

Unit-II

I. AI Applications and Tools Introduction to AI Tools (ChatGPT, Google AI tools, AI in MS Excel & Business Analytics)

Artificial Intelligence (AI) is transforming the way people learn, communicate, manage information, and make business decisions. AI tools are software applications that use techniques such as machine learning, natural-language processing, pattern recognition, and generative AI to perform tasks that traditionally require human intelligence. These tasks include understanding language, creating content, analysing data, forecasting trends, recognising patterns, and automating repetitive work. In education and business, AI should be seen not merely as a replacement for human effort, but as an assistant that improves speed, accuracy, creativity, and decision-making.

▪ AI applications

AI has applications across almost every field. In education, it supports personalised learning, automated feedback, language practice, content development, assessment design, and accessibility tools such as speech-to-text and text-to-speech. A teacher, for example, can use an AI tool to prepare comprehension questions at different difficulty levels, while students can use it to obtain explanations, summaries, vocabulary support, and practice exercises.

In business, AI helps organisations analyse customer behaviour, automate customer service, predict demand, detect fraud, manage inventory, and improve marketing. Chatbots can answer routine customer questions at any time. Recommendation systems suggest products based on customer interests. In finance, AI can identify unusual transactions and support risk analysis. In healthcare, it can assist professionals in examining medical images, organising records, and identifying patterns that may require attention.

AI is also increasingly used in communication and office productivity. It can draft e-mails, summarise long documents, translate text, produce meeting notes, organise schedules, generate presentations, and improve grammar and style. However, its outputs must always be reviewed by a knowledgeable human being, especially where accuracy, privacy, ethics, or high-stakes decisions are involved.

a. ChatGPT

ChatGPT is a conversational AI assistant. Users communicate with it in ordinary language by typing questions or instructions, often called prompts. It can answer questions, explain concepts, draft and rewrite content, summarise material, translate between languages, generate ideas, and assist with logical problem-solving. It can also help with coding, planning, analysing structured information, and creating learning materials.

For students and teachers, ChatGPT can be especially useful in the following ways:

- Explaining difficult concepts in simple language.

- Creating outlines for essays, reports, speeches, and presentations.
- Generating grammar exercises, quizzes, discussion questions, and examples.
- Improving the clarity, tone, grammar, and organisation of a draft.
- Summarising lengthy passages or converting notes into a structured format.
- Supporting language learning through dialogues, vocabulary lists, and writing practice.
- Brainstorming research topics, survey questions, case-study ideas, and project plans.

For example, a commerce student may ask: “Explain the difference between gross profit and net profit with a simple numerical illustration.” ChatGPT can provide a clear explanation and an example. An English student may ask it to identify the theme, setting, characters, and literary devices in a passage. A faculty member may use it to prepare a classroom activity, but should modify the output according to the syllabus, students’ proficiency level, and learning outcomes.

The quality of ChatGPT’s response depends greatly on the quality of the prompt. A vague instruction such as “Write about marketing” may result in a general answer. A specific prompt such as “Write a 300-word introduction on digital marketing for first-year B.Com students, using simple academic language and two examples” is more likely to give a useful result.

ChatGPT has limitations. It can sometimes produce incorrect, incomplete, biased, or invented information. Therefore, it should not be treated as an unquestionable authority. Students must verify facts from reliable textbooks, peer-reviewed journals, official reports, and credible websites. It should be used to support learning and drafting, not for plagiarism or submitting unedited AI-generated work as one’s own.

b. Google AI tools

Google provides AI-powered tools through its Gemini application and integration with Google Workspace. Gemini can assist users in generating, organising, summarising, and analysing information. In Google Workspace, AI features can be available within applications such as Gmail, Google Docs, Google Sheets, Google Slides, Google Drive, and Google Meet.

In Gmail, Gemini can help users’ draft replies, summarise e-mail threads, and retrieve information from messages. In Google Docs, it can help brainstorm ideas, prepare outlines, summarise long documents, improve writing, and suggest revisions in grammar, tone, and formatting. Google’s documentation specifically describes uses such as outlining a sales pitch, generating campaign ideas, summarising lengthy documents, and polishing written work.

Google Sheets is particularly important for students and business users because AI can assist with spreadsheet tasks. Gemini in Sheets can create tables and formulas, generate data insights,

build charts and graphs, apply formatting, create pivot tables, sort data, add filters, and perform optimisation-related tasks. For example, a small business owner can enter monthly sales data and ask the tool to identify the best-selling products, prepare a chart, or highlight months with unusually low revenue.

Google AI tools promote collaboration because several users can work on shared documents, spreadsheets, and presentations. They are useful for group projects, meeting preparation, report writing, and data organisation. Nevertheless, users must understand the organisation's privacy policies before entering confidential information such as examination records, financial data, employee details, or personal customer information.

c. AI in MS Excel

Microsoft Excel has become more powerful with AI-supported features, especially through Microsoft Copilot. Copilot in Excel enables users to interact with spreadsheet data in natural language rather than relying only on complex formulas and manual analysis. It can help generate formulas, explain formulas, summarise data, identify trends and outliers, and create new rows or columns based on existing information.

For instance, instead of manually writing a formula, a user may ask: "Calculate the total sales for each region," "Find the average expenditure by department," or "Explain the formula in cell G12." Copilot can suggest an appropriate formula and explain how it works. It can also assist with calculations across multiple worksheets and generate summaries from data.

Common applications of AI in Excel include:

- Creating formulas such as SUM, AVERAGE, IF, VLOOKUP, XLOOKUP, and percentage calculations.
- Cleaning and organising raw data.
- Detecting duplicate entries, missing values, and unusual figures.
- Classifying data into categories.
- Preparing summary tables and pivot tables.
- Identifying sales trends, correlations, and performance patterns.
- Creating charts and visualisations.
- Explaining existing formulas to beginners.
- Supporting budget preparation, inventory management, and financial analysis.

Excel users should still learn fundamental spreadsheet concepts. AI can suggest a formula, but the user must check whether it refers to the correct cells, uses the correct assumptions, and

produces a sensible answer. Blindly accepting a formula can lead to costly mistakes in accounting, budgeting, payroll, or financial reporting.

d. Business analytics and responsible use

Business analytics is the process of using data to understand business performance and support decision-making. AI strengthens business analytics by helping users process large datasets quickly, identify patterns, forecast possible outcomes, and convert numerical data into understandable insights. In Excel, a manager can ask a question such as, “Which product category showed the fastest sales growth?” Copilot can analyse the workbook and respond using the relevant data; Microsoft notes that such analysis can use natural-language questions and Python-based analysis without requiring the user to write formulas or code.

AI-supported analytics is valuable in sales forecasting, customer segmentation, inventory planning, expense control, financial reporting, human-resource analysis, and marketing performance evaluation. Yet AI does not eliminate the need for human judgment. Managers must interpret results in the context of market conditions, customer needs, ethical standards, and organisational goals.

Responsible AI use requires four habits: verify information, protect confidential data, acknowledge AI assistance where appropriate, and retain human responsibility for final decisions. AI tools are most effective when used thoughtfully—as learning partners, productivity assistants, and analytical aids rather than substitutes for critical thinking, originality, and professional accountability.

II. Basics of Prompt Writing

Prompt writing is the skill of giving clear, specific instructions to an AI tool so that it produces useful and relevant output. A prompt is the text you type into an AI system—such as a question, instruction, or description. The quality of the AI’s response depends heavily on the quality of the prompt. Well-written prompts reduce misunderstandings, save time, and help students, teachers, and professionals get better results.

▪ What is a prompt?

A prompt is any input given to an AI tool to guide its response. It can be as simple as “Explain inflation” or as detailed as “Write a 150-word paragraph on inflation for first-year B.Com students, using simple language and two real-world examples.” The second prompt is more likely to produce a focused, classroom-ready answer because it specifies audience, length, language level, and content requirements.

▪ Basic principles of prompt writing

Effective prompts share several characteristics. They are clear, specific, and context-rich. They state the task, audience, format, and length. They avoid vague words such as “something,” “anything,” or “a lot.” They also avoid contradictory instructions.

A good prompt usually includes:

- Task: What exactly should the AI do? (Explain, summarise, compare, create, rewrite, etc.)
- Audience: Who will read or use the output? (Undergraduates, school students, managers, general readers.)
- Tone and style: Formal, simple, academic, conversational, persuasive, etc.
- Length: Word limit, number of points, or number of paragraphs.
- Content requirements: Specific topics, examples, headings, or structure.
- Output format: Paragraph, bullet points, table, dialogue, script, etc.
- **Examples of basic prompts**

Example 1: Explaining a concept

- Weak prompt: “Tell me about photosynthesis.”
- Improved prompt: “Explain photosynthesis in 120 words for Class 9 students. Use simple language, define the term in the first sentence, and include one example of where it happens in nature.”

The improved prompt gives word limit, audience, language level, and structure.

Example 2: Writing practice

- Weak prompt: “Give me a paragraph on climate change.”
- Improved prompt: “Write a 100-word paragraph on the impact of climate change on Indian agriculture. Use formal academic style and include two effects such as irregular rainfall and crop loss.”

This version guides the AI to produce a focused, syllabus-relevant paragraph.

Example 3: Creating exercises

- Weak prompt: “Make some grammar questions.”
- Improved prompt: “Create 10 multiple-choice questions on subject–verb agreement for first-year undergraduate English students. Provide four options (A–D) for each question and give an answer key at the end.”

Now the AI knows the topic, level, number of items, format, and that an answer key is required.

Example 4: Summarising text

- Weak prompt: “Summarise this.” (with no context)
- Improved prompt: “Summarise the following passage in 80 words. Highlight the main idea and two supporting points. Ignore minor examples.”

This tells the AI how long the summary should be and what to prioritise.

Example 5: Improving writing

- Weak prompt: “Correct this.”
- Improved prompt: “Improve the grammar, clarity, and flow of the following 150-word paragraph. Keep the original meaning and maintain a formal academic tone suitable for a college assignment.”
- **Tips for better prompts**
 - Be specific about what you want.
 - Mention the audience and purpose.
 - State word limits or number of points.
 - Ask for examples when needed.
 - Break complex tasks into smaller steps.
 - Review and refine the output, then ask follow-up prompts such as “Make it simpler,” “Add two more examples,” or “Convert this into bullet points.”

Prompt writing is a learnable skill. With practice, students and teachers can use AI more effectively for studying, teaching, research, and professional communication.

III. AI for Report Writing and Presentations

AI can make report writing and presentation preparation faster, more organised, and more creative. It is most useful as a drafting and support tool: the student or professional must verify facts, add original thinking, cite credible sources, and make the final decisions.

AI for report writing

A report is a formal document that presents information, analysis, findings, and recommendations on a particular topic. Typical reports include academic project reports, business reports, internship reports, survey reports, financial reports, and research reports.

AI tools such as ChatGPT, Gemini, and Microsoft Copilot can assist at different stages of report writing:

- Selecting and narrowing a topic.
- Generating a report outline.
- Explaining concepts in simple language.
- Brainstorming research questions and objectives.
- Drafting an introduction or conclusion.
- Summarising notes, articles, interview responses, or survey data.
- Improving grammar, coherence, tone, and paragraph structure.
- Converting informal writing into formal academic or business language.
- Creating tables, headings, questionnaires, and presentation points.
- Identifying repeated ideas or unclear sentences.

For example, instead of using a vague prompt such as “Write a report on social media,” a student may write:

“Prepare a detailed outline for a 1,000-word undergraduate report on the impact of social media marketing on small businesses in India. Include introduction, objectives, findings, challenges, recommendations, and conclusion.”

The AI can produce a logical structure. The student should then conduct research, add reliable evidence, and write the final report in their own words.

Another useful prompt is:

“Rewrite the following paragraph in formal academic language. Keep the meaning unchanged, reduce repetition, and use clear transitions between ideas.”

This is helpful during editing. However, AI-generated references must be checked carefully. It may provide inaccurate, incomplete, or even non-existent sources. Students should locate and verify every source through library databases, reputable journals, government reports, or official organisational websites.

AI for presentations

A presentation communicates key ideas visually and orally through slides. AI can support the planning, design, content development, and revision of presentations. Google Gemini in Slides can generate new slides, write or rewrite slide content, summarise presentations, use Drive files as reference material, and generate images. Microsoft Copilot in PowerPoint can draft content, create a presentation from a topic or an existing Word file, generate summaries, and answer questions about a presentation.

AI can help users to:

- Create a presentation outline from a topic.
- Divide content into suitable slides.
- Suggest titles, subtitles, and section headings.
- Convert a report into concise slide points.
- Create speaker notes for oral delivery.
- Suggest charts, diagrams, tables, and visuals.
- Rewrite slides in simple, formal, or persuasive language.
- Summarise lengthy content into key takeaways.
- Check whether slides contain too much text.
- Generate likely audience questions for practice.

For example, a student preparing a presentation on “Artificial Intelligence in Education” may use this prompt:

“Create an eight-slide presentation outline on Artificial Intelligence in Education for first-year undergraduate students. Use simple academic language. Include definition, applications, advantages, challenges, ethical issues, one classroom example, conclusion, and three key takeaways.”

The AI may create an appropriate sequence:

1. Title slide
2. Meaning of AI in education
3. Major applications
4. Benefits for students and teachers
5. Challenges and limitations
6. Ethical and privacy concerns
7. Classroom example
8. Conclusion and key takeaways

Google's Gemini can create a fully editable Google Slides presentation from a single prompt; users can review the generated plan before approving it and then edit the resulting slides. Similarly, Copilot can begin with an outline or ask clarifying questions before generating a PowerPoint presentation for review.

Effective use process

A responsible workflow has five stages:

1. **Plan the task.** Decide the topic, audience, purpose, word limit, or number of slides.
2. **Give a clear prompt.** Mention the desired tone, level, format, and specific sections.
3. **Review the output.** Check accuracy, relevance, language, and completeness.
4. **Add original contribution.** Include your own analysis, examples, data, observations, and conclusions.
5. **Edit and practise.** Correct errors, cite sources, reduce unnecessary text, and rehearse the presentation.

Important precautions

AI should support not replace critical thinking and academic integrity. Do not submit unedited AI-generated reports as original work. Do not upload confidential institutional, student, financial, or personal data into public AI tools without permission. Avoid copying AI content word for word; revise it, check it against trustworthy sources, and acknowledge AI use if your institution requires disclosure.

A strong report depends on evidence and analysis, while a strong presentation depends on clarity, visual balance, and confident delivery. AI can help create the first draft, but human judgment makes the final work accurate, ethical, original, and effective.

IV. AI Ethics and Responsible AI

AI ethics refers to the moral principles that guide the design, development, use, and governance of artificial intelligence. Responsible AI means creating and using AI systems in ways that are safe, fair, transparent, privacy-conscious, accountable, and beneficial to people and society. The aim is not to stop innovation, but to ensure that AI supports human well-being, rights, dignity, and sustainable development.

Why AI ethics matters

AI systems are increasingly used in education, banking, healthcare, recruitment, policing, transport, customer service, and social media. Their decisions or recommendations can influence who receives a loan, which job applicant is shortlisted, what information a person sees online, or how a student's work is assessed.

If AI is designed or used carelessly, it can cause real harm. For example, a recruitment system trained mainly on data from past male employees may unfairly rate women applicants lower. A facial-recognition system may work less accurately for certain groups if its training data is not diverse. A chatbot may provide incorrect medical, legal, or financial advice. Therefore, AI systems must be tested, monitored, and used with human judgment.

UNESCO's global guidance on AI ethics places human rights and human dignity at the centre of AI governance. It highlights principles such as fairness, transparency, privacy, safety, environmental sustainability, and meaningful human oversight.

Key principles

The following principles are central to responsible AI:

- **Fairness and non-discrimination:** AI should not unfairly disadvantage people because of gender, caste, race, religion, disability, age, language, income, or location. Developers must use diverse and representative data and test systems for bias.
- **Transparency:** Users should know when they are interacting with AI and understand, as far as possible, how an AI system reaches its outputs. AI-generated content should not be presented dishonestly as human-created work.
- **Accountability:** A human person or organisation must remain responsible for the effects of an AI system. If an AI tool makes an error or causes harm, responsibility cannot simply be shifted to "the machine."
- **Privacy and data protection:** Personal data such as names, marks, medical records, financial details, addresses, photographs, and contact information must be collected and used carefully. UNESCO states that privacy should be respected and protected throughout the entire AI life cycle.
- **Safety and security:** AI systems should be designed to avoid physical, psychological, social, and financial harm. They must also be protected from hacking, manipulation, misuse, and malicious prompts.
- **Human oversight:** AI should assist humans rather than make uncontrolled high-stakes decisions. In healthcare, education, recruitment, law, and finance, qualified people should review important AI recommendations before action is taken.
- **Accuracy and reliability:** AI outputs must be checked because generative AI can produce false information, biased statements, outdated facts, or fabricated references. The NIST AI Risk Management Framework identifies trustworthy AI as valid and reliable, safe, secure and resilient, accountable and transparent, explainable, privacy-enhanced, and fair with harmful bias managed.

Responsible AI in education

In educational institutions, responsible AI use is especially important. Teachers can use AI to create lesson plans, quizzes, examples, summaries, and differentiated learning materials. Students can use it for brainstorming, grammar support, explanations, practice questions, and feedback on drafts.

However, students should not submit AI-generated assignments as their own original work. This is academically dishonest and prevents genuine learning. A responsible student may use AI to obtain an outline for a report, but should independently research the topic, verify facts, write the report in their own words, cite sources correctly, and acknowledge AI assistance if institutional policy requires it.

For example, a student may ask AI to explain “sustainable business practices.” The student should then compare that explanation with textbooks and credible sources, add examples from Indian companies or local businesses, and form an independent conclusion.

Practical guidelines

Responsible AI use can be practised through simple habits:

1. Verify important facts from reliable sources before using them.
2. Do not upload confidential personal, institutional, financial, or examination-related data into public AI tools.
3. Check AI output for bias, stereotypes, offensive language, and unfair assumptions.
4. Keep humans responsible for final decisions, especially in high-risk situations.
5. Cite original sources and follow academic-integrity rules.
6. Disclose significant AI assistance when required by a teacher, employer, or institution.
7. Use AI to enhance creativity and learning, not to replace critical thinking.

AI ethics is ultimately about ensuring that technological progress remains human-centred. Responsible AI builds trust because it combines innovation with fairness, safety, accountability, and respect for human values.

V. Data Privacy and Security Issues

AI tools such as ChatGPT, Gemini, Microsoft Copilot, image generators, AI chatbots, and data-analysis applications can improve learning, writing, communication, and business decision-making. However, they also create important **data privacy and security** concerns. Users often enter text, documents, images, spreadsheets, customer details, student records, or business information into AI tools without fully understanding where the data goes, how long it is stored, who can access it, or whether it may be used to improve the system.

Data privacy refers to the right of individuals to control how their personal information is collected, used, stored, shared, and deleted. Data security refers to the measures used to protect information from unauthorised access, theft, alteration, loss, or misuse. In AI use, both privacy and security must be considered together.

Major privacy concerns

One major concern is the exposure of personal data. Users may unintentionally paste confidential information into a public AI tool. This may include names, mobile numbers, addresses, passwords, Aadhaar or PAN details, bank-account information, medical records, student marks, attendance data, examination papers, research data, or employee information. Such data should never be entered into an AI tool unless the organisation has approved it and the tool provides appropriate privacy safeguards.

For example, a teacher should not upload a spreadsheet containing students' names, marks, contact numbers, and performance details into a public chatbot and ask it to identify weak students. The data should first be anonymised by removing names and other identifying details. Labels such as "Student A," "Student B," and "Student C" can be used instead.

AI can also create **re-identification risks**. Even when direct identifiers are removed, AI may combine different pieces of information and infer the identity of a person. NIST notes that AI creates new re-identification risks because of its ability to analyse information across multiple datasets; it can also increase behavioural tracking and surveillance.

Another issue is unclear data retention. Some AI platforms may store prompts, uploaded files, conversations, or usage records for a certain period. Users may not know whether their data is retained, shared with service providers, or used for system improvement. Before using an AI tool, users should read its privacy policy, data-retention policy, and institutional guidelines.

Major security risks

AI tools can be targeted by cybercriminals. Attackers may try to steal user data, manipulate AI outputs, gain access to organisational accounts, or use AI to create convincing phishing e-mails and fake messages. AI-generated text can make fraudulent messages appear professional and personalised, increasing the risk of financial fraud and identity theft.

Weak passwords and shared accounts are another common security problem. If an AI account is linked to e-mail, cloud storage, documents, or organisational files, unauthorised access can expose a large amount of sensitive information. Users should create strong, unique passwords and enable multi-factor authentication wherever possible.

AI systems can also produce inaccurate or manipulated outputs. A malicious person may provide misleading data or craft instructions that cause an AI tool to reveal sensitive content, ignore safeguards, or generate harmful responses. Therefore, AI output should be reviewed carefully before it is used in reports, decisions, official communications, or public documents.

Safe practices for users

Responsible use of AI requires practical precautions:

1. Do not enter confidential, personal, financial, medical, institutional, or examination-related data into public AI tools.
2. Remove names, identification numbers, contact details, and other personal identifiers before uploading documents or datasets.

3. Use only institutionally approved AI tools for official work involving student, employee, customer, or organisational data.
4. Read the tool's privacy policy and understand whether conversations and files are stored or used for training.
5. Use strong passwords and enable multi-factor authentication for AI-linked accounts.
6. Verify AI-generated facts, references, calculations, and recommendations before using them.
7. Restrict access to sensitive files and share them only with authorised persons.
8. Report suspected data leaks, suspicious AI-generated messages, or unauthorised account activity immediately.

Responsible institutional use

Educational institutions and businesses should create clear AI-use policies. These policies should explain what data may be used with AI tools, which tools are approved, who can access sensitive information, and how users should report security incidents. UNESCO emphasises that privacy and data protection must be safeguarded throughout the entire life cycle of an AI system, supported by proper data-governance frameworks.

In conclusion, AI tools are useful only when used carefully. Protecting privacy, securing accounts, anonymising data, verifying outputs, and maintaining human oversight can help individuals and organisations gain the benefits of AI without placing people's information, rights, or safety at risk.

VI. Bias and Fairness in AI

Bias in AI occurs when an artificial-intelligence system produces unfair, inaccurate, or discriminatory results for certain individuals or groups. Fairness in AI means designing and using AI systems so that they treat people equitably, avoid harmful discrimination, and provide benefits without excluding or disadvantaging particular communities. Responsible AI must be fair, transparent, accountable, and subject to human oversight.

AI systems learn patterns from data. If the data contains social inequalities, stereotypes, historical discrimination, or incomplete representation, the AI may learn and reproduce those patterns. For this reason, AI is not automatically neutral simply because it uses technology or mathematics.

Sources of bias in AI

Bias can enter an AI system at many stages. The first source is biased training data. AI systems are trained using large amounts of past data. If this data mainly represents one group of people, language, region, gender, or income category, the system may perform less accurately for others.

For example, an AI voice-recognition tool trained mainly on American or British English may struggle to understand Indian English accents. Similarly, an AI image-recognition system trained mostly on photographs of lighter-skinned people may identify darker-skinned people less accurately.

A second source is historical bias. Historical data may reflect unfair practices that already existed in society. If a company hired mostly men for technical roles in the past, an AI recruitment tool trained on those previous hiring decisions may incorrectly learn that male applicants are more suitable. As a result, qualified women candidates may receive lower rankings.

A third source is human bias. Developers, data collectors, managers, and users make choices about which data to collect, what questions to ask, which categories to use, and how to interpret results. Their conscious or unconscious assumptions can influence the AI system. NIST identifies three broad categories of bias that need to be managed: systemic bias, computational and statistical bias, and human-cognitive bias.

Examples of AI bias

Consider an AI-based loan approval system. If it is trained using past loan data from a bank that historically gave fewer loans to people from certain low-income areas, the AI

may treat applicants from those areas as high-risk, even if they are financially eligible. This can deepen existing inequalities.

Another example is an online recruitment system. The tool may analyse resumes and rank candidates according to previous hiring patterns. If those patterns favour graduates from certain institutions, cities, or social backgrounds, talented candidates from rural areas or lesser-known colleges may be unfairly excluded.

In education, an AI writing evaluator may score students unfairly if it has been trained mostly on one variety of English. It may wrongly treat expressions common in Indian English or second-language writing as poor-quality language. Teachers should therefore review AI-generated scores and feedback rather than relying on them completely.

Bias can also appear in generative AI. When asked to generate images of a “doctor,” “engineer,” or “CEO,” an AI system may repeatedly show men. When asked to generate images of nurses, teachers, or domestic workers, it may show women. Such output can reinforce gender stereotypes.

Ensuring fairness

Fairness is not achieved merely by removing sensitive information such as gender, caste, religion, or race from a dataset. AI may still use related information, such as location, language, income level, school attended, or type of employment, to infer those characteristics indirectly. Therefore, AI systems need continuous testing and monitoring.

The following practices can improve fairness:

- ✓ Use diverse, representative, and high-quality data.
- ✓ Check datasets for missing groups, stereotypes, and historical inequalities.
- ✓ Test AI performance across different genders, languages, regions, age groups, and social backgrounds.
- ✓ Involve people from diverse disciplines and communities in AI design.
- ✓ Explain how important AI decisions are made.
- ✓ Provide ways for people to challenge, correct, or appeal an AI-based decision.
- ✓ Retain human review for high-stakes areas such as healthcare, recruitment, admissions, loans, and law enforcement.
- ✓ Monitor systems regularly after they are deployed because bias can emerge or change over time.

The OECD AI Principles state that AI should respect human rights and democratic values, including non-discrimination, equality, diversity, fairness, social justice, and privacy. They also call for ongoing risk management to address AI-related harms, including harmful bias.

Human responsibility

AI can support decision-making, but it should not become the final authority in matters that affect people's opportunities, rights, and dignity. If an AI system rejects a job candidate, denies a loan, grades a student, or identifies a person, a qualified human should be able to review the decision and correct errors.

Fair AI is therefore a shared responsibility. Developers must build and test systems carefully; organisations must establish ethical policies; teachers and students must use AI critically; and users must question outputs that seem unfair. According to NIST, trustworthy AI should be fair with harmful bias managed, alongside being reliable, safe, secure, transparent, explainable, privacy-enhanced, and accountable.

VII. Legal Issues of AI in Business

AI creates important legal issues for businesses because it affects personal data, intellectual property, contracts, employment, consumer protection, and liability. A business using AI must ensure that its tools comply with applicable laws, protect stakeholders' rights, and do not cause unlawful harm. Since AI laws differ across countries, organisations should seek legal advice before deploying AI in high-risk or large-scale business activities.

Data protection and privacy

Businesses frequently use AI tools to analyse customer data, employee records, sales figures, financial information, location data, and online behaviour. This creates legal obligations

relating to privacy and data protection. A company must have a lawful reason for collecting and processing personal data. It should inform people about what data is collected, why it is used, how long it will be stored, and whether it will be shared with third parties.

For example, an online shopping company may use AI to recommend products to customers. If the company collects purchase history, browsing behaviour, addresses, or payment-related information, it must protect this data from unauthorised access and misuse. It should not share personal information with an AI provider without appropriate consent, safeguards, and contractual arrangements.

Businesses should avoid entering confidential data into public AI tools. Customer lists, employee salaries, examination data, trade secrets, product designs, passwords, bank details, and unpublished reports should not be uploaded unless the organisation has approved the tool and ensured adequate security controls.

Intellectual property rights

AI raises complex questions about ownership and copyright. Businesses may use AI to generate text, logos, software code, images, advertisements, music, designs, or reports. However, AI-generated content may resemble existing copyrighted material, trademarks, or designs. A company that uses such content without checking it could face legal claims.

For example, if an advertising agency asks an AI image generator to create a logo for a client, the output may unintentionally look like another company's trademark. Before using the logo publicly, the agency should check for trademark conflicts and ensure that the final design is original.

Copyright concerns also arise when businesses use AI tools trained on external material. Organisations should understand the licensing conditions of the tool, the ownership rights over generated outputs, and whether the material can be used commercially. The European Union's AI Act includes obligations for general-purpose AI providers concerning transparency and respect for copyright.

Discrimination and employment law

Businesses may use AI in recruitment, performance evaluation, promotion, customer profiling, credit assessment, and pricing. These uses can create legal problems if the AI system discriminates unfairly based on gender, caste, religion, race, age, disability, language, nationality, or other protected characteristics.

For instance, an AI recruitment tool may rank applicants based on past hiring data. If the earlier data reflects gender or social bias, the tool may unfairly reject capable candidates. The business can still be held responsible even if the discrimination was produced by an external AI vendor.

The EU AI Act treats several employment-related AI applications as high-risk. For high-risk systems, organisations must meet requirements relating to risk management, data quality, documentation, transparency, human oversight, accuracy, cybersecurity, and robustness.

Businesses using such systems should ensure that humans review important decisions and that applicants or employees have a way to question or appeal unfair outcomes.

Liability and consumer protection

If AI gives incorrect advice, produces harmful content, makes a defective recommendation, or causes financial loss, legal responsibility may arise. Liability may fall on the AI developer, software vendor, business using the tool, employee operating it, or more than one party. Contracts with AI vendors should clearly state responsibilities for data protection, security, intellectual property, errors, service interruptions, and compensation for losses.

Consumer-protection law is also relevant. Businesses should not use AI to deceive customers through fake reviews, misleading advertisements, deepfakes, hidden chatbots, or manipulative pricing. Customers should be informed when they are interacting with an AI system, especially if the interaction affects their decisions or rights. From 2 August 2026, the EU AI Act's transparency rules require certain AI systems to tell users when they are interacting with AI and when content has been generated or altered by AI.

Contracts and governance

AI use should be covered by clear internal policies and external contracts. Businesses should define which AI tools employees may use, what information can be uploaded, who approves

high-risk AI uses, and how outputs are reviewed. Vendor agreements should address ownership of data and AI output, confidentiality, cybersecurity, audit rights, legal compliance, and incident reporting.

The OECD states that organisations and individuals who develop, deploy, or operate AI systems should be accountable for their proper functioning and for respecting relevant values and rights. These risks include harmful bias, safety, privacy, security, labour rights, and intellectual-property rights.

In conclusion, AI can improve efficiency and innovation in business, but it cannot remove legal responsibility. Businesses must use AI with careful governance, human oversight, legal compliance, transparent communication, and respect for customer and employee rights.

VIII. Future Careers in AI

Artificial Intelligence (AI) is reshaping the world of work by automating routine tasks, enhancing human capabilities, and creating entirely new job categories. The World Economic Forum's *Future of Jobs Report 2025* estimates that by 2030, about 170 million new roles will be created while 92 million roles may be displaced, resulting in a net gain of 78 million jobs. AI, big data, cybersecurity, and digital skills are among the fastest-growing areas. Future careers in AI will not only include highly technical roles but also interdisciplinary positions that combine AI knowledge with domain expertise in fields such as business, healthcare, education, law, finance, agriculture, and creative industries.

Technical and specialised AI careers

Technical careers in AI focus on designing, developing, deploying, and maintaining AI systems. Common roles include:

- **AI and machine-learning engineers:** They build algorithms, train models, and integrate AI into software products and business systems.
- **Data scientists and data analysts:** They collect, clean, analyse, and interpret large datasets to support decision-making and AI training.
- **AI research scientists:** They work on advanced algorithms, new architectures, and innovations in areas such as natural-language processing, computer vision, and robotics.
- **MLOps and AI infrastructure engineers:** They manage the deployment, monitoring, scaling, and security of AI models in production environments.
- **AI security and privacy specialists:** They ensure that AI systems are secure, resistant to attacks, and compliant with data-protection and ethical standards.

These roles usually require strong foundations in mathematics, statistics, programming, and computer science. However, employers increasingly value problem-solving skills, domain knowledge, and the ability to work in teams.

AI-augmented and hybrid careers

Many future careers will not be purely “AI jobs” but roles where AI significantly enhances human work. In these hybrid careers, professionals use AI tools to improve productivity, accuracy, and creativity. Examples include:

- **AI-assisted doctors and healthcare professionals:** They use AI for diagnosis support, treatment planning, medical imaging, and patient monitoring while retaining final clinical judgment.
- **AI-enabled teachers and instructional designers:** They use AI to create learning materials, generate quizzes, personalise instruction, and provide feedback, while focusing on mentoring, motivation, and critical thinking.
- **AI-supported lawyers and legal analysts:** They use AI to search case law, draft documents, summarise contracts, and identify legal risks, while providing strategic advice and courtroom representation.
- **AI-driven marketers and content creators:** They use AI for audience analysis, content generation, campaign optimisation, and performance tracking, while shaping brand strategy and creative direction.
- **AI-enhanced accountants and financial analysts:** They use AI for data entry, reconciliation, forecasting, fraud detection, and reporting, while focusing on interpretation, advisory services, and compliance.

The World Economic Forum notes that AI and big data skills are the fastest-growing skills, but human skills such as creative thinking, resilience, flexibility, and collaboration remain critical.

Emerging and interdisciplinary roles

As AI matures, new specialised roles are emerging:

- **AI ethicists and governance specialists:** They design ethical guidelines, assess risks, ensure compliance with laws such as the EU AI Act, and promote fairness, transparency, and accountability.

- **AI policy analysts and regulators:** They develop national and organisational policies on AI use, data protection, intellectual property, and labour rights.
- **AI trainers and prompt engineers:** They design high-quality prompts, curate training data, refine AI outputs, and teach others to use AI effectively.
- **Human–AI interaction designers:** They create user interfaces and workflows that enable smooth collaboration between humans and AI systems.
- **AI sustainability specialists:** They assess the environmental impact of AI, promote energy-efficient models, and align AI projects with Sustainable Development Goals.

Skills for future AI careers

To prepare for AI-related careers, students and professionals should develop both technical and human skills. Key technical skills include AI and big data, networks and cybersecurity, and overall technological literacy. At the same time, analytical thinking, creative thinking, resilience, and adaptability are increasingly important.

Practical steps include learning basic programming and data analysis, using AI tools responsibly, understanding AI ethics and legal issues, and gaining domain expertise in a chosen field. Continuous learning, upskilling, and reskilling will be essential because employers expect about 39% of workers' core skills to change by 2030.

Future careers in AI will reward those who can combine technical understanding with critical thinking, ethical awareness, communication skills, and the ability to work effectively alongside intelligent machines.