
Postgraduate Certificate in AI Strategies for NGOs

Artificial Intelligence Fundamentals

Artificial Intelligence (AI) – related terms: machine learning, automation. The simulation of human intelligence processes by computers, including learning, reasoning, and self-correction. In NGOs, AI can streamline data analysis, predict humanitarian needs, and automate routine tasks. Example: An AI system analyses satellite imagery to identify flood-affected regions. Challenges include data bias, limited technical expertise, and ethical concerns about decision-making transparency.

Artificial Neural Network (ANN) – related terms: deep learning, perceptron. A computational model inspired by the brain's network of neurons, consisting of layers of interconnected nodes that process inputs to produce outputs. NGOs use ANNs for pattern recognition, such as classifying disaster-related social media posts. Practical issues involve the need for large labeled datasets and high computational resources, which may be scarce in low-budget settings.

Algorithm – related terms: procedure, code. A step-by-step set of instructions for solving a problem or performing a task. In AI, algorithms define how models learn from data. Example: The k-nearest neighbors algorithm groups similar beneficiary profiles for targeted aid. Challenges include selecting appropriate algorithms for specific data characteristics and ensuring they do not reinforce existing inequities.

Bias (Data Bias) – related terms: fairness, representation. Systematic errors that cause a model to produce prejudiced outcomes due to skewed training data. For NGOs, biased data can lead to unequal resource allocation. A practical mitigation is to audit datasets for demographic coverage and apply re-weighting techniques. The main challenge is obtaining diverse, high-quality data in crisis contexts.

Classification – related terms: categorization, supervised learning. The task of assigning inputs to predefined categories. NGOs often classify incoming assistance requests into "food", "medical", or "shelter" needs. Techniques include logistic regression and support vector machines. Challenges involve imbalanced class distributions and evolving categories as crises develop.

Clustering – related terms: unsupervised learning, segmentation. Grouping similar data points without pre-labeled categories. NGOs may cluster households based on vulnerability indicators to prioritize interventions. Common algorithms are K-means and hierarchical clustering. Practical difficulties include determining the optimal number of clusters and interpreting results for policy decisions.

Computer Vision – related terms: image processing, object detection. A field of AI that enables machines to interpret visual information from images or video. Humanitarian organizations use computer vision to assess damage from drone footage. Example: An object-detection model identifies collapsed buildings. Challenges include limited labeled imagery for rare disaster types and the need for robust models that function under varying lighting and weather conditions.

Cross-Validation – related terms: model evaluation, k-fold. A technique for assessing how a predictive

model will generalize to an independent dataset. NGOs often split limited data into training and validation folds to avoid overfitting. Example: A 5-fold cross-validation of a disease-prediction model. The main obstacle is the computational overhead when resources are constrained.

Data Augmentation – related terms: synthetic data, oversampling. Methods for artificially expanding a dataset by creating modified versions of existing data. In NGOs, augmenting limited medical images (e.g., Rotating X-rays) can improve model robustness. Challenges involve ensuring augmented data reflect realistic variations and do not introduce artifacts that mislead the model.

Data Governance – related terms: policy, stewardship. The overall management of data availability, usability, integrity, and security. NGOs need clear governance to protect beneficiary privacy while enabling AI analytics. Practical steps include establishing data access controls and consent procedures. Challenges arise from cross-border data transfers and varying legal regimes.

Data Privacy – related terms: confidentiality, GDPR. The right of individuals to control how personal information is collected and used. AI projects in NGOs must anonymize personal data, use secure storage, and obtain informed consent. Example: Masking names in a refugee-registration dataset before training a predictive model. The difficulty lies in balancing analytical utility with strict privacy safeguards.

Data Preprocessing – related terms: cleaning, normalization. The series of steps to transform raw data into a suitable format for modeling. Common tasks include handling missing values, scaling numeric features, and encoding categorical variables. NGOs often receive heterogeneous data from field surveys, requiring extensive preprocessing. Challenges include time-consuming manual cleaning and the risk of inadvertently discarding valuable signals.

Deep Learning – related terms: neural networks, representation learning. A subset of machine learning that uses multi-layer neural networks to automatically learn hierarchical feature representations. NGOs apply deep learning to speech-to-text transcription of crisis calls. While powerful, deep learning demands large datasets and GPU resources, which may be out of reach for many humanitarian actors.

Decision Tree – related terms: random forest, CART. A flowchart-like model that splits data based on feature thresholds to make predictions. Decision trees are interpretable, making them attractive for NGOs needing transparent rationale for aid allocation. Example: A tree predicts the likelihood of food insecurity based on household size, income, and location. Overfitting is a common challenge, often mitigated by pruning or ensemble methods.

Ethical AI – related terms: responsible AI, fairness. The practice of designing, deploying, and monitoring AI systems in ways that respect human rights, avoid harm, and promote equity. NGOs must embed ethical considerations such as bias mitigation, explainability, and community consent. Practical frameworks include AI impact assessments and stakeholder workshops. The challenge is operationalizing abstract principles into concrete project steps.

Feature Engineering – related terms: feature extraction, transformation. The process of creating informative variables from raw data to improve model performance. In humanitarian contexts, features might include distance to the nearest health clinic or frequency of past aid requests. Effective feature engineering can

compensate for limited data volume. However, it requires domain expertise and can be time-intensive.

Feature Selection – related terms: dimensionality reduction, Lasso. Techniques for identifying the most relevant variables for a model, reducing complexity and overfitting risk. NGOs may use recursive feature elimination to narrow down socioeconomic indicators that best predict displacement risk. Challenges include maintaining predictive power while simplifying the model and avoiding the removal of subtle but important variables.

Gaussian Process – related terms: probabilistic modeling, kernel methods. A non-parametric approach that defines a distribution over functions, providing predictions with uncertainty estimates. In NGOs, Gaussian processes can model the spread of disease with confidence intervals, aiding risk communication. The method scales poorly with large datasets, posing computational challenges for extensive geographic analyses.

Generative Adversarial Network (GAN) – related terms: synthetic data, adversarial training. A class of deep learning models where two networks (generator and discriminator) compete, leading to realistic data generation. NGOs may use GANs to create synthetic satellite images of flood-affected areas for training other models when real images are scarce. Ethical concerns include potential misuse of generated data and difficulty in verifying authenticity.

Gradient Descent – related terms: optimization, learning rate. An iterative algorithm for minimizing a loss function by moving parameters in the direction of steepest descent. All neural-network training in NGOs relies on variants of gradient descent, such as stochastic gradient descent. Selecting appropriate learning rates and handling plateaus are practical challenges that affect convergence speed and model quality.

Hyperparameter – related terms: tuning, model configuration. Settings that govern the behavior of learning algorithms but are not learned from data (e.g., Number of hidden layers, regularization strength). NGOs often perform grid or random search to find optimal hyperparameters for a disaster-prediction model. The main difficulty is the computational expense and the need for systematic experimentation.

Interpretability – related terms: explainability, model transparency. The degree to which a human can understand the reasoning behind an AI output. For NGOs, interpretability is crucial for stakeholder trust and compliance with donor requirements. Tools such as SHAP values or LIME provide local explanations for complex models. Trade-offs arise when highly accurate black-box models sacrifice transparency.

Kernel Method – related terms: SVM, feature space. Techniques that implicitly map data into higher-dimensional spaces to make it linearly separable. Support Vector Machines (SVM) with radial basis function kernels are used by NGOs to classify crisis-related texts. Kernel methods can be computationally intensive for large datasets, requiring approximation strategies.

K-Nearest Neighbors (KNN) – related terms: instance-based learning, distance metric. A simple algorithm that classifies a data point based on the majority label among its k closest neighbors. NGOs may apply KNN to recommend nearby shelters based on a beneficiary's location. The method suffers from high memory usage and sensitivity to irrelevant features, necessitating careful preprocessing.

Labeling (Data Labeling) – related terms: annotation, ground truth. The process of assigning meaningful tags to raw data to create supervised learning datasets. Humanitarian NGOs often rely on volunteers to label images of damaged infrastructure. Quality control is essential; inter-annotator agreement metrics help identify inconsistencies. Labeling is labor-intensive and can become a bottleneck for rapid AI deployment.

Latent Variable – related terms: hidden state, factor analysis. Variables that are not directly observed but inferred from observed data, often capturing underlying structure. Topic modeling in NGOs uses latent variables to uncover hidden themes in donor feedback. Interpreting latent variables requires domain expertise, and model misspecification can lead to misleading conclusions.

Logistic Regression – related terms: binary classification, odds ratio. A statistical model that predicts the probability of a binary outcome using a logistic function. NGOs frequently employ logistic regression to estimate the likelihood of disease outbreak based on environmental predictors. While easy to interpret, the model assumes linear relationships between predictors and log-odds, which may not hold in complex humanitarian settings.

Machine Learning (ML) – related terms: AI, predictive modeling. A subset of AI focused on algorithms that improve automatically through experience. NGOs use ML to forecast food price volatility, allocate resources, and detect fraud. Key steps include data collection, preprocessing, model training, and evaluation. Challenges specific to NGOs involve limited data, fluctuating contexts, and the need for explainable outcomes.

Model Overfitting – related terms: generalization, regularization. When a model captures noise instead of underlying patterns, performing well on training data but poorly on new data. In humanitarian projects, overfitting can misguide resource distribution. Techniques such as cross-validation, early stopping, and dropout help mitigate overfitting. Continuous monitoring on fresh data is essential to detect drift.

Natural Language Processing (NLP) – related terms: text mining, sentiment analysis. The field of AI that enables computers to understand, interpret, and generate human language. NGOs leverage NLP to analyze crisis-related tweets, extract beneficiary needs, and automate translation of relief documents. Challenges include multilingual corpora, informal language, and limited labeled data for low-resource languages.

Neural Architecture Search (NAS) – related terms: autoML, model design. Automated methods for discovering optimal neural network structures. While powerful, NAS requires substantial computational resources, often beyond the reach of small NGOs. Some lightweight NAS frameworks can suggest efficient models for edge devices used in remote monitoring. The trade-off between search time and model performance must be carefully managed.

One-Hot Encoding – related terms: categorical variables, dummy variables. A technique to represent categorical data as binary vectors, where each category gets a separate dimension. NGOs encode donor types (individual, corporate, government) before feeding them into a predictive model. The method can increase dimensionality dramatically when categories are numerous, leading to sparsity issues.

Outlier Detection – related terms: anomaly detection, robust statistics. Identifying data points that deviate markedly from the majority. In humanitarian logistics, outliers may indicate fraudulent expense claims or

sensor errors. Algorithms such as Isolation Forest or Local Outlier Factor are commonly applied. Distinguishing true anomalies from rare but legitimate events remains a practical difficulty.

Precision (Metric) – related terms: positive predictive value, accuracy. The proportion of true positive predictions among all positive predictions. For an NGO's disease-alert system, high precision reduces false alarms, preserving trust. However, focusing solely on precision may lower recall, missing some true cases. Balancing precision and recall is essential for effective alerting.

Probabilistic Model – related terms: Bayesian inference, uncertainty. Models that output probability distributions instead of single point estimates. NGOs use probabilistic models to forecast the probability of displacement under different climate scenarios. These models allow decision-makers to assess risk levels and allocate contingency funds. Computational complexity and the need for prior distributions can hinder adoption.

Recall (Metric) – related terms: sensitivity, true positive rate. The proportion of actual positives correctly identified by the model. In a refugee-identification system, high recall ensures most at-risk individuals are flagged. Yet, increasing recall often reduces precision, generating more false positives that strain resources. Optimizing the trade-off depends on the humanitarian cost of missed cases versus over-allocation.

Reinforcement Learning (RL) – related terms: agent, reward function. A learning paradigm where an agent interacts with an environment, receiving rewards for actions that achieve a goal. NGOs could employ RL to optimize the routing of mobile health clinics, learning policies that minimize travel time while maximizing coverage. The main challenges are defining appropriate reward signals and ensuring safe exploration in real-world deployments.

Regularization – related terms: L1, L2, overfitting. Techniques that add penalty terms to the loss function to discourage overly complex models. L1 regularization induces sparsity, useful for feature selection in limited-data NGO projects. L2 regularization smooths weights, improving generalization. Selecting regularization strength requires validation, and excessive regularization may underfit the data.

ResNet (Residual Network) – related terms: deep learning, skip connections. A neural-network architecture that uses shortcut connections to alleviate vanishing gradient problems, enabling very deep models. Humanitarian organizations use ResNet variants for high-resolution satellite-image analysis to detect building damage. While powerful, ResNets are computationally heavy, demanding careful hardware planning for field deployments.

Sentiment Analysis – related terms: opinion mining, NLP. The computational determination of emotional tone behind textual data. NGOs monitor social media sentiment during crises to gauge public fear or trust in relief efforts. Rule-based or machine-learning approaches can be employed. Language nuances, sarcasm, and multilingual texts pose challenges to accurate sentiment classification.

Support Vector Machine (SVM) – related terms: kernel, margin. A supervised learning algorithm that finds the hyperplane maximizing the margin between classes. NGOs use SVMs for classifying textual reports into "urgent" vs. "Non-urgent". SVMs perform well on small to medium datasets and can handle high-dimensional feature spaces. However, they are less scalable to very large corpora and may require

careful kernel selection.

Transfer Learning – related terms: pre-trained model, fine-tuning. The practice of adapting a model trained on one task to a related task with limited data. NGOs can fine-tune a pre-trained language model (e.G., BERT) on a small set of local crisis reports to improve extraction accuracy. Benefits include reduced training time and higher performance. Potential pitfalls involve domain mismatch and the risk of inheriting biases from the source model.

Unsupervised Learning – related terms: clustering, dimensionality reduction. Learning patterns from data without explicit labels. NGOs apply unsupervised techniques to discover emerging topics in field notes or to compress high-dimensional sensor data. Methods include autoencoders and principal component analysis. Evaluation is less straightforward than supervised learning, often relying on expert interpretation.

Validation Set – related terms: hold-out, model selection. A subset of data used to tune hyperparameters and decide between competing models, distinct from training and test sets. NGOs allocate a validation set to avoid over-optimistic performance estimates when calibrating a flood-risk model. The challenge is preserving enough data for training while maintaining a representative validation sample, especially when overall data volume is limited.

Variable Importance – related terms: feature importance, permutation. Metrics indicating how much each predictor contributes to a model's predictions. In a donor-retention model, variable importance helps NGOs focus on key factors like communication frequency. Techniques such as permutation importance or SHAP values provide insights. Interpreting importance scores can be misleading if features are correlated, necessitating careful analysis.

Word Embedding – related terms: vector representation, GloVe. Dense vector representations of words that capture semantic relationships. NGOs use embeddings to improve text classification of aid requests, enabling the model to understand synonyms (e.G., "Food" and "nutrition"). Pre-trained embeddings accelerate development but may not capture local dialects, requiring domain-specific fine-tuning.

Zero-Shot Learning – related terms: transfer learning, generalization. The ability of a model to correctly predict classes it has never seen during training, based on auxiliary information. Humanitarian NGOs could use zero-shot techniques to identify novel disaster types from textual descriptions without prior examples. While promising for rare events, performance is typically lower than supervised methods, and reliable auxiliary descriptors are essential.