AI TERMINOLOGY / LEVEL 1 / GROUP 1

Generative AI

From Basics to Advanced Research



Presenter guide with slide visuals and notes. The slide numbers match the PowerPoint.

BASICS 02–25

ADVANCED 26–36

RESEARCH 37–47

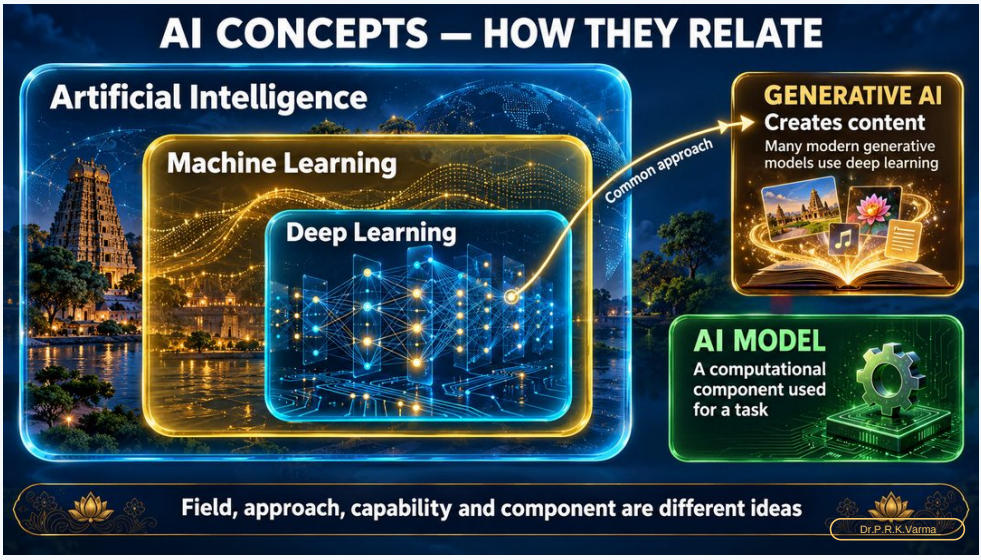

AI Concepts: How They Relate

PRESENTER NOTES

Machine learning is part of AI, and deep learning is part of machine learning. Generative AI describes content creation, while an AI model is a working computational component, not another nested field.



Generate or Retrieve?

PRESENTER NOTES

Retrieval selects material already stored. Generation constructs an output from learned patterns and the current input. One system can retrieve supporting information and then generate a response from it.



Recognising a Style versus Creating a Passage

PRESENTER NOTES

Two students studying Kalidasa illustrate different tasks: recognising a style or category and composing a new passage. Generation does not guarantee originality, truth or quality, and the analogy does not imply human understanding.



Conditional Generation

PRESENTER NOTES

A conditional model produces an output influenced by a supplied condition, such as a text request, reference image, or class label. The condition guides the output without guaranteeing an exact match.



Variation and Seeds

PRESENTER NOTES

Sampling can produce different outputs for the same prompt. A seed can help reproduce the sequence of random choices, but exact repeatability can also depend on the model, settings, and serving system.



Prompt Context and Constraints

PRESENTER NOTES

A prompt communicates what to create, what context to use, and what constraints to respect. Specify the expected format or limits when they matter, then verify that the output follows them.



Tokens and Decoding

PRESENTER NOTES

A language model operates on tokens, which can be words, pieces of words, or other units. A decoding procedure selects the next token using model scores, then repeats with the enlarged sequence.



Image Denoising

PRESENTER NOTES

In many diffusion-style generators, a learned model repeatedly guides a noisy state toward a structured image. This is generation, not the restoration of a hidden original photograph.



From Ingredients to a Recipe Proposal

PRESENTER NOTES

Ingredients and a request lead to a proposed apple-and-orange dessert recipe. Generative AI produces a candidate description; a person must check the recipe and prepare the food.

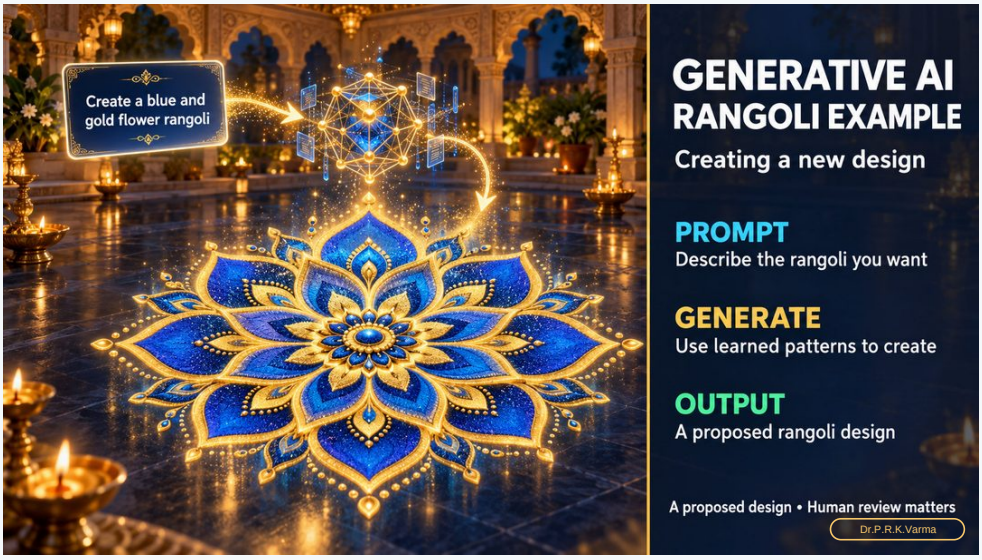

Generating a Rangoli Design

PRESENTER NOTES

A prompt describes the desired blue-and-gold flower rangoli, and a model proposes a new design. Generation differs from classifying an existing rangoli, and human review remains necessary.



Generating a Fabric Design

PRESENTER NOTES

A request for a blue-and-gold floral fabric leads to a proposed pattern. The design is an output, not finished cloth; a weaver or designer checks feasibility before manufacture.



Generating a Melody Proposal

PRESENTER NOTES

A request for a gentle flute melody illustrates content generation rather than identifying an instrument. This is a static illustration, not an audio sample or proof of musical quality.



Generating a Garden Proposal

PRESENTER NOTES

A model proposes a garden layout from a request involving roses and sunflowers. The gardener still checks plant suitability, dimensions and practical constraints before implementing it.



Generating a New Story Draft

PRESENTER NOTES

A writing prompt leads to a proposed story rather than simply retrieving a book from a shelf. Generated text needs review for coherence, accuracy and suitability for its intended audience.

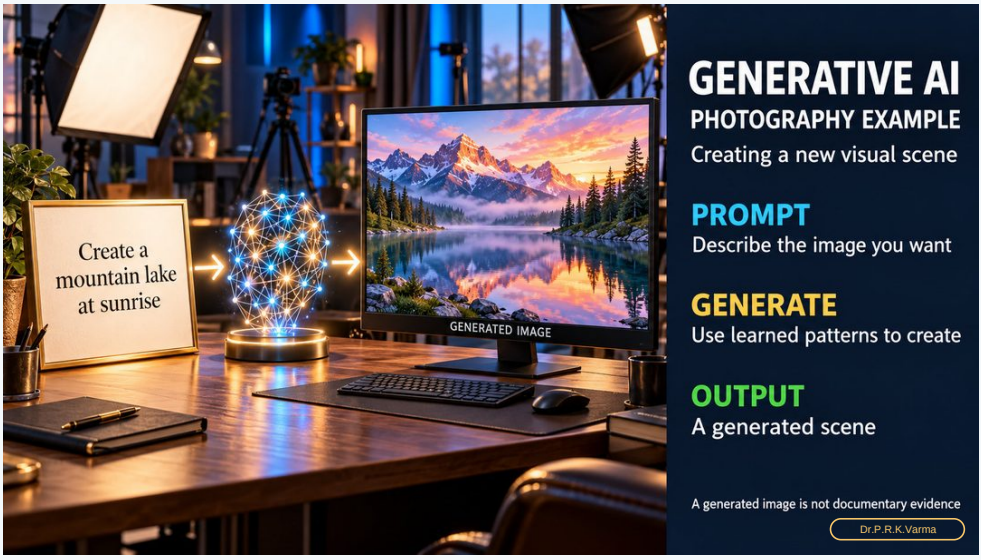

Generating an Imagined Photograph

PRESENTER NOTES

A mountain lake at sunrise is generated from a request. A realistic-looking result is not documentary evidence that the scene or event actually occurred.



Generating a Clothing Design

PRESENTER NOTES

A request for a blue dress with gold floral details produces a proposed design. A designer must still assess fit, construction and feasibility; the image itself is not a manufactured garment.

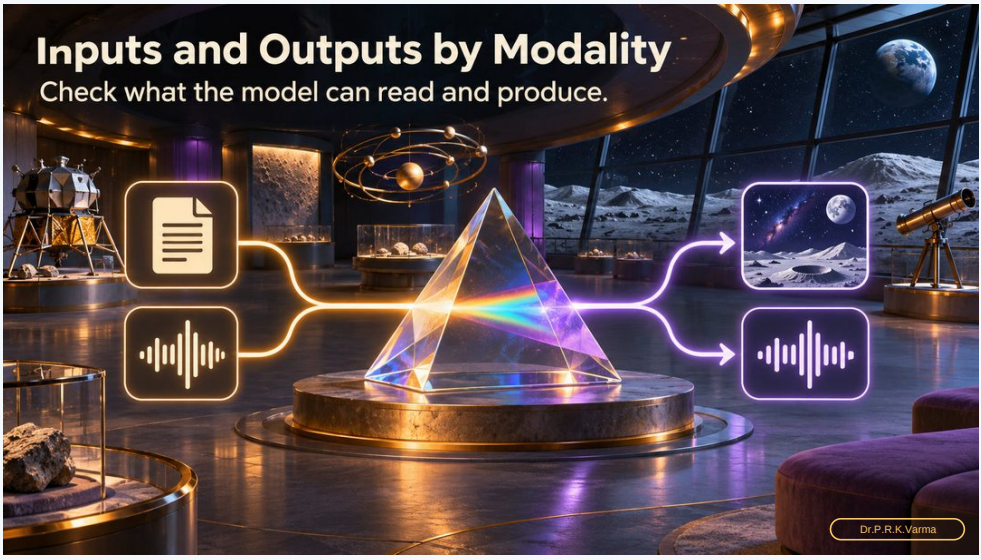

Inputs and Outputs by Modality

PRESENTER NOTES

A model may accept text, images, audio, or some combination, and its supported output types can differ from its inputs. A multimodal system does not necessarily handle every input-output pairing.

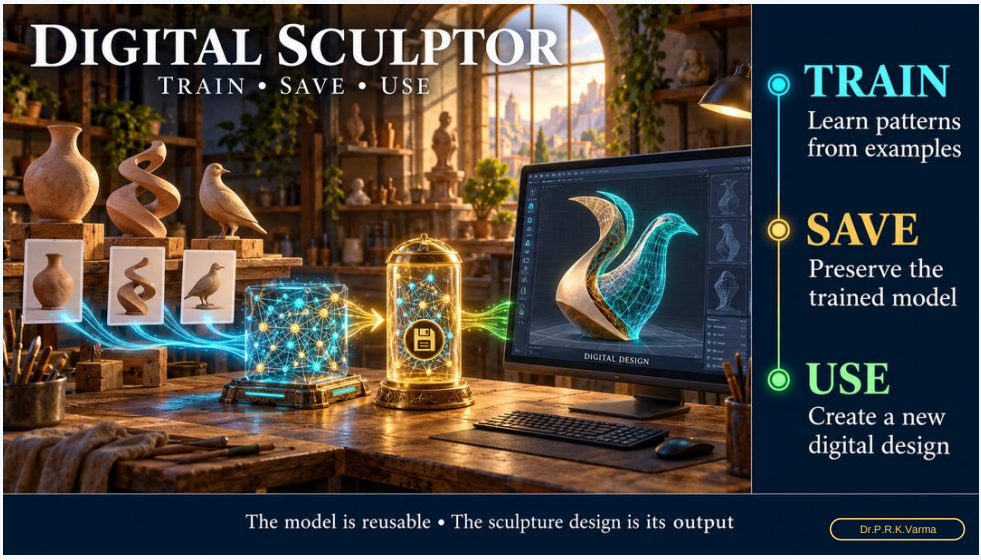

Digital Sculptor: Learn, Preserve and Generate

PRESENTER NOTES

Examples help adjust a model; saving preserves the trained system; a new request produces a digital sculpture design. The model is reusable, while the design is its output. It does not physically sculpt an object or guarantee perfect skill.

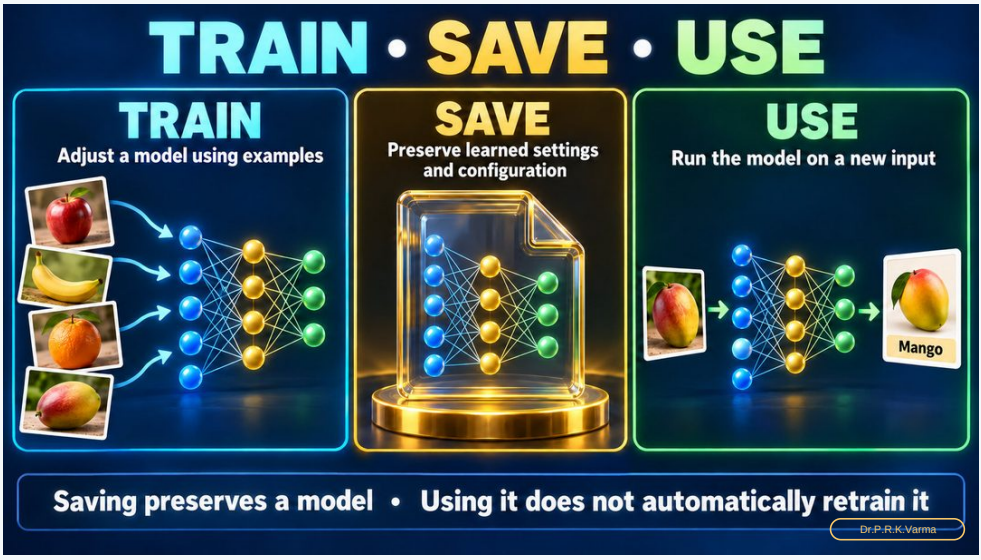

Train, Save, Use

PRESENTER NOTES

Training adjusts a model using examples. Saving preserves learned settings, and use applies the model to a new input. Saving does not turn an output into a model, and ordinary use does not retrain it.



Human Iteration and Review

PRESENTER NOTES

Review the first output against the goal, correct errors, and refine the request or the artifact. For consequential use, an appropriate human reviewer should sign off before publication or deployment.



Plausibility and Factual Grounding

PRESENTER NOTES

A fluent response may sound correct without being supported. Where accuracy matters, compare the claims with reliable evidence and check whether the cited material actually supports them.



Attribution and Resemblance

PRESENTER NOTES

A generated output can resemble an existing work without revealing which examples influenced it. Do not infer a source from appearance alone. Use reliable provenance records where available and verify attribution independently.



Dataset and Rights Awareness

PRESENTER NOTES

Record where training or reference material and outputs came from. Check applicable licenses, consent, privacy, and intended use before release. When the rules are unclear, get qualified guidance rather than assuming permission.



Five Ms: A Memory Aid

PRESENTER NOTES

Machines, Many Examples, Many Layers, Making Content and Model for Reuse recall five related concepts. They are memory hooks, not five training stages or a chain of nested categories.

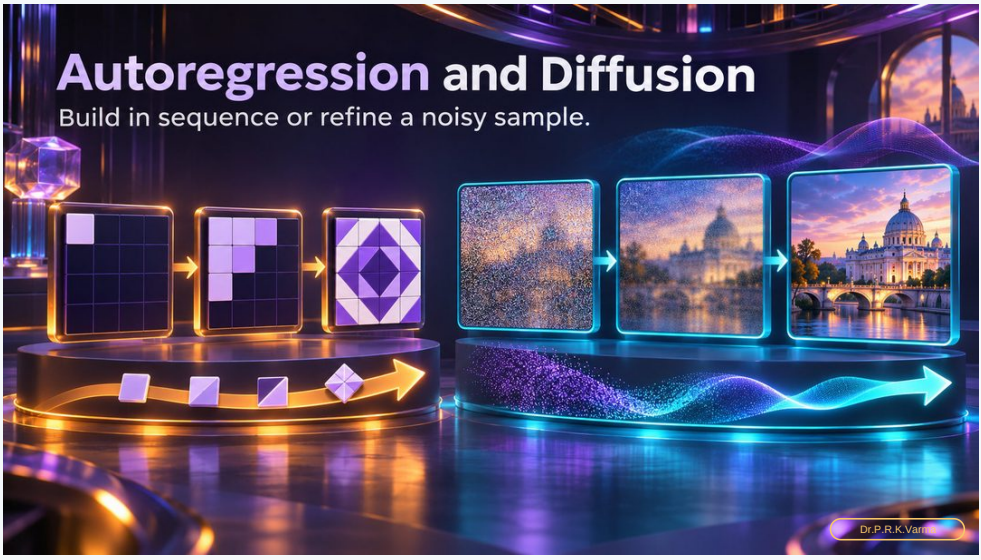

Autoregression and Diffusion

PRESENTER NOTES

Autoregressive models generate a sequence by conditioning on preceding tokens. Diffusion models iteratively reverse noise under a learned model. Both can support multiple modalities, so do not equate architecture with output format.

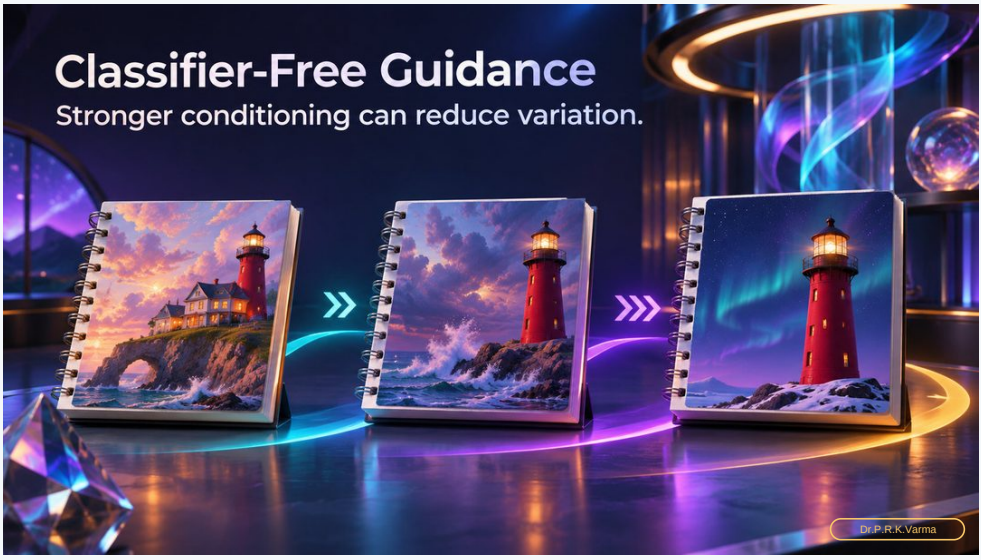

Classifier-Free Guidance

PRESENTER NOTES

Classifier-free guidance blends conditional and unconditional or negative-condition predictions in common implementations. Try multiple strengths and compare adherence, diversity and visual artifacts rather than assuming stronger is better.

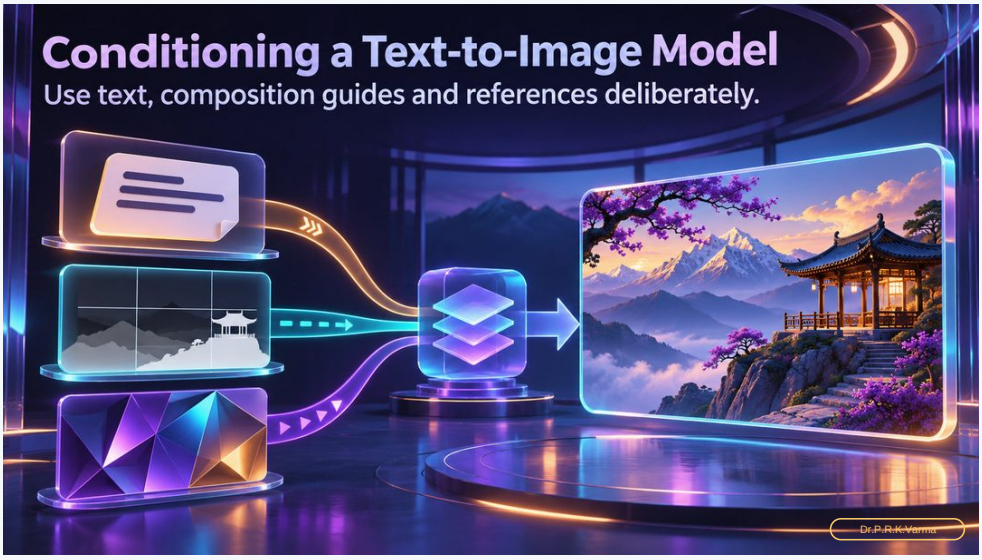

Conditioning a Text-to-Image Model

PRESENTER NOTES

Text and supported image conditions can direct different attributes such as subject, spatial arrangement or style. Add one control at a time and check for contradictions and unintended carryover.



Inpainting Mask Boundaries

PRESENTER NOTES

A mask nominates an editable region, but precise outside-mask preservation depends on the pipeline. Inspect seams, illumination and texture, and composite against the original if an exact surrounding region must be preserved.

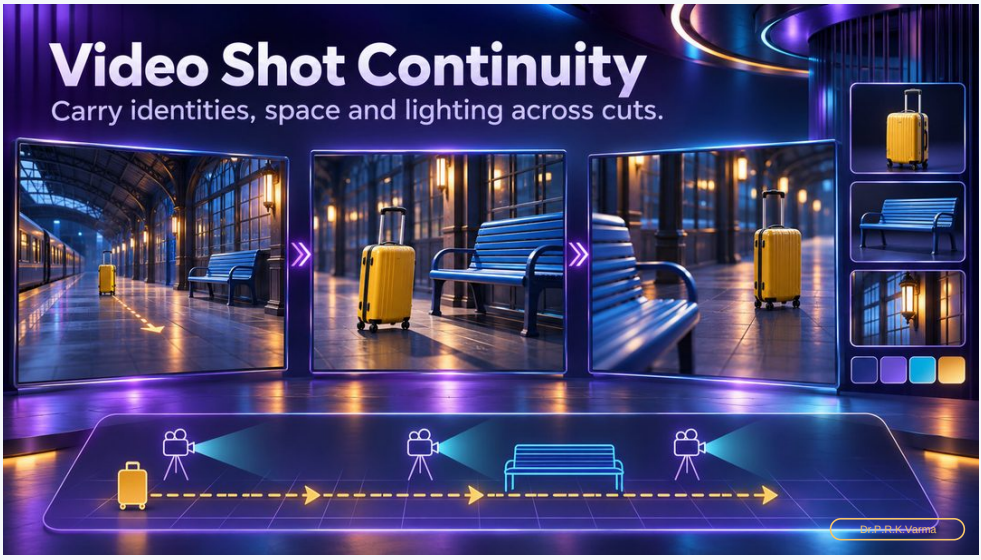

Video Shot Continuity

PRESENTER NOTES

Plan anchors such as the subject, camera direction, prop positions, lighting and timing. Review cuts in sequence and check continuity in both motion and framing; a convincing single frame does not guarantee a coherent shot.



Speaker Consistency and Consent

PRESENTER NOTES

Obtain permission for a reference voice and agree on its permitted scope. Evaluate identity consistency and intelligibility across segments. Keep the recording and generated asset within those permissions and disclose synthesis as appropriate.

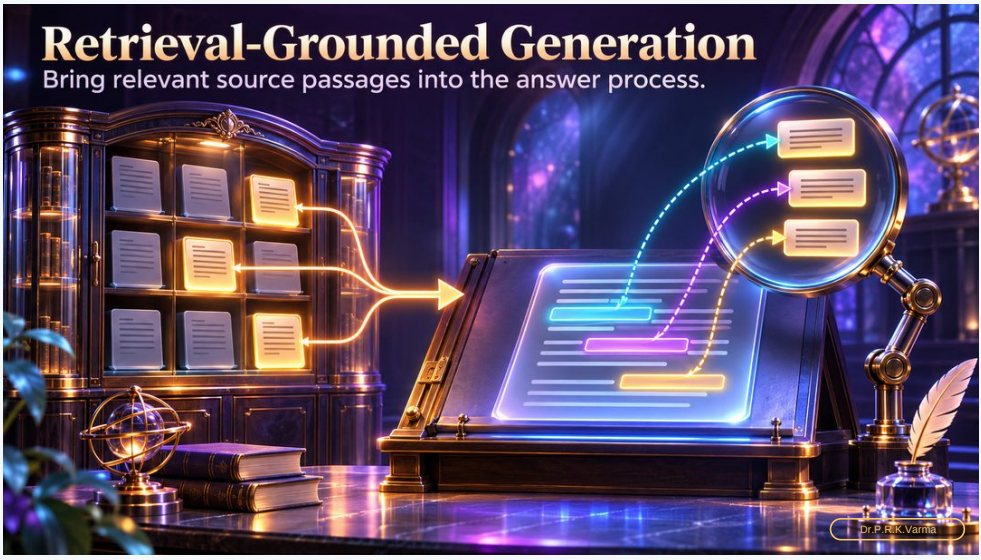

Retrieval-Grounded Generation

PRESENTER NOTES

Retrieve relevant passages, compose using that evidence, and verify each material claim against the cited source. Retrieval does not guarantee that a result is accurate, current or authorized to use.

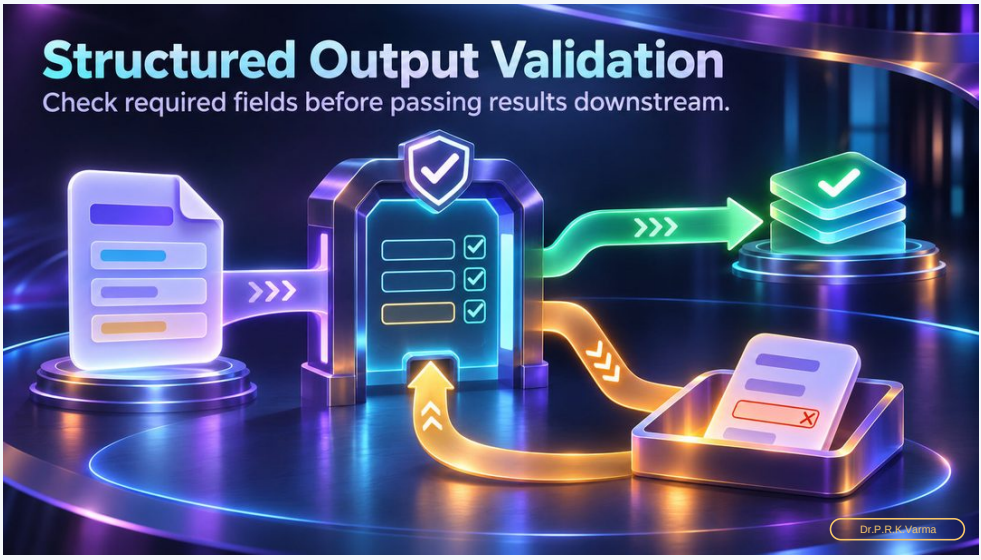

Structured Output Validation

PRESENTER NOTES

Use a schema or constrained decoder where supported, then validate both syntax and domain rules outside the model. Treat invalid output as an explicit failure with a controlled retry or fallback.

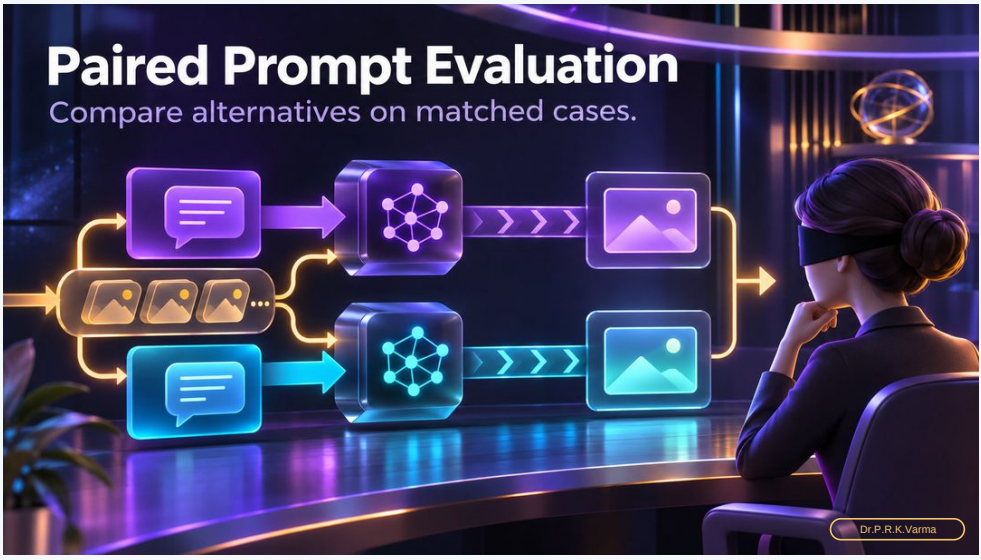

Paired Prompt Evaluation

PRESENTER NOTES

Run competing prompts on the same representative tasks with comparable sampling settings. Use a predetermined rubric and review difficult slices. Paired comparisons help distinguish a prompt effect from a change in test mix.



Diversity and Fidelity

PRESENTER NOTES

Assess diversity across repeated generations separately from individual output quality, adherence or realism. Document seeds, configurations and evaluation criteria; higher diversity need not imply higher fidelity.

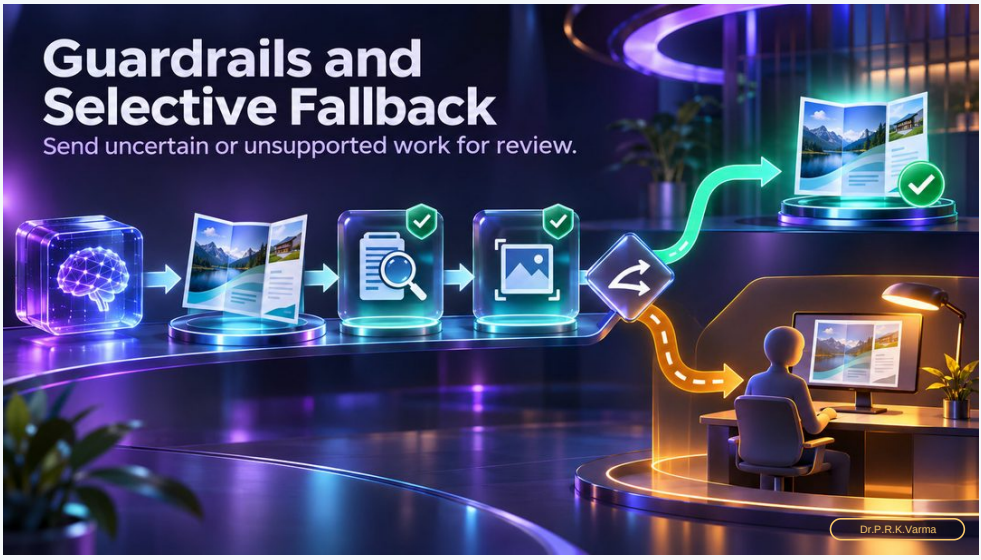

Guardrails and Selective Fallback

PRESENTER NOTES

Check policy, factual requirements and output constraints within a defined scope. If the system lacks evidence or returns invalid content, it can abstain, ask for clarification or route the task to a reviewer rather than silently proceeding.



Flow Matching and Diffusion

PRESENTER NOTES

Flow matching learns a time-dependent velocity field for chosen probability paths. Some path choices connect to diffusion; the comparison concerns objectives and paths, not a universal ranking of their sample quality.

Original research

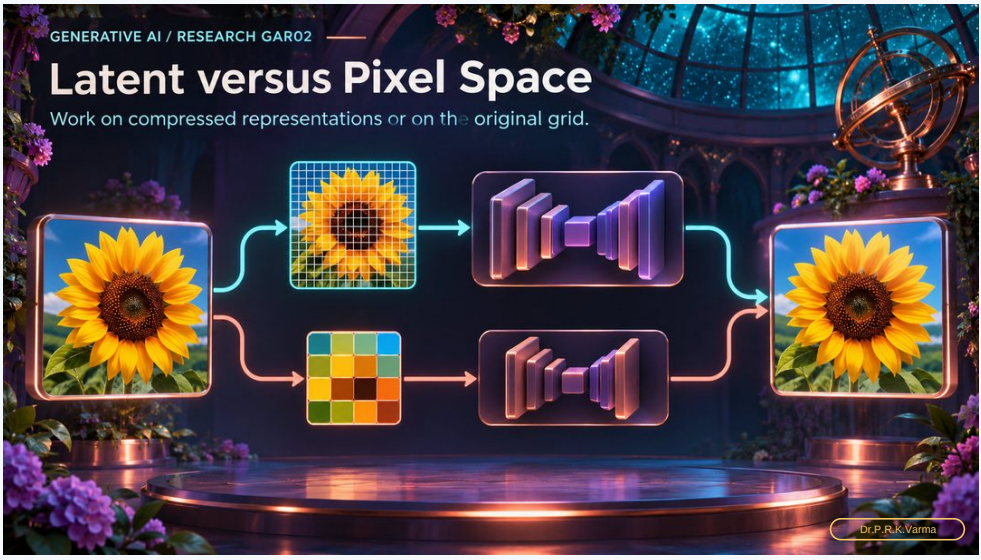

Latent versus Pixel Space

PRESENTER NOTES

Latent diffusion denoises a compressed representation and then decodes it. That can reduce compute, but the encoder and decoder also affect which detail can be recovered. Compare reconstruction and generation quality separately.

Original research

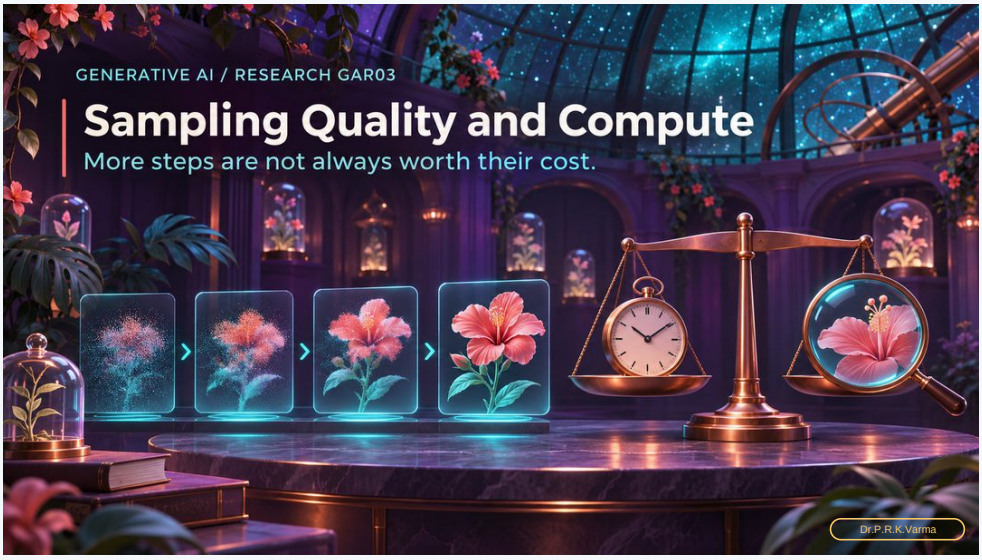

Sampling Quality and Compute

PRESENTER NOTES

Compare samplers over several step budgets while keeping the model, conditioning, and evaluation suite fixed. Report both quality and measured latency. A sampler can improve one point on the trade-off without winning everywhere.

Original research

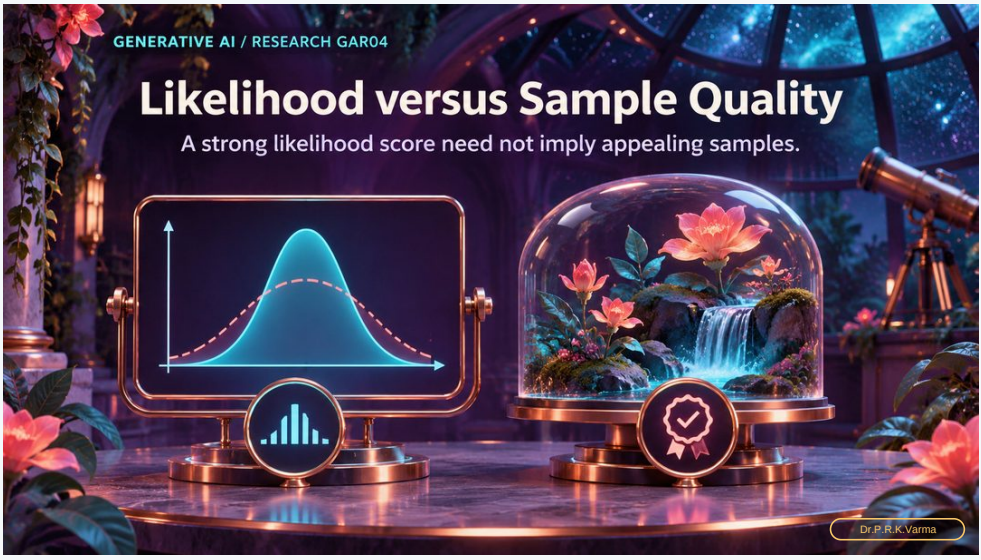

Likelihood versus Sample Quality

PRESENTER NOTES

Likelihood quantifies a model’s assigned probability density under a particular setup. Apparent sample fidelity, likelihood, and downstream usefulness can disagree, so choose evaluation measures that address the intended application.

Original research



Generative Evaluation Bias

PRESENTER NOTES

Image preprocessing and metric implementation can alter rankings; changes such as resizing can affect FID. Specify these settings and compare multiple metrics and human judgments before treating one number as definitive.

Original research

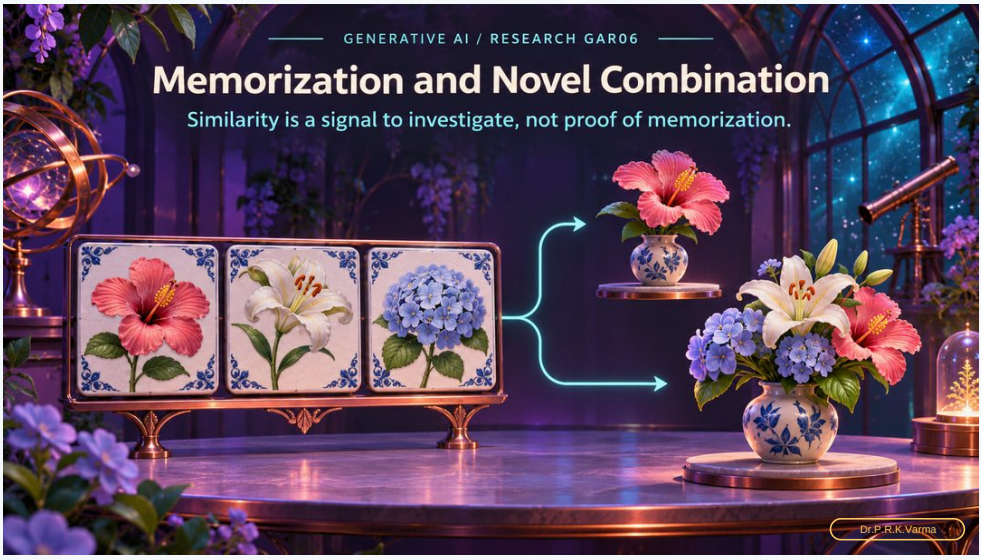

Memorization and Novel Combination

PRESENTER NOTES

Generated examples may resemble known examples through copying or shared structure. A similarity score alone cannot distinguish them. Audit with authorized references and report uncertainty, especially when the original training data cannot be inspected.

Original research

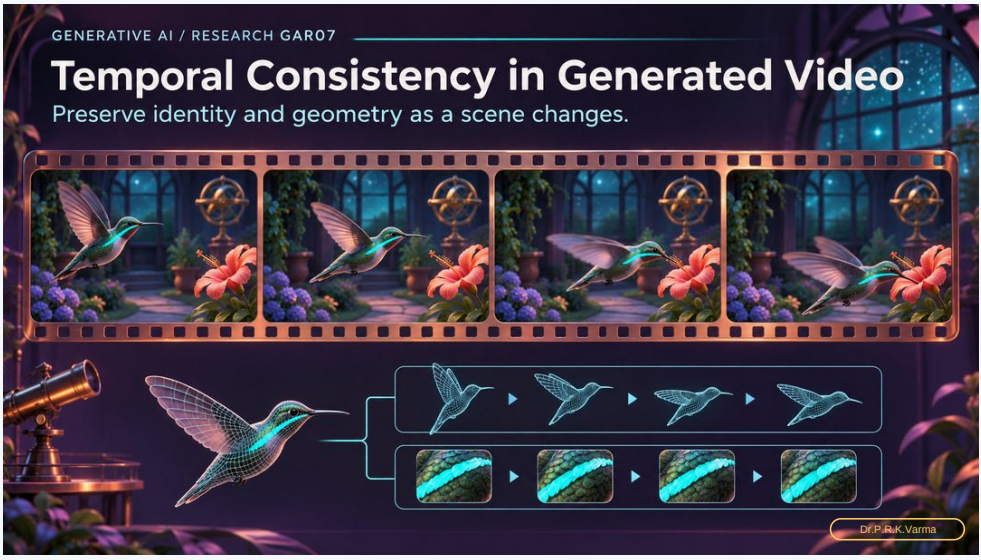

Temporal Consistency in Generated Video

PRESENTER NOTES

Video quality includes temporal properties such as identity consistency and motion smoothness. Assess them separately from individual frame quality, and include tests across motion speeds, viewpoints, and clip lengths.

Original research

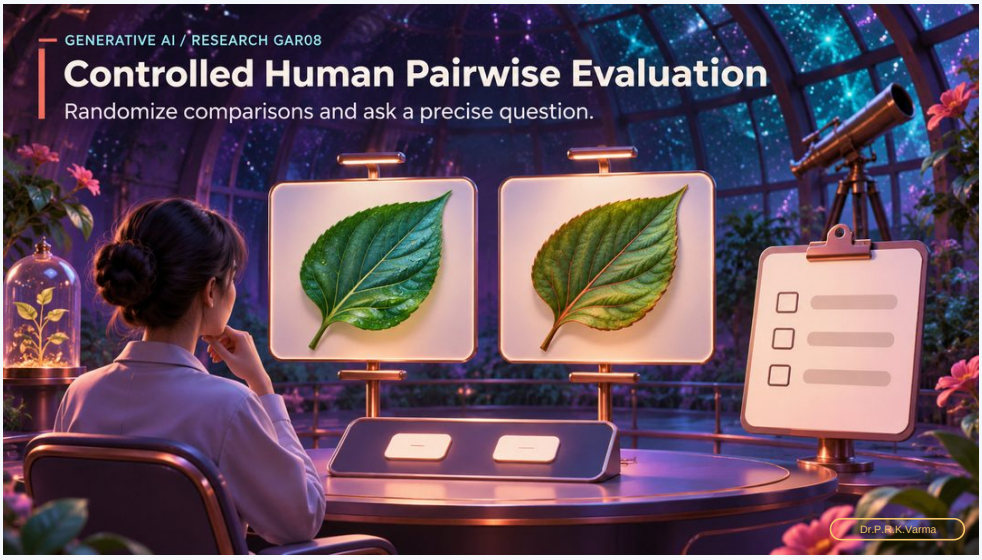

Controlled Human Pairwise Evaluation

PRESENTER NOTES

For paired judgments, keep prompts and presentation conditions comparable, randomize order, and define the criterion before rating. Report ties, annotator variation, and uncertainty rather than one universal preference score.

Original research



Diversity-Collapse Diagnostics

PRESENTER NOTES

High apparent quality can coexist with limited variety. Evaluate coverage across conditions and compare generated and reference distributions with diagnostics that distinguish fidelity from coverage; no one diversity measure is definitive.

Original research

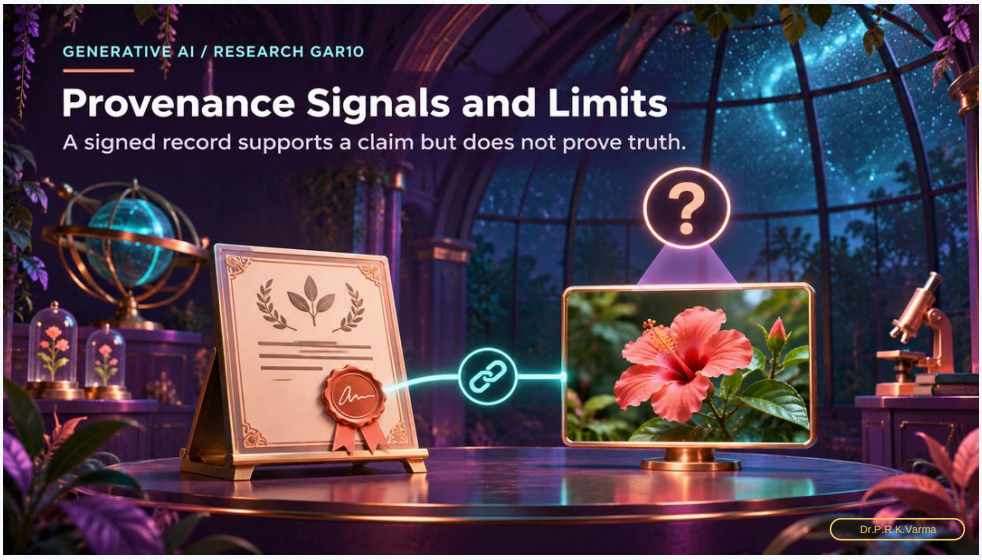

Provenance Signals and Limits

PRESENTER NOTES

A verified Content Credential can authenticate that the signed assertions are tied to the asset and intact under the relevant trust assumptions. It does not prove that the depicted event is true, and missing credentials do not prove fabrication.

Original research



Reproducible Generation Stress Tests

PRESENTER NOTES

Freeze the generation and scoring protocols before evaluating. Save prompt sets, model version, seeds when available, settings, scoring criteria, and exclusion rules so another team can attempt to reproduce the findings.

Original research