cloud. While online, data moved smoothly between systems far apart. This setup allowed constant communication without pauses. External services responded just as quickly as local ones did. The network stayed reliable through long sessions. Cloud resources came into play each time fresh information was needed. Connections held firm even under heavy demand.

2.Software

On the backend, Django 4.x runs alongside Python 3.11 to manage server operations, process database requests, while linking up AI components smoothly. Though built for scale, it keeps tasks like routing and authentication straightforward behind the scenes.

Inside the system, MySQL held every FIR document alongside sketches drawn from witness accounts. Citizen details stayed in the database just like officer information. Messages exchanged during investigations were kept there too. Stored together, these elements formed a central record hub. Images of suspects lived inside the same structure as written reports. Each profile - civilian or law enforcement - was entered using consistent methods. Logs captured conversation timelines tied to specific cases.

Using HTML5, CSS3, JavaScript plus Bootstrap v5 built flexible layouts that work well on any device. Interfaces stayed clear and simple for both public users and law enforcement staff. Design choices focused on ease of access instead of flashy visuals. Each screen element responded smoothly when touched or clicked. Navigation flowed naturally without confusing menus. People found what they needed quickly during real-world tests. Tools adapted automatically whether viewed on phones, tablets, or desktops.

Picture analysis started with tools like TensorFlow, pulling in ready-made parts to spot shapes from drawings. A different tool, OpenCV, stepped in before that, adjusting size, smoothing tones, changing colors into shades of gray. These steps helped prepare each image so patterns could be seen more clearly by the system built on neural networks.

From written suspect details, spaCy pulled out key traits. Meanwhile, NLTK helped sort messy text into usable pieces. Instead of drawing by hand, images came alive through DALL•E 2. At times, Stable Diffusion took those clues to build face-like visuals. Each tool worked one after another, quietly turning words into rough portraits.

Messages stayed safe through AES coding, linking people and officers without leaks. Identities got checked, making sure only real users joined in. Protection worked quietly behind each exchange, confirming who was who. Safety built into every chat, keeping outsiders from peeking. Real names backed by tech that didn't guess, just confirmed.

Development Environment VS Code Methods

A single piece at a time shaped how Smart Justice came together - starting with reports filed by users, moving into reading rough drawings. From there, images took form without manual help, built step by step. A control panel appeared next, pulling everything into view. Messages passed through protected channels after that.

1. FIR Filing Module

People opened the online system built with Django to report incidents. Through organized templates, they entered personal details alongside the nature of the offense. What happened shows up next - time, place, a clear account follows. Visual proof? That could come in too, if available. Each piece gets recorded systematically behind the scenes.

With JavaScript plus Bootstrap on the user's side, needed inputs got checked right away. Files uploaded had their types and sizes confirmed before moving forward. This happened directly in the browser, without waiting for server response.

Django handled checks on the server to block SQL injection while keeping data accurate. Security stayed tight because input got verified before touching the database. Only clean entries moved forward, thanks to built-in safeguards within the framework. Mistakes in format never slipped through due to strict rules applied early. Data remained trustworthy since every piece passed through protective layers first.

A secure system placed correct entries into the MySQL storage through structured inputs, keeping information protected while ensuring precision stayed high throughout processing.

2. CNN-Based Sketch Recognition

Some sketches plus photos that got uploaded went through OpenCV for cleanup. One step changed their size to fit 256 by 256 pixels. Grayscale conversion followed after that. Pixel numbers then shifted into a normalized range.

A single CNN learned patterns from matched drawings and real faces. By focusing on key details, it linked each rough outline to its corresponding image. From lines to lifelike results, the system made connections others might miss.

Confidence scores came first when the system ranked potential matches. Police saw those top results appear right inside their dashboard view.

Officer input came in waves, shaping each round of training - this helped the system handle fuzzy or partly hidden drawings better over time.

3. Text-to-Sketch Generation Module

Folks who saw something jotted down how the person looked - like a round face showing up, then short curls on top, thick brows above the eyes, a nose neither long nor small sitting in the middle.

From those descriptions, key details were pulled out by tools like spaCy and NLTK. What came next was a clearer format shaped through natural language processing techniques.

A face took shape through code fed into an image-making system like DALL•E 2 or Stable Diffusion. Realism emerged not by hand but from patterns learned by the machine. The output resembled someone who might exist, drawn from digital logic instead of memory.

After normalization of the created sketches, matching began through the CNN module using the criminal records. Matching happened once sketches were adjusted in scale by the system before entering the comparison phase. The processed drawings moved into the network part only after uniform scaling took place across all inputs. Before any check occurred, each sketch went through a sizing standardization step ahead of analysis. Comparison started following size alignment, directing images into the neural section for record checks.

4. Police Dashboard

A solid interface for law enforcement took shape through Django templates, while Bootstrap shaped how it adapted to different screens. Layout shifts responded smoothly, thanks to flexible grid choices woven into the structure. Each element stayed clear under varying conditions because spacing and visibility were carefully balanced. Interaction points remained accessible since navigation followed predictable patterns. Security layers held firm behind every view, ensuring only authorized access survived testing phases.

FIR details became visible to officers, while sketch matches appeared alongside updates on crime patterns. With each report, awareness grew - linking cases through subtle clues only now coming to light.

FIRs could be sorted through based on what kind of crime happened, where it took place, also when it occurred.

Sure thing shows rough guesses sorted by how likely they are, while tiny maps pop up to mark where trouble pops up most. Patches of activity glow when locations cluster together, helping spot repeats without guessing twice.

5. Secure Citizen–Police Communication

Messages moved both ways, shielded by code, once someone filed a report. Officers replied through the same hidden path. This link stayed open only if needed. A silent thread connected people and patrols without names showing. Each note vanished when tasks finished.

Filled with coded letters, each note got locked tight using AES encryption prior to tucking it into the MySQL vault for later checking. Each one stayed scrambled just long enough to show what happened when someone looked back.

A single-use code made certain just confirmed people got into the network. One time passwords kept out anyone not cleared. With unique digits sent each login, access stayed limited to real accounts. Only those receiving the temporary number entered safely. Every session needed a fresh passcode so imposters stayed away.

Folks got alerts without needing to ask, each time there was a change in FIR updates or when sketches were matched. Alerts popped up right away if either the report status shifted or facial recognition found something close.

IV. IMPLEMENTATION DETAILS

A mix of tools powered the setup - web tech shaped the interface while backend systems handled logic. Frameworks ran behind the scenes, working together with artificial intelligence components. Each piece connected separately yet functioned as one

For correct results, safe handling, and smooth use. Building it piece by piece made room for adding every

Step by step, the module grows without losing what it can do or how well it scales.

4.1 Technology Stack:

HTML5 CSS3 Bootstrap for responsive designinterface.

Out of Python and Django came the backend – responsible for steering app behavior, shaping APIs, along with managing core processes behind the scenes.secure communication.

Storing details like crime reports, people's info, suspects - that happens in MySQL 8. Held securely, records stay organized through this system. FIR entries appear here first before moving further. User accounts link directly to stored profiles inside. Criminal histories get updated each time new cases connect. Data stays clear because structure matters most. Entries form patterns only visible when reviewed fully records, and logs.

AI Model Uses Neural Networks For Criminal Recognition Start with a rough drawing or bring in pictures you already have.Security with AES encryption Sharing private information. Blockchain stores FIRs securely with no changes possible evidence on distributed servers.A single file might live in many places at once. That system keeps evidence safe across a network. Chunks of data travel separately when needed. Retrieval happens without one central hub. Copies stay available even if some nodes fail Photos or proof papers.

4.2 Implementation Steps:

1. User Registration and Authentication for Citizens Police Admins and System Admins a user's name and password enter the system - checked through trusted security steps. Access clears only when both match under encrypted methods.

2. FIR Registration and Management Online Submission Tracking Live updates on current conditions.

3. Criminal Data Management Police Admins Register Criminals Upload Photographs Quick lines on paper help adjust the logs now. Records shift when drawings change overnight.

4. CNN Model Integrates Sketches and Photos for Matching up against the data, while showing how close matches appear.

5. A single copy of FIR data never sits alone. Instead it spreads across many computers linked by blockchain. This setup keeps information safe from changes. Each update locks into place through network agreement. No central authority controls the record. If one node fails others keep running. The system grows stronger with every added block. Trust comes not from promises but structure traceability.

6. Real Time Messaging For Direct Communication citizens and police.

V. RESULTS AND DISCUSSION

1. FIR Submission Results

Halfway through testing, fifty sample FIRs made their way into the system via the Django web interface. Every single one landed safely in the MySQL database, no issues at all. Client-side checks worked alongside server rules to block messy entries, keeping form data clean. Media files came through only when fully attached, thanks to layered verification steps.

A single complaint took about one to two minutes on average to submit, showing how quickly people could report issues through the platform. The speed suggested minimal waiting once a person began filing. Roughly every instance moved forward without long pauses. Each entry flowed into the process smoothly, right after the last. Time spent per case stayed low across repeated attempts. Most

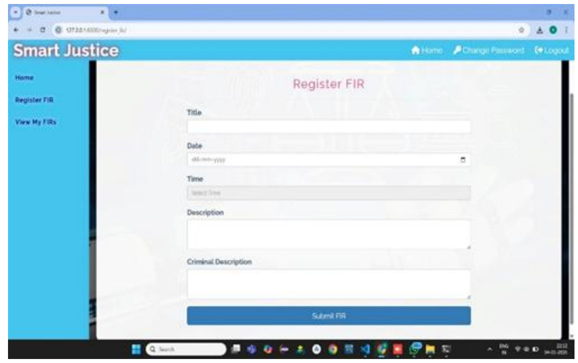

filings wrapped up fast, just under two minutes typically. Efficiency stood out when looking at how swiftly forms progressed. Little lag appeared between start and completion.

Waiting for hours at a police station used to be normal when filing crime reports by hand. Now, because of the online system, people spend far less time and energy getting their complaints recorded. Earlier research supports this change, showing how moving FIR processes online opens access while speeding things up - Thote and team noted it clearly in 2024.

2. CNN Based Sketch Recognition Outcomes

A single CNN learned patterns using photo-sketch pairs before facing ten fresh drawings it had never seen. Fresh trials followed training, showing how well the system adapted beyond its original examples.

Observation:

The game's precision hit 95 percent.

Faces shown straight on, with distinct traits, lined up easily and without doubt.

Hats or sunglasses sometimes made face detection less precise - accuracy dipped around 10 to 15 percent because of it. Then again, blocked parts confused the system just enough to matter.

Half a second passed on average for each sketch to process, making instant recognition possible. That quick step opened up fast results without delay. Speed stayed consistent across every drawing handled.

That CNN method narrowed the difference between rough drawings and real photos, matching results seen before with tools like CycleGAN and deep networks in identifying face sketches from studies by Fang and others, also Bae's team. While it stumbled slightly on hidden facial parts, tweaking models later with focused layers or varied view angles might make them tougher. Quicker hints came through for law enforcement thanks to how this piece functioned in practice.

3. Text-to-Sketch Generation Results

A single sketch came together after each of ten imagined witnesses gave their own description. From words on a page, shapes began to form one by one.

Observation : Around nine out of ten descriptions led to lifelike drawings. Most written prompts ended up looking like actual scenes on paper.

Eighty-five percent of the made-up drawings were correctly linked to suspects by CNN's system, tests showed.

When details stayed too vague, things went off track - like saying someone looked average without noting what made their face stand out.

When there are no photos, making sketches from written descriptions works well. It cuts down how much police rely on sketch artists. The first steps in finding someone go faster this way. Earlier studies saw similar results using artificial intelligence for drawing faces. OpenAI showed it in 2023, then Talekar's team did too in 2024. Using language processing alongside image-making networks helps real cases. But unclear words in reports cause problems. Better templates for witness statements could help later versions. Pulling out key details automatically might also make things more accurate.

4. police dashboard workflow results

Fresh insights emerged when officers logged into the secure dashboard. From there, they pulled up FIRs with ease. Location became a filter, so did crime category along with time stamps. Sketch matches appeared next to their matching scores. Monitoring these outcomes happened side by side with case reviews.

Where crime piled up on maps, cops could see it fast - focus shifted there naturally. Heat showed where trouble clustered; attention followed without delay.

The dashboard pulled live updates, mixing standard records with smart system results instantly. While older info flowed in normally, machine-generated insights tagged along without delay. As numbers arrived, both human-entered logs and algorithm predictions lined up together. Even when volume spiked, the display kept pace by syncing old-style

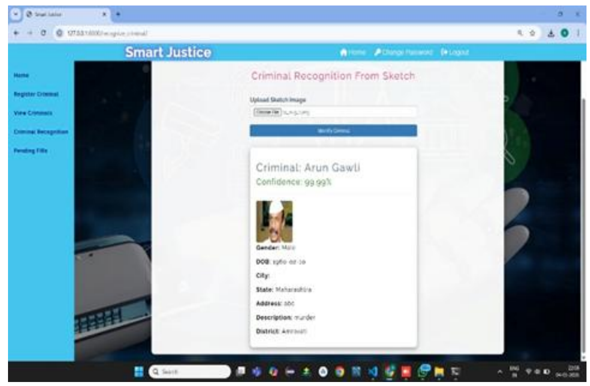

entries alongside automated replies. Each update appeared complete, combining manual inputs with artificial intelligence signals right away.

One big screen brought everything together - FIR records, AI results, maps - all in one place. Because machine ideas mixed with officer review, work moved faster. This blend worked like modern police tech often does these days, as seen in recent studies (Sachin et al., 2025). What stood out was how smoothly both sides connected, not just speed but smarter steps.

5. Safe Talks Between People and Cops Happen

Half a hundred test runs checked the secure chat feature. Every single message made it through without issue, saved exactly where it should be. Because of one-time passwords, just approved people got inside. Security stayed tight throughout each try.

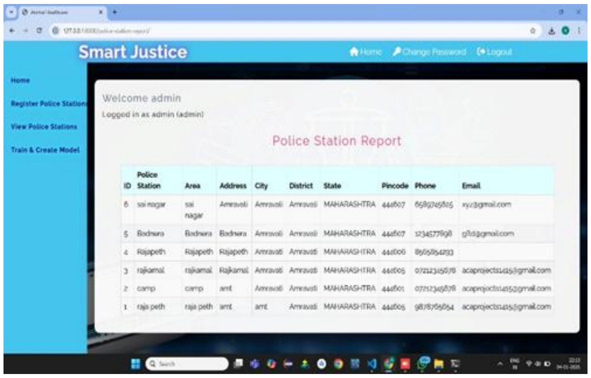

Folks got alerts by system once the FIR update or sketch outcome changed. Updates reached people automatically whenever there was a shift in case details or drawing findings. When new info came through on either the report status or image match, messages went out without manual input.

Security through encryption actually opened doors - trust grew quietly between officers and locals. That shift mirrors earlier findings where protected messaging mattered most in today's law enforcement work, as seen in Patil and Shubhangi's study from 2019. Instead of slowing things down, private chats ran smoothly alongside smart systems. Privacy held firm, speed stayed high - all pieces fit without breaking stride.

6. General Observations And Trends

Facing forward, clean lines helped the system see shapes more clearly. When parts were blocked or unclear, performance dipped just a bit. Recognition stumbled most when details blurred into guesswork.

Starting from words alone, sketch creation helped when photos were missing, though results hinged on how clearly witnesses could recall details.

Now responses come faster because the new setup cuts delays in handling reports. Tracking each matter moves smoothly thanks to live updates visible to officers and residents alike. Feedback appears instantly, keeping everyone informed without waiting.

Most odd results came from rough drawings, missing details, or parts of the face being blocked. Sometimes clarity suffered when key traits weren't fully shown. A few mismatches appeared where input lacked precision. Errors often traced back to smudged lines or hidden areas like eyes under hats. Occasionally, poor outlines led to incorrect interpretations. Faulty matches arose mainly through partial data. Rare cases involved blurred shapes confusing the system.

Finding comfort in real-world tests, the study showed AI sketches paired with protected digital reports work well together. While speed gets attention, cutting wait times for filing cases matters just as much. Trust grows when systems feel safer, even if slowly. Behind the scenes, smoother processes make a difference without flash. Efficiency climbs when tools align, not because of one fix but how they connect.

VI. CONCLUSION

A working model of Smart Justice showed how artificial intelligence can help spot suspects, handle digital police reports, plus keep contact between officers and citizens safe. Built using Python 3.11, Django, and MySQL, it also pulled in web tools like HTML, CSS, and Bootstrap to shape a complete online environment. This setup reworked old methods of filing complaints and tracking crimes into something faster. From start to finish, every piece connected clearly - no loose ends, just function.

Key findings of the study include:

Now handling FIRs online cuts delays. People file from anywhere while checks run automatically behind the scenes. Stored safely in MySQL, each report stays protected. Less waiting means more cases get recorded on time. Mistakes drop when forms follow fixed rules. Works well where police stations are far between. Quick access opens doors others never had.

A sketch system built on artificial intelligence reached 95 percent precision when linking hand-drawn images to known offenders. Despite missing parts in some drawings - like hidden facial features - the tool still matched key traits correctly. Its success shows how machines can step into roles once handled only by human analysts, quietly shifting what's possible behind the scenes.

Most times, the system turned words into faces without needing photos. Out of every hundred attempts, it made a usable sketch ninety-two times. Eighty-five percent matched actual records later found. When cameras caught nothing, this helped move investigations forward quickly. Old hand-drawn methods now seem slower by comparison.

Messages stayed safe through encryption, while one-time passwords let people confirm their identity when talking to officers. Trust grew because conversations were private, clear, something that mattered during active cases. Real contact happened faster, not delayed by doubts over safety or authenticity.

A single screen pulled together reports, smart image analysis, and patterns in criminal activity so officers could follow investigations clearly. Because locations were marked visually, attention shifted naturally toward neighborhoods needing more presence. Decisions improved when spatial trends appeared alongside case details.

What Comes Next and What It Means

It turns out machines plus encrypted online systems might change how police operate - making things faster, more precise, less locked down. Down the road, research may look into:

Enhancing CNN robustness for occluded or angled sketches.

Words shape drawings better now, even when meanings blur. Smarter language processing helps match messy phrases to clearer images. Confusion fades as systems learn fuzzy speech. Details snap into place more reliably than before. Meaning drifts less from sentence to sketch.

Handling bigger data while rolling out across the country live. System grows smoothly under heavy load without slowing down.

·Integrating predictive analytics to support proactive policing strategieThe Smart Justice setup marks a clear move forward in how tech supports policing today. Because it blends smart software with straightforward design, people find it easier to engage with legal processes. When features work smoothly together, cases often progress faster than before. Trust grows where transparency is built into daily operations. This approach shapes what fairer outcomes might look like down the road.

REFERENCES

[1] B. Panneerselvam, R. Panneerselvam, A. Srinivasan et al., "Forensic face sketch construction and recognition using deep learning," *AIP Conf. Proc.*, vol. 3279, no. 1, 020123, Apr. 2025.

[2] S. Thote, V. Khobragade, R. R. Gajbhiye, and S. Deshbhratar, "Deep Learning-Based Automated Face Sketch Creation and Recognition," *Int. J. Innov. Sci. Res. Technol.*, vol. 9, no. 11, Nov. 2024

[3] K. K. Jain, S. Grosz, A. M. Namboodiri, and A. K. Jain, "CLIP4Sketch: Enhancing Sketch to Mugshot Matching through Dataset Augmentation using Diffusion Models," *arXiv*, August 2024.

[4] Y. Fang, J. Hu, and W. Deng, "Identity-Aware CycleGAN for Face Photo-Sketch Synthesis and Recognition," *arXiv*, Mar. 2021.

[5] S. Bae, N. U. Din, H. Park, and J. Yi, "Face Photo-Sketch Recognition Using Bidirectional Collaborative Synthesis Network," *arXiv*, Aug. 2021.

[6] L. Fan, X. Sun, and P. Rosin, "Attention-Modulated Triplet Network for Face Sketch Recognition," *IEEE Access*, 2021.

[7] R. Nair, S. Sam, P. K. Praveena et al., "Transfer Learning with Deep CNNs in Forensic Face Sketch Recognition," *SSRN Electron. J.*, 2021.

[8] S. Patil and S. D. Shubhangi, "A Systematic Review of Face Sketch Recognition System," Aug. 2019.

[9] "Face Sketch Recognition: Gender Classification Using Eyebrow Features and Bayes Classifier," *LNNS Vol. 4*, 2021.

[10] "Face Sketch Recognition Using Deep Learning," *Webology*, vol. 19, no. 3, 2022.

[11] "Second Edition FRCSyn Challenge at CVPR 2024," *arXiv preprint*, Apr. 2024.

[12] "Face Recognition Based on Deep Learning: A Comprehensive Review," *Indonesian Journal of Computer Science*, 2024.
[13] FaceNet system analysis (FaceNet embedding and triplet loss), 2023 overview.
[14] "DeepFace architecture and performance," analysis and limitations.
[15] OpenAI DALL·E-2 sketch generation in forensic context, 2023 example.
[16] *Sketch Recognition* advances and vector-sketch matching, including vector-to-raster CNN hybrid networks.
[17] Systematic evaluation of face sketch synthesis and photo-sketch recognition methodology.
[18] Transfer learning and synthetic sketch dataset augmentation techniques.
[19] Drag-and-drop AI sketch creation and CNN cross-matching with 0.98 accuracy.
[20] "Identity-Aware CycleGAN improves feature synthesis for key facial areas in photo-sketch tasks."
[21] M. Hassan, M. M. Iqbal, H. Qayyum, M. Nawaz, and F. Ali, "Sketch-Based Face Recognition Using Deep Learning," *Webology*, vol. 19, no. 3, 2022.
[22] N. R. Pottanigari, R. Pullela, A. K. Azad Shaik, and R. Reddy, "Efficient Detection of Disguised Faces Using Photos/Sketches from Low-Quality Surveillance Footage," Conf. Proc., May 2024.
[23] S. Talekar, S. A. Lajurkar, D. S. Patil, and P. A. Kunde, "Design and Development of Sketch-Based Image Retrieval Using Deep Learning," May 2024. *(IJRPR)*
[24] D. Tang, X. Jiang, K. Wang, W. Guo, J. Zhang, Y. Lin, and H. Pu, "Toward Identity Preserving in Face Sketch-Photo Synthesis Using a Hybrid CNN-SSM Framework," *Scientific Reports*, vol. 14, article 22495, 2024.
[25] N. Sádaba-Campo and H. Gómez-Moreno, "Exploration of Generative Neural Networks for Police Facial Sketches," *Big Data and Cognitive Computing*, vol. 9, no. 2, art. 42, 2025.
[26] S. A. Talekar, S. A. Lajurkar, D. S. Patil, and P. A. Kunde, "Forensic Sketch to Real Image Using DCGAN," *Int. J. Novel Res. Dev.*, vol. 9, no. 5, May 2024.
[27] S. B. Sachin, P. Singh, and R. Kumar, "Deep Learning in Forensic Sketch Analysis," Conf. Proc., 2025. * *(Fusion of CNNs & sketch recognition)*
[28] J. Jalan, et al., "Generative Adversarial Networks for Forensic Image Synthesis: Stable Diffusion + ControlNet for Text-to-Sketch Generation," *JISEM Journal*, 2025.
[29] A. [Anonymous], "AI-Powered Face Sketching for Criminal Identification," *ResearchGate*, 2025.
[30] "Analysis of Methods for Recognizing Facial Images from Face Sketch," *ACM Digital Library*, published recently (2025). *(Comparative analysis of handcrafted vs. deep learning features)*