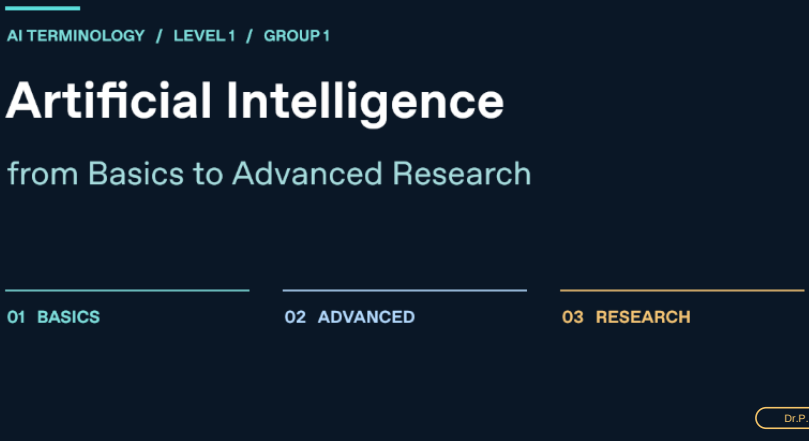
AI TERMINOLOGY / LEVEL 1 / GROUP 1

Artificial Intelligence

From Basics to Advanced Research

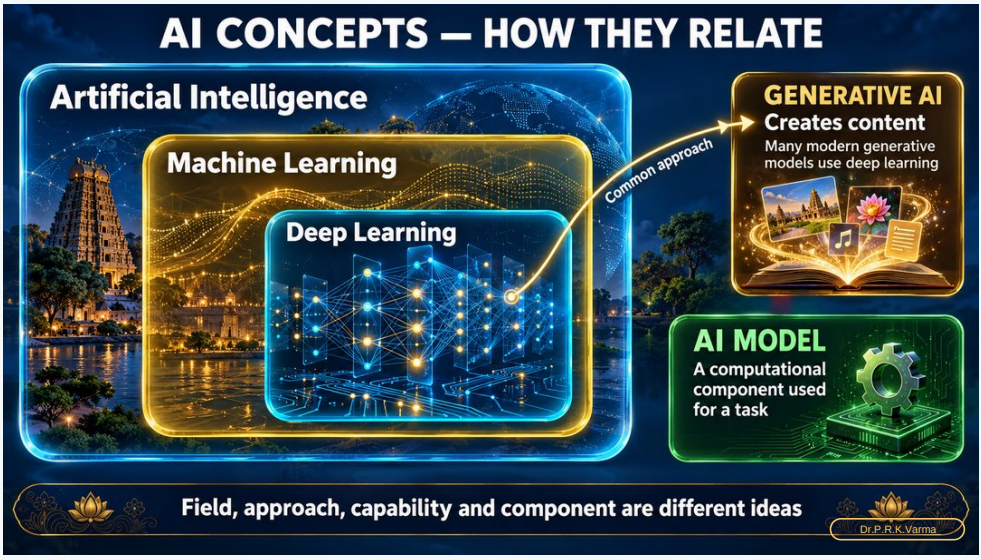


Presenter guide with slide visuals and notes. The slide numbers match the PowerPoint.

BASICS 02–23

ADVANCED 24–34

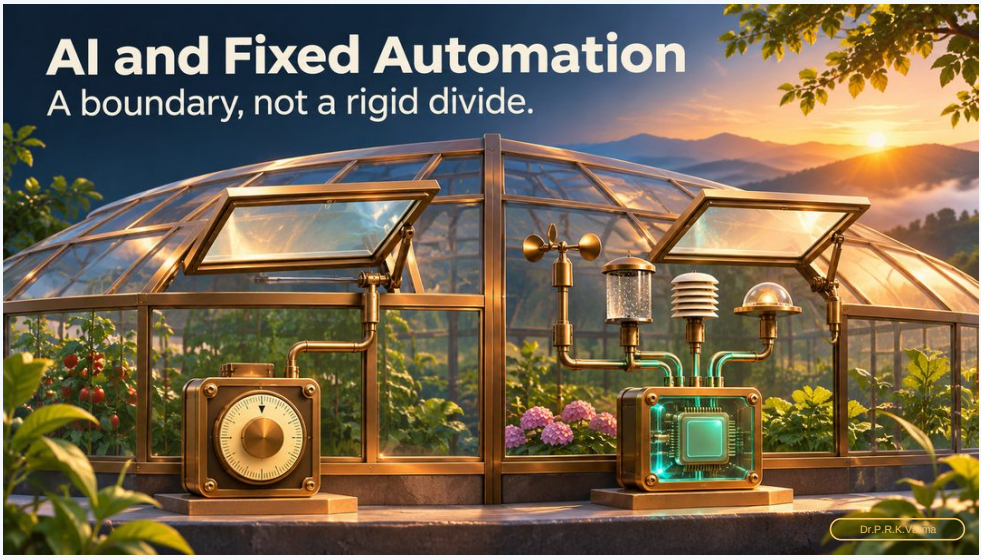
RESEARCH 35–45



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.



AI and Fixed Automation

PRESENTER NOTES

A fixed rule repeats prescribed steps. Some AI systems infer or choose what to do under varying conditions. The boundary is not exact, and practical systems often combine both.



Goals and Task Scope

PRESENTER NOTES

Describe the task, the allowed inputs and outputs, and the conditions where the system should work. A system that sorts books does not necessarily know how to repair them.



AI System: Kitchen

PRESENTER NOTES

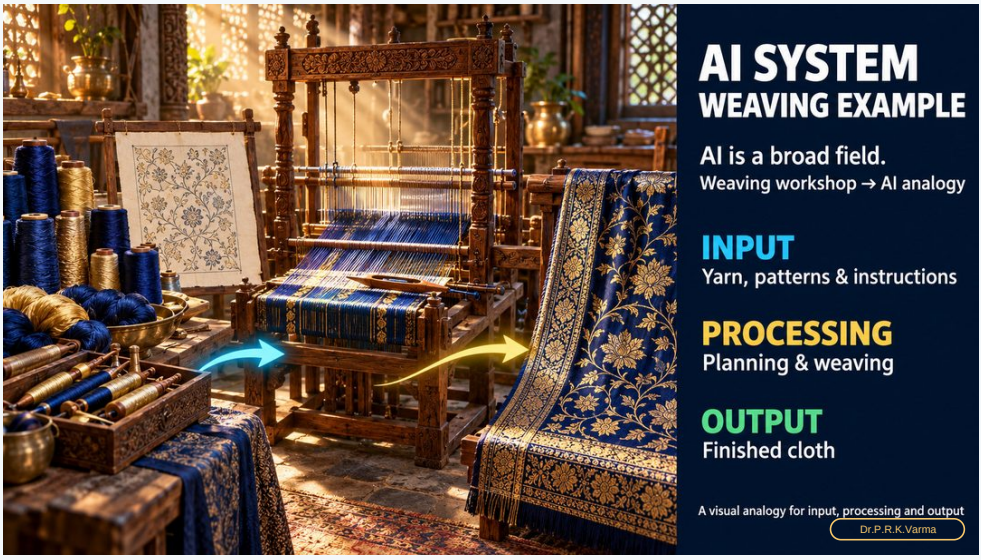
Ingredients and instructions are inputs; planning and cooking represent processing; the prepared dish represents an output. The kitchen is an analogy for a whole system, not a literal explanation of computation.



AI System: Rangoli

PRESENTER NOTES

Colours, patterns and instructions become inputs to a planning and design process, leading to a rangoli result. Use the example to explain input, processing and output without equating AI with human artistic understanding.



AI System: Weaving

PRESENTER NOTES

The weaving example follows input, processing and output. Materials and instructions enter a coordinated process that produces cloth; the whole workflow illustrates an AI system rather than only its model.



Different Routes to Intelligent Tasks

PRESENTER NOTES

Like different routes to a destination, AI can use rules, search or learning. The aim is to find a useful plan for the conditions, not to claim every AI system learns in the same way.



Perception and Uncertainty

PRESENTER NOTES

Perception systems interpret signals such as images, sounds, or sensor readings. When the signal is incomplete, several interpretations can fit. An uncertain answer is not the same as a known fact.



Language Needs Context

PRESENTER NOTES

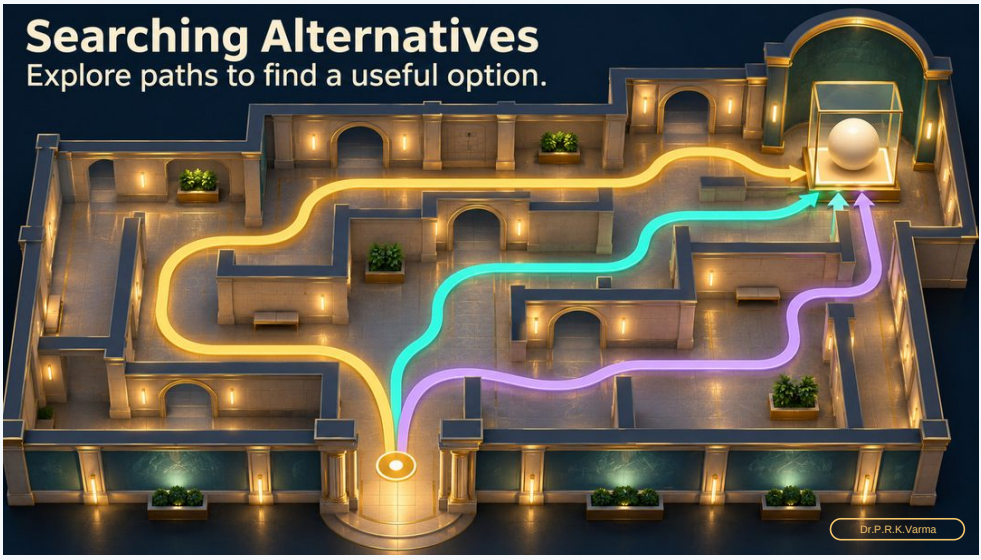
The surrounding words and situation help a system choose among possible meanings. A clear question or an example can resolve ambiguity, but inferred intent can still be wrong.



Knowledge Representation

PRESENTER NOTES

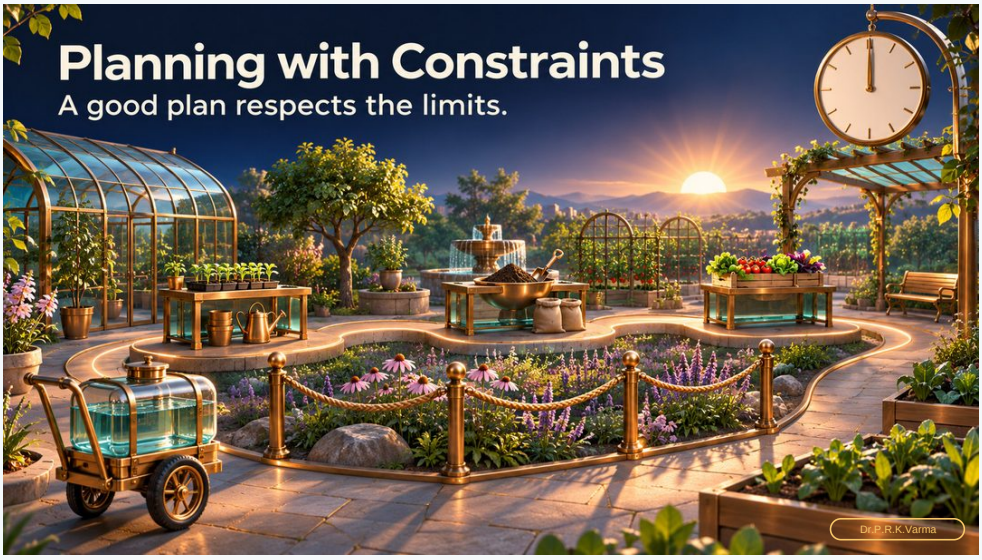
A representation gives information a usable structure. A knowledge graph, for example, links entities and relationships so a system can find relevant facts or follow chains of relationships.



Searching Alternatives

PRESENTER NOTES

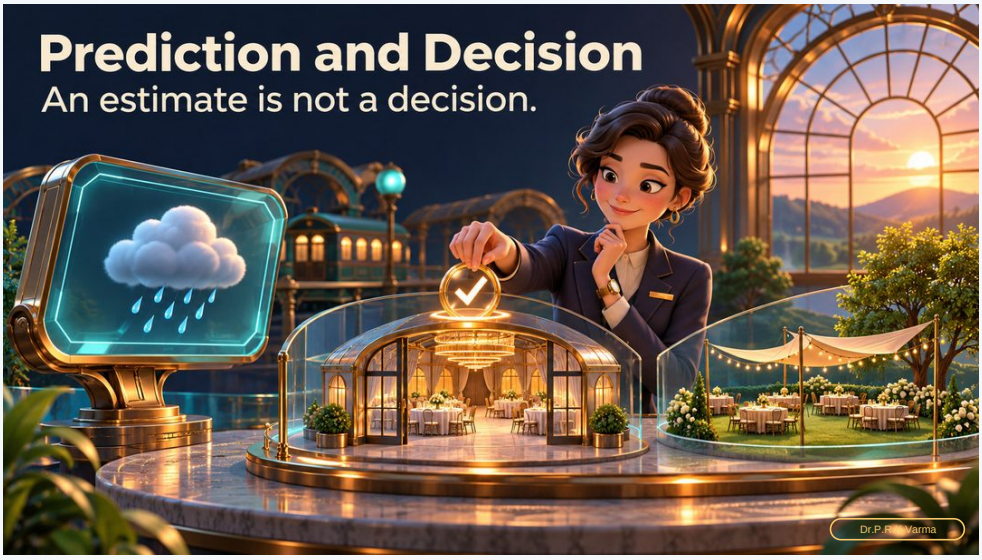
Search explores alternatives rather than committing to the first one. A route planner can compare candidate museum paths for distance or accessibility, depending on the task.



Planning with Constraints

PRESENTER NOTES

Planning arranges actions to reach a goal. A garden plan might schedule watering before opening while respecting limited water, staff availability, and areas where carts are not allowed.



Prediction and Decision

PRESENTER NOTES

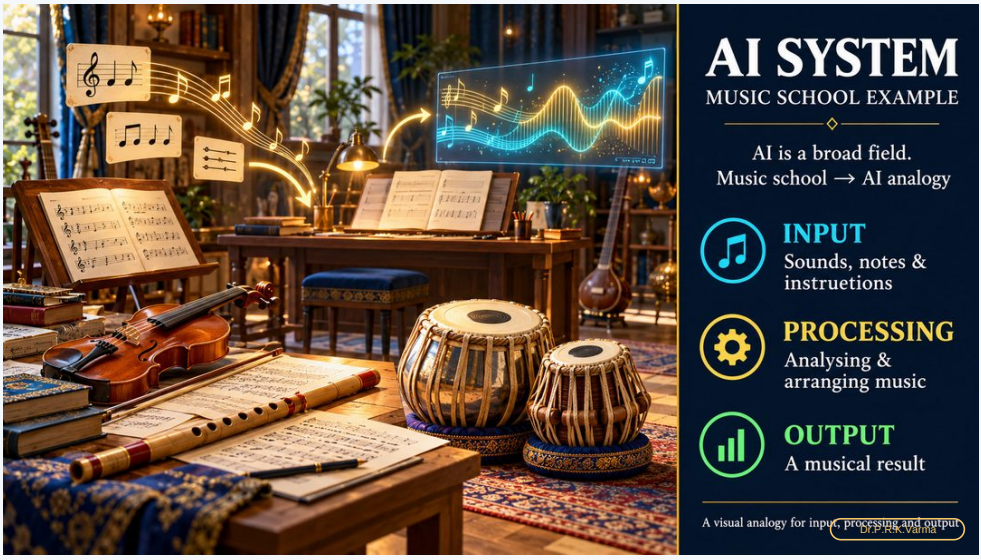
Prediction estimates what may happen. Decision-making chooses what to do while considering the estimate, goals, costs, or rules. A rain prediction does not by itself decide whether to move an event indoors.



Human Judgment and Escalation

PRESENTER NOTES

Set rules for when the system should stop, ask for clarification, or route an unusual case to a person. A human reviewer needs enough context and authority to make or correct the final decision.



AI System: Music School

PRESENTER NOTES

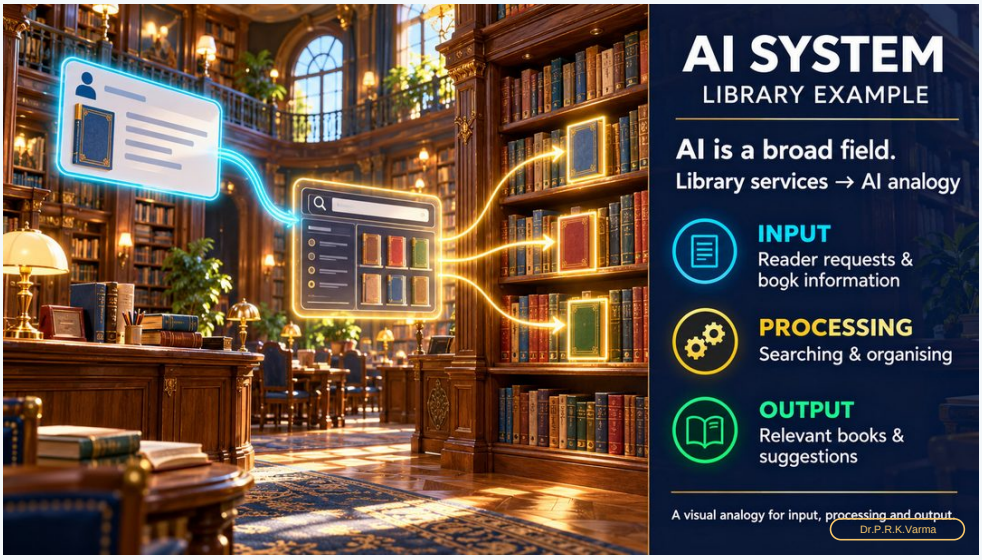
Sounds, notes and instructions are analysed and arranged to produce a musical result. The example separates what enters a system, what it does and what it returns.



AI System: Garden

PRESENTER NOTES

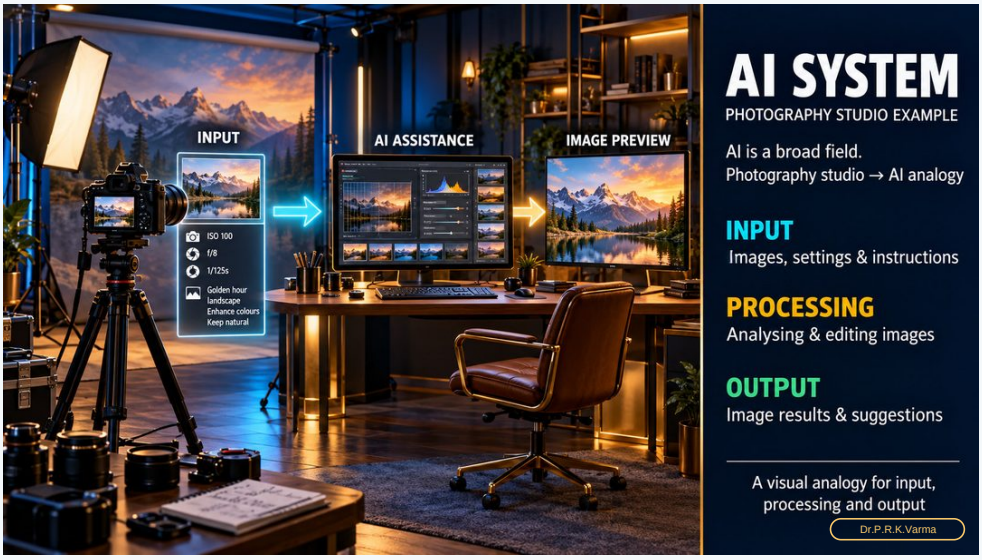
Plants, conditions and instructions guide planning and tending, leading to a garden result. Garden work is an input–processing–output analogy, not a claim that plant growth itself is AI.



AI System: Library

PRESENTER NOTES

Reader requests and book information enter a system that searches and organises them, returning relevant books and suggestions. This distinguishes a complete service from the individual methods it uses.



AI System: Photography Studio

PRESENTER NOTES

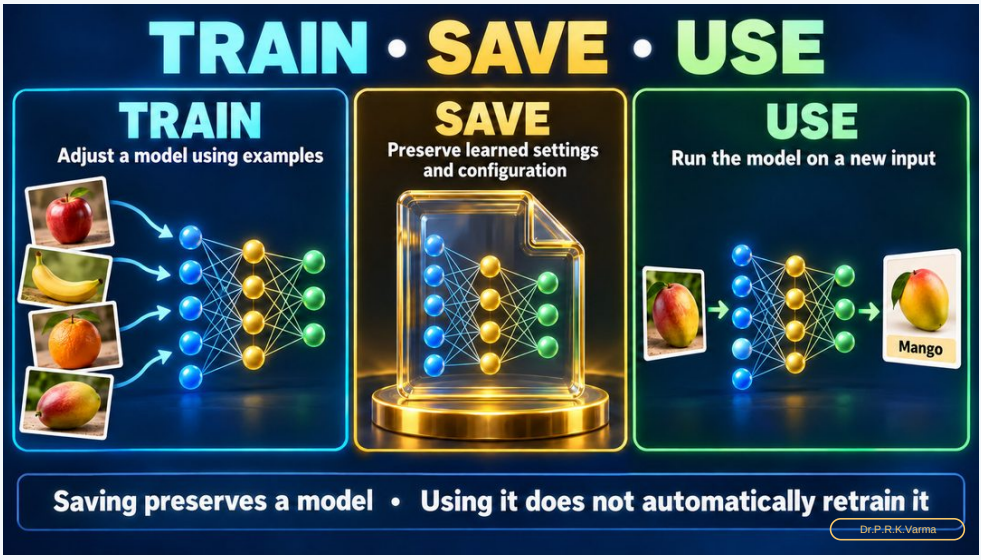
Images, settings and instructions are analysed and edited to produce image results and suggestions. Input quality and the intended task affect how useful those results are.



AI System: Dress Design Studio

PRESENTER NOTES

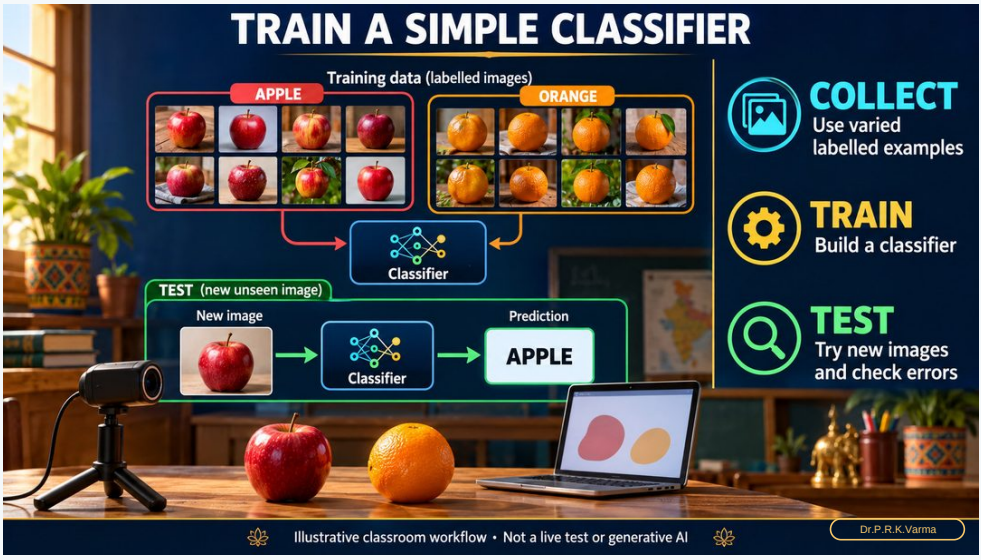
Fabric, measurements and instructions guide planning and designing to produce a clothing design. The example makes the system's inputs, processing and outputs visible.



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.



A Practical Classification Demonstration

PRESENTER NOTES

Train with labelled apple and orange images, then classify a new image. This is a teaching illustration, not the measured result of an experiment that was actually run.



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.



Model and System Boundaries

PRESENTER NOTES

Draw the boundary around the model, data interfaces, decision rules, tools and people. Assign ownership and test input, output and failure assumptions at the boundaries. A model improvement is not automatically a system improvement.



Rules and Learned Perception

PRESENTER NOTES

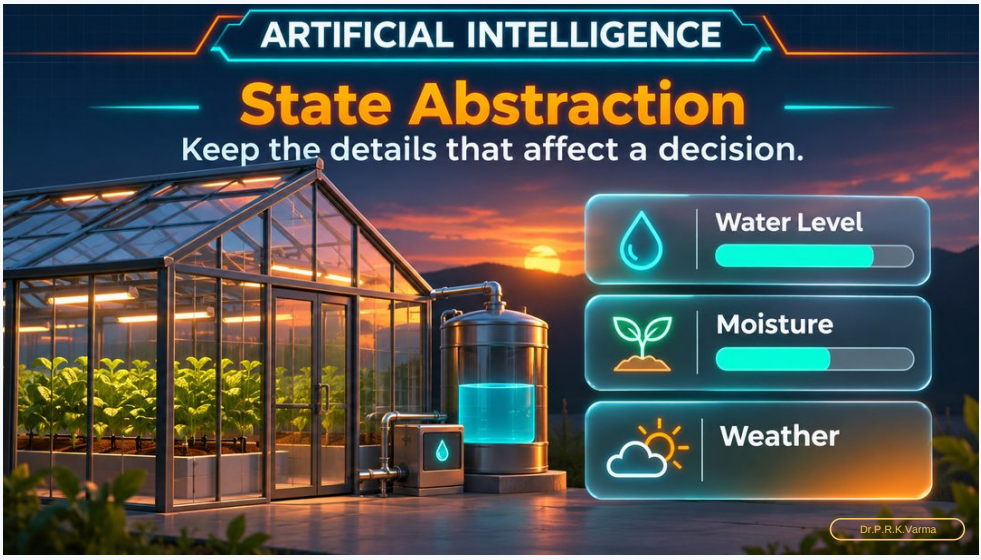
Use learned perception for ambiguous signals, then apply explicit rules for deterministic requirements. If confidence is low or a rule cannot be applied, send the item to review rather than treating a model guess as a verified fact.



Multimodal Fusion Tradeoffs

PRESENTER NOTES

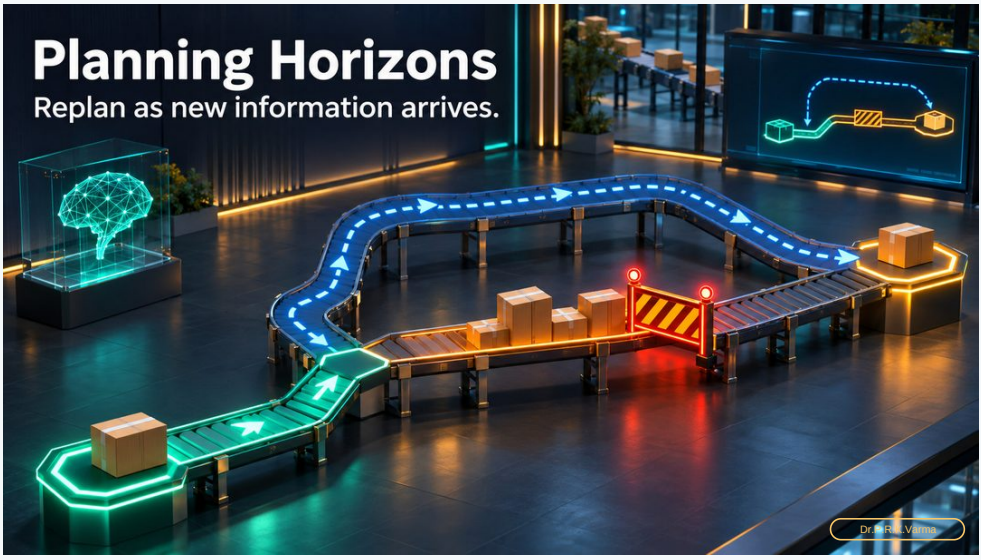
Early fusion combines inputs, intermediate fusion combines features, and late fusion combines predictions. Evaluate time alignment, absent modalities, compute and robustness. Check whether each additional modality improves performance on the intended cases.



State Abstraction

PRESENTER NOTES

A state abstraction omits detail to make decisions more manageable. Test whether discarded information changes downstream decisions, especially in unusual conditions.



Planning Horizons

PRESENTER NOTES

A receding-horizon system plans ahead, takes the first action, observes the new state and plans again. A longer lookahead can improve anticipation but increases computation and dependence on the forecast.



Permission Gates for Tools

PRESENTER NOTES

Give tools limited permissions. For a calendar assistant, reading availability and drafting a change can be separated from committing it. Require authorization at the appropriate boundary, log the result and distinguish a proposed action from a completed one.



Recovery After an Action

PRESENTER NOTES

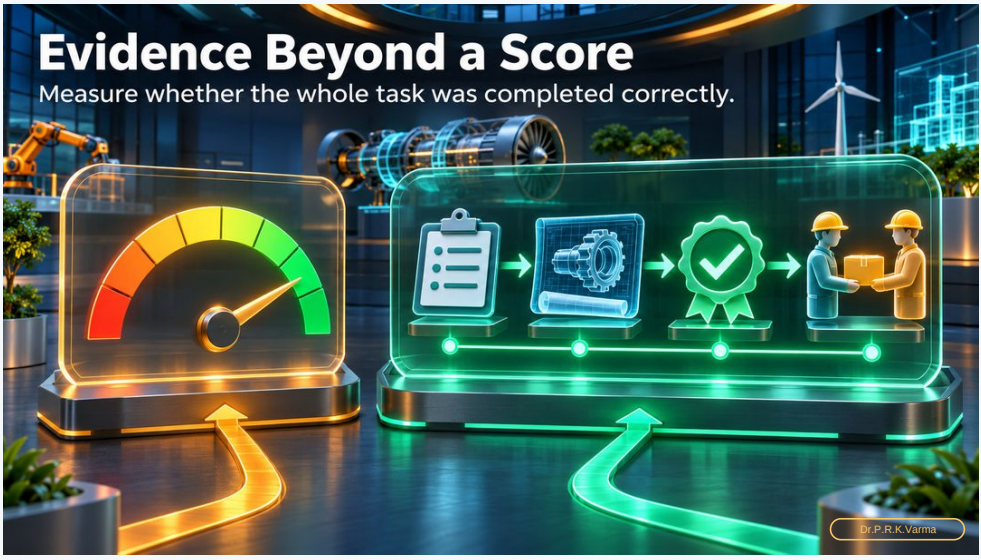
Check the observed result rather than assuming an action succeeded. Treat partial completion and duplicate actions explicitly, with human review when the system cannot establish the outcome.



Verification and Fallback

PRESENTER NOTES

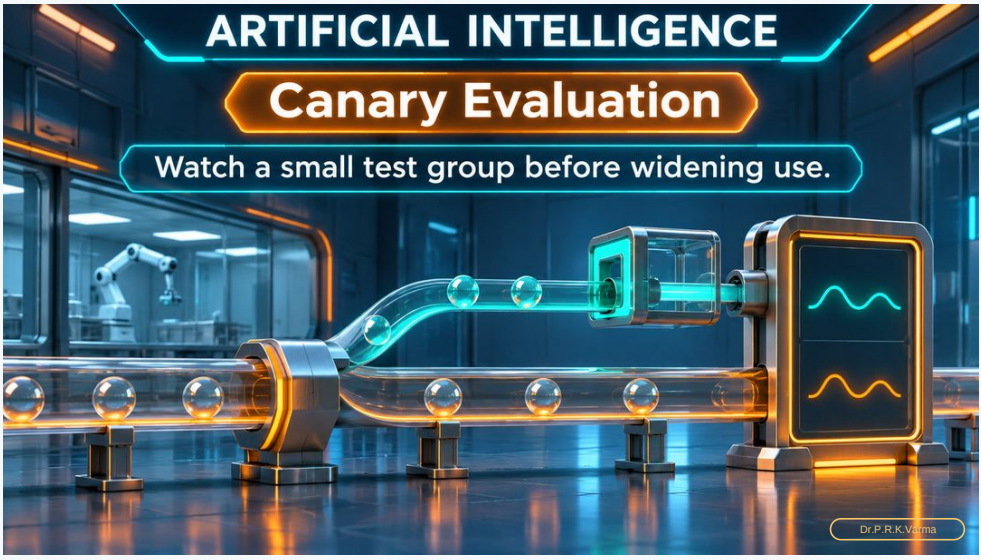
Validate structured outputs, inspect supporting evidence and test known constraints before use. When checks fail or evidence is inadequate, choose a safe fallback, ask for clarification, or route to human review.



Evidence Beyond a Score

PRESENTER NOTES

Use model scores as diagnostic signals, not substitutes for end-to-end evidence. Test representative tasks and inspect correct outputs, constraint satisfaction, latency, human rework and failed-case patterns.



Canary Evaluation

PRESENTER NOTES

Start with a bounded trial and compare against the current workflow. Define what would pause or reverse the rollout, and do not interpret a tiny sample as proof of reliability.



Deployment Monitoring

PRESENTER NOTES

Monitor failures, input changes, output checks and user impact. Define alert ownership and rollback triggers, preserve useful diagnostic records and review the underlying causes before reintroducing a changed system.



Search Under a Compute Budget

PRESENTER NOTES

A heuristic helps choose which states to explore first. Compare answer quality and work done under the same time or memory budget; a useful shortcut is not automatically guaranteed to find the best solution.

Original research

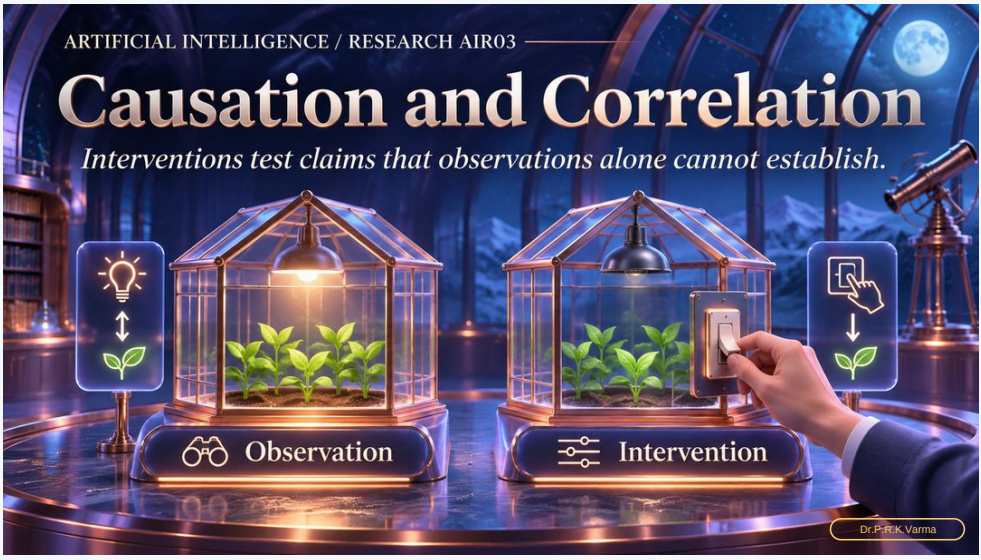


Probabilistic Reasoning

PRESENTER NOTES

A belief network models dependencies so that partial observations update plausible hypotheses. The inference is conditional on modeled relationships; uncertainty remains when evidence is sparse or the model is misspecified.

Original research



Causation and Correlation

PRESENTER NOTES

Co-occurrence alone does not show what would happen after an intervention. Use a controlled comparison when possible; otherwise name the causal assumptions that support an observational inference.

Original research

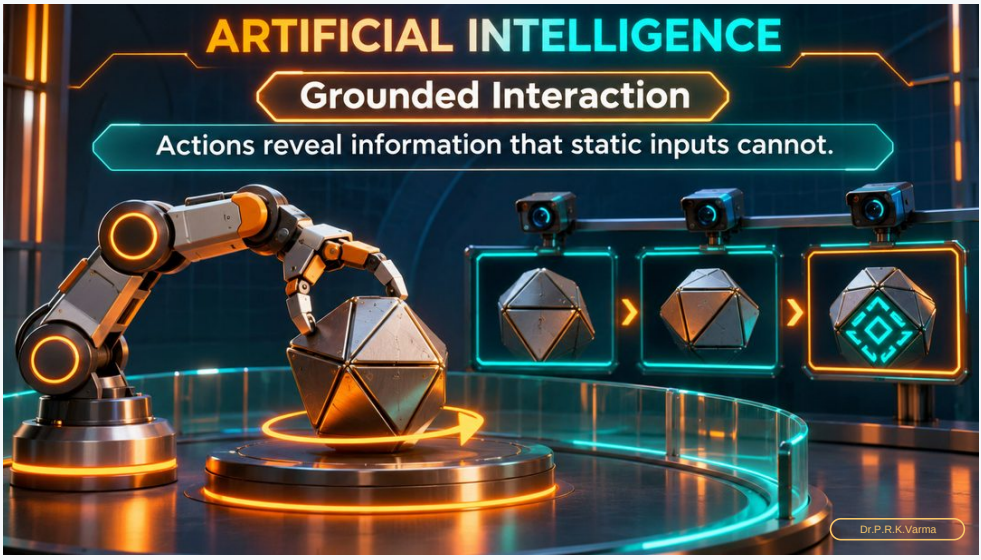


Neuro-Symbolic Generalization

PRESENTER NOTES

A neuro-symbolic system can extract learned concepts and apply structured operations to them. Hold out novel combinations to test composition rather than merely recognizing combinations already seen during training.

Original research



Grounded Interaction

PRESENTER NOTES

Embodied evaluation can test how an agent acquires task-relevant information through action. Hold the task and simulator assumptions explicit; success need not transfer to a real environment.

Original research

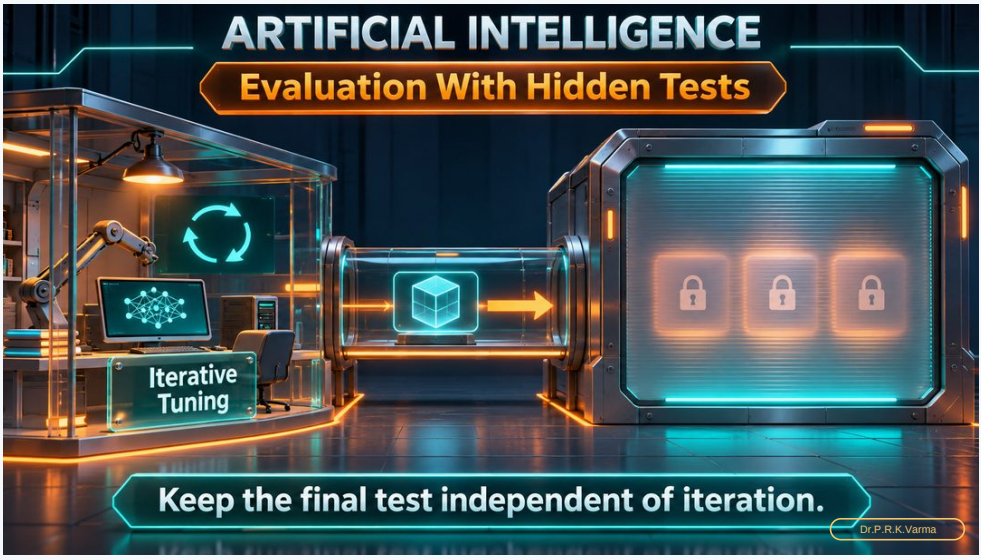


Benchmarks for Changing Environments

PRESENTER NOTES

For stream benchmarks, preserve time order and state when the environment changes. Test before any adaptive update on that example, then separate initial degradation from recovery to avoid hiding drift with an average.

Original research



Evaluation With Hidden Tests

PRESENTER NOTES

Repeated feedback from an evaluation can invite adaptation to the test. Lock down a confirmatory set and record access, tuning trials and evaluation assumptions.

Original research

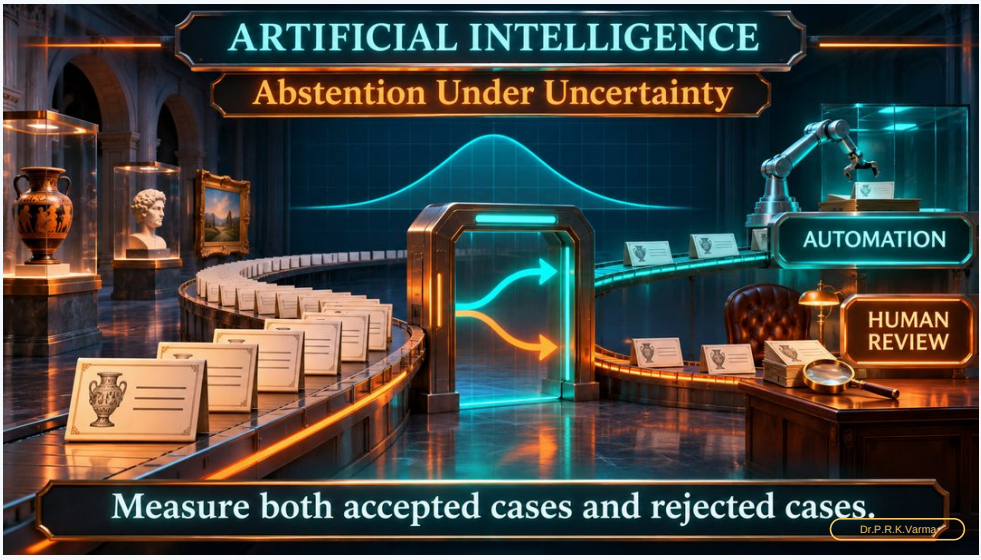


Scalable Oversight and Verification

PRESENTER NOTES

Intermediate verification can surface errors earlier than checking only a final answer. A checker can also fail; use independent test cases and compare verified outputs with trusted references before generalizing.

Original research



Abstention Under Uncertainty

PRESENTER NOTES

Selective prediction lets a system abstain. Evaluate the trade-off between coverage and error among accepted cases and check whether rejected cases overload review.

Original research



Human–AI Team Complementarity

PRESENTER NOTES

Compare the combined workflow against a person alone and AI alone on the same cases. A real collaboration benefit beats the stronger baseline, and should be checked across tasks and for failure recovery.

Original research



Cost–Quality Frontier

PRESENTER NOTES

Use the same task workload and reporting rules, then vary allowable resources and show quality across budgets. Include selection and tuning costs where relevant; a single headline score conceals trade-offs.

Original research