

CertNexus Certified Ethical Emerging Technologist Exam CET-110

Exam Information

Candidate Eligibility

The *Certified Ethical Emerging Technologist (CEET)* exam requires no application fee, supporting documentation, or other eligibility verification measures in order to take the exam. An exam voucher may come bundled with your training program, or can be purchased separately [here](#) or directly from Pearson VUE. Once purchased, you will receive more information about how to register for and schedule your exam. Once you have obtained your voucher information, you can register for an exam time [here](#). By registering, you agree to our Candidate Agreement included [here](#).

Exam Prerequisites

While there are no formal prerequisites to register for and schedule an exam, we strongly recommend you first possess the knowledge, skills, and abilities to do the following:

- Demonstrate an understanding of the fundamental/foundational concepts related to ethics in data-driven technologies.
- Identify common ethical principles and frameworks and select the appropriate framework to understand and/or address ethical issues.
- Identify regulations, standards, and best practices utilized in the industry and identify the ethical challenges that may conflict with or compromise their implementation.
- Identify and mitigate the myriad risks that arise within the development, utilization, and/or implementation of data-driven technologies.
- Communicate about ethical risks and ethical practices internally within the organization and externally to relevant third parties.
- Create, implement, and evaluate ethical policies and governance regarding data-driven technology throughout the organization.

You can obtain this level of skill and knowledge by taking the following online course, which is available through Coursera, or by attending an equivalent third-party training program:

- *Certified Ethical Emerging Technologist*

Exam Specifications

Number of Items: 80

Passing Score: 62%

Duration: 120 minutes (**Note:** Published exam times include the 10 minutes you are allotted for reading and signing the Candidate Agreement and reviewing exam instructions.)

Exam Options: In person at Pearson VUE test centers or online via Pearson OnVUE online proctoring

Item Formats: Multiple Choice/Multiple Response

Exam Description

Target Candidate:

This certification exam is designed for individuals seeking to demonstrate a vendor neutral, cross-industry, and multidisciplinary understanding of applied technology ethics that will enable them to navigate the processes by which ethical integrity may be upheld within emerging data-driven technology fields (such as artificial intelligence (AI), Internet of Things (IoT), and data science).

Exam Objective Statement:

This exam will certify that the successful candidate has the knowledge, skills, and abilities required to apply foundational ethical principles, follow industry-standard frameworks, identify and mitigate risks, and navigate ethical organizational governance in order to devise and maintain ethical, trusted, and inclusive data-driven technologies.

To ensure exam candidates possess the aforementioned knowledge, skills, and abilities, the *Certified Ethical Emerging Technologist (CEET)* exam will test them on the following domains with the following weightings:

Domain	% of Examination
1.0 Fundamental Concepts for Data-Driven Technology Ethics	17%
2.0 Ethical Frameworks	23%
3.0 Risk Identification and Mitigation	30%
4.0 Communication	12%
5.0 Organizational Policy and Governance	18%
Total	100%

The information that follows is meant to help you prepare for your certification exam. This information does not represent an exhaustive list of all the concepts and skills that you may be tested on during your exam. The exam domains, identified previously and included in the objectives listing, represent the large content areas covered in the exam. The objectives within those domains represent the specific tasks associated with the job role(s) being tested. The information beyond the domains and objectives is meant to provide examples of the types of concepts, tools, skills, and abilities that relate to the corresponding domains and objectives. All of this information represents the industry-expert analysis of the job role(s) related to the certification and does not necessarily correlate one-to-one with the content covered in your training program or on your exam. We strongly recommend that you independently study to familiarize yourself with any concept identified here that was not explicitly covered in your training program or products.

Objectives

Domain 1.0 Fundamental Concepts for Data-Driven Technology Ethics

Objective 1.1 Identify and describe common terminology or concepts important to data-driven technology ethics

- AI-related concepts
 - Narrow AI
 - General AI
 - Superintelligence
 - Ambient intelligence
 - Black box
 - Model training
- Data science-related concepts
 - Pipeline
 - Ground truth dataset
- Legal-related concepts
 - EULA
 - ToS
 - Click-through agreement
 - SLA
 - Visual contracts
 - Smart contracts
 - Data sharing agreement
 - Liability
 - Legal responsibility

- Privacy-related concepts
 - PII
 - Pseudonymization
 - Anonymization
 - Differential privacy
 - Protected attributes
 - Passive listening
 - Opt-in/Opt-out
 - Informed consent
 - Facial recognition
 - De-identification
 - Privacy-by-Design
 - Data ethics
- Ethics
 - Autonomy
 - Justice
 - Ethical risk
 - Techno-social engineering
 - Personhood
 - Engineering activism
 - Ethics-by-Design
 - Beneficence
 - Non-maleficence
 - Universalizability
 - Moral relativism
 - Moral agency
 - Moral psychology
 - Moral reasoning
 - Moral judgments
- Bias
 - Implicit bias
 - Complacency bias
 - Automation bias
 - Data collection bias
 - Reinforcement bias
 - Temporal bias
 - Human-cognitive bias
 - Cultural bias
 - Statistical bias
 - Sociological bias

- Equity
 - Demographic parity
 - Equal opportunity
 - Equal accuracy
- Evaluation metrics
 - Accuracy and specificity
 - Reliability
 - Goodhart's Law

Objective 1.2 Identify and describe common ethical theories

- Moral philosophy
 - Utilitarianism/Consequentialism
 - Determinism
 - Virtue ethics
 - Deontology
- Applied ethics
 - Bioethics
 - Business ethics
 - Engineering ethics
 - Professional ethics
 - Social ethics
 - Environmental ethics

Objective 1.3 Identify when it is appropriate to conduct an ethical risk review

- Ideation and innovation management activities
- What-if scenario planning sessions
- New product/service development
- All along the data science/AI development lifecycle, from ideation and design through deployment and maintenance
- Stage gates and other points as appropriate within an organization's Ethics-by-Design practices
- When legal and/or regulatory non-compliance has occurred
- When an ethical violation or incident has occurred

Domain 2.0 Ethical Frameworks

Objective 2.1 Identify common ethical principles cited by major ethical frameworks

- Privacy
- Accountability
- Safety and Security
- Transparency and Explainability
 - Auditability
 - Explicability
 - Interpretability

- Fairness and Non-Discrimination
- Human Control of Technology
- Professional Responsibility
- Promotion of Human Values

Objective 2.2 Given an ethically challenging dilemma, identify and select an ethical framework to understand the issue

- Initiatives
 - Academia
 - Technical communities
 - Business
 - Civil society
 - Intergovernmental organizations
 - Trade unions
- Examples
 - The Montreal Declaration for a Responsible Development of Artificial Intelligence
 - Ethics Guidelines for Trustworthy AI
 - Beijing AI Principles
 - Government AI Readiness Index
 - Universal Guidelines for Artificial Intelligence
 - Top 10 Principles for Ethical Artificial Intelligence
 - OECD Principles on Artificial Intelligence
 - Ethically Aligned Design (IEEE)
 - G20 AI Principles
 - The Asilomar AI Principles
 - The Toronto Declaration

Objective 2.3 Follow applicable regulations, standards, and best practices

- Regulations
 - The Code of Fair Information Practices
 - OECD Privacy Guidelines
 - GDPR
 - CCPA
 - PIPEDA
 - POPI
 - LGPD
 - HIPAA
 - COPPA
 - Algorithmic Accountability Act
 - FERPA
 - PCI DSS

- BIPA
- Standards/best practices
 - NIST CIS (NIST 800.53, SANS CSC, ISO 27001)
 - ISO 27017
 - NISTIR 8288
 - MITRE ATT&CK Framework
 - IEEE 7000 series

Objective 2.4 Identify ethical challenges that may conflict or require compromise with regulatory and/or business constraints or demands

- Data minimization principle vs. need for data
- Performance vs. explainability
- Compliance vs. cost
- Transparency/explainability vs. intellectual property rights
- Company/stakeholder needs vs. ethical decision making
- Ethics washing
 - Marketing vs. integrity
 - Use of collected data vs. integrity
- Efficiency vs. the risk of collateral damage
 - Militarization/weaponization of AI
- Proliferation of AI to unscrupulous actors vs. democratization of AI (e.g., open source)
- Efficiency of development vs. cultural/contextual sensitivity
- Availability of datasets for ML algorithms vs. privacy protection
- Big data generated through devices (cloud, IoT) vs. concentration of power in big tech
- Fair competition vs. corporate hegemony (data assets)
- Efficiency/streamlining experience vs. enabling human agency/autonomy
- Moral relativism vs. evidence-based policy

Domain 3.0 Risk Identification and Mitigation

Objective 3.1 Identify and mitigate privacy risks

- Source
 - Data use without informed consent
 - Cross-correlation of combined data sources
 - Third-party data
 - Secondary use of data
 - Use/integration of third-party products
- Methods of identification
 - Check for presence of PII
 - Check for presence of consent/legitimate interest/appropriate use
 - Validate that legal and regulatory compliance has been met

- Persona modeling with external input for human/machine actors
- Put in place and maintain records of relevant personal data protection policies and processes
- Track information about customer data, such as when it was collected and the terms governing its collection; accessing and using that data; and auditing access and use
- Mitigation strategies
 - Communicate and verify intent
 - Manage consent over time
 - Use of synthetic data
 - Avoid collecting PII or associated metadata
 - Preserve data provenance
 - Minimize amount of data shared
 - Renew informed consent agreements
 - Differential privacy
 - Opt-in/opt-out
 - User inspection
- Tools for identification/mitigation
 - Privacy legislation databases
 - World Legal Information Institute Database for International Privacy Law
 - Trusted party/SSO/MFA
 - Blockchain
 - SDKs
 - HealthKit
 - ResearchKit
 - Anonymization and pseudonymization
 - Homomorphic encryption
 - Zero-Knowledge Protocols

Objective 3.2 Identify and mitigate accountability risks

- Source
 - Lack of transparency/explainability
 - Use of third-party components (open source libraries, etc.)
 - Lack of paper trails and black box records
 - Task delegation to autonomous systems
 - Data collected via federated learning/integrated systems
 - Lack of guiding principles/governance
 - Complacency/automation bias
 - Extrajudicial judgment (eg., drone attacks based on metadata)
 - Use/integration of third-party products

- Methods of identification
 - Be familiar with types of algorithms/models that are more “black box”
 - Check for governance structure outlined by your organization
- Mitigation strategies
 - Document company policies clearly and provide them to all design and development teams
 - Be aware of the parameters of your organization’s, other products, or vendor’s responsibilities
 - Document and record design processes and decisions
 - Follow the business conduct guidelines/governance provided by your organization
 - Pilot testing
 - RACI matrices/RAM
 - SOPs
 - Internal review boards
 - Ongoing collaboration and mutual accountability between data sharing partners
 - Document processes for auditing operations of AI systems
- Tools for identification/mitigation
 - Algorithmic Impact Assessment
 - Data visualization and dashboard reporting

Objective 3.3 Identify and mitigate transparency and explainability risks

- Source
 - Self-learning models
 - Black box nature of the system
 - Intellectual property rights
 - Shadow banning (eg., content/users named from networks without knowing why)
 - Use/integration of third-party products
- Methods of identification
 - Identify algorithmic decisions
 - Deconstruct specific decisions
 - Explainable AI
- Mitigation strategies
 - Provide contextual information about how an AI system works and interacts with data
 - Publish the algorithms or algorithmic principles underlying AI systems
 - Establish a channel by which an individual can seek an explanation
 - Human-in-the-loop methods
 - Account for downstream uses of data sets

- Ensure data uses are consistent with intentions of data disclosers
- Explain methods for analysis and marketing to data disclosers to maximize transparency at the point of data collection
- Share information about potential inadequacies in training data
- Tools for identification/mitigation
 - SHAP
 - Alibi
 - ELI5
 - LIME
 - What If tool

Objective 3.4 Identify and mitigate fairness and non-discrimination (bias) risks

- Source
 - Implicit bias
 - Data collection and sampling bias
 - Automation bias
 - Reinforcement bias
 - Validation/verification of data labels
 - Overfitting to training data
 - Temporal bias
 - Edge cases, outliers, and fitting to the mean
 - Use/integration of third-party products
- Methods of identification
 - Analytical techniques – systematically assess data used to train AI systems for appropriate representativeness and document its origins and characteristics
 - Persona modeling
 - Analyze model behavior in different environments to identify harmful bias or discrimination
- Mitigation strategies
 - Differentiate between useful algorithmic pattern matching and bias/prejudice
 - Inclusive design
 - Foreseeability
 - Appropriate analysis
 - PESTLE
 - STEEPV
 - Conduct representative user testing at appropriate points in the development lifecycle
 - Model how a product or application functions in different environments or situations (information collection, engineering requirements, etc.)

- Engage with appropriate external stakeholders to solicit their input during the development lifecycle
- Tools for identification/mitigation
 - Bias and safety bounties
 - The AI Fairness Project by IBM
 - Radioactive data tracing

Objective 3.5 Identify and mitigate safety and security risks

- Source
 - Abnormal system behaviors/unintended system interactions
 - Bad actors
 - Malicious data
 - Groupthink
 - Complacency bias
 - Cyber attacks
 - DoS
 - Malware and ransomware
 - Passive Wiretapping
 - SQLi
 - Wardriving
 - Zero-day exploits
 - Adversarial machine learning
 - Use/integration of third-party products
- Methods of identification
 - Threat modeling
 - Engage SMEs in design and testing in order to utilize human judgment to identify blind spots and inference bias
 - Evaluate quality/suitability of data and models used to train and operate AI
 - Vetting liability in negligence law (consumer safety)
 - Attack trees
 - VAST modeling
 - Penetration testing
 - Red teams
 - Forensic analysis
 - Threat intelligence
 - Quantitative analysis – probability and impact
- Mitigation strategies (Security-by-Design practices)
 - Verify system is behaving as intended under actual operating conditions and can respond safely to unanticipated situations

- Ensure more rigorous standards in design, testing, operations documentation for AI systems are used to make consequential (life/death) decisions about people
- Implement controls for when/how AI should seek human input during critical situations and easily understood handover
- Designate rapid response team members for emergency reaction
- Feedback mechanism for users to report performance issues
- Revisit data security practices on a regular basis
- Set up secure environment for handling static data
- Protect integrity and security of data throughout networks and supply chains
- Destroy temporary databases that may contain aggregated data
- Adhere to license agreements associated with APIs
- Be conscious of the lack of control over streamed data
- Tools for identification/mitigation
 - Breach and Attack simulation tools
 - Threat modeling and analysis tools
 - Vulnerability scoring tools
 - Challenger models
 - SIEM
 - Black Box Multi-vector Testing
 - Threat Libraries
 - Risk Pattern Libraries
 - Cognitive security

Domain 4.0 Communication

Objective 4.1 Effectively communicate with key stakeholders and/or team members (internal communication)

- Identified ethical risks
- Business impacts
 - Social impacts
 - Organizational reputation
 - Consumer trust
 - Liability
 - Legal/regulatory obligations
- Business incentives

Objective 4.2 Effectively communicate about the ethical practices of the organization to outside parties (external communication)

- Marketing/Public Relations
- Brand awareness/value
- Media inquiries

- Corporate reporting
- Organizational philosophy
- Disclosure statements
 - Rationale explanation
 - Fairness explanation
 - Safety explanation

Domain 5.0 Organizational Policy and Governance

Objective 5.1 Identify the elements that can help foster an ethical organizational culture

- Training
- Leadership championing
- Incentive structures
- Culture-building workshops
- Creation of an ethics board
- Organizational resourcing

Objective 5.2 Identify and describe the ethical considerations that shape policies regarding the development, use, and governance of technology

- Fair competition
- Open data
- Privacy
- Intellectual property
- Fairness
- Non-discrimination
- Legal and regulatory requirements
- Human rights
- Accountability
- Transparency
- Animal rights/welfare
- Safety and reliability
- Environmental concerns
- Economic impacts
- Workforce impacts

Objective 5.3 Follow recommended guidelines for developing a code of ethics

- Identify internal and external stakeholders who should review or contribute
- Determine your organization's memberships in ecosystems, industry, and professional groups
- Collect codes of ethics from the above groups and aggregate their codes as a minimum baseline for your own code
- Consult with process owners to understand any factors which may frustrate or impede the adoption of more ethical practices

- Publish draft among stakeholders who will participate in a pilot for a determined length of time
- During the pilot period, interview stakeholders at predetermined intervals to understand impacts of change
- Upon completion of the pilot, update and ratify the code of ethics

Objective 5.4 Follow recommended guidelines for developing an ethical organizational policy

- Identify the need for the policy
 - Assess organizational activities, responsibilities, and the external environment to identify the need (e.g., anticipated need and in response to need) for policies and procedures
- Identify the owner(s) of the policy
 - Delegate responsibility to an individual, working group, staff, etc. according to expertise required
- Gather information about the policy
 - Leverage existing guidance, templates, or examples from internal and external sources
- Draft the policy
 - Ensure that verbiage, length, and complexity of policy are appropriate for those who will be tasked with implementation
- Consult appropriate stakeholders
 - Provide the appropriate affected parties with the opportunity to review and discuss potential implications of the policy
- Approve and publish the policy
 - Resolve any open issues
 - Obtain appropriate approval(s)
 - Publish the final policy
 - Make the policy available to the required parties
- Establish procedures to support the policy
 - Put existing procedures that support the policy into place
 - Provide guidance regarding how the policy will be implemented and by whom
- Implement the policy
 - Coordinate implementation
 - Determine how the policy will be communicated and by whom, if training is required, and if PR/external communications is required
- Monitor and refresh the policy at regular intervals
 - Establish monitoring and measurement/reporting systems to assess policy implementation, usage, and compliance
 - Review feedback from key stakeholders
 - Revise policy based on periodic refresh cycle

Objective 5.5 Evaluate the effectiveness of internal and external ethical policies

- Sentiment analysis of public discussion
- Surveys and focus groups
- Periodic health check of ethical policies
- Number and severity of ethical violations
- Industry best practices/leading best practices
 - Peer benchmarking

Recertification Requirements

The *Certified Ethical Emerging Technologist (CEET)* certification is valid for 3 years from the time certification is granted. In order to maintain a continuously valid certification, candidates can recertify via one of the following options:

1. Retake the most recent, up-to-date version of the exam before their certification expires.
2. Earn and submit enough continuing education credits (CECs) to recertify without retaking the exam.

Certified Ethical Emerging Technologist (CEET) Acronyms

Acronym	Expanded Form
AI	artificial intelligence
API	application program interface
BIPA	Illinois Biometric Information Privacy Act
CCPA	California Consumer Privacy Act
COPPA	Children's Online Privacy Protection Act
DoS	Denial of Service
ELI5	Explain Like I'm 5
EULA	End-user license agreement
FERPA	Family Educational Rights and Privacy Act
GDPR	General Data Protection Regulation
HIPAA	Health Insurance Portability and Accountability Act
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of Things
ISO	International Organization for Standardization
LGPD	Lei Geral de Proteção de Dados (General Data Protection Law)
LIME	Local Interpretable Model-agnostic Explanations
MFA	multi-factor authentication
ML	machine learning
NIST	National Institute of Standards and Technology
NISTIR	National Institute of Standards and Technology Interagency or Internal Report
OECD	Organisation for Economic Co-operation and Development
PCI DSS	Payment Card Industry Data Security Standard

PII	personally identifiable information
PIPEDA	Personal Information Protection and Electronic Documents Act
POPI	Protection of Personal Information Act
RACI	Responsible, Accountable, Consulted, and Informed
RAM	responsibility assignment matrix
SDK	software development kit
SHAP	SHapley Additive exPlanations
SIEM	Security Information Event Management
SLA	service-level agreement
SME	Subject Matter Expert
SOP	standard operating procedure
SQLi	Structured Query Language Injection
SSO	single sign-on
ToS	Terms of Service
VAST	Visual, Agile, and Simple Threat