
Undergraduate Certificate in Artificial Intelligence for Indirect Tax Management

Ethical and Legal Issues in AI

Algorithmic Bias: Systematic prejudice or unfairness in automated systems and decision-making processes, often due to biased data or biased decision-making rules. For example, if an AI system used to screen job applications was trained on data from a company that historically favored hiring male candidates, the system might learn to discriminate against female applicants.

Artificial Intelligence (AI): The ability of machines to perform tasks that would typically require human intelligence, such as visual perception, speech recognition, decision-making, and language translation. AI can be categorized into two main types: narrow or weak AI, which is designed to perform a specific task, and general or strong AI, which can perform any intellectual task that a human can do.

Automated Decision-Making: The use of AI systems to make decisions based on predefined rules and algorithms, without human intervention. For example, an AI system might be used to automatically approve or reject loan applications based on various factors, such as credit score, income, and debt-to-income ratio.

Big Data: Large and complex sets of data that can be analyzed computationally to reveal patterns, trends, and associations, especially relating to human behavior and interactions. Big data is often characterized by its volume, variety, velocity, and veracity.

Computer Vision: A field of AI that deals with the automatic processing and interpretation of visual information, such as images and videos. Computer vision algorithms can be used for tasks such as object recognition, image classification, and facial recognition.

Data Mining: The process of discovering patterns and insights in large datasets, using various statistical and machine learning techniques. Data mining can be used for various purposes, such as market research, fraud detection, and predictive maintenance.

Deep Learning: A subset of machine learning that uses artificial neural networks with multiple layers to learn and represent complex patterns and relationships in data. Deep learning algorithms can be used for various tasks, such as image and speech recognition, natural language processing, and autonomous driving.

Discrimination: The unfair or unequal treatment of individuals or groups based on certain characteristics, such as race, gender, age, or religion. Discrimination can take various forms, such as direct discrimination, indirect discrimination, and systemic discrimination.

Explainable AI (XAI): The ability of AI systems to provide clear and understandable explanations of their decisions, actions, and recommendations, especially in high-stakes domains such as healthcare, finance, and criminal justice. Explainable AI is important for building trust, ensuring accountability, and avoiding bias and discrimination.

Facial Recognition: A type of computer vision technology that uses algorithms to identify and verify

individuals based on their facial features, such as the distance between their eyes, nose, and mouth. Facial recognition has various applications, such as access control, surveillance, and law enforcement.

Fairness: The quality of treating individuals or groups equally and without bias or discrimination, based on relevant and objective criteria. Fairness is an important principle in AI design, development, and deployment, as it can help to promote trust, accountability, and social responsibility.

General Data Protection Regulation (GDPR): A comprehensive European Union (EU) law on data protection and privacy in the EU and European Economic Area (EEA), which came into force on May 25, 2018. GDPR aims to give individuals more control over their personal data, and to impose strict obligations and penalties on organizations that process or handle personal data without proper consent or safeguards.

Machine Learning: A subset of AI that deals with the automatic learning and improvement of algorithms, based on data and feedback. Machine learning algorithms can be supervised, unsupervised, or reinforcement learning, depending on the type of data and task.

Natural Language Processing (NLP): A field of AI that deals with the automatic processing and interpretation of human language, such as text and speech. NLP algorithms can be used for various tasks, such as language translation, sentiment analysis, and question-answering.

Neural Networks: A type of AI model that is inspired by the structure and function of the human brain, and that consists of interconnected nodes or neurons that can process and transmit information. Neural networks can be used for various tasks, such as image and speech recognition, natural language processing, and autonomous driving.

Predictive Analytics: The use of statistical and machine learning techniques to predict future events or outcomes, based on historical data and patterns. Predictive analytics can be used for various purposes, such as risk assessment, fraud detection, and demand forecasting.

Privacy: The right of individuals to control their personal information and to protect it from unauthorized access, use, or disclosure. Privacy is an important principle in AI design, development, and deployment, as it can help to build trust, prevent harm, and comply with legal and ethical norms.

Robotic Process Automation (RPA): The use of software robots or agents to automate repetitive and routine tasks, such as data entry, document processing, and customer service. RPA can help to improve efficiency, accuracy, and productivity, and to reduce costs and errors.

Sentiment Analysis: A type of natural language processing that deals with the automatic detection and interpretation of emotions, opinions, and attitudes in text or speech. Sentiment analysis can be used for various purposes, such as market research, customer feedback, and social media monitoring.

Supervised Learning: A type of machine learning that uses labeled data, where the input and output variables are known and the algorithm is trained to predict the output based on the input. Supervised learning can be used for various tasks, such as image classification, speech recognition, and fraud detection.

Transparency: The quality of providing clear, accurate, and complete information about the design,

development, and deployment of AI systems, including their goals, methods, and limitations. Transparency is an important principle in AI design, development, and deployment, as it can help to build trust, ensure accountability, and avoid bias and discrimination.

Unsupervised Learning: A type of machine learning that uses unlabeled data, where the input variables are known but the output variables are unknown and the algorithm is trained to discover patterns and relationships in the data. Unsupervised learning can be used for various tasks, such as clustering, dimensionality reduction, and anomaly detection.

Virtual Assistants: Software agents or bots that use natural language processing and machine learning to provide various services, such as information search, recommendation, and automation. Virtual assistants can be integrated into various platforms, such as websites, mobile apps, and smart speakers.

These are some of the key terms and concepts related to ethical and legal issues in AI in the course Undergraduate Certificate in AI for Indirect Tax Management. These terms are important for understanding the opportunities and challenges of AI in tax management, and for designing and implementing AI systems that are ethical, legal, and socially responsible.