Quest Journals
Journal of Research in Humanities and Social Science
Volume 14 ~ Issue 11 (November 2026) pp: 25-38
ISSN(Online):2321-9467
www.questjournals.org

Research Paper

Artificial Intelligence to Map Cybercrime Hotspots: Applications, Challenges, and Policy Implications

Dr.S. Krishnan[1] and Dr.Soniya Dashehariya[2]

Abstract
Cybercrime has emerged as a significant challenge to public safety, economic stability, national security, and digital governance. The expansion of digital payments, cloud computing, social media, artificial intelligence (AI), and interconnected information systems has created new opportunities for criminal activity while complicating the identification of offenders and the geographical analysis of cybercrime. Traditional crime-mapping methods rely primarily on physical locations, police records, victim reports, and identifiable crime scenes. Cybercrime, by contrast, frequently crosses national boundaries, involves multiple jurisdictions, and employs anonymisation technologies that obscure the physical locations of offenders. Artificial intelligence, geospatial analysis, machine learning, and big-data analytics offer promising approaches to identifying cybercrime hotspots, detecting emerging patterns, and supporting evidence-based crime prevention.
This article examines the conceptual foundations, technological applications, methodological approaches, and policy implications of using AI to map cybercrime hotspots. It distinguishes between locations where victims report cybercrime, locations associated with attack infrastructure, and locations where offenders are believed to operate. The article reviews international research, including the World Cybercrime Index, and considers the relevance of cybercrime mapping to India. It proposes an integrated analytical framework combining official complaint records, verified incident data, financial fraud reports, threat intelligence, and geospatial information. It also discusses the limitations of predictive models, including underreporting, biased datasets, inaccurate geographical attribution, privacy risks, and the possibility of discriminatory policing. The article argues that AI-assisted cybercrime mapping should be treated as a decision-support mechanism rather than an autonomous instrument for identifying offenders or determining criminal responsibility. Effective implementation requires transparent methodology, reliable data, independent evaluation, human oversight, and cooperation among law enforcement agencies, financial institutions, technology companies, and international organisations.

Keywords: *Artificial Intelligence; cybercrime; hotspot mapping; geospatial intelligence; machine learning; predictive analytics; cybercrime geography; digital forensics; India; crime prevention.*

Received 25 Oct., 2026; Revised 04 Nov., 2026; Accepted 06 Nov., 2026 © The author(s) 2026.
P**ublished with open access at **www.questjournals.org

I. Introduction

The rapid development of information and communication technologies has transformed the ways in which individuals, businesses, governments, and criminal organisations interact. Digital platforms have improved communication, expanded access to financial services, facilitated electronic commerce, and supported the delivery of public services. At the same time, these developments have increased exposure to phishing, identity theft, online financial fraud, ransomware, unauthorised system access, data breaches, digital extortion, and other forms of cybercrime.

Cybercrime is no longer adequately understood as a collection of isolated technical incidents. It is also a social, economic, organisational, and geographical phenomenon. Criminal networks operate through combinations of digital infrastructure, financial systems, communication platforms, intermediaries, and locally embedded networks. Although an attack may be executed remotely, its planning, financing, recruitment,

[1] The Author is an Academician and an Experienced Journalist based in Jaipur, Rajasthan.
[2] The Author is an Advocate in High Court, Rajgarh, Madhya Pradesh.

infrastructure management, and proceeds may involve people and organisations situated in different geographical locations.

Consequently, understanding the geography of cybercrime has become an important research and policy objective. Authorities need to know where victims are concentrated, where attacks originate, which locations repeatedly appear in verified investigations, and where preventive resources might produce the greatest benefit. Artificial intelligence offers analytical tools that can assist with these questions by processing large datasets, identifying recurring patterns, classifying reported incidents, detecting unusual activity, and presenting complex information through interactive maps.

The application of AI to crime mapping has developed from broader advances in geographic information systems (GIS), statistical hotspot analysis, predictive modelling, natural-language processing, network analysis, and anomaly detection. Conventional GIS can display the geographical distribution of reported offences. Machine-learning methods can extend this capability by identifying nonlinear relationships, combining multiple information sources, and estimating changing levels of risk. Natural-language processing can help classify complaint descriptions, while network analysis can reveal connections among accounts, transactions, devices, and reported incidents.

Nevertheless, cybercrime mapping presents a fundamental difficulty: the geographical location of a cyberattack does not necessarily identify the location of its perpetrator. An internet protocol (IP) address may identify a server, virtual private network, compromised computer, proxy service, or other intermediary rather than the person responsible. Similarly, a victim's residential address identifies the location of victimisation, not necessarily the location from which the offence was coordinated. The geographical interpretation of cybercrime therefore requires careful distinctions between observed events, technical indicators, investigative attribution, and statistical prediction.

The present article explores how artificial intelligence can support a more comprehensive understanding of cybercrime geography. It develops a conceptual and methodological framework for identifying cybercrime hotspots, examines the opportunities and limitations of AI-assisted analysis, and proposes practical measures for implementing such systems responsibly. Particular attention is given to India, where digital payments, online public services, mobile connectivity, and internet-based commerce make cybercrime prevention an important component of digital governance.

The article addresses five principal questions:

1. How should a cybercrime hotspot be defined and measured?
2. Which AI and geospatial techniques can support cybercrime mapping?
3. What data are required to identify reliable geographical patterns?
4. What methodological, legal, and ethical challenges affect AI-based hotspot analysis?
5. How can such systems support cybercrime prevention, investigation, and public policy?

The central argument is that AI can improve the identification and interpretation of cybercrime patterns, but only when analytical results are grounded in reliable evidence, appropriate geographical definitions, transparent modelling, and accountable institutional practices.

II. Conceptualising Cybercrime Hotspots

2.1 Meaning of cybercrime

Cybercrime broadly refers to unlawful activities involving computers, digital networks, communication technologies, or online services. The term encompasses offences in which digital systems are the direct targets of criminal activity and offences in which technology facilitates conventional forms of crime.

Cyber-dependent offences include certain forms of unauthorised system access, malware deployment, denial-of-service attacks, and interference with digital infrastructure. Cyber-enabled offences include many forms of financial fraud, impersonation, identity misuse, extortion, and deceptive online transactions that may also occur through non-digital channels.

This distinction matters for hotspot mapping because different offences have different spatial characteristics. A ransomware incident affecting a business network may be associated with the location of the affected organisation, the location of compromised infrastructure, and the location of investigators handling the case. An online payment scam may involve a victim in one state, a bank account in another, an intermediary receiving funds elsewhere, and an offender operating across national boundaries.

Cybercrime should therefore not be represented as a single, uniform geographical phenomenon. A useful analytical system should distinguish offence categories and model their spatial patterns separately before combining them into an overall assessment.

2.2 Defining a hotspot

In conventional criminology, a crime hotspot is a geographical area in which crime occurs at a comparatively high concentration. Hotspots may be identified through counts, rates, spatial clustering, or statistically significant deviations from an expected distribution.

For cybercrime, a hotspot can be defined as a geographical area exhibiting a high concentration of a clearly specified cybercrime-related outcome during a defined period.

However, the outcome must be explicitly identified. At least four types of cybercrime hotspot should be distinguished.

First, victimisation hotspots. These are areas in which comparatively high numbers of victims report cybercrime. They may reflect population size, internet usage, digital-payment adoption, economic activity, or the accessibility of reporting services.

Second, incident hotspots. These are areas in which verified cyber incidents repeatedly affect organisations, networks, businesses, or individuals. Their interpretation depends on the geographical variable used, such as the location of an affected organisation or the location where an incident was detected.

Third, offender-associated hotspots. These are locations linked to suspected or established offender activity through corroborated investigative evidence. Such assessments require stronger attribution standards than those used to map victim complaints or technical indicators.

Fourth, infrastructure hotspots. These are areas in which particular technical resources, such as malicious hosting infrastructure or compromised devices, are observed. They identify characteristics of digital infrastructure rather than necessarily identifying the people who control it.

A fifth category, the *predicted-risk hotspot*, can be used for locations where a model estimates elevated future victimisation or incident risk. Such locations represent forecasts, not established concentrations of criminal activity.

These categories should never be merged without explanation. A district reporting a high number of complaints may be experiencing extensive victimisation, better public awareness, or more effective reporting mechanisms. A country associated with many malicious servers may host infrastructure used by foreign offenders. Neither observation, by itself, establishes that residents of the location are responsible for criminal activity.

2.3 The geographical dimensions of cybercrime

The geography of cybercrime can be examined at several levels:

- **Individual level:** Victim and incident characteristics, handled with appropriate privacy safeguards.
- **Local level:** Neighbourhoods, municipal wards, police-station jurisdictions, or districts.
- **Regional level:** States, provinces, or other administrative divisions.
- **National level:** Country-level differences in victimisation, infrastructure, enforcement capacity, and offender activity.
- **Transnational level:** Connections among victims, infrastructure, financial intermediaries, offenders, and investigative authorities across countries.

Each level serves a different purpose. Local mapping may assist police departments in organising victim-support services and public-awareness programmes. Regional analysis may help allocate forensic resources and specialised investigators. National and transnational analysis may support strategic planning and cross-border cooperation.

The appropriate geographical scale depends on the research question, the reliability of the available data, and the intended policy application. Highly detailed maps are not automatically more informative: they can expose individuals, exaggerate small-area differences, or imply a level of precision unsupported by the underlying records.

2.4 Why geographical analysis matters

Geographical analysis can help policymakers identify recurring patterns, allocate limited investigative resources, understand differences among offence types, and assess whether prevention programmes are reaching the communities most affected.

For example, a concentration of digital-payment fraud complaints may justify targeted consumer education, improved bank-fraud escalation procedures, or additional victim-support capacity. Repeated ransomware incidents affecting a particular industrial sector may indicate a need for improved security controls and incident-response preparedness.

At the same time, hotspot identification should not be confused with causal explanation. A map can show where recorded incidents are concentrated without establishing why that concentration exists. Explanations require further investigation into social, economic, technological, institutional, and behavioural factors.

III. The Global Geography of Cybercrime

3.1 Challenges in identifying offender locations

Cybercrime presents a substantial attribution problem because digital communications can traverse multiple jurisdictions and use technical services designed to obscure the source of activity.

A single fraudulent transaction may involve a victim's device, a compromised account, a payment intermediary, a receiving account, and communication infrastructure in several countries. The apparent location of any one component may differ from the location of the person directing the offence.

Traditional cyberthreat reports frequently describe the geographical distribution of detected attacks, malicious infrastructure, or suspicious network traffic. These observations are valuable for technical defence but do not necessarily provide a representative picture of offender locations.

This distinction is identified as a central methodological problem in measuring the global geography of cybercrime. Their research explains why technical indicators alone cannot establish a reliable global measure of where cybercriminals operate.

3.2 The World Cybercrime Index

The World Cybercrime Index is an important contribution to the study of cybercrime geography. The survey, conducted from March to October 2021, asked respondents to identify countries they considered significant sources of five forms of profit-driven cybercrime and assess the impact, professionalism, and technical skill associated with offenders in those locations.

The five categories were:

1. Technical products and services.
2. Attacks and extortion.
3. Data and identity theft.
4. Scams.
5. Cashing out and money laundering.

The resulting index illustrates that cybercrime is not geographically uniform and that locations may be associated with different types of criminal activity. It also provides a framework for considering the local dimensions of a phenomenon commonly described as borderless.

The findings should nevertheless be interpreted within the limits of the research design. The index is based on expert judgements rather than a comprehensive census of offenders. Its authors acknowledge that the expert sample was relatively small and geographically uneven, with underrepresentation of some regions, including Africa and South America. It should therefore be treated as an important research instrument rather than a definitive measurement of all cybercrime worldwide.

The index also demonstrates why a country-level cybercrime ranking should not be interpreted as a ranking of the criminality of its population. The presence of organised criminal networks in a jurisdiction does not establish collective responsibility among its residents, and rankings do not measure the prevalence of offending among ordinary citizens.

3.3 International threat assessments

International organisations provide additional evidence about the changing nature of cybercrime.

Europol's *Internet Organised Crime Threat Assessment 2024* examines cyberattacks, online and payment fraud, and other forms of internet-enabled organised crime. It describes an environment in which criminal networks adapt to law-enforcement operations, reorganise their activities, and exploit the opportunities created by digital services and illicit online markets.

INTERPOL's 2025/2026 Asia and South Pacific Cyberthreat Assessment Report discusses the expansion of cybercrime in the region and the growing use of AI, sophisticated social engineering, and organised criminal business models. Its findings reinforce the need for stronger information-sharing arrangements, analytical capacity, and coordinated responses across jurisdictions.

These assessments are relevant to AI-based hotspot mapping because they show that cybercrime patterns change over time and differ across regions and offence types. A static map or single national ranking cannot adequately represent such changes.

A further implication is that cybercrime intelligence should be analysed temporally as well as geographically. A location that exhibits an unusual increase in reported incidents over a short period may deserve attention even if its total number of cases remains below that of a much larger urban centre.

3.4 Implications for AI-assisted mapping

The international evidence suggests that cybercrime mapping requires several complementary forms of analysis. Victim reports can identify the distribution of harm, technical telemetry can identify suspicious infrastructure, and verified investigative records can help establish offender associations. Financial intelligence may reveal patterns of transfers and cashing out, while expert assessments can contribute to understanding criminal ecosystems that remain difficult to observe directly.

Artificial intelligence can help integrate these sources, but it cannot remove the underlying uncertainties merely by processing more data. Its value depends on whether the information being combined

measures the same phenomenon, whether geographical attribution is reliable, and whether uncertainty is explicitly represented in the resulting maps.

IV. Artificial Intelligence and Geospatial Crime Analysis

4.1 From conventional mapping to intelligent analysis

Geographic information systems have long supported the visualisation and analysis of spatial information. Conventional GIS can display reported cybercrime cases by district, compare regional counts, identify clusters, and show changes across reporting periods.

AI extends these capabilities by enabling automated classification, anomaly detection, forecasting, and analysis of complex relationships across large datasets. A system can, for example, classify incoming complaint descriptions into offence categories, detect sudden increases in particular types of reported fraud, and estimate which areas may require additional victim-support resources.

The distinction between AI and GIS is important. GIS provides the spatial framework for storing, displaying, and analysing geographical data. AI provides methods for detecting patterns or making predictions. Their integration can improve analysis, but neither technology independently guarantees valid conclusions.

4.2 Machine learning

Machine learning involves algorithms that identify statistical relationships from data and use those relationships to classify observations, estimate outcomes, or detect unusual patterns.

For cybercrime mapping, supervised machine learning can be trained on labelled incident records to classify reports into categories such as phishing, identity theft, ransomware, or payment fraud. Where reliable historical labels are available, models can also estimate the likelihood that a new report belongs to a particular category.

Unsupervised learning can help identify groups of incidents with similar characteristics when reliable labels are unavailable or incomplete. For instance, clustering may reveal groups of reports that share a suspected fraud pattern, transaction characteristics, or technical indicators.

These methods can reduce the burden of manually reviewing large datasets. However, their output depends on the quality and representativeness of the training data. If historical records systematically underrepresent particular communities, languages, offence categories, or reporting channels, the model may reproduce those gaps.

Machine-learning predictions should therefore be evaluated against appropriate baselines and tested on data that were not used during training.

4.3 Natural-language processing

Cybercrime complaints frequently contain unstructured text, including descriptions of fraudulent messages, suspicious calls, impersonation attempts, payment instructions, and unauthorised account activity.

Natural-Language Processing (NLP) can assist analysts by extracting relevant information from these descriptions and classifying reports according to offence type. Language models may also help summarise long narratives, identify recurring themes, and support the translation or categorisation of multilingual complaints.

In a multilingual environment such as India, this capability may be particularly useful because complaints can be submitted in different languages and writing styles. Nevertheless, language models may misinterpret colloquial expressions, regional terminology, sarcasm, or incomplete narratives.

An NLP classification should therefore be treated as an initial analytical assessment rather than a final determination that an offence occurred. Where a classification affects investigation or resource allocation, appropriate human review and confidence reporting are essential.

4.4 Spatial statistics

Spatial statistics provides a means of determining whether the observed distribution of incidents differs meaningfully from a random or expected pattern.

Kernel density estimation, for example, produces a smoothed representation of incident concentration. Spatial autocorrelation measures whether nearby geographical units tend to have similar values. Local cluster statistics can help distinguish areas of relatively high values surrounded by other high-value areas from isolated observations.

These techniques can be used with cybercrime complaint data, but their interpretation requires care. A cluster of complaints may result from a genuine concentration of victimisation, a population concentration, a reporting campaign, or the location of a major reporting institution.

AI can complement spatial statistics by incorporating additional variables, such as population, internet access, digital-payment activity, reporting delays, and historical incident patterns. However, the choice of variables and geographical scale must be justified by the research question.

4.5 Deep learning and anomaly detection

Deep-learning models can analyse complex data, including sequences of transactions, communication patterns, or large volumes of network telemetry. Anomaly-detection systems can identify observations that differ substantially from established patterns.

These techniques may support the identification of unusual transaction clusters, coordinated phishing campaigns, or unexpected changes in incident reporting. They can also help security teams prioritise events for further review.

However, unusual activity is not necessarily criminal activity. A sudden increase in legitimate transactions, an unusual pattern of internet traffic, or a concentrated reporting campaign can resemble a threat signal. Automated alerts should consequently be followed by verification before they inform consequential decisions.

4.6 Network analysis

Cybercrime often involves relationships among multiple entities rather than isolated offenders. Network analysis represents these relationships as nodes and connections.

Depending on the lawful purpose and available evidence, nodes may represent reported incidents, accounts, devices, organisations, transactions, or verified entities. Connections may represent documented transfers, shared indicators, or other established relationships.

Network analysis can help investigators identify recurring links across incidents and examine how suspected criminal activity spans geographical boundaries. AI-assisted methods may identify groups of related events that would be difficult to discover through individual case reviews.

Nevertheless, a connection between two records does not automatically establish common ownership, intent, or criminal responsibility. Shared infrastructure may be used by unrelated parties, and common characteristics may arise from ordinary business practices. Analytical links must be corroborated before they are interpreted as evidence of offending.

4.7 Temporal analysis

Cybercrime patterns evolve with technological changes, public events, economic conditions, and criminal adaptation. Time-series analysis can help identify seasonal variation, emerging increases, or sudden changes in the distribution of reported incidents.

For example, analysts might examine whether online shopping fraud complaints increase during major sales periods, whether particular forms of impersonation increase during public benefit announcements, or whether ransomware incidents rise within a specific industrial sector.

Temporal analysis can also help distinguish persistent hotspots from short-lived clusters. This distinction is important for deciding whether a response should involve long-term investment, a temporary public-awareness campaign, or an immediate investigative intervention.

AI models should be periodically reassessed because relationships learned from historical data may become unreliable when criminal methods or reporting practices change.

V. Proposed Integrated AI Architecture

An effective AI-based cybercrime mapping platform can be organised into six connected layers.

Data layer: Collects authorised complaint records, verified incident information, relevant contextual statistics, and appropriately sourced threat intelligence.

Data-quality layer: Standardises classifications, removes duplicates, documents missing values, and separates observations according to their evidentiary status.

Geospatial layer: Assigns appropriate geographical references, stores different types of location separately, and supports spatial aggregation.

Analytical layer: Applies statistical hotspot detection, machine learning, anomaly detection, temporal analysis, and network analysis according to the stated objective.

Visualisation layer: Presents historical patterns, estimated risk, confidence information, and relevant contextual variables through maps, dashboards, and analytical reports.

Governance layer: Controls access, records model changes, supports audits, protects personal information, and establishes procedures for human review and correction.

The governance layer should operate across the entire platform rather than functioning as an optional final check.

A simplified workflow is:

Authorised data collection → validation and classification → geographical attribution → spatial and temporal analysis → AI-assisted assessment → human verification → proportionate intervention → outcome evaluation.

The architecture should preserve the distinction between evidence and inference at every stage. For example, a confirmed complaint is an observed report; a spatial cluster is an analytical finding; a forecast is an estimate; and an offender attribution is an investigative conclusion requiring its own evidentiary basis.
AI should not automatically convert the first three into the fourth.

VI. Applications of AI in Cybercrime Hotspot Mapping

6.1 Online financial fraud

Financial fraud is an important application because fraudulent schemes often generate multiple reports involving similar narratives, payment methods, or transaction patterns.
AI can classify complaint descriptions, identify repeated fraud indicators, and help analysts detect emerging concentrations of victimisation. When lawfully available and properly protected, financial-sector information may provide additional context about the distribution of reported harm.
Geographical analysis can help distinguish areas with large absolute numbers of complaints from areas with unusually high reporting rates. It may also identify locations where additional consumer education, reporting assistance, or victim-support services are required.
However, the location of a victim or receiving account should not automatically be interpreted as the location of the offender. Financial intermediaries, money mules, compromised accounts, and legitimate account holders may occupy different positions within a fraudulent transaction.

6.2 Phishing and digital impersonation

Phishing and impersonation campaigns exploit deceptive messages, websites, social-media accounts, or other communication channels to obtain information or induce victims to transfer money.
NLP can help classify complaints and extract recurring themes, while clustering techniques can identify reports that appear to be associated with similar campaigns. Temporal analysis can help establish whether a particular form of impersonation is increasing during a given period.
Mapping the locations of victims can reveal the distribution of harm and help target public-awareness campaigns. Mapping malicious infrastructure can support technical response, provided that infrastructure location is not confused with offender location.

6.3 Ransomware and attacks on organisations

Ransomware incidents can affect businesses, public institutions, hospitals, educational organisations, and critical infrastructure.
AI-assisted systems can combine verified incident records with sectoral and geographical information to identify concentrations of reported organisational harm. They can also support the prioritisation of incident-response assistance and resilience assessments.
The location of an affected organisation is particularly useful for understanding operational impact, while information about attacker infrastructure and suspected operators belongs in separate analytical layers.
Because ransomware incidents may be underreported or disclosed according to different institutional practices, comparisons across sectors and regions should account for differences in reporting obligations and incident-detection capacity.

6.4 Identity theft and account compromise

Identity theft and account compromise may be associated with compromised credentials, impersonation, unauthorised access, or misuse of personal information.
AI can support the classification of reported incidents, the detection of repeated patterns, and the analysis of changing complaint volumes. It can also help identify locations where people may benefit from targeted information about account security and recovery procedures.
The analysis should avoid exposing sensitive personal details or treating the repeated appearance of a person's account in different records as proof of wrongdoing. Victims may themselves be connected to multiple fraudulent events because their information has been stolen or misused.

6.5 Cybercrime affecting children and vulnerable groups

AI-assisted mapping may help authorities understand where particular forms of reported online exploitation or digital abuse are generating demand for specialised support.
Such analysis requires particularly strong safeguards. Sensitive case information should not be placed on public maps, and detailed geographical displays should not permit the identification of children, victims, or protected witnesses.
Aggregate analysis can support the planning of specialist services, training, public education, and reporting channels without exposing individual cases.

6.6 Organised cybercrime networks

Organised cybercrime may involve specialised roles, including the provision of technical services, access to compromised systems, financial facilitation, and laundering of proceeds.

Network analysis and AI-assisted entity resolution can help investigators identify relationships across verified cases and examine how different activities are distributed geographically.

These applications require strong provenance controls. A shared telephone number, account, technical indicator, or infrastructure resource may be relevant to an investigation but is not necessarily conclusive evidence that all associated entities are part of the same criminal organisation.

6.7 Resource allocation and public policy

Cybercrime mapping can assist institutions in deciding where to place specialised investigators, improve reporting facilities, expand forensic capacity, and deliver prevention programmes.

For example, an authority may use complaint rates and estimated workload to identify districts that need additional victim-support staff. A separate analysis of verified ransomware incidents may inform regional investment in incident-response capabilities.

Resource allocation should not rely exclusively on predicted risk. Authorities must also consider population needs, service accessibility, existing capacity, and the possibility that low reporting levels reflect barriers to reporting rather than low levels of harm.

6.8 Measuring the effectiveness of interventions

AI-based mapping can support the evaluation of cybercrime prevention programmes by comparing patterns before and after an intervention.

A public-awareness campaign might be evaluated using changes in reporting rates, victim knowledge, repeat victimisation, and the time between an incident and its report. A new incident-response team might be evaluated through response times, case outcomes, and the quality of investigative support.

A reduction in reported incidents is not automatically proof that crime has declined. It may also reflect reduced awareness, changes in reporting channels, or barriers to accessing assistance. Multiple indicators are needed to assess whether an intervention has achieved its intended result.

VII. Cybercrime Hotspot Mapping in India

7.1 Digital expansion and cybercrime exposure

India's growing digital economy has increased the importance of reliable cybercrime prevention and investigation. Digital payments, mobile banking, online marketplaces, social-media platforms, and internet-based public services provide substantial benefits, but they also create opportunities for financial fraud, impersonation, identity misuse, and other offences.

The geographical distribution of these harms is likely to reflect differences in population, economic activity, connectivity, digital-service adoption, reporting access, and enforcement capacity. It should not be assumed that the areas with the largest number of complaints necessarily have the highest underlying prevalence of cybercrime.

India's federal administrative structure adds another dimension. Cybercrime reporting and investigation involve institutions operating across state, district, and national levels. An effective mapping framework should therefore support coordination while respecting the responsibilities and legal authorities of the relevant agencies.

7.2 Official data and reporting infrastructure

The National Cyber Crime Reporting Portal is a relevant source of official information about India's cybercrime reporting framework. It provides an online mechanism for submitting complaints, including complaints concerning financial fraud, identity theft, hacking, and offences affecting women and children.

The National Crime Records Bureau (NCRB) compiles national crime statistics through its publication *Crime in India*. The Ministry of Home Affairs' March 2026 parliamentary information release identifies the 2023 edition as the latest published NCRB report at that time and provides state- and union-territory-level details of registered cybercrime cases for 2021–2023.

These sources serve different purposes. Complaint records can help researchers examine reported demand and emerging patterns, whereas registered-case statistics reflect offences recorded through the relevant law-enforcement processes. The two datasets should not be combined without reconciling their definitions, coverage, time periods, and units of observation.

73 A proposed Indian hotspot-mapping model

A national research programme could begin with aggregate district-level data and gradually incorporate more detailed information where access is lawful, data quality is sufficient, and privacy protections are in place.

The model could include the following components:

1. **Victimisation mapping:** Display reported cybercrime incidents and appropriate rates by district and state.
2. **Offence classification:** Separate financial fraud, phishing, identity theft, ransomware, unauthorised access, and other categories where the data support those distinctions.
3. **Temporal monitoring:** Compare monthly or quarterly patterns to identify significant changes in reporting and recorded incidents.
4. **Contextual analysis:** Examine relationships with population, digital connectivity, economic activity, and access to reporting facilities.
5. **Verified incident analysis:** Use appropriately authorised investigative data to examine recurring patterns without treating unverified indicators as confirmed offender locations.
6. **Service planning:** Identify locations where additional cybercrime awareness, reporting support, forensic capacity, or victim assistance may be needed.
7. **Evaluation:** Measure whether interventions improve reporting access, response times, victim outcomes, and investigative efficiency.

This framework should begin as a research and decision-support system, not as an automated mechanism for assigning criminality to geographical areas.

7.4 State and district comparisons

Comparing Indian states and districts requires careful attention to differences in scale and data collection.

A state with a large population and extensive digital-payment activity may record more complaints than a smaller state even when the underlying rate of victimisation is lower. Similarly, a district that has improved reporting access may show a rise in recorded cases because more victims are coming forward.

Researchers should therefore report absolute counts alongside appropriate rates, explain the denominator used, and examine changes over time. Where data are sparse, estimates should be presented with suitable uncertainty rather than converted into definitive rankings.

District-level maps can be useful for planning services, but the results should not be interpreted as proof that residents of a particular district are more likely to commit cybercrime.

7.5 Institutional coordination

An effective Indian system would require cooperation among law-enforcement agencies, cyber-forensic laboratories, financial institutions, telecommunications providers, technology companies, academic researchers, and relevant government bodies.

Data-sharing arrangements should specify lawful purposes, permitted uses, security requirements, retention periods, access rights, and procedures for correcting inaccurate information.

Common definitions and interoperable reporting formats would improve the comparability of records across jurisdictions. Independent research access to appropriately anonymised aggregate data could also strengthen methodological evaluation and public accountability.

7.6 Public-interest applications

The most immediate benefits of cybercrime mapping may arise from better prevention and victim support rather than from attempts to predict individual offending.

For example, an authority could use verified reporting patterns to identify areas where people need greater awareness of financial fraud, where businesses require additional cybersecurity guidance, or where victims face delays in obtaining assistance.

The results could support targeted educational programmes, improved complaint-handling procedures, and more equitable distribution of investigative resources. These applications use geographical analysis to respond to demonstrated needs without assuming that an area's residents are responsible for the crimes reported there.

VIII. Limitations and Risks of AI-Based Mapping

8.1 Underreporting and incomplete data

Cybercrime is often underreported. Victims may be unaware that an offence has occurred, may not know how to report it, or may avoid reporting because of embarrassment, fear, financial concerns, or doubts about the likelihood of recovery.

Businesses may also vary in their willingness to disclose incidents because of reputational, commercial, or legal considerations.

As a result, complaint-based maps describe reported cybercrime rather than the full underlying distribution of victimisation. AI cannot correct this problem automatically. In some circumstances, it may amplify reporting differences by treating the observed data as a complete representation of reality.

Researchers should examine reporting coverage, missingness, and possible differences between recorded incidents and the underlying phenomenon.

8.2 Geographical misattribution

The physical location of a technical resource is not necessarily the location of the person responsible for its use.
Cloud infrastructure, proxy services, compromised devices, and other intermediaries can complicate attribution. Geographical databases may also be inaccurate or outdated.
An AI system trained on unverified technical locations may consequently generate misleading offender maps. The problem is especially serious if such maps are used to justify enforcement action or stigmatise particular areas.
Separate geographical layers, confidence assessments, and independent corroboration are essential safeguards.

8.3 Algorithmic bias

Historical crime data reflect not only underlying events but also the practices and priorities of institutions that record them.
If some communities have greater access to reporting channels, their incidents may be more visible in administrative records. If enforcement resources are concentrated in particular areas, more offences may be detected there. Models trained on these records may learn institutional patterns rather than actual differences in criminal behaviour.
Bias can arise through data selection, feature design, classification errors, geographical aggregation, and the choice of performance measures.
Mitigation requires examining model performance across relevant groups and regions, documenting data limitations, and ensuring that high-risk predictions do not become self-reinforcing justifications for disproportionate surveillance.

8.4 False positives and false negatives

A false positive occurs when a model flags an event or location as suspicious when the relevant harmful activity is absent. A false negative occurs when the model fails to identify a genuine event or elevated risk.
Both have practical consequences. False positives may divert investigators, generate unnecessary scrutiny, or cause reputational harm. False negatives may leave victims without timely support or allow serious incidents to remain undetected.
Evaluation should consider the relative costs of these errors and the intended use of the model. A tool for prioritising educational outreach can tolerate different error trade-offs from a system used to initiate an investigation.

8.5 Privacy and data protection

Cybercrime datasets may contain names, telephone numbers, financial information, account identifiers, communication records, and sensitive descriptions of personal experiences.
Combining these records with geographical information increases the risk of re-identification. Even data that appear anonymous may expose individuals when joined with external datasets or displayed at a sufficiently detailed spatial scale.
Responsible systems should apply data minimisation, access controls, secure storage, appropriate retention limits, and privacy-preserving aggregation. Personal information should not be collected or retained merely because it might improve model performance.
The purpose of mapping should be defined before data are collected, and secondary uses should be subject to appropriate legal and governance review.

8.6 Transparency and explainability

Some AI systems provide complex outputs that are difficult to explain to investigators, policymakers, or affected communities.
A hotspot score may combine multiple features, historical patterns, and model assumptions. Without documentation, users may interpret the result as an objective measurement even when it contains substantial uncertainty.
Systems should disclose their data sources, definitions, analytical methods, evaluation results, known limitations, and intended uses. Where a prediction informs a consequential decision, the responsible authority should be able to explain why the result was considered relevant and how it was independently assessed.

8.7 Changing criminal behaviour

Cybercriminals adapt to security measures, reporting systems, and law-enforcement interventions. A model trained on historical patterns may become less accurate when offenders change methods, exploit new platforms, or move activity to different channels.

Continuous monitoring and periodic retraining may help, but frequent updates also create risks if the model begins learning from poor-quality or manipulated data.

Model changes should be documented, tested, and independently evaluated. The system should retain enough historical information to determine whether its performance has improved or deteriorated.

8.8 Overreliance on predictive policing

Predictive policing raises concerns when model outputs are treated as authoritative evidence of future offending or individual culpability.

An area-level risk estimate is a statistical statement about a defined outcome under specified assumptions. It does not establish that a particular person is likely to commit a crime.

Using hotspot maps to justify intrusive action against residents solely because they live in a designated area would be both analytically unsound and ethically problematic. Geographical predictions should inform proportionate, evidence-based interventions and should not substitute for individualised evidence where the law requires it.

8.9 Resource and capacity constraints

Advanced AI systems require reliable data infrastructure, trained analysts, cybersecurity controls, technical maintenance, and institutional coordination.

Organisations with limited resources may find it difficult to sustain complex models or validate their outputs. A simpler statistical system with transparent assumptions may be more useful than an advanced model that cannot be audited or maintained.

Implementation should therefore be guided by demonstrated operational needs, available expertise, and the expected benefits relative to cost.

8.10 Limited evidence of causal effectiveness

A model may accurately identify historical clusters without reducing future cybercrime.

To establish practical value, researchers must determine whether interventions informed by the model improve relevant outcomes compared with alternative approaches. They must also examine whether crime is displaced, reporting changes, or benefits arise from factors unrelated to the AI system.

This distinction between prediction and intervention is central to responsible evaluation. Accurate prediction is a technical achievement; effective prevention requires additional evidence.

IX. Challenges and Limitations

Despite its potential, AI-based cybercrime mapping has several limitations.

- **Underreporting:** Many cybercrimes are never reported. A location with more complaints may have better reporting systems rather than a higher actual crime rate.
- **Geographic uncertainty:** The location of a victim, attacker, server, and financial account may differ. IP addresses and other digital indicators do not always establish an offender's physical location.
- **Data quality:** Incomplete, outdated, duplicated, or inconsistent records can produce misleading maps.
- **Algorithmic bias:** Models may overemphasize places that generate more recorded data and overlook less-reported forms of cybercrime.
- **Privacy concerns:** Precise incident locations and personal details can expose victims to further harm if sensitive data is disclosed.
- **Changing criminal behavior:** Cybercriminals continuously adapt their methods, so models trained on historical data can become less accurate.
- **Limited explainability:** Some AI models make predictions that are difficult to interpret, potentially complicating their use in investigations.

The National Institute of Standards and Technology (NIST) recommends managing AI risks throughout the system lifecycle, with attention to reliability, transparency, privacy, fairness, and security. These principles are relevant when designing cybercrime mapping systems.

X. Ethical and Legal Considerations

10.1 Presumption of innocence

A hotspot map must not be treated as proof of criminal conduct. Statistical association, geographical proximity, and technical linkage are not equivalent to evidence sufficient to establish responsibility for an offence.

Systems should clearly distinguish reported allegations, confirmed incidents, investigative hypotheses, and established findings. This distinction is especially important when maps are shared beyond specialist investigative teams.

10.2 Proportionality and non-discrimination

Interventions should be proportionate to the evidence and legitimate public objective. The use of AI should not result in collective suspicion based on residence, socioeconomic characteristics, language, or other unrelated personal attributes.

Public-awareness programmes, reporting assistance, and cybersecurity support may be appropriate responses to elevated victimisation risk. Intrusive investigative measures require their own lawful basis and evidentiary justification.

10.3 Data governance in India

AI-assisted cybercrime mapping in India should be designed with regard to applicable criminal procedure, electronic-evidence requirements, data-protection law, and relevant government policies.

The Digital Personal Data Protection Act, 2023, is an important part of India's personal-data governance framework. Its applicability to a particular processing activity, together with the status of relevant rules, exemptions, and commencement provisions, should be assessed against the law in force at the time of implementation.

Authorities should obtain appropriate legal review before combining personal data from different institutions or using such data for new analytical purposes. Research and operational access should be governed by clear authorisation and accountability arrangements.

10.4 Accountability and human oversight

Responsibility for decisions must remain with accountable institutions and authorised personnel.

Human review should not be a merely formal approval of automated outputs. Reviewers need sufficient information, time, expertise, and authority to challenge a prediction, correct errors, and request additional verification.

A documented process should allow affected individuals to seek correction where inaccurate data or inappropriate use has caused harm, subject to applicable legal and investigative restrictions.

10.5 Responsible public communication

Publishing a map of cybercrime hotspots can influence public perceptions of safety, business investment, and the reputation of communities.

Public-facing maps should explain what is being measured, when the data were collected, and what conclusions cannot be drawn from them. They should avoid language that equates high complaint volumes with the criminality of local populations.

Where the evidence is incomplete, the publication should state that limitation directly rather than concealing uncertainty behind precise-looking scores or colour-coded rankings.

XI. Policy Recommendations

The effective use of AI in cybercrime hotspot mapping requires coordinated improvements in data quality, analytical capability, institutional governance, and evaluation.

11.1 Develop standardised data definitions

Authorities should establish common definitions for major cybercrime categories, reporting periods, geographical fields, and incident status. Standardisation would make it easier to compare records across jurisdictions and reduce errors caused by inconsistent classification.

Data dictionaries should specify whether records represent complaints, unique incidents, affected accounts, registered cases, or confirmed offences. These units must not be treated as interchangeable.

11.2 Establish secure data-sharing arrangements

Relevant agencies should develop lawful and secure procedures for sharing appropriately limited information. Agreements should define permitted purposes, access controls, retention periods, audit requirements, and responsibility for data accuracy.

Where detailed personal information is unnecessary, aggregate or pseudonymised data should be preferred, subject to a realistic assessment of re-identification risks.

11.3 Invest in analytical capacity

Cybercrime mapping requires expertise in GIS, statistics, machine learning, digital forensics, criminology, and data governance.

Training programmes should help investigators and analysts understand the limits of geographical attribution, the difference between observed and predicted patterns, and the risks of biased data.

Universities and research institutions can contribute through methodological evaluation, independent audits, and interdisciplinary research.

11.4 Begin with limited, measurable pilot projects

Rather than deploying a nationwide predictive system immediately, authorities should begin with clearly defined pilot projects.

A pilot might examine aggregate financial-fraud reporting patterns across a selected group of districts. Its success criteria could include data completeness, classification quality, usefulness to service planning, privacy protection, and measurable improvements in response.

The pilot should be independently evaluated before expansion. If the model does not outperform simpler alternatives or introduces unacceptable risks, it should be revised or discontinued.

11.5 Separate victimisation maps from offender intelligence

Public-service planning and offender investigation should use distinct analytical products.

Victimisation maps can support education, assistance, and resource allocation. Offender-associated intelligence requires stricter access controls, corroboration, and evidentiary review.

The separation reduces the risk that a geographical concentration of reported harm will be misinterpreted as evidence of criminal responsibility.

11.6 Introduce independent audits

AI systems should be reviewed for predictive performance, data quality, geographical stability, privacy risks, and disparate impacts.

Audits should examine whether the model performs differently across districts, demographic groups, languages, and offence categories. They should also assess whether changes in reporting practices are being mistaken for changes in underlying crime.

Audit findings and corrective actions should be documented, with public reporting where appropriate and consistent with legitimate confidentiality requirements.

11.7 Strengthen international cooperation

Cybercrime frequently involves multiple jurisdictions, so national hotspot mapping should be complemented by appropriate international cooperation.

Information-sharing arrangements can support the investigation of transnational fraud, ransomware, and criminal financial networks. INTERPOL's regional initiatives illustrate the importance of coordinated operations, shared threat assessments, and capacity-building efforts.

International cooperation should include common standards for evidence handling, lawful information exchange, and the interpretation of technical indicators.

11.8 Prioritise prevention and victim support

AI-derived findings should lead to practical interventions where evidence supports them.

Possible interventions include targeted digital-literacy campaigns, improved reporting access, support for affected businesses, stronger bank-fraud response processes, and expanded specialist assistance for victims.

These measures can address demonstrated needs without relying on speculative assumptions about the identities or intentions of people living in particular areas.

11.9 Evaluate outcomes rather than technological sophistication

Institutions should assess systems by their contribution to public safety and service quality, not by the complexity of their algorithms.

Relevant measures include improved response times, better classification of complaints, more efficient allocation of support, reduced repeat victimisation, and demonstrable gains in investigative effectiveness.

Where simpler statistical methods achieve comparable results with lower cost and greater transparency, they may be preferable to complex AI models.

XII. Conclusion

Artificial Intelligence, machine learning, and Geographic Information Systems can provide valuable tools for analyzing cybercrime patterns and mapping areas with elevated reported risk. By combining incident classification, anomaly detection, spatial statistics, and interactive visualization, these technologies can support cybersecurity planning, fraud prevention, and law enforcement investigations.

However, cybercrime mapping is more complex than conventional geographic crime mapping because digital offenses frequently cross national borders and involve multiple locations. Underreporting, inaccurate geolocation, data quality, privacy, and algorithmic bias must therefore be addressed before the results are used for operational decisions.

Ultimately, AI should serve as a decision-support tool rather than an unquestionable authority. Reliable data, transparent methods, human oversight, and independent evaluation are essential for developing useful and responsible cybercrime hotspot mapping systems.

References

[1]. Federal Bureau of Investigation. (2025). *2024 Internet Crime Report*. Internet Crime Complaint Center (IC3).
[2]. United Nations Office on Drugs and Crime. (2013). *Comprehensive Study on Cybercrime: Draft, February 2013*. United Nations Office on Drugs and Crime. .
[3]. Tabassi, E. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. National Institute of Standards and Technology, NIST AI 100-1. .
[4]. Hu, Y., Wang, F., Guin, C., & Zhu, H. (2018). A spatio-temporal kernel density estimation framework for predictive crime hotspot mapping and evaluation. *Applied Geography, 99*, 89–97.
[5]. Nakaya, T., & Yano, K. (2010). Visualising crime clusters in a space-time cube: An exploratory data-analysis approach using space-time kernel density estimation and scan statistics. *Transactions in GIS, 14*(3), 223–239. .
[6]. National Institute of Standards and Technology. (n.d.). *AI Risk Management Framework*. .
[7]. NIST
[8]. Hana, Y., Hu, Y., Zhu, H., & Wang, F. (2023). A cyclically adjusted spatio-temporal kernel density estimation method for predictive crime hotspot analysis. *Annals of GIS, 29*(2), 177–191.