



A PRIMER ON **Artificial Intelligence** *IN FINANCIAL MARKETS*

This primer format is intended to be an educational tool regarding emerging AI/FinTech innovations. It is not intended to state the official policy or position of the Commodity Futures Trading Commission (CFTC), or to limit the CFTC's current or future positions or actions. The CFTC does not endorse the use or effectiveness of any of the financial products or technologies in this presentation.

The CFTC's jurisdiction over any particular area will depend on specific facts and circumstances. Any examples included in this Primer are illustrative only and do not indicate a determination by the CFTC or its staff of jurisdiction.

Overview of AI

- What is “Artificial Intelligence” or “AI”?
- Key Catalysts
- Applications of AI in Financial Markets
- Potential Benefits

The Role of the CFTC

- Promote Responsible AI and Innovation
- Understand Applications of AI for Markets and Participants
- Identify Trends and Impacts

Moving AI Forward

- Risks and Challenges
- Responding to AI
- Key Considerations for AI Development



Overview of AI and Machine Learning

- *What is “Artificial Intelligence” or “AI”?*
- *Key Catalysts*
- *Applications of AI in Financial Markets*
- *Potential Benefits*

What is “Artificial Intelligence” or “AI”?



“The simulation of intelligent behavior in computers.”

- Merriam-Webster

Machine Learning is:
“A field of study that gives computers the ability to learn without being explicitly programmed.”

– Arthur Samuel, 1959

Putting AI in Context

For centuries, philosophers and mathematicians have considered the potential of machine-based intelligence.

Since the early 20th Century, AI has featured prominently in popular culture and captured human imagination. Futurama at the 1939 World's Fair Exhibits, films and other media featured self-driving cars, thinking machines and other AI applications.

In his seminal 1950 paper *Computing Machinery and Intelligence*, Alan Turing helped lay the groundwork for modern AI, posing a test of when machines could be said to think: "A computer would deserve to be called intelligent if it could deceive a human into believing that it was human."

In the 1950s, research on machine intelligence intensified and became more focused. The term "AI" was coined by John McCarthy in 1956.

Approaches to AI have evolved. Early calculating machines gave rise to "expert systems."

Subsequent waves of AI have focused on enabling machines to acquire new understanding based on experience.

Such "machine learning" has become a key element of AI.

Current applications of AI vary widely and are growing exponentially.




- “Every aspect of learning or other feature of intelligence can, in principle, be so precisely described that a machine can be made to simulate it.”
 - *Mission statement of the Dartmouth Conference, 1956*
- “Artificial intelligence would be the ultimate version of Google. It would understand exactly what you wanted, and it would give you the right thing.”
 - *Larry Page, Google Co-Founder, 2000*
- “Just as electricity transformed everything almost 100 years ago, today I actually have a hard time thinking of an industry that I don’t think AI will transform in the next several years.”
 - *Andrew Ng, Professor, Stanford University, 2017*

AI Approaches: Two End Points

Two end points define a continuum of AI research and applications.

GENERAL INTELLIGENCE

- Theme: Machines behaving like humans.
- Adaptable to any job. Not domain specific.
- Able to perceive and react to broad range of contexts.
- The pop culture perception of AI.
- Does not presently exist.
- High transferability within and between domains.



SPECIALIZED INTELLIGENCE

- Theme: Machines exceeding human capability at specific tasks.
- Narrow focus. Domain specific.
- Task or job focused.
- Use cases are proliferating.
- Limited transferability across domains.

General Intelligence is much more difficult to develop. For the present and foreseeable future, efforts are focused on developing use cases for specialized intelligence.

Applications of specialized intelligence are often designed to **augment** human intelligence and capabilities, but not act independently of human oversight.

AI is a set of technologies and approaches that has evolved over time.

Expert Systems

- Rules-based, hard-coded algorithms.
- Developers provide the machine with a roadmap of anticipated input and directed output.
- Examples: First generation chess-playing machines and “domain expert” programs.

Machine Learning

- Machine learning may take a variety of approaches that include learning algorithms, pattern recognition, graphical and statistical modeling, and decision trees.
- Examples: Natural language processing, facial recognition, and robotic process automation are each applications of machine learning.

Deep Learning

- A type of machine learning in which the machine incorporates context sensitivity and machine-driven pattern discovery. Can make use of reinforcement learning to extract progressively higher level features from data and master complex topics in a short time frame.
- Examples: Gaming machines like Alpha Zero, language translation, medical diagnosis, and object detection by self-driving vehicles.

AI efforts have developed in waves. Now, a convergence of technologies coupled with advances in data science, mass adoption, and increased investment is driving AI forward.

Advanced Software

- Open Source Software and Libraries that focus on and facilitate AI (e.g., TensorFlow, Python)
- Deep Learning
- Vision, Natural Language Processing, Robotic Process Automation

Big Data

- Vast troves of accessible data
- Internet of Things (IoT) (e.g., sensors, satellites)
- Cloud storage


Big Compute

- Greatly increased compute power
- Specialized Hardware (e.g., application-specific integrated circuits)
- Cloud computing

Catalyst - The Rise of Big Data

*“In a world where data is coin of the realm...
Data is becoming the new raw material of business.”*

— *The Economist, “Data, Data Everywhere,” Feb 27, 2010*



- With digitization and greater storage capability, big data has become ubiquitous.
- Data has proliferated through Internet of Things (IoT), sensors, streaming, text, audio, and video files.
- The use of cloud storage has massively increased the accessibility of large-scale storage and processing of big data.
- Data is the raw material that enables new AI systems to train.
- Big data fuels AI. AI's power is analyzing vast pools of data to detect patterns and generate insights.

Catalyst – Big Compute

“Now, what used to be thought of as supercomputers are inside smartphones. They cost a million times less, are a million times faster, and have a million times as much memory.”

-Richard Mark Soley, Ph.D., Executive Director, Industrial Internet Consortium

- Growth of AI – particularly Deep Learning processes – requires significant compute power.
- As it has increased access to data, the development of cloud storage has facilitated widespread access to high powered computing, enabling substantial advances in AI.
- Computing power has grown exponentially, while costs have declined.
- AI-tailored hardware solutions such as Programmable Systems on a Chip (PSoC) and parallel core architecture are being developed.
- Going forward, the advent of quantum computing is expected to further expand AI’s potential.



Case Study – Four Hours to Master Chess




Case Study: Playing Games from “Deep Blue” to “Alpha Zero”

- IBM’s Deep Blue supercomputer defeated world grand master Gary Kasparov at chess in 1997. Deep Blue began a progression of AI game-playing master machines.
- AlphaZero is a recent incarnation that was developed by DeepMind, which is now part of Google. Starting from scratch, with only the game rules, AlphaZero became expert in chess, shogi and go. In each case, the system beat a human world champion.
- In 2016, AlphaZero also beat the world’s top chess-playing computer “Stockfish 8” after only four hours of learning.
- Within the domain of each game, the AI system demonstrated an ability to learn “context” – which move makes sense when – in a way that approximates human grand masters.



Applications of AI to financial markets are wide-ranging.



- Trading (e.g., market intelligence, robo-advisory, sentiment analysis, algorithmic trading, smart routing, and transactions)
- Risk Management (e.g., margin and capital requirements, trade monitoring, fraud detection)
- Risk Assessments and Hedging
- Resource Optimization (e.g., energy and computer power)
- RegTech – Applications that enhance or improve compliance and oversight activities (e.g., surveillance, reporting)
- Compliance (e.g., identity and customer validation, anti-money laundering, regulatory reporting)
- Books and Records (e.g., automated trade histories from voice / text)
- Data Processing and Analytics
- Cybersecurity and Resilience
- Customer Service

Potential Benefits of AI



Mitigate Human Error: AI systems should not fatigue, commit computational errors, or make ad hoc or subjective judgments.

Economy and Speed: Automation reduces transaction times and unnecessary manual processes.

Leverage Crowd Wisdom: AI systems are able to identify trends from vast pools of data – essentially to discern insights from the crowd.

Productivity and Efficiency: AI systems can leverage computational power to increase productivity and free humans to focus on higher level tasks.

Improve Human Experience: In numerous work and personal contexts, AI has the potential to improve human experience.

Facilitate Market-Enhancing Technologies: AI will enable new business models far beyond present capabilities.

Inclusion: AI has the potential to counter human biases, enabling a more inclusive financial system.

Improve Access: AI can streamline customer onboarding and account setup.

Enhance Regulatory Compliance: AI systems can enhance identity, know-your-customer, and anti-money laundering procedures.

Improve Risk Management: AI can help individual participants identify and manage risk as well as broader systemic risk monitoring and mitigation.

Improve Market Integrity: AI can help detect and minimize fraud, manipulation, and other illicit activity.



The Role of the CFTC

- *Promote Responsible AI and Innovation*
- *Understand Applications of AI for Markets and Participants*
- *Identify Trends and Impacts*

“Artificial Intelligence (AI) promises to drive growth of the United States economy, enhance our economic and national security, and improve our quality of life.... It is the policy of the United States Government to sustain and enhance the scientific, technological, and economic leadership position of the United States in AI R&D and deployment through a coordinated Federal Government strategy, the American AI Initiative....”

Source: “Executive Order on Maintaining American Leadership in Artificial Intelligence,” February 11, 2019, available at <https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/>

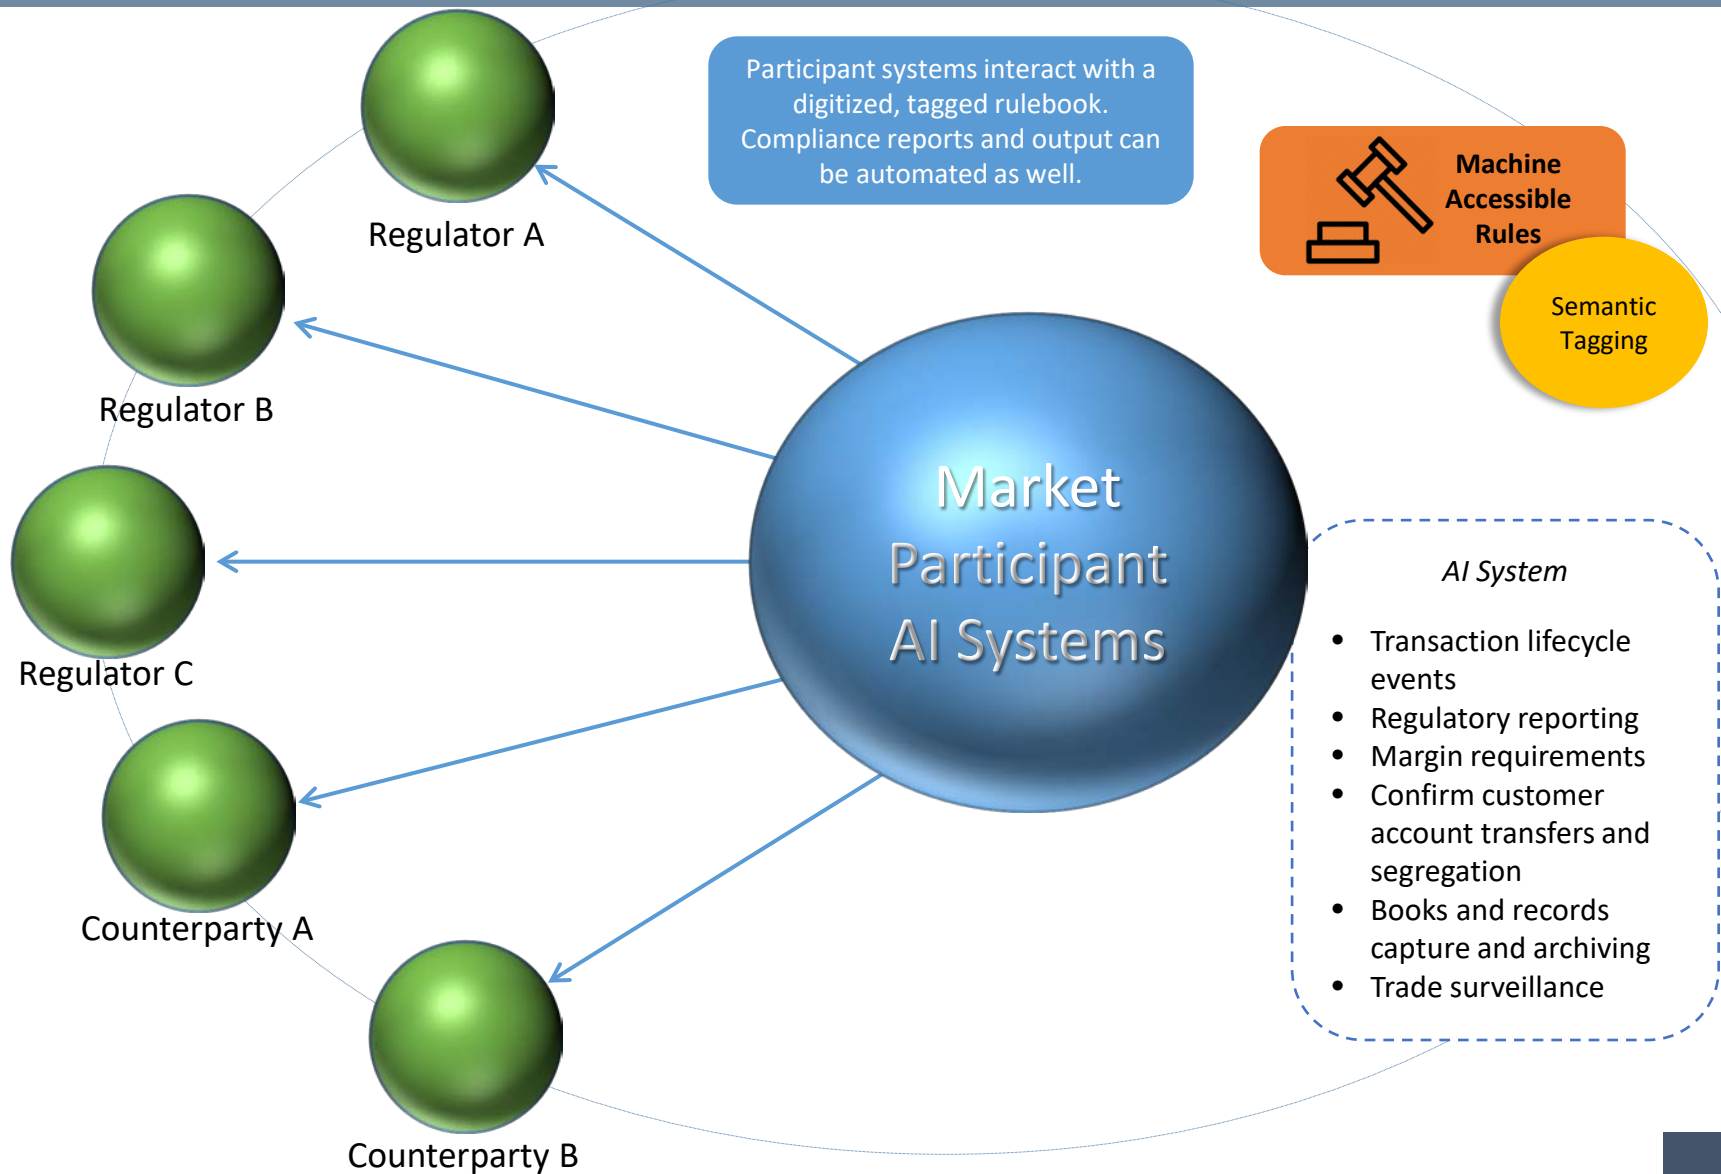


The CFTC's Mission and Approach



- The CFTC's mission is to promote the integrity, resilience, and vibrancy of the U.S. derivatives markets through sound regulation.
- Applications of AI and related technologies have the potential to enhance our markets. Accordingly, the CFTC proactively engages with relevant stakeholders and seeks to promote promising new technologies.
- Dedicated FinTech Stakeholder: LabCFTC
 - Maintain open door with innovator community.
 - Track developments to understand market impacts and regulatory implications.
 - Identify opportunities to leverage new technology, including AI, at the agency as well as in markets that the CFTC oversees.

Potential Application of AI: Market Participant Compliance



Potential Application of AI: Market Oversight

Modern markets are fast-moving, complex networks with numerous participants carrying out various functions. Effective oversight supports strong, well-working markets. As a market regulator, the CFTC could leverage AI to:

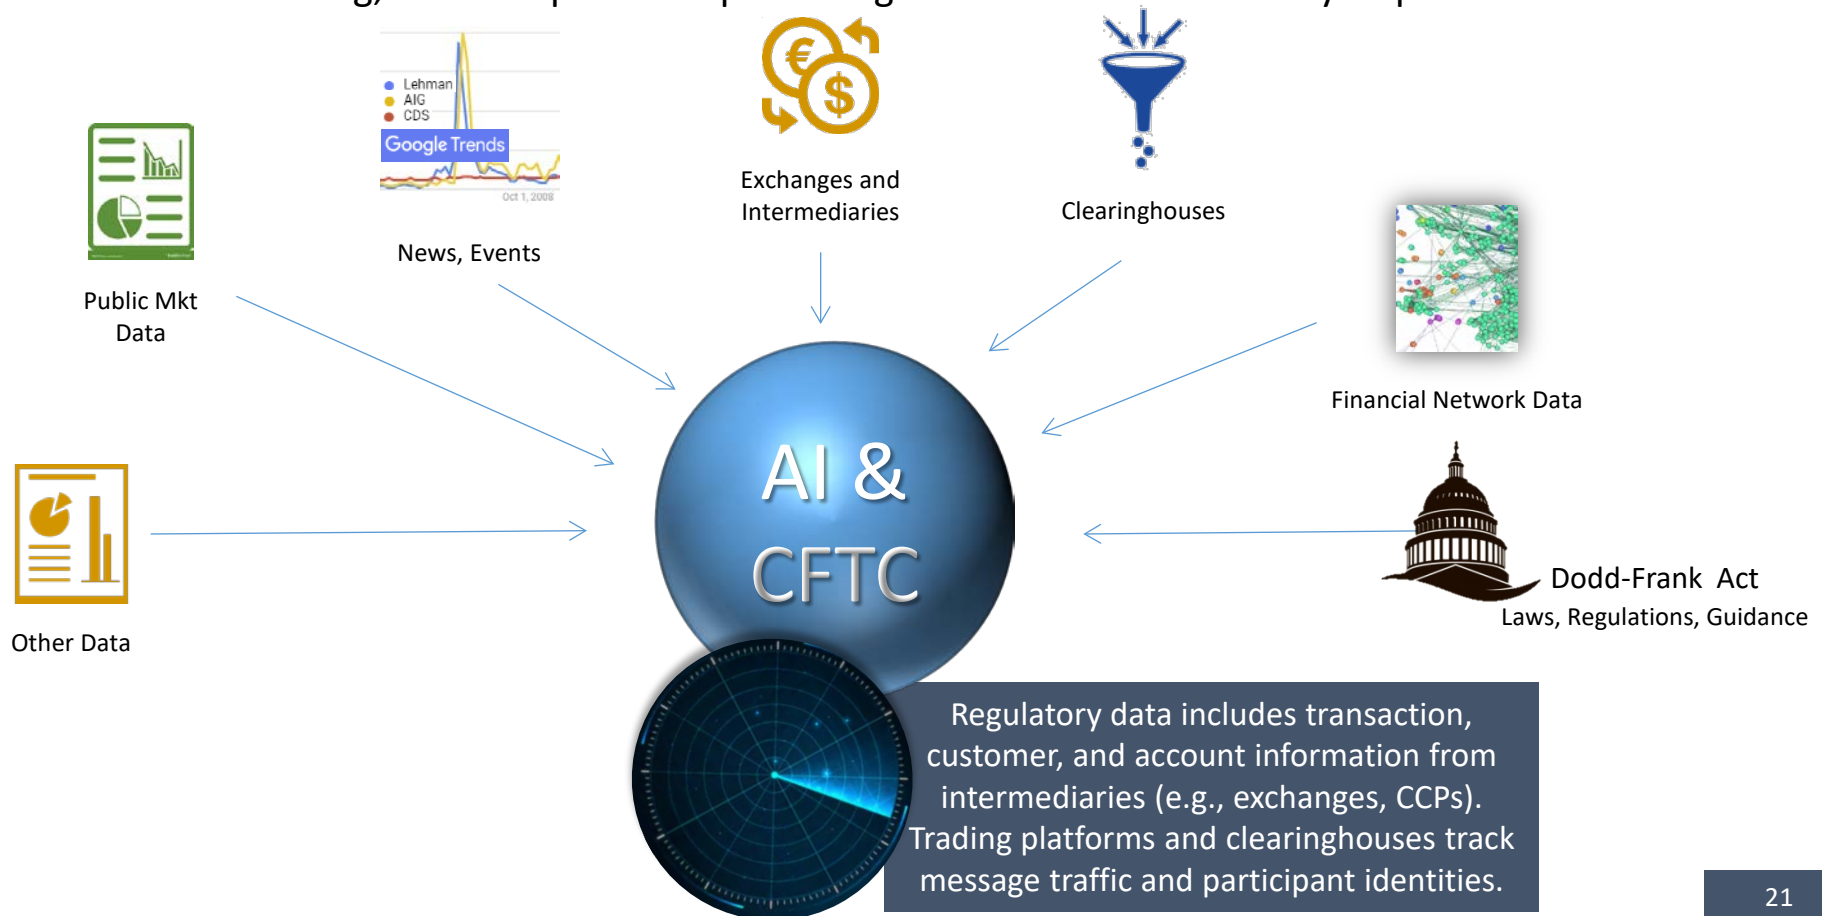
- “Read” a complex marketplace and distinguish salient activity.
- Use data to develop market models and identify risk factors.
- Apply the models to conduct ongoing market and risk surveillance.
- Help identify market manipulation, abusive trading, and fraud.



Potential Application of AI: Systemic Risk Monitoring

Predicting catastrophic market events, such as the cascading defaults of 2008, is like predicting the weather. There are many variables, which can generate diverging predictions, and some key information may be overlooked or not available. This can impede corrective action.

A strength of AI is its ability to identify correlations in vast data sets. Such correlations can be helpful in systemic risk monitoring, which depends on predicting when circumstances may require intervention.



Regulatory Implications of AI



- How might AI and related technology impact market structure? How can it enhance the quality of our markets?
 - New business models.
 - New mechanisms for market oversight.
- What safeguards are necessary or appropriate?
 - Regarding data acquisition, use, management, distribution, and access.
 - Regarding systems resiliency and disaster recovery.
 - Regarding undesirable outcomes or performance of AI-based tools.
- Who is accountable and legally responsible for the use of AI in markets the CFTC oversees?
 - The CFTC's current regulatory framework is based around registrants – a defined set of market participants.
 - Traditionally, third-party service providers provide services to CFTC-registered entities, who are legally responsible for compliance. Third-party service providers are not separately regulated.
 - The proliferation of AI systems with increasing functionality has the potential to blur traditional boundaries.

- The CFTC regulates derivatives market activity, not particular technologies.
- Uses of AI may include regulated activities (e.g., customer onboarding, KYC and AML) or regulatory requirements (e.g., transaction and position reporting).
- CFTC registrants that use AI, or engage with firms that do, remain responsible for compliance with all applicable requirements.
- Where an AI system engages in regulated activity, such work must be conducted on behalf of and subject to the direction of a CFTC registrant.
- Use of AI systems by CFTC registrants should comply with all applicable legal rules and standards.
- Entities using AI should consider consulting with competent counsel.

Data
mining

ms

**MACHINE
LEARNING**

Artificial
Intelligence

Prediction

Moving AI Forward

- *Risks and Challenges*
- *Responding to AI*
- *Key Considerations for AI Development*

Key Challenges and Risks



AI presents opportunity, but also certain risks. Addressing the challenges will facilitate AI's ongoing growth and help enable its value to be fully realized.

- **Operational Risk:** To be effective, AI systems must operate as intended. What approaches can help ensure that they do?
- **Economic Risk:** AI systems have economic impact. How can less desirable impacts on workers or other aspects of the economy be mitigated?
- **Trust:** To provide value, AI systems interact with vast amounts of data, including personal information. What can be done to maintain public trust?
- **Accountability:** AI systems raise novel questions about accountability. If something goes awry, who is responsible and for what? Regarding data, what rules apply?
- **Security Risk:** AI systems are subject to various types of attack and must be resilient. What safeguards can mitigate this risk?



While AI has the potential to enhance human capabilities and experiences, in their use of AI, financial market participants should work to ensure:

- AI systems are working as intended.
- AI systems are safe and do not generate or exacerbate systemic risk.
- AI systems serve the needs of all relevant stakeholders and operate fairly.
- AI systems protect privacy appropriately and as required.

Further development of AI can benefit from collaboration and the open exchange of ideas between market participants, regulators, and other stakeholders. Key elements of a response include:

- Development Standards.
- Governance and Oversight.
- Security and Resiliency.

A broad spectrum of principles may be useful as AI continues to develop.

AI Systems

- **Transparency:** Degree to which mechanisms providing “explainability” are built into systems.
- **Performance:** Metrics to assess how the AI system performs and quantify output in terms of accuracy, false positives, and false negatives.
- **Interoperability:** How should systems interface with the world and each other?
- **Testing Standards:** AI systems should be rigorously tested before they are implemented and as they operate over time.
- **Security and Resiliency:** Safeguards to protect against attacks and manipulation of AI systems, and contingency plans to respond.

AI System Data

- **Data Sourcing and Integrity:** Appropriate practices and principles for sourcing and safeguarding data, including disclosure, rights to review and correct, privacy, and anonymity.
- **Data Access:** Is necessary data accessible? How can it be effectively used and managed?
- **Resiliency:** Appropriate practices and principles to ensure trustworthiness, safety, and resiliency of data storage facilities and systems.

AI System Governance and Oversight



Effective governance and oversight is vital for AI systems.

Key elements include:

- **Accountability:** Who is responsible and for what?
- **Sound Data Policies:** Who has the right to capture and use data? To what extent may AI systems be deployed on such data?
- **Compliance:** What types of licensing or certification may be appropriate?
- **Oversight:** How will new AI systems be reviewed? How should AI systems be supervised as they learn and evolve? How will complaints be addressed?
- **AI Development:** What principles should guide the design and development of AI systems to reflect social standards and support reliable, robust, and trustworthy systems?
- **AI Economy:** How will a diverse, growing set of AI systems interact? How will unintended consequences of interactions between multiple AI systems and operators be addressed? What common standards may be helpful?

Security and Resiliency

- AI systems develop and change over time, and may not function properly or as intended. There could be a design fault or an unanticipated error. Data the system uses may be corrupt, biased, or incomplete.
 - What guardrails can prevent AI systems from running amok? How can we ascertain when AI systems may not be working properly? What metrics or approaches may be helpful to evaluate AI system operation and outcomes?
- AI systems and data centers can be hacked. AI systems can be impacted by adverse data attacks.
 - What approaches can mitigate these risks? What types of deterrence may be most effective? How can stakeholders coordinate effectively?



Key Considerations for AI Development

Implemented responsibly, AI has great potential to enhance our markets.



- AI systems depend on good design and ongoing monitoring.
- To build effective AI systems, careful consideration must be given to the choice of algorithms, the sourcing of data, and the evolution of AI models.
- Appropriate governance and controls are vital for AI to succeed.
- AI systems assist and augment, but cannot replace, human judgment.
- For long-term success, AI systems must be reliable, resilient, and trustworthy.



Questions?

Contact us at
LabCFTC@cftc.gov
www.cftc.gov/LabCFTC