In brief about the project: **B.R.A.I.N. – EmpowerAI: Artificial Intelligence and Citizen Science for Inclusion and Participation** is a project of the Ion Creangă State Pedagogical University (UPSC) dedicated to supporting Ukrainian refugees in Chișinău through accessible digital education and civic engagement, project funded under the "Sergio Vieira de Mello Academic Chair (SVMAC)" initiative, with support from UNHCR Moldova, and administered by the Law Centre of Advocates (CDA).

Guide in Romanian

Guide in Ukrainian

Vitalie MAMOT Tudor LAPP

Introduction to Artificial Intelligence

GUIDE

RO | EN | UKR



Chișinău, 2025

UNHCR The UN Refugee Agency

CSVM CATEDRA SÉRGIO VIEIRA DE MELLO

CDA Centrul de Drept al Avocaților

UNIVERSITATEA PEDAGOGICĂ DE STAT ION CREANGĂ DIN CHIȘINĂU

Vitalie MAMOT **Tudor LAPP**

Introduction to Artificial Intelligence

Guide

Chişinău, 2025

Autori: **Vitalie Mamot**, Associate Professor, PhD
Tudor Lapp, AI coach, Master's in Educational Sciences

Reviewer: Andrei Braicov, Professor, PhD

Translated and adapted into English by Tatiana Yavuz, PhD

The "Introduction to Artificial Intelligence" guide briefly explains what AI is, where it works well, and what its limits are, so it can be integrated responsibly into education and research. It presents definitions, types (narrow, general, generative), a history from symbolic systems to deep learning, the role of data quality and metrics, as well as the gap between laboratory performance and real-world performance. It offers practical applications for teaching, feedback, and analysis, rules for prompting, ethical principles (transparency, privacy, human-in-the-loop), and an inventory of accessible tools. It includes a glossary and bibliography, aiming to make AI a credible ally.

This material contains informational sections augmented with the help of generative artificial intelligence tools.

DESCRIEREA CIP A CAMEREI NATIONALE A CĂRȚII DIN REPUBLICA MOLDOVA

Mamot, Vitalie.

Introduction to Artificial Intelligence : Guide / Vitalie Mamot, Tudor Lapp ; translated and adapted into english by Tatiana Yavuz ; Universitatea Pedagogică de Stat "Ion Creangă" din Chişinău, CDA : Centrul de Drept al Avocaților. – Chişinău : [S. n.], 2025 (CEP UPSC). – 28 p. : il. color, tab.

Cerințe de sistem: PDF.

Referințe bibliogr.: p. 30 (18 tit.). – Cu sprijinul UNHCR Moldova.
ISBN 978-9975-48-283-7 (PDF).
004.8:37.01
M 19

©Vitalie Mamot, 2025

Editorial and Printing Center of the Ion Creangă State Pedagogical University, Chișinău, 1 Ion Creangă Street, MD-2069

CUPRINS

	pag.
B.R.A.I.N. – EmpowerAI: Artificial Intelligence and Citizen Science for Inclusion and Participation	6
What is Artificial Intelligence?	7
Why this guide?	8
Artificial Intelligence: promise, limitations, and responsibility	9
Anticipating frequently asked questions	10
A brief history of AI: from symbolic to deep learning	11
Types of AI: Narrow Artificial Intelligence, General Artificial Intelligence, Generative Artificial Intelligence	12
Artificial Intelligence in practice (education and research)	13
Basic prompting & prompt engineering	14
Ethics, safety, and regulation	15
Artificial intelligence tools	16
Runway — AI-powered video generation & editing	17
InVideo — fast template-based editing	18
Synthesia — virtual presenters (avatars)	19
Lumen5 — text to video (for social)	20
HailuoAI — AI-assisted content creation (text/images/video)	21
Midjourney — generative art on Discord	22
DALL·E 3 — friendly image generation, variations, and edits	23
Stable Diffusion — control maxim, local sau în cloud	24
Adobe Firefly — generative AI integrated into the Adobe ecosystem	25
Leonardo AI — templates, models, and assets for design	26
Suno — voice-based song generation (text-to-music & text-to-song)	27
Stable Audio — flexible audio generation (sounds, loops, music)	28
Glossary of Artificial Intelligence Terms	29
Selected Bibliography	30

B.R.A.I.N. – EmpowerAI: Artificial Intelligence and Citizen Science for Inclusion and Participation

"B.R.A.I.N. – EmpowerAI: Artificial Intelligence and Citizen Science for Inclusion and Participation" is a project funded under the "Sergio Vieira de Mello Academic Chair (SVMAC)" initiative, with support from UNHCR Moldova, and administered by the Law Centre of Advocates (CDA).

In brief about the project: B.R.A.I.N. – EmpowerAI: Artificial Intelligence and Citizen Science for Inclusion and Participation is a project of the Ion Creangă State Pedagogical University of Chișinău (UPSC) dedicated to supporting Ukrainian refugees in Chișinău through accessible digital education and civic engagement. The initiative combines introductory workshops on artificial intelligence with Citizen Science activities to build skills, confidence, and a sense of belonging to the host community.

The curriculum includes 6 beginner-oriented AI workshops: fundamentals and everyday applications; an introduction to machine learning; hands-on exercises with visual tools; data literacy; and ethics for responsible use. The Citizen Science module prepares participants to observe and document Chișinău's urban space (indicators of accessibility, pedestrian comfort, cultural landmarks), to collect data ethically, and to create digital maps and posters suitable for public presentation.



For emotional well-being and family involvement, the project offers 10 psychosocial support sessions for adults and youth, as well as 2 career-guidance sessions for parents focused on educational opportunities and socio-economic integration.

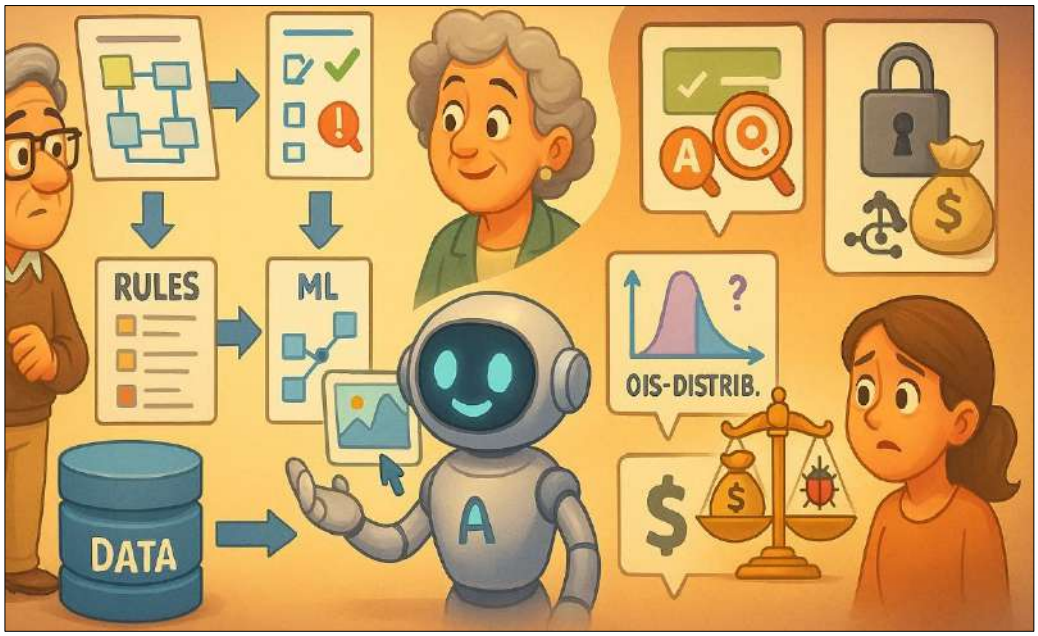
The analytical component includes an exploratory study on the dynamics of Ukrainian refugees with temporary protection/refugee status, with an emphasis on access to services, education, psychosocial support, and urban mobility.

What is Artificial Intelligence?

Definitions. Artificial Intelligence is the field that creates systems capable of performing tasks that typically require human intelligence: perception, language understanding, approximate reasoning, planning, and content generation. Practically, such a system ingests data, frames the problem through rules or by learning from examples, and optimizes a target function within a clearly defined context. A distinction is made between narrow artificial intelligence, specialized for a single task, and general artificial intelligence, a theoretical ideal of broad flexibility. Machine learning is the subfield that extracts patterns from data, while deep learning uses multi-layer neural networks.

Goals. The set of task spans prediction, classification, object or anomaly detection, planning and decision-making under constraints, as well as the generation of text, images, audio, or code. Success is measured by metrics aligned with the real cost of error.

Limits. Performance depends on data quality and representativeness; generalization degrades outside the known distribution; correlation does not imply causation; models can be opaque, vulnerable, and costly; and generative systems may produce content that is plausible yet incorrect.





Why this guide?

Artificial Intelligence (AI) has moved from the realm of spectacular experiments into the daily infrastructure of life: we search for information, translate texts, classify images, and explore datasets—all with tools that embed AI. Precisely because it is already "invisible" in many applications, AI is at once overhyped and misunderstood. This guide aims to provide a clear, practical, and responsible framework for anyone who wants to understand what AI is, what it is genuinely useful for, what its limits are, and how it can be judiciously integrated into professional practice.

Why now? Because the pace of new developments (models, services, integrations) can generate both confusion and hasty decisions. The guide offers a set of stable reference points: how to correctly formulate a problem, how to choose the appropriate metric, what "good" data means and the risks involved, how to document limitations, and how to put ethical guardrails in place when implementing intelligent solutions.

Our ambition is simple: to turn AI from a "buzzword" into a reliable ally. If, after reading, you can formulate problems more effectively, assess solutions clear-eyedly, and integrate AI responsibly into your daily work, then the guide has achieved its purpose. In a world that is changing rapidly, the same guiding markers remain essential: clarity of the problem, quality of the data, rigor in evaluation, respect for people, and transparency of the process.

Artificial Intelligence: promise, limitations, and responsibility

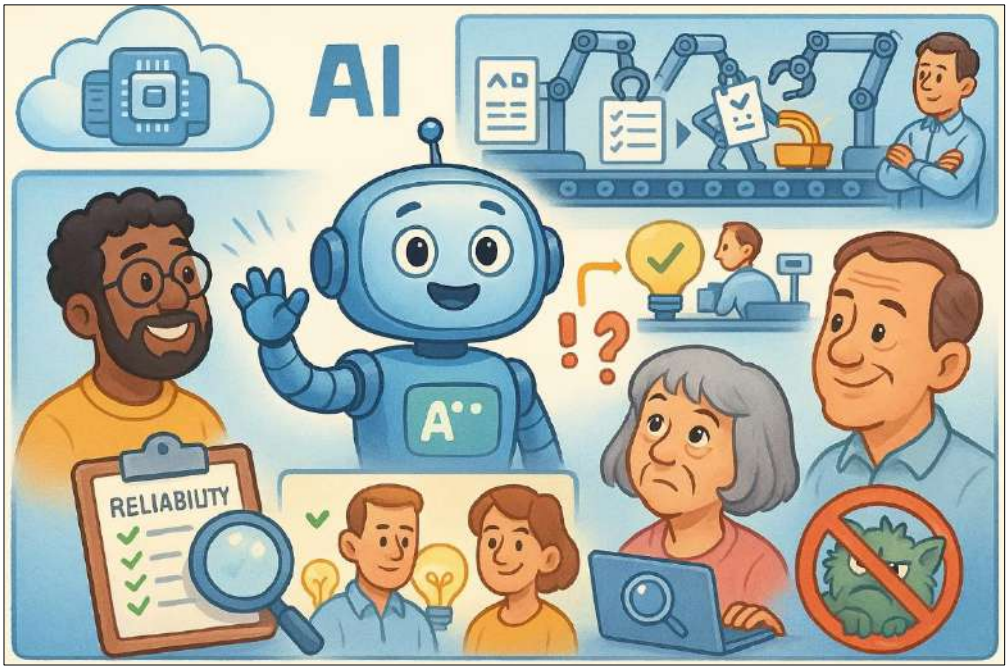
Artificial Intelligence excels at modeling patterns in data. This means its performance depends on the quality, representativeness, and quantity of the training data, as well as on how well the architecture and metrics fit the real problem. In this guide, we emphasize the difference between **performance in the laboratory** (on controlled test sets) and **performance in the real world**.



Another central aspect is the **clarity and applicability of the explanation**. Not every model needs to be fully "explainable," but the process must be transparent: what data were used, how they were processed, which model versions and which parameters were employed.

Data protection: in many scenarios, security, and informed consent are not merely advisable but mandatory.

There is also **pedagogical and social responsibility**. In education, access to AI should not replace learning but enhance it. If a student uses an assistant to write an essay, real competence requires understanding the content, evaluating it critically, being able to defend the ideas, and acknowledging the sources.



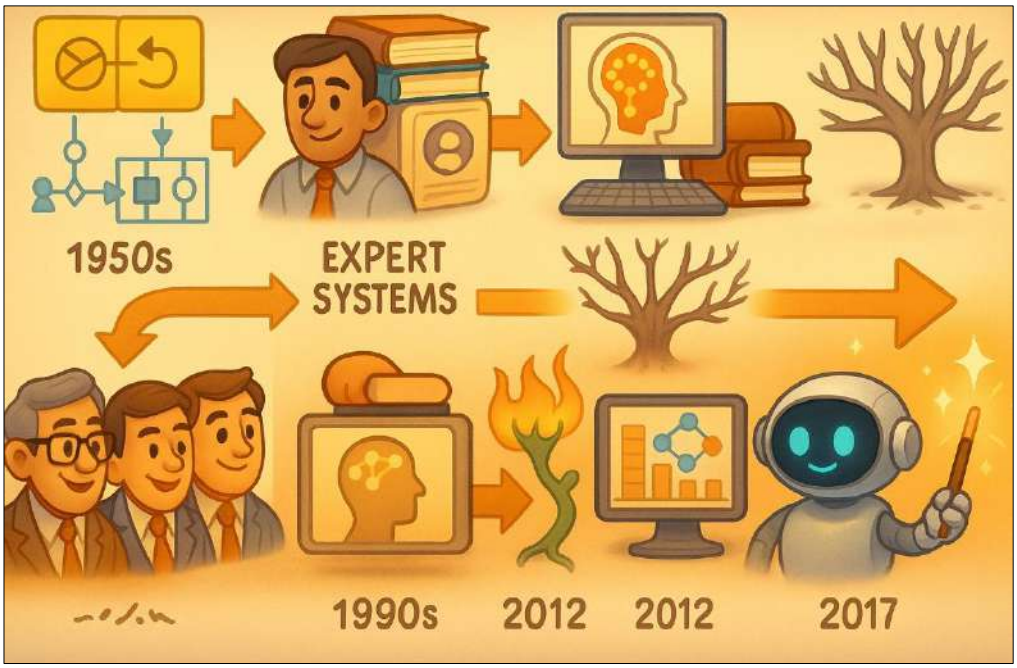
Anticipating frequently asked questions

"Is AI just a fad?" No. Although waves of enthusiasm come and go, the broad adoption of AI tools in digital infrastructure—from optimized processors to standard library stacks and cloud services—shows we're dealing with a structural transformation.

"Will AI take my job?" AI will automate parts of tasks, especially those that are repetitive and well-structured. At the same time, it will create demand for skills in orchestration, verification, interpretation, communication, and process design. Jobs don't "disappear" monolithically; they are reconfigured.

"Can I trust AI results?" Trust is not a given; it's the outcome of a process: correctly defining the problem, using quality data, conducting rigorous evaluations, monitoring over time, and having clear intervention protocols when drift occurs.

"What do I do about errors and hallucinations in generative models?" Anticipate and manage them: validate with independent sources, craft prompts that request citations/verifications, set usage boundaries (e.g., "assistant for ideas, not the final verdict"), and include human review for sensitive outcomes.



A brief history of AI: from symbolic to deep learning

Modern roots trace back to the 1950s, when researchers proposed that reasoning could be represented through rules and symbols. In 1956, at the Dartmouth conference, the "artificial intelligence" research program was launched. The first decades relied on symbolic systems: logic, search in state spaces, planning, and later expert systems that encoded specialists' knowledge. After initial waves of optimism, "AI winters" followed when promises failed to materialize on the available hardware. In parallel came the perceptron and neural networks, reinvigorated in the 1980s by backpropagation. The 1990s brought a shift to statistical methods and learning from data, against a backdrop of growing compute capacity. The turning point arrived in 2012, when deep learning decisively surpassed classical methods in image recognition, and then in speech and language.

Since 2017, attention-based architectures have propelled generative models, ushering in an era of systems capable of producing text, images, and audio at scale.



Types of AI: Narrow Artificial Intelligence, General Artificial Intelligence, Generative Artificial Intelligence

Narrow artificial intelligence refers to systems specialized in a single well-defined task: machine translation, image classification, anomaly detection, content recommendation. They achieve high performance in their domain but do not easily transfer knowledge to other problems.

General artificial intelligence is a theoretical target: an intelligence capable of learning, reasoning, and acting effectively across a wide range of situations, approaching human cognitive flexibility. There is no consensus that it has been achieved; discussions are mainly conceptual and ethical.

Generative artificial intelligence encompasses models that can produce new content—text, images, audio, code—based on instructions or examples. Useful for quick ideas, prototyping, and creative support, it requires human verification because it can generate information that is plausible yet incorrect.

Artificial Intelligence in practice (education and research)

Artificial Intelligence can support teaching, learning, and research work when used with a clear purpose and human verification.

In education:

- lesson design: generating ideas, examples tailored to the class level, variants of exercises, and assessment rubrics;
- rapid feedback: formative comments on essays, pointing out common mistakes, suggestions for improvement;
- accessible materials: summaries at different difficulty levels, simplified texts, translations, visual aids;
- balanced assessment: generating varied items for the same competency, analyzing results to identify gaps;
- personalized learning: recommending resources and pathways based on progress.

In research:

- literature review: semantic search, article summarization, extraction of tables and concepts;
- data analysis: cleaning, initial classification, visualizations, detection of patterns or anomalies;
- assisted writing: article skeletons, section lists, checks for clarity and coherence;
- prototyping: basic code for experiments, small-scale hypothesis testing;
- project management: timelines, task lists, templates for reports.

Best practices: transparency about tool usage, data protection, source citation, critical verification of results, and clear rules regarding the originality of work.

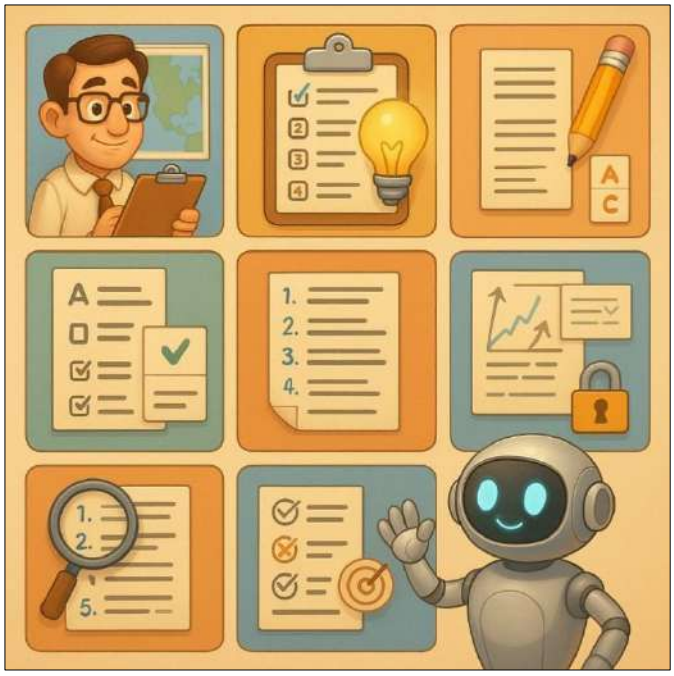


Basic prompting & prompt engineering

A "prompt" is the request you send to the model. A good prompt is like a clear brief: it states **who you are, what you want, for whom**, and **in what form**.

Simple rules:

1. **Context:** briefly describe the situation ("Geography teacher preparing a worksheet for 8th grade").
2. **Task:** state the action precisely ("generate 5 single-answer multiple-choice questions").
3. **Constraints:** length, tone, level, format ("max 120 words, friendly tone, numbered list").
4. **Examples:** show a sample output ("Example: Question… Correct answer…").
5. **Delimitation:** use clear markers for data ("Source text: '…'").
6. **Steps:** ask for the solution in stages ("first an outline, then the final version").
7. **Criteria:** say how you will evaluate ("no fabricated data; include sources when available").
8. **Iterate:** refine the prompt after the result; add missing details, reduce ambiguity.
9. **Verify:** confirm important facts from sources; don't use raw outputs for sensitive decisions.
10. **Good practices:** avoid confidential data; ask for alternatives ("one simple version and one advanced"); specify the target audience and purpose. **Clear prompt = more relevant results.**



Ethics, safety, and regulation

The ethics of using artificial intelligence means designing and deploying systems that respect the dignity, rights, and interests of the people they affect. We start with a clear, legitimate purpose: we define beneficiaries, possible side effects, and situations in which the system must not be used. We ensure informed consent and data minimization: we collect only what is necessary, explain how it will be used, and keep access under control. We **practice transparency**: we explain in plain language what the system can and cannot do, what data it used, how it was evaluated, and what its limits are. We keep a human in the loop for sensitive decisions: results are recommendations, not final verdicts.
Accountability is explicit: we designate system owners, feedback channels, and **remediation procedures** for when errors or complaints arise. We prevent **harm through impact** assessments, testing on "edge cases," policies against misuse, and **incident response plans** (shutdown, notification, correction).

How we apply ethics in practice:

- Establish institutional principles (fairness, confidentiality, explainability).
- Implement technical controls (filters, logging, periodic reviews).
- Communicate decisions and users' rights (appeals, data deletion).
- Continuously measure real-world effects, not just lab scores.



Artificial intelligence tools

Domain	Primary purpose	Tools (examples)
General assistants	Chat, ideation, summarization, writing	ChatGPT; Claude; Gemini; Perplexity
Education – lesson design	Plans, worksheets, rubrics, items	Khanmigo; Magic School AI; Canva Magic Write; Copilot for Education
Assessment & quizzes	Items, results analysis	Quizizz AI; Kahoot! AI; Google Forms + add-ons; Socrative
Personalized learning	Progress-based recommendations	Century Tech; Squirrel AI; Notion Q&A pentru curs
Scholarly search	Literature discovery	Elicit; Semantic Scholar; Connected Papers; Research Rabbit
PDF scraping / summarization	Extract tables, summarize	Scholarcy; Paperpile Notes; Zotero + ZoteroGPT
Classical NLP	Tokenization, NER, classification	spaCy; NLTK; Hugging Face Transformers; BERTopic
AI-assisted coding	Suggestions, completions, tests	GitHub Copilot; Codeium; Tabnine; Replit AI
Data analysis	Cleaning, classical models	scikit-learn; XGBoost; LightGBM; Dataiku; RapidMiner
Visualization & BI	Dashboards, insights	Power BI (Copilot); Tableau (Pulse/GPT); Looker Studio
Forecasting/time series	Forecasts, scenarios	Prophet; statsmodels; Nixtla (statsforecast); Orbit
Computer vision	Detection/segmentation	Segment Anything; Detectron2; MMDetection; OpenCV
Remote sensing / GIS	Classification, change detection	Google Earth Engine; eo-learn; Raster Vision; Sentinel Hub
OCR & documents	Scan-to-text, tables	Tesseract; docTR; Adobe Acrobat AI; Azure Form Recognizer
Speech (ASR/TTS)	Transcription, synthesis	Whisper; Vosk; Coqui TTS; Azure/Google Cloud Speech
Translation	RO↔EN/DE/FR etc.	DeepL; Google Translate; ModernMT
Automation & agents	Workflows, bots	Zapier AI; Make + LLM; CrewAI; AutoGen
Ethics, bias & safety	Evaluation, audit	Fairlearn; AIF360; Model Cards Toolkit; Great Expectations
Project management	Planning, notes, templates	Notion AI; Atlassian Intelligence (Jira); Asana AI
References & citation	Bibliography management	Zotero + plugins; EndNote; Paperpile

Runway — AI-powered video generation & editing

https://runwayml.com/

What it is: Runway is a video creation/editing app that uses artificial intelligence to generate short video clips from text (text-to-video), turn images into clips (image-to-video), and quickly edit footage.

What you can do as a beginner:

1. Write a brief description of the desired scene and get a short clip (10–20 sec).
2. Upload a video and "erase" objects with inpainting; the tool automatically reconstructs the background.
3. Use "green screen without a green screen" to separate the subject from the background.
4. Assemble several short clips into a simple edit, add music, and export.

Strengths: quick visual ideation for social media, stylized B-roll, animated backgrounds; lots of one-click effects.
Watch out for: quality and coherence can vary over longer durations; you'll usually work with short clips that you then edit together; results depend on how clear your prompt (description) is.
Tips: start with a style + mood ("night, rain, neon, urban"), request 2–3 variants, pick the best-looking one, then fine-tune (color, cuts). Keep on-screen text to a minimum (if adding titles) and check the export format for your target platform (9:16 for Reels/Shorts).
Who it's for: teachers, NGOs, and small businesses that want short, memorable materials without filming equipment.

InVideo — fast template-based editing

https://invideo.io/

What it is: InVideo is an online video editor built around templates. Instead of starting from scratch, you pick a template (promo, announcement, educational), insert your text, select images/clips from the built-in library, and get a video ready to publish.

What you can do as a beginner:

1. Choose a template for your purpose (promo, event, summary).
2. Write or paste your text; the tool distributes it across scenes.
3. Replace images/footage with your own or stock; add logo, colors, and fonts (brand kit).
4. Generate subtitles and voice-over (if desired), then export in your preferred format.

Strengths: very fast; no editing experience needed; perfect for 30–90 sec clips.
Watch out for: if you don't customize, the video may look "templated." Choose relevant visuals, keep text large and minimal, and check readability on mobile.
Tips: structure the message into 3–4 ideas, start with a clear hook (first 2–3 seconds), add a call to action (what you want people to do after watching). Export different versions for 1:1, 9:16, 16:9.
Who it's for: education, quick marketing, event announcements, NGOs, freelancers.



Synthesia — virtual presenters (avatars)

https://www.synthesia.io/

What it is: Synthesia creates videos with an avatar (digital presenter) that "speaks" your input text in many languages. It's like a virtual studio: you provide the script, choose the avatar and voice, and the platform does the rest.

What you can do as a beginner:

1. Write a short script (30–120 sec).
2. Choose the avatar (digital person), language, and voice.
3. Add images, screens, or lists; set subtitles; export the video. You can also tweak pacing with per-scene durations, upload a brand kit, and store a pronunciation dictionary for names. Start from templates, drag-drop screen recordings, and swap backgrounds with stock b-roll. Generate captions automatically, adjust styles for accessibility, and export multiple aspect ratios (16:9, 1:1, 9:16) for each platform.

Strengths: fast and consistent for tutorials, announcements, short lessons; excellent for multilingual use (the same lesson in multiple languages).
Watch out for: on very long clips it can feel "artificial." Keep sentences short, use a conversational tone, add pauses. For emotional topics, a real presenter may land better.
Tips: use simple structures (introduction → 3 points → conclusion), add screenshots and subtitles for clarity, and make 60–120 sec versions for each subtopic.
Who it's for: internal training, onboarding, video FAQs, micro-learning, rapid content localization.

Lumen5 — text to video (for social)

https://www.lumen5.com/

What it is: Lumen5 turns articles, scripts, or bullet points into short, social-ready videos. Paste text or a link, and it auto-suggests scenes with matching stock footage, images, and captions.

What you can do as a beginner:

1. Paste a paragraph or URL; auto-storyboard generates slides.
2. Pick a template (square, vertical, landscape) and a style pack.
3. Swap visuals from the built-in library or upload your own; add logo/colors/fonts.
4. Auto-caption, tweak timing, add background music, and export.

Strengths: very fast for repurposing blogs/newsletters into 30–90s clips; good templates, layouts, and auto-captioning; easy brand kits.
Watch out for: over-reliance on stock can feel generic—customize visuals and keep text concise; check readability on mobile and contrast for captions.
Tips: open with a strong hook (first 2–3s), structure as problem → 3 points → CTA, keep captions ≤12–14 words per scene, export variants for 9:16 / 1:1 / 16:9.
Who it's for: educators, NGOs, and small teams that need quick, on-brand social videos from existing text.

HailuoAI — AI-assisted content creation (text/images/video)

https://hailuoai.video/

What it is: Under the HailuoAI name you'll find AI-powered content generation solutions (typically text, images, sometimes video), often with templates for posts, articles, scripts, and visual elements. Ideal for beginners who want quick results without complex setup.

What you can do as a beginner:

1. Choose a template (social post, short article, video script).
2. Fill in the fields: topic, audience, tone, length.
3. Receive a draft text and, if needed, generated images; edit and export.

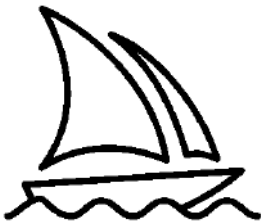
Strengths: speed and simplicity; a solid starting point for ideas, summaries, and promotional copy; sometimes includes basic graphic tools.

Watch out for: fact-check the content, rewrite in your own style, add local examples; generated images may need retouching.

Tips: start with a clear brief (purpose, audience, tone), request 2 variants, verify facts, and adapt the language to your audience.

Who it's for: teachers, NGOs, and small businesses that need quick text/images for social, blogs, or presentations.

It also supports brand kits, collaborative editing, and export presets, so teams can keep outputs consistent across channels. Use built-in plagiarism checks, citation helpers, and tone analyzers to reduce risk and maintain voice. Batch generation accelerates campaigns; API access enables integration with CMS or automation tools and custom workflows too.



Midjourney — generative art on Discord

https://www.midjourney.com/

What it is: Midjourney is an AI image generator that runs via Discord. You describe what you want in a prompt, and it produces stylized images; you can upscale, vary, and refine results directly in the chat.

What you can do as a beginner:

1. Join the Discord server and type /imagine followed by your prompt.
2. Pick one of the four initial images; upscale (U) or create variations (V).
3. Refine with short edits (add/remove concepts, change style or aspect ratio).
4. Download the final image or generate more variations.

Strengths: striking, artistic visuals; fast iterations; rich style control (photoreal, cinematic, watercolor, anime, etc.); great for moodboards, covers, posters.
Watch out for: prompts are public by default in many channels; consistency across many images can be tricky; hands/text and fine details may need retries; mind licensing/use policies.
Tips: start with style + subject + mood + composition (e.g., "portrait of a geography teacher, soft light, 35mm, editorial style"); add aspect ratio --ar 9:16/16:9/1:1; request 2–3 re-rolls and pick the best; keep prompts concise, then iterate.
Who it's for: educators, NGOs, small teams, and creators needing quick concept art, thumbnails, posters, and illustrative visuals without a full design workflow.



DALL·E 3 — friendly image generation, variations, and edits

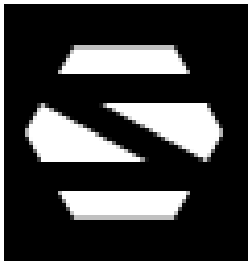
https://openai.com/index/dall-e-3/

What it is: DALL·E 3 generates images from text prompts and supports easy variations and in-painting/out-painting (adding, removing, or extending content). It's designed for natural language prompts—no complex syntax required.

What you can do as a beginner:

1. Write a plain-English prompt describing subject, style, and mood.
2. Ask for 2–3 variations and pick the best.
3. Use editing: select an area to add/remove objects or extend the canvas.
4. Download in common formats and sizes for posts, slides, or print.

Strengths: very good prompt understanding; coherent compositions with readable text more often than many models; quick variations and simple edits for posters, slides, thumbnails.
Watch out for: fine text or logos can still be imperfect; photoreal faces or trademarked characters may be restricted; always check usage rights and avoid sensitive/unsafe topics.
Tips: specify subject + style + mood + lighting + framing (e.g., "isometric classroom infographic, soft light, minimal palette"); keep prompts concise, then iterate; provide short constraints; export multiple aspect ratios (1:1, 16:9, 9:16).
Who it's for: teachers, NGOs, and small teams that need quick, on-brand visuals for lessons, social posts, covers, and presentations without a full design workflow.





Stable Diffusion — control maxim, local sau în cloud

https://stablediffusion.com/

What it is: Stable Diffusion is an open-source text-to-image model that you can run locally (on your own GPU/PC) or via cloud services and hosted UIs. It supports fine control through negative prompts, seeds, samplers, LoRA/Checkpoint models, ControlNet, and image-to-image.

What you can do as a beginner:

1. Use a hosted UI (e.g., web demos or managed platforms) to avoid setup; type a simple prompt and generate.
2. Try image-to-image: upload a sketch/photo and guide the style with a prompt and strength slider.
3. Experiment with negative prompts ("no watermark, no text, no extra fingers") and a fixed seed for reproducibility.
4. Export multiple aspect ratios and upscale with built-in upscalers.

Strengths: high customizability; runs offline for privacy; vast ecosystem of checkpoints, LoRA/embeddings, and ControlNet modules; great for consistent styles and pipelines.

Watch out for: local setup can be technical (GPU/VRAM requirements); model quality varies by checkpoint; prompts may need iteration; respect licenses and safety filters.

Tips: start from a well-rated checkpoint; set style + subject + mood + composition; add a short negative prompt; use a fixed seed to fine-tune; for structure control, add ControlNet; save presets for fast iteration.

Who it's for: educators, NGOs, researchers, and small teams needing controllable, repeatable visuals; creators who want privacy/offline use or to build custom styles and workflows.

Adobe Firefly — generative AI integrated into the Adobe ecosystem

https://www.adobe.com/products/firefly.html/

What it is: Firefly is Adobe's suite of generative AI models built into Creative Cloud apps (Photoshop, Illustrator, Premiere Pro, Express, Acrobat). It focuses on design-safe features like text-to-image, generative fill/expand, text effects, vector recoloring, and style variations—directly inside familiar Adobe workflows.

What you can do as a beginner:

1. Photoshop (Generative Fill/Expand): remove/add objects, extend canvases, change backgrounds with natural-language prompts.
2. Illustrator (Text to Vector / Recolor): generate vector motifs/icons and instantly try palette variations.
3. Express (all-in-one): create social posts, posters, and thumbnails using templates + Firefly text-to-image/text effects.

Strengths: tight integration with Adobe apps; non-destructive edits; brand kits and libraries; licensing and safety emphasis; high-quality print-ready outputs.

Watch out for: advanced controls may still be simpler than power-user tools like Stable Diffusion; results depend on prompt clarity and post-editing; asset licensing/credits vary by plan—review terms before commercial use.

Tips: start from a template or existing PSD/AI file; prompt with subject + action + style + lighting.

Who it's for: designers, educators, NGOs, and small teams already in the Adobe ecosystem who want fast, on-brand visuals and edits without leaving Photoshop/Illustrator/Express.



Leonardo AI — templates, models, and assets for design

https://leonardo.ai/

What it is: Leonardo AI is a creative platform for generating and editing images with ready-made templates, community models, and an asset library (textures, icons, backgrounds). It supports text-to-image, image-to-image, background removal, upscaling, and style consistency via trained models/LoRA.

What you can do as a beginner:

1. Pick a template (poster, thumbnail, product shot) and replace text/images.
2. Use text-to-image for concept art or social posts; generate 2–3 variations.
3. Try image-to-image to restyle sketches/photos (e.g., "watercolor map, minimalist").
4. Remove/change backgrounds, upscale, and export multiple aspect ratios.

Strengths: fast start with templates; good style presets and community models; asset library for quick compositions; useful for consistent branding across batches.

Watch out for: some models produce varied quality—test a few; fine details (hands/text) may need retries; check licensing if using community assets in commercial work.

Tips: define subject + style + mood + composition; add a short negative prompt (e.g., "no watermark, no text artifacts"); save favorite models/settings as presets; keep captions minimal and export 1:1, 9:16, 16:9 variants.

Who it's for: teachers, NGOs, and small design/marketing teams that want fast, on-brand visuals, thumbnails, and product/lesson imagery without a full pro design workflow.



Suno — voice-based song generation (text-to-music & text-to-song)

https://suno.com/

What it is: Suno is an AI music generator that creates full songs—instrumentals and vocals—from plain text prompts. You can specify genre, mood, instruments, tempo, and lyrical themes; it outputs short tracks with sung vocals and a mixdown ready for sharing.

What you can do as a beginner:

1. Write a simple prompt: *genre* + *mood* + *topic* (e.g., "upbeat pop, hopeful, about climate-smart farming").
2. Add a short lyric snippet or let Suno auto-write lyrics.
3. Choose a style/voice and generate 1–2 minutes of music; request 2–3 variants.
4. Regenerate sections or extend the track; download the audio.

Strengths: rapid song drafts with vocals; good for jingles, intros/outros, classroom projects, and quick demos; minimal setup.

Watch out for: coherence over long durations can vary; melodies/lyrics may feel generic—edit or iterate; review licensing/usage terms before commercial use; avoid trademarked names or sensitive content.

Tips:

- Start with genre + mood + theme + reference vibes ("lo-fi hip-hop, chill, study focus, sparse drums").
- Provide a chorus hook (1–2 lines) and let the tool fill verses.
- Generate multiple takes; pick the best parts and stitch externally if needed.
- Keep prompts concise; specify tempo/keys only if important.

Who it's for: teachers, NGOs, and small teams needing quick, catchy music with vocals for lessons, social clips, podcasts, promos, or prototyping song ideas.



Stable Audio — flexible audio generation (sounds, loops, music)

https://stability.ai/stable-audio/

What it is: Stable Audio is an AI tool for generating audio from text prompts and/or reference audio—ranging from one-shots and sound effects to seamless loops and full music beds.

What you can do as a beginner:

1. Describe the target audio ("warm lo-fi drum loop, 85 BPM, vinyl hiss" or "rainforest ambience with distant birds").
2. Choose duration (e.g., 5–45s for SFX/loops, longer for music) and, if available, set BPM/looping.
3. Generate 2–3 takes, pick the best, and download stems/loops if the mode provides them.

Strengths: great for quick SFX, atmospheres, and genre-specific loops; prompt-friendly; useful for podcasts, lessons, reels, and prototypes.

Watch out for: long-form musical structure can be less coherent; exact melodic control is limited; check licensing/attribution for commercial use; levels may need gain/EQ tweaks.

Tips:

- Be concrete: genre + mood + instrumentation + tempo ("cinematic pad, evolving, 70 BPM, minor, no drums").
- For loops, request "seamless loop" and specify bar length/BPM.
- Generate multiple variations and comp the best parts.
- Do light post-processing (normalize, EQ, fade/crossfade loop points).

Who it's for: educators, NGOs, small teams, and creators who need fast, customizable sound effects, backgrounds, or short music loops for videos, slides, games, and social posts.

Glossary of Artificial Intelligence Terms

Accuracy — The proportion of correct results out of the total.
Algorithm — A step-by-step procedure a computer follows to solve a problem.
AI Ethics — Principles for proper use (fairness, transparency, safety).
Big Data — Very large/variable datasets that require special tools.
Chatbot — A program that converses in natural language (e.g., a virtual assistant).
Classification — Assigning a label (e.g., "cat" vs. "dog").
Computer Vision — Models that "see" and understand images/video.
Context (in a prompt) — Background information you give the model so it works correctly.
Conversational Interface — Interaction via text or voice messages.
Data — Information a model learns from (text, images, audio).
Data Protection — Rules and technologies to keep data safe.
Data/Model Drift — When the environment or data patterns change over time, reducing performance.
Deep Learning — Neural-network–based models with many layers.
Differential Privacy — Techniques that add controlled noise to protect data.
Embeddings (semantic vectors) — Numeric representations of text/images that capture meaning.
Evaluation — Measuring model quality with metrics (e.g., F1).
Generalization — Ability to perform well on new, unseen data.
Generative AI — Models that create text, images, audio, or code.
Hallucination — When a generative model invents plausible but incorrect information.
Label — The correct answer associated with an example.
Latency — A system's response time.
LLM (Large Language Model) — A large model for language tasks (chat, summarization, translation).
Machine Learning (ML) — Techniques that let computers learn patterns from data.
Metric — A way to measure performance.
Model — A trained program that makes predictions or generates content.
Neural Network — A structure inspired by the brain that learns complex patterns.
Parameters — The internal numbers a model learns.
Pipeline — A chain of steps: data collection → cleaning → model → evaluation → deployment.
Prompt — The request you send to a model.
Prompt Engineering — Techniques for crafting prompts to get better results.
RAG (Retrieval-Augmented Generation) — A model plus an internal/external search to answer with sources.
Relevance — How well an answer matches the question/purpose.
Rule(s) — Explicit human-written instructions (as opposed to learning from examples).
Token — A chunk of text (word/subword) used internally by models.
Training — The process by which a model "learns" from data.
Vectorization — Converting data (text, images) into numbers so the model can "understand" it.
Zero-shot / Few-shot — Solving a task with no examples or with a few examples in the prompt.

Selected Bibliography

1. Barocas, S., Hardt, M., & Narayanan, A. (2023). *Fairness and machine learning* (ediție actualizată, carte online).
2. Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.
3. Chollet, F. (2021). *Deep learning with Python* (ed. a 2-a). Manning.
4. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., … Houlsby, N. (2021). An image is worth 16×16 words: Transformers for image recognition at scale. În *International Conference on Learning Representations (ICLR 2021)*.
5. Eisenstein, J. (2019). *Introduction to natural language processing*. MIT Press.
6. European Union. (2024). *Artificial Intelligence Act (AI Act)*. Official Journal of the European Union.
7. Géron, A. (2022). *Hands-on machine learning with scikit-learn, Keras, and TensorFlow* (ed. a 3-a). O'Reilly.
8. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
9. Hartley, R., & Zisserman, A. (2004). *Multiple view geometry in computer vision* (ed. a 2-a). Cambridge University Press.
10. Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning* (ed. a 2-a). Springer.
11. Jurafsky, D., & Martin, J. H. (2023). *Speech and language processing* (ed. a 3-a, manuscris în lucru).
12. Kearns, M., & Roth, A. (2019). *The ethical algorithm*. Oxford University Press.
13. Manning, C. D., Raghavan, P., & Schütze, H. (2008). *Introduction to information retrieval*. Cambridge University Press.
14. Murphy, K. P. (2022). *Probabilistic machine learning: An introduction*. MIT Press.
15. Murphy, K. P. (2023). *Probabilistic machine learning: Advanced topics*. MIT Press.
16. National Institute of Standards and Technology. (2023). *AI risk management framework*. NIST.
17. Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (ed. a 4-a). Pearson.
18. Szeliski, R. (2022). *Computer vision: Algorithms and applications* (ed. a 2-a). Springer.