Undergraduate Certificate in AI for Public Policy and Governance:

Algorithm – A step-by-step procedure for solving a problem or performing a computation. Related terms: deterministic algorithm, heuristic, complexity. Explanation: Algorithms are the core of AI systems; they dictate how data is processed to produce outputs. Example: A sorting algorithm arranges data records by date. Practical application: Search engines use ranking algorithms to order results. Challenges: Ensuring efficiency, avoiding bias, and maintaining transparency.

Artificial Intelligence (AI) – The field of study focused on creating machines that can perform tasks requiring human intelligence. Related terms: machine learning, deep learning, neural networks. Explanation: AI encompasses a range of techniques from rule-based systems to data-driven models. Example: An AI chatbot answers citizen inquiries about public services. Practical application: Predictive analytics for traffic congestion. Challenges: Ethical considerations, data privacy, and interpretability.

Artificial Neural Network (ANN) – A computational model inspired by the structure of biological neurons. Related terms: deep neural network, backpropagation, activation function. Explanation: ANNs consist of layers of interconnected nodes that transform inputs into outputs. Example: A feed-forward network classifies images of public infrastructure. Practical application: Detecting anomalies in utility consumption patterns. Challenges: Overfitting, high computational demand, and explainability.

Bias (Algorithmic) – Systematic error that leads to unfair outcomes for certain groups. Related terms: fairness, discrimination, bias mitigation. Explanation: Bias can arise from skewed training data or flawed model design. Example: An AI hiring tool that favors candidates from a particular university. Practical application: Auditing public-policy models for equitable impact. Challenges: Identifying hidden biases, balancing trade-offs, and regulatory compliance.

Big Data – Extremely large and complex datasets that exceed traditional processing capabilities. Related terms: volume, velocity, variety, veracity. Explanation: Big data provides the raw material for AI-driven policy analysis. Example: Real-time sensor feeds from smart city infrastructure. Practical application: Analyzing crime patterns across metropolitan areas. Challenges: Storage costs, data quality, and privacy protection.

Classification – The task of assigning items to predefined categories. Related terms: supervised learning, label, confusion matrix. Explanation: Classification models learn from labeled examples to predict categories for new data. Example: Categorizing public comments as supportive, neutral, or opposed. Practical application: Automating the triage of service requests. Challenges: Imbalanced classes, label noise, and interpretability.

Clustering – Grouping similar data points without pre-assigned labels. Related terms: unsupervised learning, k-means, hierarchical clustering. Explanation: Clustering reveals hidden structures in data, useful for exploratory analysis. Example: Segmenting neighborhoods by socioeconomic indicators. Practical application: Targeted outreach for community programs. Challenges: Determining the optimal number of clusters and handling high-dimensional data.

Computational Ethics – The study of moral implications of algorithmic decision-making. Related terms: algorithmic accountability, responsible AI, ethical frameworks. Explanation: It guides the design of AI systems that align with societal values. Example: Designing a policy-impact model that respects privacy. Practical application: Drafting guidelines for AI use in welfare eligibility. Challenges: Translating abstract principles into concrete technical constraints.

Confidentiality – The principle of protecting personal data from unauthorized disclosure. Related terms: privacy, data protection, GDPR. Explanation: Confidentiality is a core requirement in public-policy data pipelines. Example: Encrypting health records before analysis. Practical application: Secure sharing of citizen feedback with policymakers. Challenges: Balancing transparency with privacy, and managing consent.

Correlation vs. Causation – Distinguishing statistical association from a direct cause-effect relationship. Related terms: confounding variable, counterfactual, regression analysis. Explanation: AI models may uncover correlations that do not imply policy-relevant causality. Example: Increased bike-share usage correlates with lower traffic accidents, but may be driven by weather. Practical application: Informing evidence-based policy decisions. Challenges: Designing experiments or quasi-experiments to infer causality.

Data Governance – The set of policies, standards, and processes that ensure data quality and responsible use. Related terms: data stewardship, metadata, data lineage. Explanation: Effective governance is essential for trustworthy AI in the public sector. Example: A municipal data portal that tracks provenance of traffic datasets. Practical application: Establishing data sharing agreements between agencies. Challenges: Coordinating across jurisdictions, and maintaining up-to-date documentation.

Data Literacy – The ability to read, interpret, and critically evaluate data. Related terms: statistical reasoning, visualization, numeracy. Explanation: Policymakers need data literacy to understand AI outputs and their implications. Example: Interpreting a heat map of pollution levels. Practical application: Training workshops for elected officials. Challenges: Overcoming technical jargon and fostering a culture of evidence-use.

Data Privacy – Protecting personal information from misuse while allowing legitimate analysis. Related terms: anonymization, differential privacy, consent. Explanation: Public-policy AI must comply with legal frameworks and ethical norms. Example: Publishing aggregated crime statistics without revealing individual identities. Practical application: Deploying privacy-preserving analytics for health data. Challenges: Balancing data utility with privacy guarantees.

Data Quality – The degree to which data accurately represents the real-world phenomenon it intends to capture. Related terms: accuracy, completeness, consistency. Explanation: Poor data quality undermines AI model reliability and policy relevance. Example: Incomplete census records leading to misallocation of resources. Practical application: Implementing validation rules for citizen-report forms. Challenges: Detecting and correcting errors, and handling missing values.

Data Scientist – A professional who extracts insights from data using statistical, computational, and domain-specific expertise. Related terms: machine learning engineer, analyst, researcher. Explanation: In the AI-policy context, data scientists collaborate with policymakers to translate findings into action. Example: Building a model that predicts the impact of a new tax policy. Practical application: Conducting impact assessments for social programs. Challenges: Communicating technical results to non-technical stakeholders.

Decision Support System (DSS) – A computer-based tool that assists decision-makers by providing relevant data, models, and visualizations. Related terms: dashboard, scenario analysis, what-if modeling. Explanation: AI-enhanced DSS can simulate policy outcomes under varying assumptions. Example: A platform that forecasts housing affordability based on zoning changes. Practical application: Real-time allocation of emergency resources. Challenges: Ensuring model validity, user adoption, and avoiding over-reliance on automation.

Deep Learning – A subset of machine learning that employs multi-layer neural networks to learn hierarchical representations. Related terms: convolutional neural network, recurrent neural network, autoencoder. Explanation: Deep models excel at processing unstructured data such as images, audio, and text. Example: Analyzing satellite imagery to detect illegal dumping sites. Practical application: Automated transcription of public hearing recordings. Challenges: Large data requirements, opacity of decisions, and energy consumption.

Diffusion of Innovation – A theory describing how new ideas and technologies spread through societies. Related terms: adoption curve, early adopters, network effects. Explanation: Understanding diffusion helps policymakers design interventions that accelerate beneficial AI uptake. Example: Tracking how municipalities adopt AI-driven traffic-management tools. Practical application: Crafting incentive programs for AI pilots in local governments. Challenges: Measuring adoption rates and addressing resistance.

Ethical AI – The practice of developing and deploying AI systems that are fair, transparent, accountable, and respect human rights. Related terms: responsible AI, AI governance, trustworthy AI. Explanation: Ethical AI frameworks guide the public sector in avoiding harms while leveraging AI benefits. Example: Publishing model documentation for a predictive policing system. Practical application: Establishing an ethics board for AI projects in a city council. Challenges: Operationalizing abstract principles and reconciling competing stakeholder interests.

Explainable AI (XAI) – Techniques that make the inner workings of AI models understandable to humans. Related terms: interpretability, model transparency, feature importance. Explanation: XAI is crucial for public-policy contexts where accountability and trust are paramount. Example: Using SHAP values to show which variables drive a welfare-eligibility score. Practical application: Providing policymakers with rationales for automated decisions. Challenges: Balancing explanatory depth with model performance.

Feature Engineering – The process of creating informative variables from raw data to improve model performance. Related terms: feature selection, dimensionality reduction, transformations. Explanation: Good features often capture domain knowledge that raw data alone cannot convey. Example: Deriving “distance to nearest hospital" from GIS coordinates. Practical application: Enhancing a model that predicts emergency-service response times. Challenges: Time-consuming, risk of leakage, and maintaining reproducibility.

Governance (AI) – The structures, policies, and processes that oversee AI development and deployment. Related terms: AI strategy, regulation, risk management. Explanation: In public policy, governance ensures AI aligns with societal goals and legal standards. Example: A municipal AI charter outlining permissible uses of surveillance data. Practical application: Conducting periodic audits of AI systems used for social benefits. Challenges: Keeping pace with rapid technological change and coordinating across agencies.

Ground Truth – Accurate, real-world data used as a benchmark to evaluate AI model performance. Related terms: labeling, validation set, reference data. Explanation: Ground truth is essential for training supervised learning models and for assessing bias. Example: Manually annotated images of road conditions used to train a classification model. Practical application: Validating an AI system that predicts flood risk. Challenges: High labeling cost, subjectivity, and ensuring representativeness.

Human-in-the-Loop (HITL) – A design approach that incorporates human judgment into AI decision processes. Related terms: oversight, collaborative AI, feedback loop. Explanation: HITL helps mitigate errors and maintain accountability, especially in high-stakes policy domains. Example: An analyst reviews AI-generated eligibility recommendations before final approval. Practical application: Real-time moderation of citizen-generated content on government portals. Challenges: Designing efficient interfaces and preventing automation bias.

Impact Assessment – Systematic evaluation of the potential effects of a policy or technology before implementation. Related terms: cost-benefit analysis, risk assessment, evaluation framework. Explanation: AI-driven impact assessments leverage predictive models to forecast outcomes. Example: Estimating job displacement from automated public-service workflows. Practical application: Guiding legislative debates on AI adoption in social services. Challenges: Data scarcity, uncertainty quantification, and stakeholder buy-in.

Inference – The process of applying a trained AI model to new, unseen data to generate predictions. Related terms: prediction, deployment, runtime. Explanation: Inference speed and scalability are critical for real-time public-policy applications. Example: Predicting traffic congestion levels during rush hour. Practical application: Triggering dynamic signal control based on live predictions. Challenges: Latency constraints, model drift, and hardware limitations.

Institutional Review Board (IRB) – A committee that reviews research involving human participants to ensure ethical standards. Related terms: ethical clearance, informed consent, risk mitigation. Explanation: AI projects using citizen data often require IRB approval to safeguard privacy. Example: A study analyzing social-media posts for sentiment on public health measures. Practical application: Securing clearance before deploying a pilot AI survey tool. Challenges: Navigating bureaucratic timelines and aligning research goals with public interest.

Integration (Systems) – Combining AI components with existing IT infrastructure to create a cohesive solution. Related terms: API, interoperability, middleware. Explanation: Seamless integration reduces duplication and maximizes the utility of AI insights. Example: Linking a predictive policing model with the city's dispatch system. Practical application: Automating data flow from sensors to a policy-analytics dashboard. Challenges: Legacy system constraints, data format mismatches, and security concerns.

Interpretability – The degree to which a human can understand the cause of a model's output. Related terms: explainability, transparent modeling, model introspection. Explanation: High interpretability is often required for regulatory compliance in government settings. Example: A linear regression where coefficients directly indicate impact of variables. Practical application: Justifying budget allocations based on model-derived insights. Challenges: Trade-offs with complex, high-performing black-box models.

Knowledge Graph – A network-based representation of entities and their relationships, often used to support reasoning. Related terms: ontology, semantic web, triples. Explanation: Knowledge graphs can encode policy domains, enabling AI to answer complex queries. Example: Linking legislation, agencies, and affected demographic groups in a graph. Practical application: Assisting officials in locating relevant statutes for a case. Challenges: Curating accurate relationships and maintaining updates.

Legislation (AI) – Laws and regulations that govern the development, deployment, and use of AI technologies. Related terms: regulatory framework, compliance, policy. Explanation: Policymakers must stay informed of emerging AI legislation to ensure lawful operations. Example: The EU's AI Act setting risk-based obligations for public authorities. Practical application: Conducting compliance reviews for municipal AI projects. Challenges: Interpreting ambiguous language and adapting to cross-jurisdictional rules.

Logistic Regression – A statistical model used for binary classification that estimates the probability of an outcome. Related terms: odds ratio, sigmoid function, maximum likelihood. Explanation: Logistic regression offers interpretability and is often a baseline for policy-focused models. Example: Predicting whether a household qualifies for a subsidy. Practical application: Screening applicants for social-housing programs. Challenges: Linear decision boundary limitations and handling multicollinearity.

Machine Learning (ML) – A subset of AI that enables computers to learn patterns from data without explicit programming. Related terms: supervised learning, unsupervised learning, reinforcement learning. Explanation: ML algorithms power predictive analytics used in public-policy decision-making. Example: Predicting the spread of an infectious disease using time-series data. Practical application: Allocating resources for disaster response based on forecasted needs. Challenges: Data bias, model drift, and ensuring robustness.

Model Drift – The degradation of model performance over time due to changes in underlying data distributions. Related terms: concept drift, monitoring, retraining. Explanation: In dynamic policy environments, regular model evaluation is essential. Example: A traffic-prediction model becomes less accurate after a new bike-lane network is introduced. Practical application: Scheduling periodic retraining of a welfare-eligibility classifier. Challenges: Detecting drift early and allocating resources for maintenance.

Natural Language Processing (NLP) – Techniques for analyzing, understanding, and generating human language. Related terms: sentiment analysis, named entity recognition, topic modeling. Explanation: NLP enables the extraction of insights from public comments, legislation texts, and social media. Example:

Classifying citizen feedback on a new zoning law as positive or negative. Practical application: Automating the summarization of city council meeting minutes. Challenges: Ambiguity, multilingual support, and bias in language models.

Neural Architecture Search (NAS) – Automated methods for designing optimal neural network structures. Related terms: hyperparameter optimization, AutoML, search space. Explanation: NAS can produce efficient models tailored to specific policy data constraints. Example: Discovering a lightweight CNN for on-device air-quality monitoring. Practical application: Deploying AI on low-power sensors in remote regions. Challenges: Computational expense and ensuring discovered architectures meet interpretability needs.

Open Data – Data that is freely available for anyone to use, modify, and share. Related terms: data portals, licensing, transparency. Explanation: Open data fuels civic AI initiatives and promotes accountability. Example: Publishing anonymized transit ridership data for public analysis. Practical application: Hackathons that develop AI tools for community planning. Challenges: Balancing openness with privacy and ensuring data quality.

Optimization – The process of finding the best solution among many feasible alternatives, often under constraints. Related terms: objective function, linear programming, gradient descent. Explanation: Optimization models support resource allocation, budgeting, and scheduling decisions. Example: Minimizing emergency-service response times while respecting budget caps. Practical application: Designing optimal routes for waste-collection trucks. Challenges: Non-convex problems, computational tractability, and data uncertainty.

Overfitting – When a model captures noise instead of the underlying pattern, leading to poor generalization. Related terms: regularization, cross-validation, bias-variance trade-off. Explanation: Overfitted models can mislead policymakers with overly optimistic predictions. Example: A model that predicts crime spikes perfectly on training data but fails on new neighborhoods. Practical application: Employing dropout and early stopping to mitigate overfitting. Challenges: Detecting subtle overfit, especially with high-dimensional data.

Policy Simulation – The use of computational models to explore the outcomes of alternative policy choices. Related terms: scenario analysis, agent-based modeling, system dynamics. Explanation: AI-enhanced simulations provide quantitative evidence for decision-makers. Example: Simulating the impact of a carbon tax on household energy consumption. Practical application: Interactive dashboards where users adjust policy levers and view projected results. Challenges: Model calibration, data availability, and communicating uncertainty.

Privacy-Preserving Machine Learning – Techniques that protect individual data while still enabling model training. Related terms: federated learning, differential privacy, secure multi-party computation. Explanation: These methods allow agencies to collaborate without exposing raw data. Example: Multiple municipalities jointly train a fraud-detection model without sharing citizen records. Practical application: Deploying federated learning across health departments for disease-trend analysis. Challenges: Communication overhead, reduced model accuracy, and regulatory compliance.

Predictive Analytics – The use of statistical techniques and ML models to forecast future events. Related terms: forecasting, risk modeling, time-series analysis. Explanation: Predictive tools assist governments in proactive planning and resource allocation. Example: Anticipating peak electricity demand during heat waves. Practical application: Pre-positioning emergency supplies based on predicted disaster pathways. Challenges: Data lag, model interpretability, and handling rare events.

Public-Private Partnership (PPP) – Collaborative arrangements between government entities and private sector firms. Related terms: contracting, risk sharing, joint venture. Explanation: PPPs often fund AI initiatives that would be unaffordable for public agencies alone. Example: A city partners with a tech company to develop an AI traffic-management platform. Practical application: Leveraging private-sector expertise for AI model deployment. Challenges: Aligning incentives, ensuring public interest, and managing intellectual property.

Qualitative Data – Non-numeric information such as text, audio, or video that describes attributes or meanings. Related terms: thematic analysis, coding, content analysis. Explanation: Qualitative insights complement quantitative AI models, especially in understanding citizen sentiment. Example: Interview transcripts from community meetings on a new zoning plan. Practical application: Using NLP to extract themes from public comments. Challenges: Subjectivity, scaling analysis, and integrating with numeric datasets.

Reinforcement Learning (RL) – A learning paradigm where agents learn optimal actions through trial-and-error interactions with an environment. Related terms: policy, reward function, Markov decision process. Explanation: RL can optimize sequential decisions such as dynamic resource allocation. Example: An RL agent that adjusts traffic-signal timings to minimize overall delay. Practical application: Real-time control of public-transport fleets. Challenges: Defining appropriate rewards, safety constraints, and convergence guarantees.

Regulation (AI) – Formal rules that dictate how AI systems may be designed, tested, and deployed. Related terms: compliance, standards, audit. Explanation: Regulatory frameworks aim to protect citizens while fostering innovation. Example: Mandatory impact assessments for AI tools that affect civil liberties. Practical application: Establishing a certification process for AI procurement. Challenges: Keeping regulations up-to-date with rapid technological change.

Responsible AI – An approach that embeds ethical, legal, and societal considerations throughout the AI lifecycle. Related terms: accountability, fairness, transparency. Explanation: Responsible AI practices are essential for public trust in government AI projects. Example: Conducting a fairness audit before launching an AI-driven welfare eligibility system. Practical application: Publishing model cards that disclose performance metrics across demographic groups. Challenges: Institutionalizing responsible-AI processes and measuring impact.

Risk Assessment – Systematic identification and evaluation of potential adverse outcomes associated with a policy or technology. Related terms: hazard analysis, mitigation strategy, contingency planning. Explanation: AI risk assessments address concerns such as bias, security, and unintended consequences. Example:

Evaluating the risk of algorithmic discrimination in loan-approval scoring. Practical application: Creating a mitigation plan that includes regular bias checks. Challenges: Quantifying intangible risks and obtaining stakeholder consensus.

Robustness – The ability of an AI model to maintain performance under diverse conditions and adversarial attacks. Related terms: adversarial robustness, stress testing, generalization. Explanation: Robust models are crucial for mission-critical public-policy functions. Example: A flood-prediction model that remains accurate despite noisy sensor inputs. Practical application: Implementing adversarial training to harden image-recognition systems used for infrastructure monitoring. Challenges: Trade-offs with model complexity and computational cost.

Scalability – The capacity of a system to handle increasing workloads without degradation. Related terms: horizontal scaling, cloud computing, load balancing. Explanation: Scalable AI solutions enable city-wide deployment and cross-agency integration. Example: Extending a pilot AI traffic-optimization tool from one district to the entire metropolitan area. Practical application: Leveraging serverless architectures for on-demand inference. Challenges: Managing data pipelines, latency, and cost control.

Sentiment Analysis – The computational determination of emotional tone behind a body of text. Related terms: opinion mining, polarity, subjectivity. Explanation: Sentiment analysis helps gauge public reaction to policy proposals. Example: Measuring citizen sentiment toward a new public-transport fare policy from social-media posts. Practical application: Real-time dashboards that track sentiment shifts during a policy rollout. Challenges: Sarcasm detection, domain adaptation, and language diversity.

Simulation Modeling – The creation of a virtual representation of a system to study its behavior under varying conditions. Related terms: Monte Carlo, agent-based model, system dynamics. Explanation: Simulations support “what-if” analyses for complex policy environments. Example: Simulating the spread of misinformation across social networks. Practical application: Testing the impact of regulation changes before enactment. Challenges: Parameter estimation, validation against real data, and computational intensity.

Smart City – An urban area that uses digital technology, including AI, to enhance the quality of life and efficiency of services. Related terms: IoT, urban analytics, digital twin. Explanation: AI drives many smart-city functions such as traffic management, energy optimization, and public safety. Example: Adaptive lighting systems that dim based on pedestrian presence. Practical application: Integrating AI sensors with municipal dashboards for real-time decision-making. Challenges: Data integration, privacy concerns, and equitable service distribution.

Social Impact Assessment (SIA) – Evaluation of the social consequences of a policy, program, or project. Related terms: stakeholder analysis, community engagement, impact metrics. Explanation: AI can augment SIAs by providing predictive insights and automating data collection. Example: Forecasting displacement effects of a new highway construction using demographic data. Practical application: Incorporating AI-generated risk maps into planning approvals. Challenges: Capturing qualitative nuances and ensuring community participation.

Supervised Learning – A machine-learning paradigm where models are trained on input-output pairs. Related terms: labelled data, regression, classification. Explanation: Supervised methods are widely used for policy-relevant prediction tasks. Example: Predicting school-performance scores from socioeconomic indicators. Practical application: Automating eligibility checks for subsidy programs. Challenges: Obtaining high-quality labels and avoiding over-reliance on historical patterns.

Supply Chain Transparency – Visibility into the origins, processes, and flows of goods and services. Related terms: traceability, blockchain, audit trail. Explanation: AI can analyze supply-chain data to detect fraud, ensure compliance, and support sustainable procurement. Example: Using anomaly detection to spot irregularities in procurement invoices. Practical application: Public dashboards that display sourcing information for government contracts. Challenges: Data standardization across vendors and protecting confidential commercial information.

Synthetic Data – Artificially generated data that mimics the statistical properties of real data. Related terms: data augmentation, generative models, privacy. Explanation: Synthetic datasets enable model development when real data is scarce or sensitive. Example: Simulated traffic patterns for training a congestion-prediction model. Practical application: Sharing synthetic health records with researchers without violating privacy. Challenges: Ensuring realism, avoiding inadvertent leakage of original data, and validating utility.

Temporal Fusion Transformer (TFT) – A deep learning architecture designed for multi-horizon time-series forecasting. Related terms: attention mechanism, sequence modeling, forecasting. Explanation: TFT handles heterogeneous inputs (static and time-varying) and provides interpretability via attention weights. Example: Forecasting weekly demand for public-housing applications. Practical application: Planning staffing levels for social-service offices. Challenges: Model complexity, need for extensive training data, and interpretability for non-technical stakeholders.

Transparency (AI) – Openness about how AI systems are built, trained, and deployed. Related terms: model documentation, open source, auditability. Explanation: Transparency builds trust and facilitates accountability in public-policy contexts. Example: Publishing the source code of an AI-driven resource-allocation tool. Practical application: Conducting third-party audits of predictive policing algorithms. Challenges: Balancing intellectual-property concerns with the need for openness.

Uncertainty Quantification – Techniques for estimating the confidence or error bounds of model predictions. Related terms: confidence interval, Monte Carlo dropout, Bayesian inference. Explanation: Quantifying uncertainty helps policymakers assess risk and make informed choices. Example: Providing a 95 % confidence band around projected air-quality improvements. Practical application: Decision dashboards that display prediction intervals alongside point estimates. Challenges: Computational overhead and communicating uncertainty to non-technical audiences.

Validation (Model) – The process of assessing a model's performance on unseen data to ensure generalizability. Related terms: test set, cross-validation, performance metrics. Explanation: Rigorous validation guards against overfitting and misleading conclusions. Example: Using a hold-out dataset of recent census records to evaluate a demographic-prediction model. Practical application: Setting performance thresholds before deploying an AI system for benefits distribution. Challenges: Data leakage, selection bias, and evolving data distributions.

Version Control (Data) – Systematic tracking of changes to datasets, code, and models. Related terms: Git, data lineage, reproducibility. Explanation: Version control ensures reproducibility, auditability, and collaborative development. Example: Maintaining a Git repository for all scripts involved in a policy-impact analysis. Practical application: Rolling back to a previous dataset version after discovering an error. Challenges: Managing large binary data files and coordinating across multiple teams.

Virtual Public Consultation – Digital platforms that enable citizens to provide feedback on policy proposals. Related terms: e-participation, crowdsourcing, online surveys. Explanation: AI can automate the analysis of large volumes of citizen input, extracting key themes and sentiment. Example: An AI-driven portal that clusters feedback on a new housing ordinance. Practical application: Real-time dashboards that inform legislators about public concerns during debates. Challenges: Ensuring representativeness, mitigating manipulation, and protecting participant privacy.

Weighted Decision Matrix – A tool that scores alternatives based on multiple criteria, each weighted by importance. Related terms: multi-criteria analysis, scoring system, prioritization. Explanation: AI can generate the underlying scores by predicting criterion outcomes from data. Example: Evaluating technology vendors for a city-wide AI procurement. Practical application: Combining cost, scalability, and ethical compliance scores to select a solution. Challenges: Assigning appropriate weights and avoiding subjectivity.

Zero-Shot Learning – A learning paradigm where a model can correctly recognize classes it has never seen during training. Related terms: transfer learning, semantic embedding, few-shot learning. Explanation: Zero-shot techniques are useful when new policy categories emerge without labeled data. Example: Classifying a novel type of infrastructure request based on textual description. Practical application: Rapidly adapting a complaint-routing AI to a newly introduced service line. Challenges: Designing robust semantic representations and handling ambiguity.

Bias Mitigation Techniques – Methods employed to reduce unfairness in AI models. Related terms: pre-processing, in-processing, post-processing. Explanation: Techniques span data rebalancing, algorithmic constraints, and outcome adjustments. Example: Re-weighting training samples to achieve demographic parity. Practical application: Deploying a fairness-constrained classifier for public-housing eligibility. Challenges: Selecting appropriate fairness metrics and avoiding unintended side effects.

Data Anonymization – The process of removing personally identifying information from datasets. Related terms: de-identification, k-anonymity, masking. Explanation: Anonymization enables secondary analysis while protecting privacy. Example: Stripping names and addresses from a health-service dataset before sharing with researchers. Practical application: Publishing open datasets for civic tech initiatives. Challenges: Re-identification risk, utility loss, and compliance with regulations.

Decision Tree – A flow-chart-like model that splits data based on feature thresholds to arrive at predictions. Related terms: entropy, gini impurity, pruning. Explanation: Decision trees are intuitive, making them suitable for policy contexts where interpretability matters. Example: A tree that determines eligibility for a subsidy based on income, age, and residency. Practical application: Interactive policy tools that let users explore decision pathways. Challenges: Prone to overfitting and sensitivity to small data variations.