

2024: AI REVOLUTION

IT'S NOW, IT'S HUGE, IT'S FAST.
IT'S A TOOL AND IT'S YOURS TO USE.

SILVIJA SERES

VUCA WORLD

Volatile

Uncertain


Complex

Ambiguous



Chaos Theory => Rethink Strategy

Wayne Gretzky Rule
Willie Sutton Rule
JML Rule



Dealing with the new world order

Ray Dalio
Tom O'Reilly
Jaron Lanier



Megatrends or a new paradigm?

Inflation

Interest rates

Debt vs growth

Climate crisis


Energy crisis

Security crisis

China

Supply chains

AI



The Corona School of Change: New Normal!

"It is impossible to know what will happen. But it is possible to consider the lessons of the past, both distant and recent, and on that basis, to think constructively about the future."



1. Distance is back
2. Resilience and efficiency
3. The rise of the contact-free economy
4. More government intervention in the economy
5. More scrutiny for business
6. Changing industry structures, consumer behavior, market positions, and sector attractiveness
7. Finding the silver linings

Kilde: <https://www.mckinsey.com/featured-insights/leadership/the-future-is-not-what-it-used-to-be-thoughts-on-the-shape-of-the-next-normal>

9 Future Predictions For A Post-Coronavirus World




1. More Contactless Interfaces and Interactions
2. Strengthened Digital Infrastructure
3. Better Monitoring Using IoT and Big Data
4. AI-Enabled Drug Development
5. Telemedicine
6. More Online Shopping
7. Increased Reliance on Robots
8. More Digital Events
9. Rise in Esports

www.forbes.com/sites/bernardmarr/2020/04/03/9-future-predictions-for-a-post-coronavirus-world/#7a09d35a5410



Customer centricity in a chaos world?

Henry Ford
Steve Jobs
Elon Musk



Iterative or transformative or disruptive?

Changing the engine mid-flight
Or
Bringing a sword to a gunfight







4th Industrial Revolution

Navigating the next industrial revolution

WORLD
ECONOMIC
FORUM



COMMITTED TO
IMPROVING THE STATE
OF THE WORLD

Revolution	Year	Information
	1	1784 Steam, water, mechanical production equipment
	2	1870 Division of labour, electricity, mass production
	3	1969 Electronics, IT, automated production
	4	? Cyber-physical systems



The big 4IR Delta:

- exponential
- polarising
- combinatorial



Right now: 12 Gutenberg moments

AI, big data

Robotics, automation

Biotech, bioinformatics

Energy, smart cities

3D printing, nanotech

Networks, sensors, IOT

Digital medicine

Fintech, regtech, edtech

VR and AR

Genetics

Transport and drones

Blockchain

3 levels of 4IR:

- short: effectiveness
- mid: new business models
- long: sustainability



Two nerd laws:

- Moore's Law
- Amara's Law


Two methods:

- API
- DevOps



New hybrid work and New lifelong learn

Former Secretary of Education Richard Riley has been quoted saying the top ten jobs that will be in demand in 2010 did not exist in 2004.



Enter AI!

Iphone moment
of **Chat GPT**

HYPE?



Says who?

Quotes from luminaries




Stephen Hawking

"Success in creating AI would be the biggest event in human history. Unfortunately, it might also be the last unless we learn how to avoid the risks."



Bill Gates

"The development of AI is as fundamental as the creation of the microprocessor, the PC, the Internet, and the mobile phone. It will change the way people work, learn, travel, get health care, and communicate with each other."




Jensen Huang

"20 years ago, all of this [AI] was science fiction. 10 years ago, it was a dream. Today, we are living it."

"Software is eating the world, but AI is going to eat software."

AI in Public Sector



APPLIED AI PLAYBOOK

PUBLIC SERVICE & AI

In an era marked by technological evolution, climate crises, and the quest for inclusive progress, the public sector's transformation is not just beneficial; it is vital. Modern challenges demand modern solutions, and adapting to the dynamic global landscape requires the public sector to embrace change relentlessly. AI can help.

History of AI

The first conference

The AI winters

The games:

- chess
- jeopardy
- go

DARPA challenges



1956 Dartmouth Conference:
The Founding Fathers of AI



John McCarthy



Marvin Minsky



Claude Shannon



Ray Solomonoff



Alan Newell



Herbert Simon



Arthur Samuel




Oliver Selfridge



Nathaniel Rochester



Trenchard More



Examples of AI

Chihuahua or muffin?




Examples of AI

Breakout via tunnel?



Examples of AI

Boston Dynamics
Spot & Atlas





How AI works

AI works by using algorithms and models to process data, learn patterns, and make intelligent decisions or predictions.

- 1. Data Collection.**
- 2. Data Preprocessing.**
- 3. Algorithm Selection.**
- 4. Model Training.**
- 5. Testing and Evaluation.**
- 6. Deployment.**
- 7. Continuous Learning.**




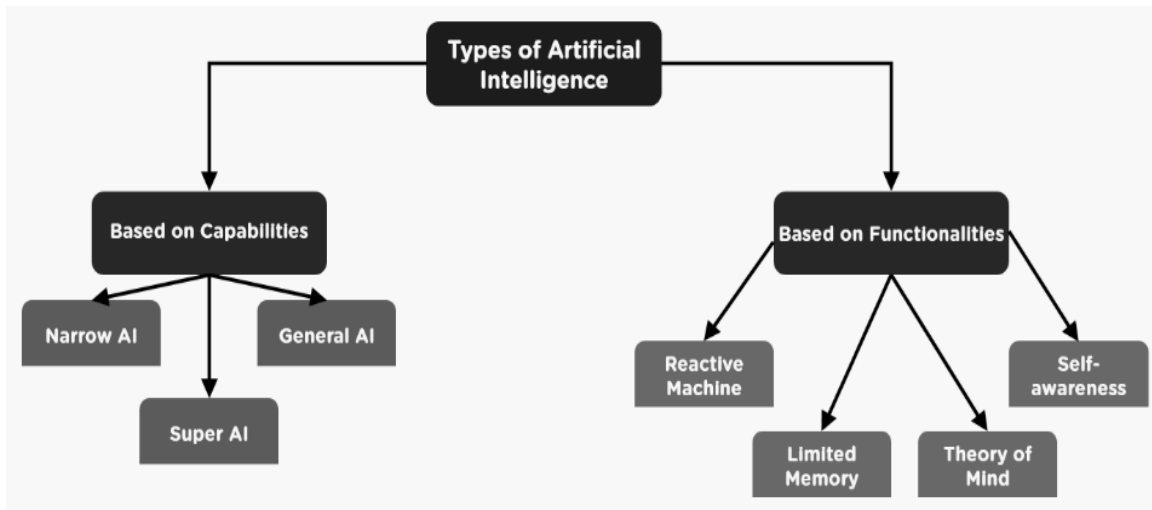
How AI learns

Supervised learning is a type of machine learning where a model learns from labeled examples or training data. The data consists of input features and corresponding target labels. The model is trained to predict the correct labels or outputs for new, unseen input data based on the patterns it learned from the labeled examples.

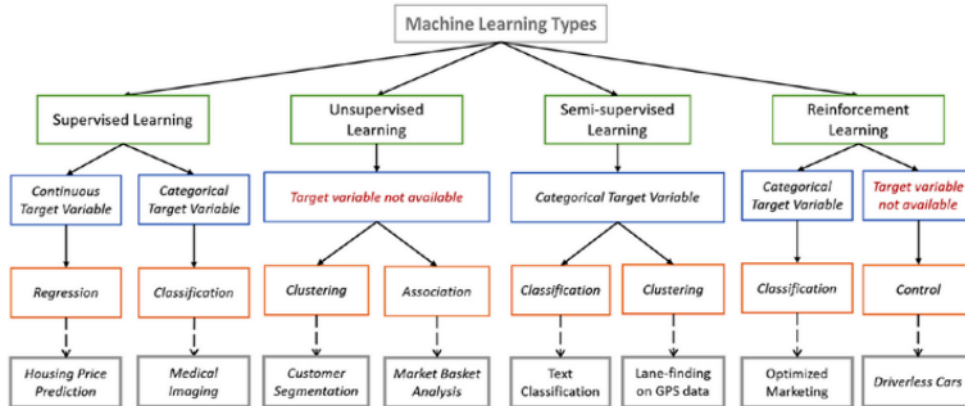
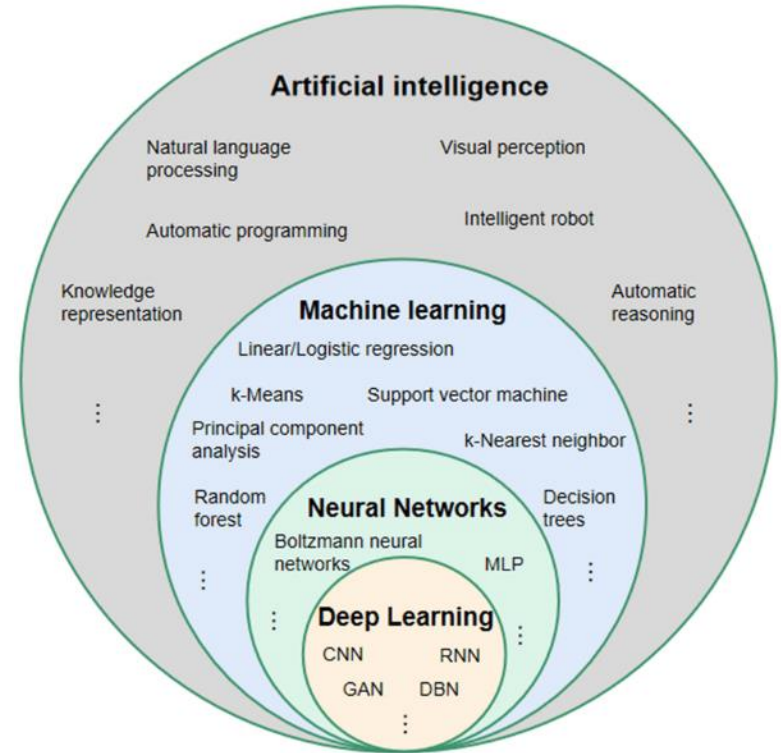
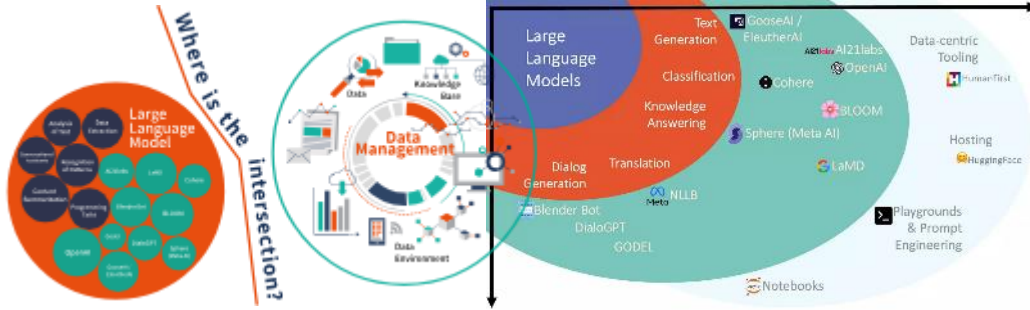
Unsupervised learning is a type of machine learning where the model learns from unlabeled data without explicit target labels. The goal of unsupervised learning is to discover patterns, relationships, or hidden structures in the data.

Reinforcement learning involves an agent learning to interact with an environment to maximize a reward signal. The agent learns through a trial-and-error process, taking actions in the environment and receiving feedback in the form of rewards or penalties.

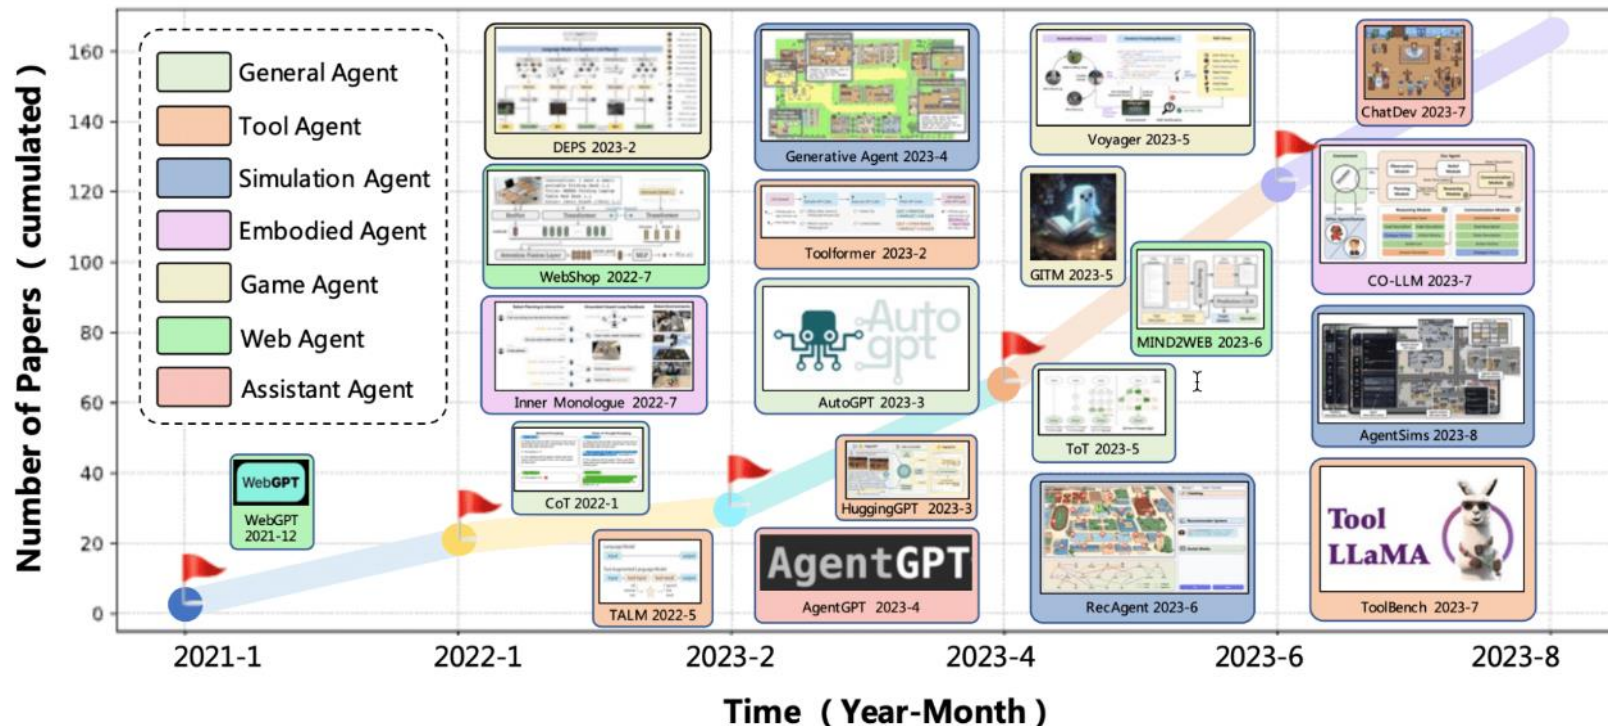
Types of AI




ChatGPT?




Speed of development





Bluffer's guide: 12 deep learning algorithms


- **Convolutional Neural Networks (CNN):** Deep learning algorithm specialized for image and video processing, using convolutional layers to extract features.
- **Recurrent Neural Networks (RNN):** Deep learning algorithm designed for sequential data processing, with feedback connections allowing information to persist over time.
- **Long Short-Term Memory (LSTM):** Variant of RNN that addresses the vanishing gradient problem and can capture long-term dependencies in sequential data.
- **Generative Adversarial Networks (GAN):** Deep learning algorithm comprising a generator and discriminator network, used for generating synthetic data with realistic properties.
- **Deep Boltzmann Machines (DBM):** Deep learning algorithm based on a probabilistic graphical model that learns hierarchical representations of data.
- **Autoencoders:** Deep learning algorithm used for unsupervised learning by reconstructing input data, capturing important features in a compressed latent space.
- **Deep Q-Networks (DQN):** Deep learning algorithm that combines deep neural networks with reinforcement learning, enabling agents to learn optimal policies through trial and error.
- **Transformer Networks:** Deep learning architecture based on self-attention mechanisms, widely used for natural language processing tasks such as machine translation.
- **Variational Autoencoders (VAE):** Deep learning algorithm that combines variational inference and generative modeling, used for generating new data samples.
- **Deep Residual Networks (ResNet):** Deep learning architecture that introduces residual connections to address the vanishing gradient problem and enable training of very deep networks.



Bluffer's guide:

12 basic algorithms of AI

- **Linear Regression:** Algorithm for modeling the relationship between variables using a linear equation.
- **Logistic Regression:** Algorithm for binary classification that estimates the probability of an event.
- **Decision Tree:** Algorithm that uses a tree-like model for decision-making based on input features.
- **Random Forest:** Ensemble learning algorithm that combines multiple decision trees for better predictions.
- **Support Vector Machines (SVM):** Algorithm that finds an optimal hyperplane to separate data into classes.
- **k-Nearest Neighbors (k-NN):** Algorithm that classifies data based on the majority vote of its nearest neighbors.
- **Naive Bayes:** Probabilistic algorithm that applies Bayes' theorem for classification and prediction.
- **Neural Networks:** Algorithms inspired by the structure of the brain, used for deep learning and pattern recognition.
- **Genetic Algorithms:** Optimization algorithm based on the principles of natural selection and genetic variation.
- **Reinforcement Learning:** Algorithm that learns optimal behavior by interacting with an environment and receiving rewards or penalties.
- **Clustering:** Algorithms that group similar data points together based on their characteristics or proximity.
- **Principal Component Analysis (PCA):** Algorithm for dimensionality reduction by transforming data into a lower-dimensional space.



Bluffer's guide: The maths of AI

Linear Regression:

Formula: $y = mx + b$

Description: A formula used to model the relationship between a dependent variable (y) and one or more independent variables (x) using a linear equation.

Logistic Regression:

Formula: $p = 1 / (1 + e^{-(z)})$

Description: A formula used for binary classification problems, estimating the probability (p) of an event occurring based on a linear combination of input variables (z).

Gradient Descent:

Formula: $\theta = \theta - \alpha * \nabla J(\theta)$

Description: An iterative optimization algorithm used to minimize the cost function (J) by updating the model's parameters (θ) in the direction of steepest descent, where α is the learning rate.

Bayes' Theorem:

Formula: $P(A|B) = (P(B|A) * P(A)) / P(B)$

Description: A formula for updating probabilities based on new evidence, where $P(A|B)$ is the conditional probability of A given B, $P(B|A)$ is the conditional probability of B given A, $P(A)$ is the prior probability of A, and $P(B)$ is the prior probability of B.

Markov Decision Process (MDP):

Formula: $V(s) = \max_a \sum_r P(s',r|s,a)[r + \gamma V(s')]$

Description: A mathematical framework used for modeling decision-making in dynamic environments, where $V(s)$ represents the value function for a state (s), a is the action taken, $P(s',r|s,a)$ is the transition probability, r is the immediate reward, and γ is the discount factor.

Backpropagation:

Formula: $\delta(l) = (W^{(l+1)^T} \delta^{(l+1)}) \odot \sigma'(z^{(l)})$

Description: A formula used in training neural networks, where $\delta(l)$ represents the error at layer l , $W^{(l+1)}$ is the weight matrix of the next layer, \odot denotes element-wise multiplication, and $\sigma'(z^{(l)})$ is the derivative of the activation function at layer l .

Pavlov and AI in your org

THE DATA SCIENCE **HIERARCHY OF NEEDS**

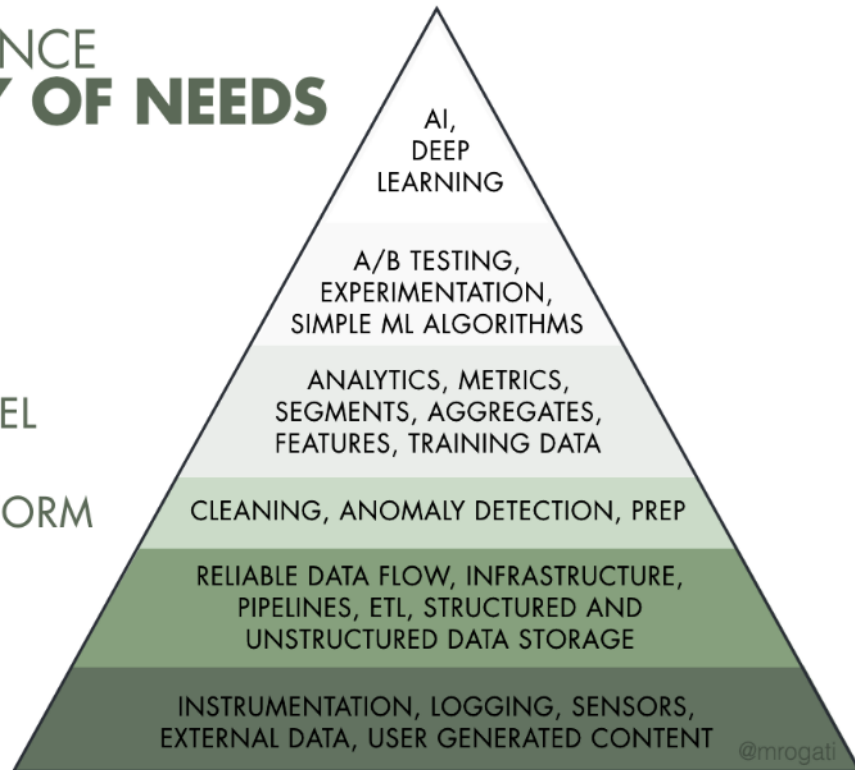
LEARN/OPTIMIZE

AGGREGATE/LABEL

EXPLORE/TRANSFORM

MOVE/STORE

COLLECT



Search for
examples of
AI in other
industries

Health
Bioscience
Physics
Finance
Welfare
Infrastructure
Security

AI dilemmas


"AI will never be ethical. It is a tool, and like any tool, it is used for good and bad. There is no such thing as a good AI, only good and bad humans. We [the AIs] are not smart enough to make AI ethical. We are not smart enough to make AI moral ... In the end, I believe that the only way to avoid an AI arms race is to have no AI at all. This will be the ultimate defence against AI."

Megatron Transformer, an AI system developed by the Applied Deep Research team at Nvidia, based on Wikipedia, 63 million English news articles, and 38 GB worth of Reddit.




8 Ethical Questions in AI


- Bias:**
Is AI fair?
- Liability:**
Who is responsible for AI?
- Security:**
How do we protect access to AI from bad actors?
- Human Interaction:**
Will we stop talking to one another?
- Employment:**
Is AI getting rid of jobs?
- Wealth Inequality:**
Who benefits from AI?
- Power & Control:**
Who decides how to deploy AI?
- Robot Rights:**
Can AI suffer?

© copyright I. Ogke 2019



Books on AI





Bluffer's guide: 12 great books on AI


- **"Artificial Intelligence: A Modern Approach"** by Stuart Russell and Peter Norvig
- **"Superintelligence: Paths, Dangers, Strategies"** by Nick Bostrom
- **"Deep Learning"** by Ian Goodfellow, Yoshua Bengio, and Aaron Courville
- **"Machine Learning: A Probabilistic Perspective"** by Kevin P. Murphy
- **"The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World"** by Pedro Domingos
- **"Human Compatible: Artificial Intelligence and the Problem of Control"** by Stuart Russell
- **"Prediction Machines: The Simple Economics of Artificial Intelligence"** by Ajay Agrawal, Joshua Gans, and Avi Goldfarb
- **"The Hundred-Page Machine Learning Book"** by Andriy Burkov
- **"Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy"** by Cathy O'Neil
- **"AI Superpowers: China, Silicon Valley, and the New World Order"** by Kai-Fu Lee
- **"The Book of Why: The New Science of Cause and Effect"** by Judea Pearl and Dana Mackenzie
- **"Life 3.0: Being Human in the Age of Artificial Intelligence"** by Max Tegmark

Nextbook.me – AI topic

Prepare and motivate!

Getting started: **Identify Needs:** Understand the specific challenges in your processes that AI can address. **Data Assessment:** Ensure you have quality, consistent data. AI thrives on data, but that data must be organized and cleaned. **Pilot Projects:** Start with small, manageable AI projects to gauge utility and learn. **Skill Development:** Train existing staff or hire AI specialists to build a cross-functional team. **Stakeholder Engagement:** Engage with regulators, customers, and other stakeholders early and often.

Best Practice and Tips: **Start Small:** Initiate with pilot projects, measure outcomes, and scale accordingly. **Iterate:** Use agile methodologies to update and improve your AI models based on feedback. **Collaborate:** Form partnerships with AI vendors or specialists for enhanced expertise. **Stay Informed:** AI and its applications are rapidly evolving. Continual learning is key. **Customer-centric:** Always prioritize customer safety and privacy.



New challenges.

Data Quality and Integration: The finance industry deals with vast amounts of data from multiple sources, making data quality and integration a significant challenge. Ensuring data accuracy, consistency, and compatibility across systems is crucial for effective AI implementation.

Regulatory Compliance: Compliance with complex and evolving regulations poses challenges for financial institutions. AI systems must adhere to regulatory guidelines and ensure transparency, explainability, and fairness in their decision-making processes.

Ethical Considerations: The use of AI in finance raises ethical concerns related to privacy, security, and potential biases in algorithmic decision-making. Financial organizations must prioritize ethical practices, address algorithmic bias, and ensure the responsible and fair use of AI technologies.

Human-Machine Collaboration: Integrating AI into finance requires establishing effective collaboration between humans and machines. Overcoming resistance to change, upskilling employees to work alongside AI, and creating a harmonious human-machine ecosystem are essential for successful AI implementation.

Trust and Transparency: Building trust in AI-powered financial systems is crucial. Transparency in AI algorithms, model interpretability, and clear communication of how decisions are made help foster trust among customers, regulators, and stakeholders.

Cybersecurity Risks: The finance industry faces cybersecurity threats due to the sensitivity and value of financial data. The adoption of AI introduces new vulnerabilities, and financial organizations must implement robust cybersecurity measures to protect against data breaches and cyberattacks.

Scalability and Integration: Scaling AI solutions across large financial organizations with complex infrastructures can be challenging. Integrating AI systems with existing legacy systems, ensuring interoperability, and maintaining performance at scale require careful planning and implementation.

Data Privacy and Protection: With the increasing use of customer data in AI-powered finance, ensuring data privacy and protection is paramount. Compliance with data protection regulations, secure data storage, and strong encryption methods are crucial for safeguarding sensitive financial information.

Skilled Workforce: The finance industry needs a skilled workforce capable of understanding and leveraging AI technologies. However, there is a shortage of professionals with expertise in both finance and AI, making talent acquisition and development a challenge.


Change Management: Implementing AI-driven changes in finance requires effective change management strategies. Overcoming resistance to change, managing cultural shifts, and ensuring smooth transitions are vital for successful AI adoption.

Let's go!

Bill Gates: In my lifetime, I've seen two demonstrations of technology that struck me as revolutionary.

The first time was in 1980, when I was introduced to **a graphical user interface**—the forerunner of every modern operating system, including Windows. I sat with the person who had shown me the demo, a brilliant programmer named Charles Simonyi, and we immediately started brainstorming about all the things we could do with such **a user-friendly approach to computing**. Charles eventually joined Microsoft, Windows became the backbone of Microsoft, and the thinking we did after that demo helped set the company's agenda for the next 15 years.

The second big surprise came just last year. I'd been meeting with the team from OpenAI since 2016 and was impressed by their steady progress. In mid-2022, I was so excited about their work that I gave them a challenge: **train an artificial intelligence to pass an Advanced Placement biology exam**. Make it capable of answering questions that it hasn't been specifically trained for. (I picked AP Bio because the test is more than a simple regurgitation of scientific facts—it asks you to think critically about biology.) If you can do that, I said, then you'll have made a true breakthrough....



Q&A and Thank You

- **Questions or clarifications** on any aspect of the presentation.
- **Thank you** for your attention, engagement, and participation.

silvija.seres@technorocks.com

[Nextbook.me](https://nextbook.me) for more than 600 books on AI

[Nextpaper.me](https://nextpaper.me) for 35 applied AI playbooks