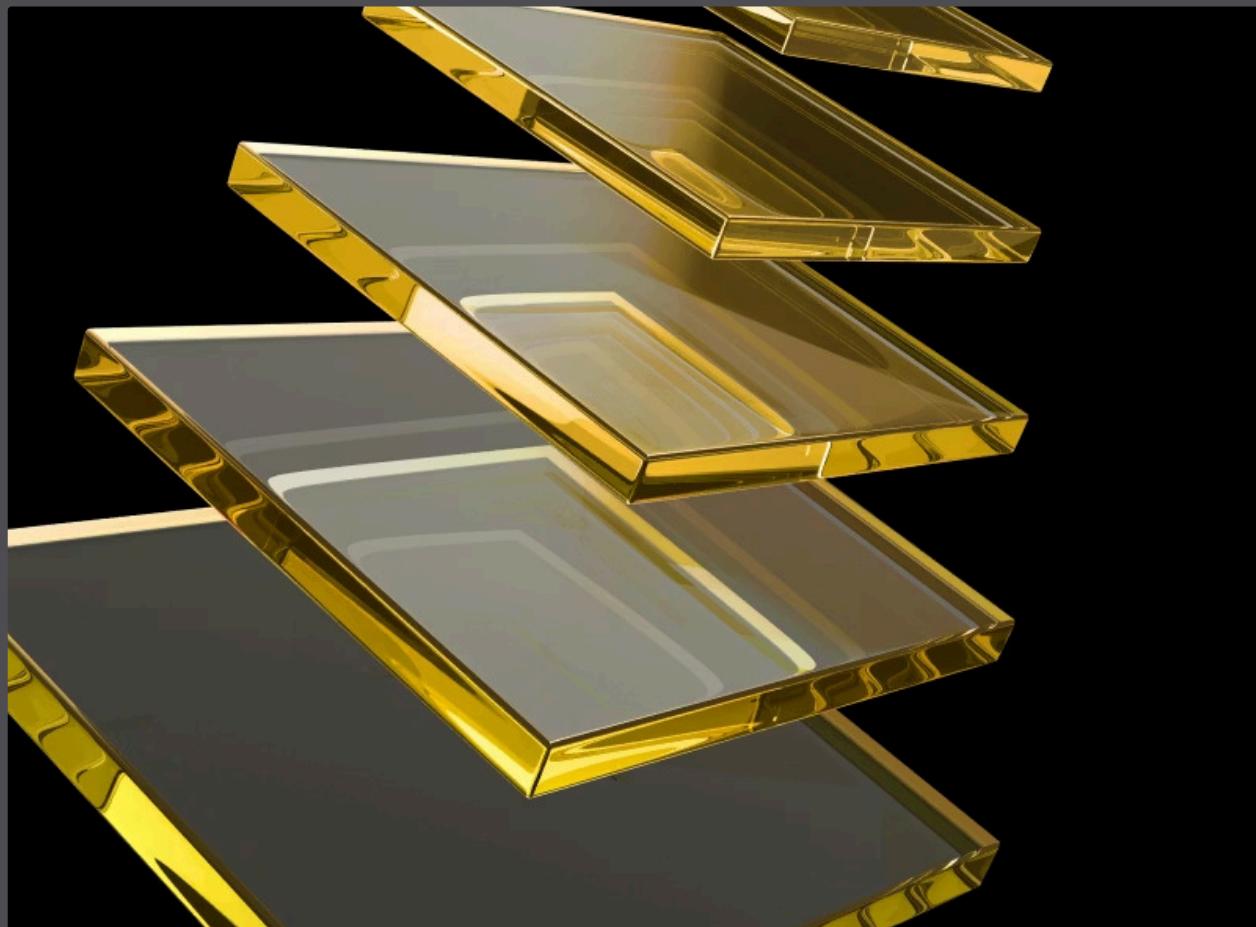


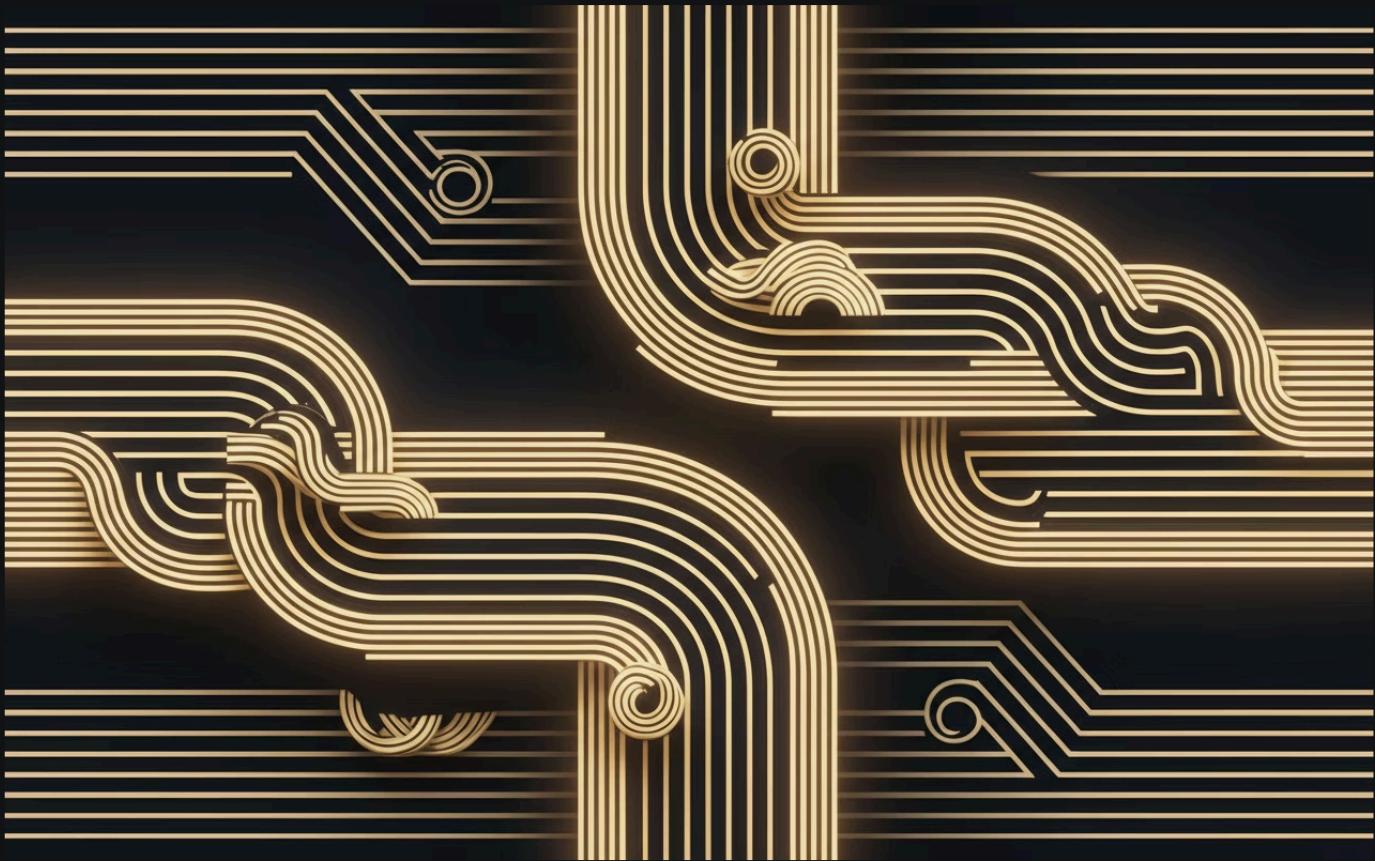
AI-Driven Verification: Accelerating Quality and Coverage in the Era of Complex Silicon



Abstract

The rising complexity of modern SoCs, AI accelerators, and advanced digital-IP subsystems has stressed traditional verification methodologies to their limits. Manual testbench creation, coverage closure, and debug consume the majority of development time and cost. This white paper explores how AI techniques—particularly generative models, reinforcement learning, and hybrid AI-formal approaches—are transforming the verification engineering landscape. We present a framework for how AI augments and automates key verification tasks, quantify the expected gains, discuss practical challenges and mitigations, and provide guidance for integration into existing flows. Finally, we introduce Moores Lab AI's VerifAgent™ as a vendor-agnostic, but production-grade example of how an AI-driven verification assistant can lower cost, shorten schedules, and elevate verification quality.

Introduction



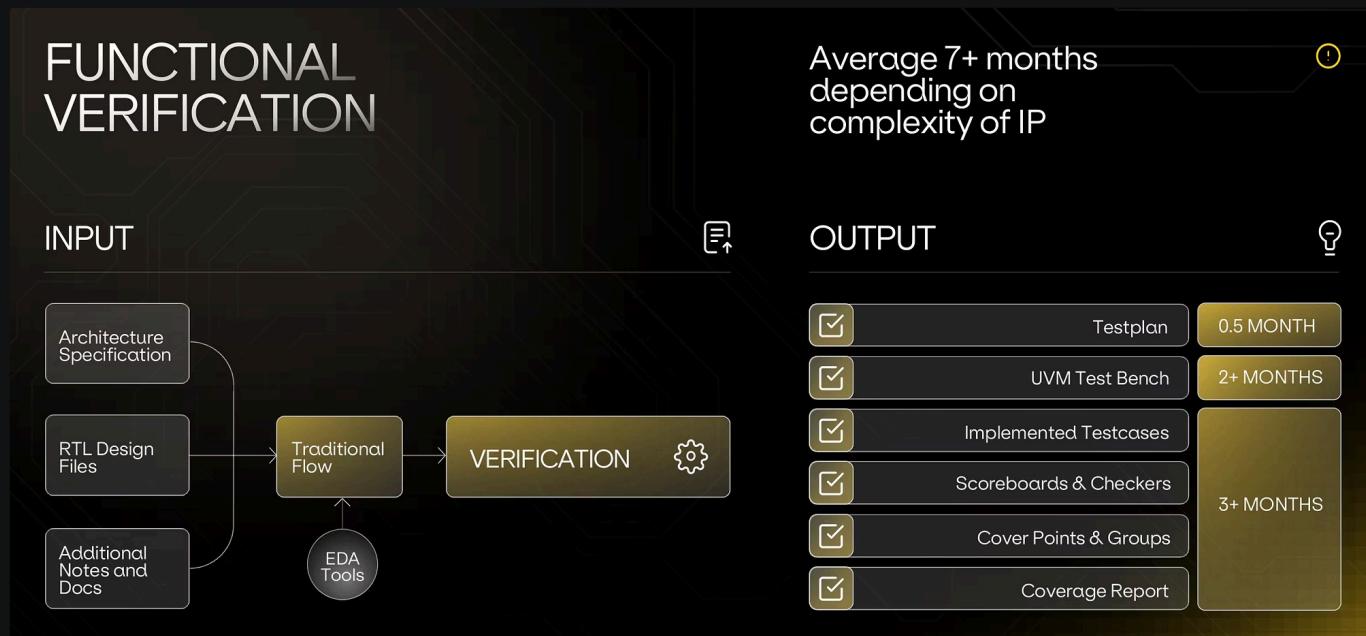
The verification burden in modern silicon design is overwhelming:

- Simulation cycles are long and regressions take hours.
- Verification teams require large headcounts to meet coverage and corner-case goals.
- Manual testbench design, stimulus writing, and debugging are error-prone and slow.
- Coverage closure (especially functional coverage) often becomes a "whack-a-bug" iteration late in the schedule.
- The talent gap is widening: it's hard to attract and retain top verification engineers at scale.
- Variation in methodology, environment, assertion coverage, and IP reuse quality means inconsistent verification quality across teams.

In many projects, verification takes more than 60% of the RTL-to-tapeout schedule, and costs can run into the tens of millions of dollars. Management often demands faster time-to-market, but verifying correctness and corner-case behavior is non-negotiable in today's safety, security, and reliability domains.

The central challenge is: How to reduce the human burden and accelerate verification without compromising coverage, quality, or debug effectiveness.

Traditional Verification Methodologies

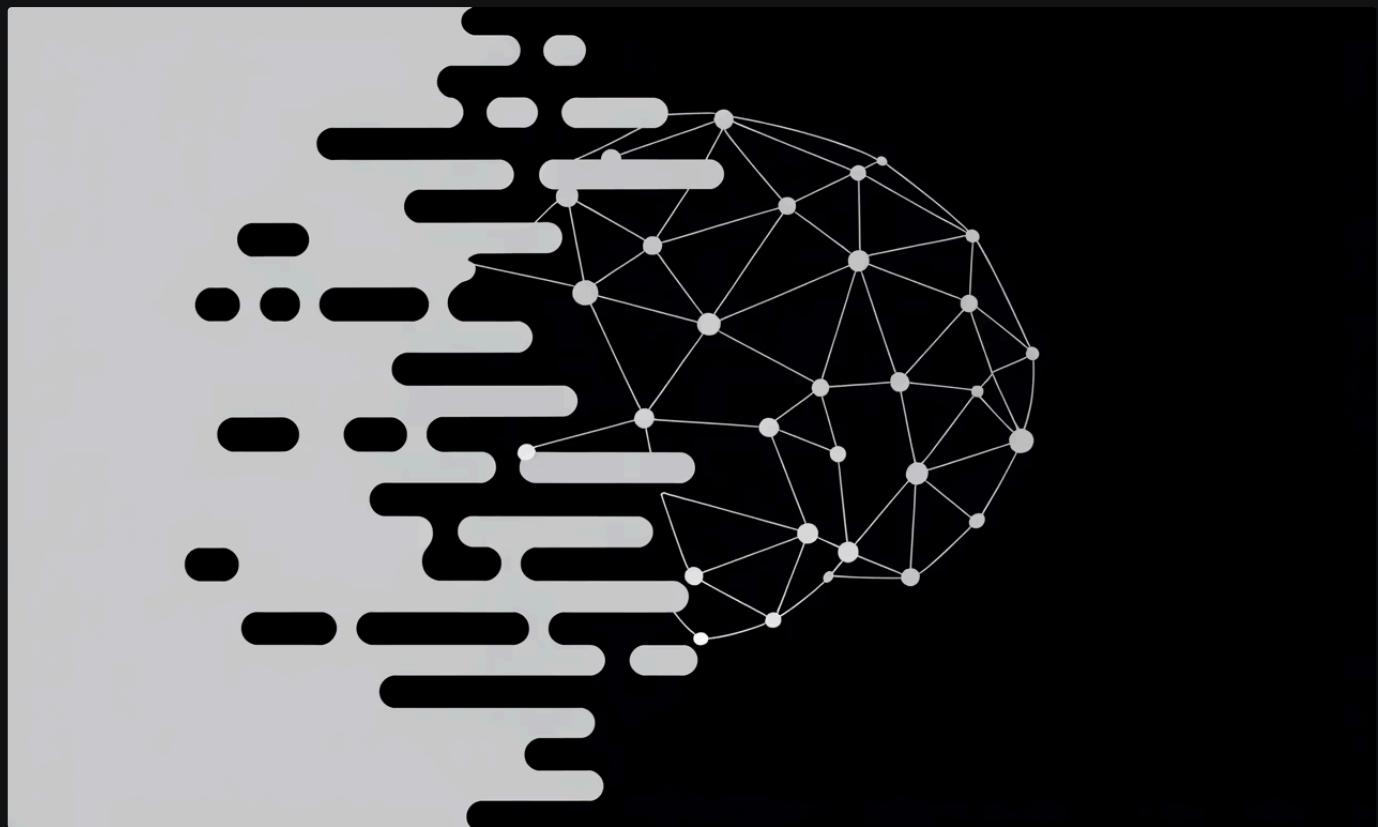


Over the past two decades, verification teams have adopted layered, methodology-driven flows:

-  Test plan » Environment » Stimulus » Coverage & assertions » Debug & refine » Signoff.
-  Methodology frameworks like UVM (Universal Verification Methodology) and its ancestors (e.g., OVM) provide structured reuse, sequence layering, transaction-level abstraction, factory overrides, and constrained-random stimulus.
-  Formal and equivalence checking are used for certain data paths, corner-case assertions, or design equivalence (RTL vs higher-level model).
-  Hardware acceleration / emulation (e.g. ZeBu) helps with system-level handoff and firmware-driven verification.
-  Incremental regression and coverage feedback loops; human engineers review coverage and decide what strategy to use to close coverage holes.

These flows have matured, but they still heavily rely on human insight: deciding stimulus strategy, writing corner-case sequences, managing coverage dependencies, triaging failures, and debugging interactions among modules.

Early AI / ML in EDA & Verification



Over the last few years, AI and ML have begun to appear in EDA tools in limited roles:

- Synopsys introduced VSO.ai and TSO.ai to accelerate coverage closure and test generation, leveraging machine learning to predict coverage holes or optimize test patterns.
- AWS published generative AI approaches to design assistants and prompting-based exploration of design alternatives.
- Academic work, such as MAVF (Multi-Agent Generative Verification Framework), demonstrates that dividing verification tasks across specialized AI agents (spec parser, strategy generator, code emitter) outperforms single-LLM approaches in generating testbenches and verification artifacts.
- Other academic approaches, such as generative induction in formal verification, show how AI can assist or accelerate proof search.
- **The Saarthi agent demonstrates** a proof-of-concept "autonomous AI formal verification engineer" that can take RTL and verify end-to-end under a constrained domain.
- Intelligent coverage closure strategies (e.g. auto-feedback of coverage holes to stimulus generation) have been used to reduce manual effort. This is sometimes referred to as "intelligent verification."

These developments suggest that AI/ML can be more than a helper; they can become an integral co-engineer in verification.

The AI-Driven Verification Solution

This section describes how AI techniques can be applied to key subdomains of verification, how they work in concert, and what benefits accrue (both technical and business). Note that these verification solutions apply exclusively to UVM testbenches.

Task Decomposition: Where AI Fits

Below, the verification problem is broken into subdomains and shows the AI augmentation opportunities for each one:

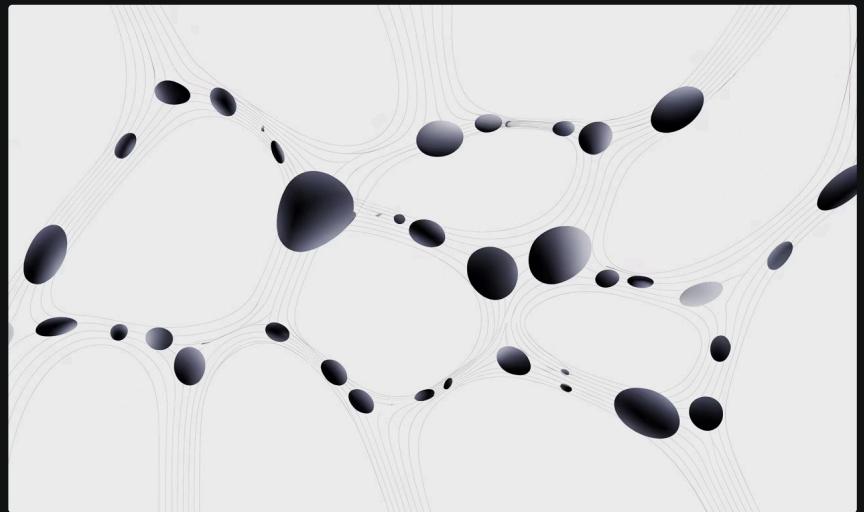
Subdomain	Traditional Process	AI-Augmented / Automated Role
Specification / Requirements Parsing	Manual reading, deriving features, corner-case identification	Use vLLMs or domain-tuned models to parse spec documents (protocols, state machines) and propose candidate coverage items, edge-case scenarios, or protocol fuzzers
Test Plan Generation / Strategy	Human crafts a plan with constrained-random, directed tests, corner-case enumeration	AI suggests prioritized test plan templates, gap-driven tests, scenario permutations, and cross dependencies
Testbench / Environment Code Generation	Manual scaffolding of monitors, drivers, scoreboard, interface infrastructure, constraints	Agent-based generative-AI emits transaction-level or UVM-style boilerplate, connectivity wrappers, assertion skeletons, environment scaffolds

Stimulus / Sequence Synthesis	Writing constrained-random sequences, directed corner-case sequences, scenario-based stimuli	AI-driven sequence generation targeting coverage holes, edge-case generation, reinforcement-learning based adaptive stimulus
Coverage Analysis & Closure Guidance	Engineers analyze coverage reports, identify unhit bins, infer dependencies, iterate	AI clusters coverage holes, suggests stimulus or constraint modifications, predicts which holes are feasible toA improve the coverage
Bug Triage / Debug / Root-Cause Analysis	Engineers trace failing tests, isolate signal-assertion interactions, collaborate with RTL designers	AI-assisted failure clustering, trace-to spec mapping, likely root-cause ranking, hints for additional assertions or exclusions
Formal Assistance & Hybrid AI-Formal	Formal proofs or bounded model checking performed with manual abstraction/hinting	AI can propose abstraction, lemmas, induction invariants, or direction heuristics to help formal engines scale
Regression Prioritization / Risk Estimation	Human picks which test regressions to run first	ML models score regressions' likely coverage yield or bug-finding power, schedule regressions accordingly

Together, these form a "stack" of AI-augmented verification capabilities.

Architecture: Multi-Agent & Feedback Loops

A practical AI verification system should be architected as a multi-agent system rather than a monolithic LLM. AI agents perform functions such as:



Specification

Processes natural-language spec, extracts state machines, transactions, legal behaviors

Planning

Translates spec features into coverage plan and test goals

Generation

Emits testbench code, sequences, environments

Stimulus

Dynamically generates or mutates stimulus during regression

Feedback

Ingests coverage/bucket data and suggests improvements

Debugging

Processes failures, signals, assertions and suggests root-cause hypotheses

These specialized agents communicate, refine each other's outputs, and loop until a verification closure target is reached (or diminishing returns kick in). MAVF, for example, shows better performance than single-LLM pipelines.

The system should include feedback loops from coverage results and debug outcomes back into the AI agents. As regressions run, new data is ingested (e.g. coverage maps, assertion hits, bug history), and the AI adapts for better future generations.

Integration into Existing Flows

To be realistic and adoptable, the AI stack should be:

Vendor-agnostic

Integrate via script interfaces, open APIs, or adapters to standard EDA tools and flows

Prompt-free / Minimal prompt engineering

Users shouldn't need to reverse-engineer prompts or micromanage the AI

Reviewable & controllable

Engineers should be able to review and override AI suggestions

Incrementally adoptable

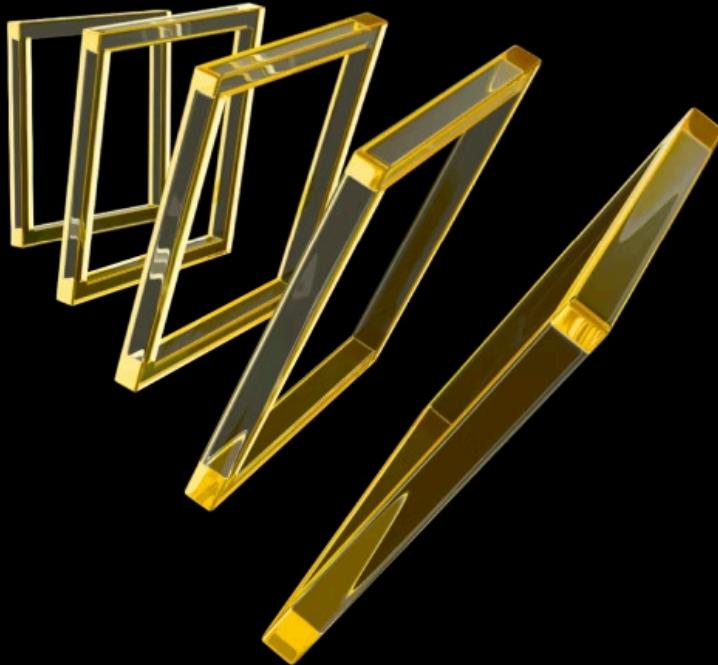
Begin with assistance in a subset (say, testbench generation or coverage closure) before scaling

Secure & confidential

Models should run on-prem or in secure enclaves, not require IP exposure to public clouds

Performance-aware

Generation should scale for large modules, with modularization, streaming, and caching



Quantified Benefits & Empirical Gains

Based on industry claims, internal studies, and academic benchmarks, here are realistic expectations:

Metric	Typical Baseline	Expected Improvement
Testbench & environment scaffolding time	Weeks to months	10×–30× faster
Coverage closure cycle count	Many nested loops	Reduction by 30–50% or more
Regression runtime utilization	Conservative ordering / serial	Better prioritization, fewer wasted runs
Bug-detection yield earlier, including documentation errors	Late-stage surprises	Shift-left coverage and more bugs caught earlier
Human headcount burden	Large teams per IP	Up to 50% fewer manual engineers needed
Cost savings	High verification spend	Reductions in cost up to 40–70% (depending on IP complexity)

Moores Lab AI VerifAgent™ customers have experienced up to 7× faster time-to-market and 86% cost reduction on IP verification. In the academic domain, MAVF shows competitive or superior results vs manual approaches for multi-module verification. Other industry sources (e.g., Tessolve) highlight manual-effort reductions and better coverage closure via AI + formal techniques.

Risks, Limitations & Mitigations

No AI system is magic; prudent engineering is required to mitigate risks:

Over-reliance / complacency



Engineers may accept suboptimal or incorrect generated tests.

Mitigation: require human-in-the-loop review, offer transparency, and ensure code is auditable.

Model generalization failures

Edge-case domains or niche IP may lie outside training experience.

Mitigation: allow domain-specific fine tuning and feedback learning.

Scalability and runtime

Model generation latency might be high for large modules.

Mitigation: modular generation, caching, parallelism, streaming generation.

Debug complexity

Generated testbench may produce failures that are hard to interpret.

Mitigation: couple debug agent, logging, and trace mapping.

False coverage confidence

There's a risk of treating AI-generated coverage as "complete" prematurely.

Mitigation: retain manual coverage goals, assertion-based guardrails, and sanity checks.

Tool & methodology compatibility

Generated code may not match local coding style or methodology.

Mitigation: provide adapters, style templates, constraint configuration.

Security / IP leakage

While not an AI risk per se, the use of cloud LLMs could expose proprietary RTL or specifications. *Mitigation:* on-prem/private deployment, federated models, encryption.

With careful safeguards (especially reviewability and feedback), these risks are manageable, and the benefits outweigh them in most IP-level flows.

Conclusion

The semiconductor industry is under intense pressure to shrink schedules, reduce headcount, and maintain high assurance in increasingly complex IPs. Verification engineering is arguably the greatest bottleneck in many tapeout flows. AI offers a transformative path — not by replacing verification engineers, but by *augmenting* them, reducing manual toil, and accelerating coverage convergence.

To those leading verification teams and their management, the message is: start small, prove value, then scale. Begin by adopting AI assistance for modules where testbench scaffolding or stimulus generation is expensive. Use the ROI to build confidence and expand the role of AI agents deeper into your flow.

Your road map might look like:

01

Pilot AI-assisted testbench generation on a mature IP.

02

Compare auto-generated vs human stimulus in functional coverage and bug yield.

03

Incorporate coverage-feedback agents to suggest missing scenarios.

04

Add debug / triage agent for failure analysis.

05

Expand to cover hybrid formal/AI for difficult corner-case properties.

06

Finally, adopt the full multi-agent AI verification stack across multiple IPs in parallel.

The urgency is real: companies that adopt AI early will out-compete in silicon delivery. Verification teams should be proactive, not reactive, in embracing this shift.

About Moores Lab AI



Moores Lab AI is dedicated to bringing AI-native automation to silicon verification. Our initial product, VerifAgent, is an agentic AI verification assistant that integrates into standard EDA flows and automates test plan generation, testbench and stimulus creation, assertion scaffolding, coverage guidance, and debug assistance.

Some of the benefits our VerifAgent™ customers have experienced include:

7x

faster time-to-market

(i.e. reducing verification cycle time)

86%

cost reduction

in IP verification expenditure

- Seamless integration with existing flows (vendor-agnostic, prompt-free)
- Automated generation of full UVM testbenches, coverage models, and assertions delivered in hours rather than weeks
- Corner-case bugs that had previously escaped human test benches caught in "preverified" IPs
- The choice of on-prem or cloud deployment

We provide not just software but an IP verification package (test plan, test bench, test cases with coverage) via a structured engagement flow. Other agentic AI silicon engineering products are under development and on our roadmap.

If you're leading a verification group and evaluating AI adoption, contact us for a pilot on a moderately complex IP and benchmark the results against your internal metrics. Consider broader deployment once validated.

Visit: www.mooreslab.ai