Unit 7: Data Mining and Ethics

2-Hour Lecture for AI Undergraduates

Sunil Regmi, Lecturer

09:41 PM +0545, Wednesday, July 02, 2025

Lecture Overview

- Outline:
  - Data Privacy and Security
  - Social Impacts of Data Mining
  - Algorithmic Bias & Biased Datasets
  - Transparency, Accountability, and Governance
  - 5 Principles of Data Mining Ethics
  - Ethical Use of Algorithms
  - Conclusion and User Protection

Data Privacy and Security: Definitions

- **Data Privacy**: Appropriate handling of personal data to ensure user control over access and use.
  - Example: User consent for cookie tracking on a website.
  - Focus: User rights and consent.
- **Data Security**: Protecting data from unauthorized access, corruption, or theft using technical measures.
  - Example: Encrypting a customer database with AES-256.
  - Focus: Data integrity and confidentiality.
- Importance: Safeguards individual rights and ensures compliance with legal standards.

Data Privacy and Security: Key Differences

- **Focus**:
  - Privacy: User control and consent.
  - Security: Protection from external threats.
- **Objective**:
  - Privacy: Empowers users with data rights (e.g., right to erasure).
  - Security: Maintains data integrity and availability.
- **Implementation**:
  - Privacy: Policies and consent mechanisms (e.g., pop-up agreements).
  - Security: Technical safeguards (e.g., firewalls, authentication).
- **Responsibility**:
  - Privacy: Shared between organizations and users.
  - Security: Primarily organizational.
- Example: A company informs users of data use (privacy) and encrypts it (security).

Data Privacy and Security: Importance

- **Data Privacy**:
  - Protects individual rights and builds trust.
  - Ensures compliance with laws (e.g., GDPR fines up to €20M).
  - Example: A healthcare firm protects patient consent to avoid legal issues.
- **Data Security**:
  - Prevents financial and reputational damage from breaches.
  - Ensures business continuity.
  - Example: Equifax breach (2017) cost $575M due to unsecured data.
- Interdependence: Privacy requires secure data handling to be effective.

Data Privacy and Security: Techniques

- **Data Privacy Techniques**:
  - Anonymization: Removes identifiers (e.g., k-anonymity with k=5).
  - Differential Privacy: Adds noise to protect individual data (e.g., Google analytics).
  - Consent Management: Opt-in/opt-out options (e.g., website pop-ups).
  - Example: A retail site anonymizes customer names before mining patterns.
- **Data Security Techniques**:
  - Encryption: Unreadable data formats (e.g., AES-256).
  - Access Controls: Passwords or MFA.
  - Firewalls/Intrusion Detection: Blocks unauthorized access.
  - Example: A bank uses MFA and encryption for transaction data.

Data Privacy and Security: Examples and Context

- **Data Privacy Example**:
  - Cambridge Analytica (2018): Misused 87M Facebook users' data without consent.
  - Lesson: Consent is non-negotiable in ethical data mining.
- **Data Security Example**:
  - Equifax Breach (2017): Exposed 147M records due to unpatched software.
  - Lesson: Security lapses undermine privacy efforts.
- **Combined Impact**: Illegal web scraping risks both privacy (no consent) and security (vulnerable storage).

Data Privacy and Security: Related Concepts

- **Overlap**: Privacy and security are interdependent.
  - Example: A secure but non-consented database still violates privacy.
- **Legal Frameworks**:
  - GDPR (2018): Privacy rights (e.g., data portability) and security obligations (72-hour breach notification).
  - AI Act (Proposed): Regulates AI data practices.
  - Nepal's Data Privacy Bill (Draft): Requires consent and secure storage.
- **Ethical Link**: Aligns with ownership and transparency principles.

Data Privacy and Security: Challenges and Best Practices

- **Challenges**:
  - Privacy: Balancing personalization with user control reduces revenue.
  - Security: Strong measures (e.g., encryption) increase costs.
  - Trade-Off: Over-securing limits usability; lax privacy erodes trust.
- **Best Practices**:
  - Privacy: Obtain consent, provide transparent policies, offer opt-out options.
  - Security: Regular updates, audits, staff training.
  - Integrated: Combine anonymization and encryption.
  - Example: A social media platform uses consent pop-ups and end-to-end encryption.

Social Impacts of Data Mining: Positive

- **Innovation**: Advances research and technology.
  - Example: Mining genomic data to develop personalized medicine.
  - Impact: Faster disease detection (e.g., cancer predictors).
- **Personalized Services**: Enhances user experience.
- Example: Amazon recommending products based on browsing history.
- **Fraud Detection**: Protects financial systems.
- Example: Banks flagging unusual transactions (e.g., $10Kwithdrawalinanewlocation$).

Social Impacts of Data Mining: Negative

- **Surveillance**: Enables mass monitoring.
  - Example: NSA's PRISM program collecting call metadata.
  - Concern: Erosion of personal freedom.
- **Job Displacement**: Automates roles.
- Example: AI chatbots replacing call center jobs.
- **Manipulation**: Influences decisions.
- Example: Targeted ads swaying voter preferences (Cambridge Analytica).

Social Impacts: Case Studies

- **Cambridge Analytica (2018)**: Used 87M Facebook profiles to target voters.
  - Outcome: Led to data protection laws (e.g., GDPR).
  - Debate: Innovation vs. privacy violation.
- **Predictive Policing**: Analyzes crime data for hotspots.
  - Example: Chicago's system over-targeted minority neighborhoods.
  - Outcome: Raised fairness concerns.

Algorithmic Bias & Biased Datasets: Sources

- **Training Data Bias**: Reflects societal inequalities.
  - Example: Hiring datasets with 80
  - Issue: Models learn from historical disparities.
- **Historical Data**: Reinforces past discrimination.
- Example: Loan data favoring high-income groups.
- Impact: Perpetuates unfair decisions in AI systems.

Algorithmic Bias: Examples

- **Lending**: Higher denial rates for minorities.
  - Example: 2019 study showed 20
  - Cause: Biased credit scoring models.
- **Hiring**: Bias in resume screening.
- Example: Amazon's AI rejected female resumes due to male-dominated training data.
- **Facial Recognition**: Accuracy disparities.
- Example: NIST (2018) reported 34

Algorithmic Bias: Mitigation

- **Fairness Metrics**: Ensure equity.
  - Example: Equal opportunity (same true positive rate across genders).
  - Metric: Demographic parity (equal selection rates).
- **Techniques**: Re-sampling, bias correction.
- Example: Oversampling minority groups in training data.
- **Ethics in AI**: Inclusive design.
- Example: Diverse teams in facial recognition development.

Transparency, Accountability, and Governance: XAI

- **Explainable AI (XAI)**: Makes models interpretable.
  - Example: SHAP values show feature importance in loan decisions.
  - Tool: LIME explains predictions locally for black-box models.
- Need: Ensures trust, compliance with laws (e.g., GDPR's right to explanation).
- Challenge: Balancing interpretability with accuracy.

Transparency, Accountability, and Governance: Responsibility

- **Data Scientists**: Ensure ethical data practices.
  - Example: Auditing datasets for bias before training.
- **Developers**: Build secure systems.
- Example: Implementing end-to-end encryption in data pipelines.
- **Institutions**: Enforce policies.
- Example: Universities mandating ethics training for AI projects.

Transparency, Accountability, and Governance: Legal Frameworks

- **GDPR (EU, 2018)**: Grants data rights (e.g., erasure).
  - Example: Users can request data deletion from social media.
- **AI Act (EU, Proposed)**: Regulates high-risk AI.
- Example: Requires transparency for AI in hiring.
- **Nepal's Data Privacy Bill (Draft)**: Local data protection.
- Example: Mandates consent for data collection.
- Role: Enforces accountability and transparency.

Ethical Principles in Data Mining

Five Core Principles:

- Transparency
- Privacy
- Accountability
- Fairness
- Consent

These principles guide ethical data collection, usage, and sharing practices in the field of data mining.

Transparency

Definition: Individuals should understand how their data is being collected, used, and shared.

Explanation:

- Ensure clear communication of data practices.
- Provide users with access to data usage policies.

Example:

- A company explains what data its fitness tracker collects and how it is used to personalize health tips.

Privacy

Definition: Respect individuals' privacy by limiting and securing personal data.

Explanation:

- Avoid collecting unnecessary personal data.
- Use encryption and anonymization techniques.

Example:

- A health app stores only non-identifiable data for research purposes.

Accountability

Definition: Organizations must be responsible for their data handling practices.

Explanation:

- Take ownership of data misuse or breaches.
- Set up audit trails and compliance protocols.

Example:

- A company issues a public notice and compensates users after a data breach.

Fairness

Definition: Ensure unbiased and equitable data handling and outcomes.

Explanation:

- Avoid discrimination in data analysis.
- Use bias-detection tools in model training.

Example:

- An HR tool is tested to ensure it does not favor candidates based on gender or race.

Consent

Definition: Users should have control over their data and grant informed permission.

Explanation:

- Allow users to opt-in/opt-out.
- Provide clear consent forms with data usage details.

Example:

- A website asks users to accept or reject cookies before tracking behavior.

Responsible Web Scraping

Ethics in Data Collection:

- Follow website terms of service.
- Avoid scraping personal or sensitive data without permission.
- Use rate limits to avoid server disruption.

Example: Scraping public product prices for comparison, not user reviews or emails.

Ethical Use of Algorithms

- **Training**: Avoid unrepresentative datasets.
  - Example: Oversampling minorities to balance gender data.
- **Code**: Check for unintentional bias.
- Example: Reviewing code for hardcoded preferences (e.g., favoring certain zip codes).
- **Feedback**: Prevent biased learning.
- Example: Filtering biased user ratings in recommendation systems.
- **Goal**: Align algorithms with data science ethics.

Conclusion and User Protection

- **Ethical Need**: Protects users from sensitive data exposure.
  - Example: Illegal scraping risking personal details.
- **Business Benefit**: Avoids legal hassles (e.g., fines under GDPR).
- **User Action**: Use proxies (e.g., Smartproxy) to mask digital footprints.
- **Takeaway**: Ethical practices balance innovation with responsibility.

Discussion Questions

- Should businesses prioritize user consent over data-driven profits? Why?
- Can predictive policing justify surveillance if it reduces crime?
- Is it ethical to use historical data with known biases? Discuss solutions.
- Who should enforce algorithm transparency: governments or companies?
- How can organizations ensure ethical outcomes without sacrificing efficiency?

Real-World Scenario Analysis

- **Scenario:** A retail company uses web scraping to collect customer reviews. The data reveals a decline in brand perception, prompting a new ad campaign. However, the scraping was done without user consent, risking legal action.
- **Task:** Analyze the ethical dilemma. Should the company halt the campaign, seek retroactive consent, or proceed? Propose a solution balancing ethics and business needs.
- **Time:** 20 minutes for group discussion.

Conclusion

- Data mining ethics is vital for responsible AI development.
- Critical evaluation of privacy, bias, and governance shapes future innovations.
- Next Steps: Investigate XAI tools and local ethical guidelines.