AI+EDA for Design Verification Engineers

(with examples for RISC-V processors)

Simon Davidmann

AI+EDA Researcher, School of Electronics and Computer Science

Southampton University, UK

DVClub, Bristol, UK

Wed 28 Jan 2026

Note: some of this data may be incorrect as I used AI for some of the research… (and it kept making things up)

“Artificial Intelligence **isn’t just another tech trend** — it’s a seismic shift redefining industries and improving lives. From energy production to healthcare breakthroughs to climate resilience, AI promises solutions to some of society’s toughest challenges”

- Article on AI at CES26

“Artificial Intelligence **isn’t just another tech trend** — it’s a seismic shift redefining industries and improving lives. From energy production to healthcare breakthroughs to climate resilience, AI promises solutions to some of society’s toughest challenges”

- Article on AI at CES26

I believe AI / ML is going to completely transform the EDA industry

Flow of this presentation

- AI?
- Where is AI in EDA
- RISC-V trends
- AI needs
- We need more data

AI (Neural Nets) since 1980's in UK with HW*

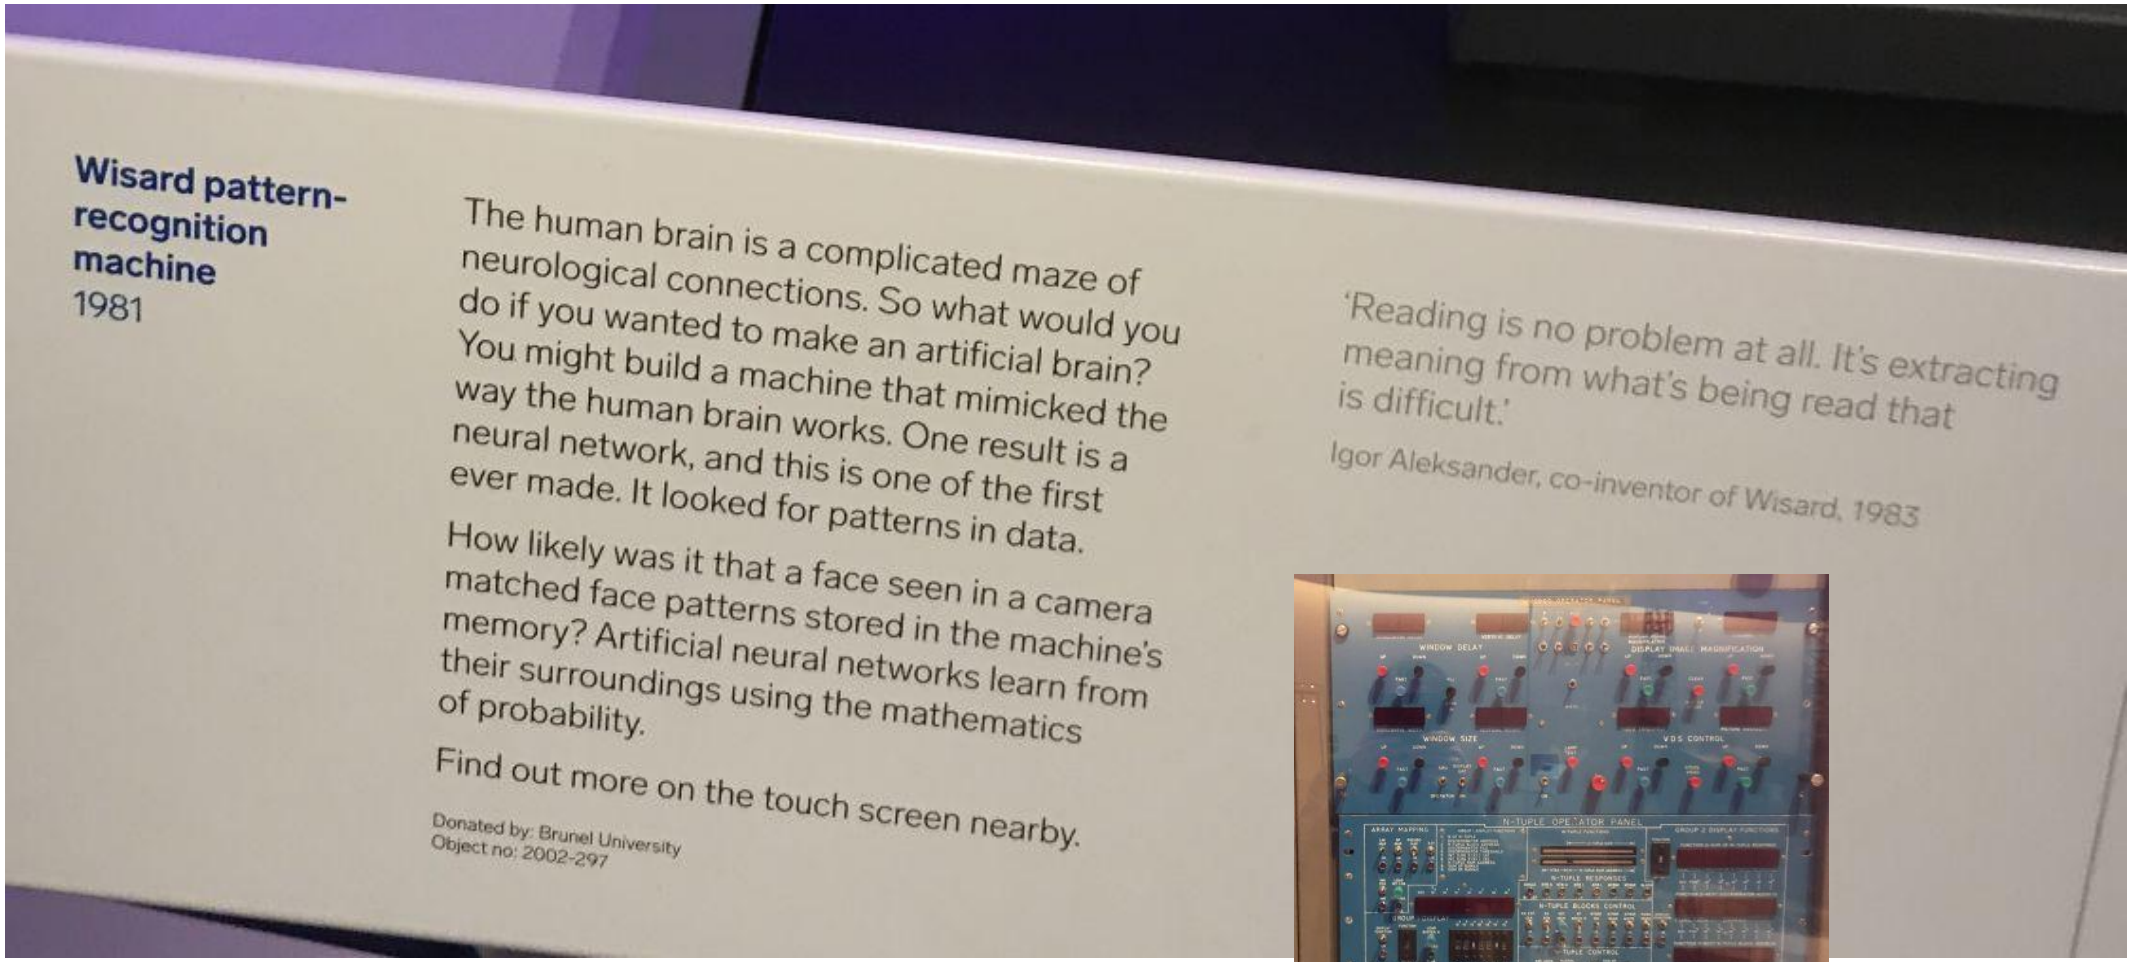

*While I was in simulation research at Brunel Univ., Prof. Aleksander had office down corridor…

In Science Museum, London

Fast Forward to 2025…
The world of AI is now moving very quickly…

‘The Era of AI Employees is Here’…

Bus stop advert outside hotel in San Francisco when visiting DAC62 June 2025

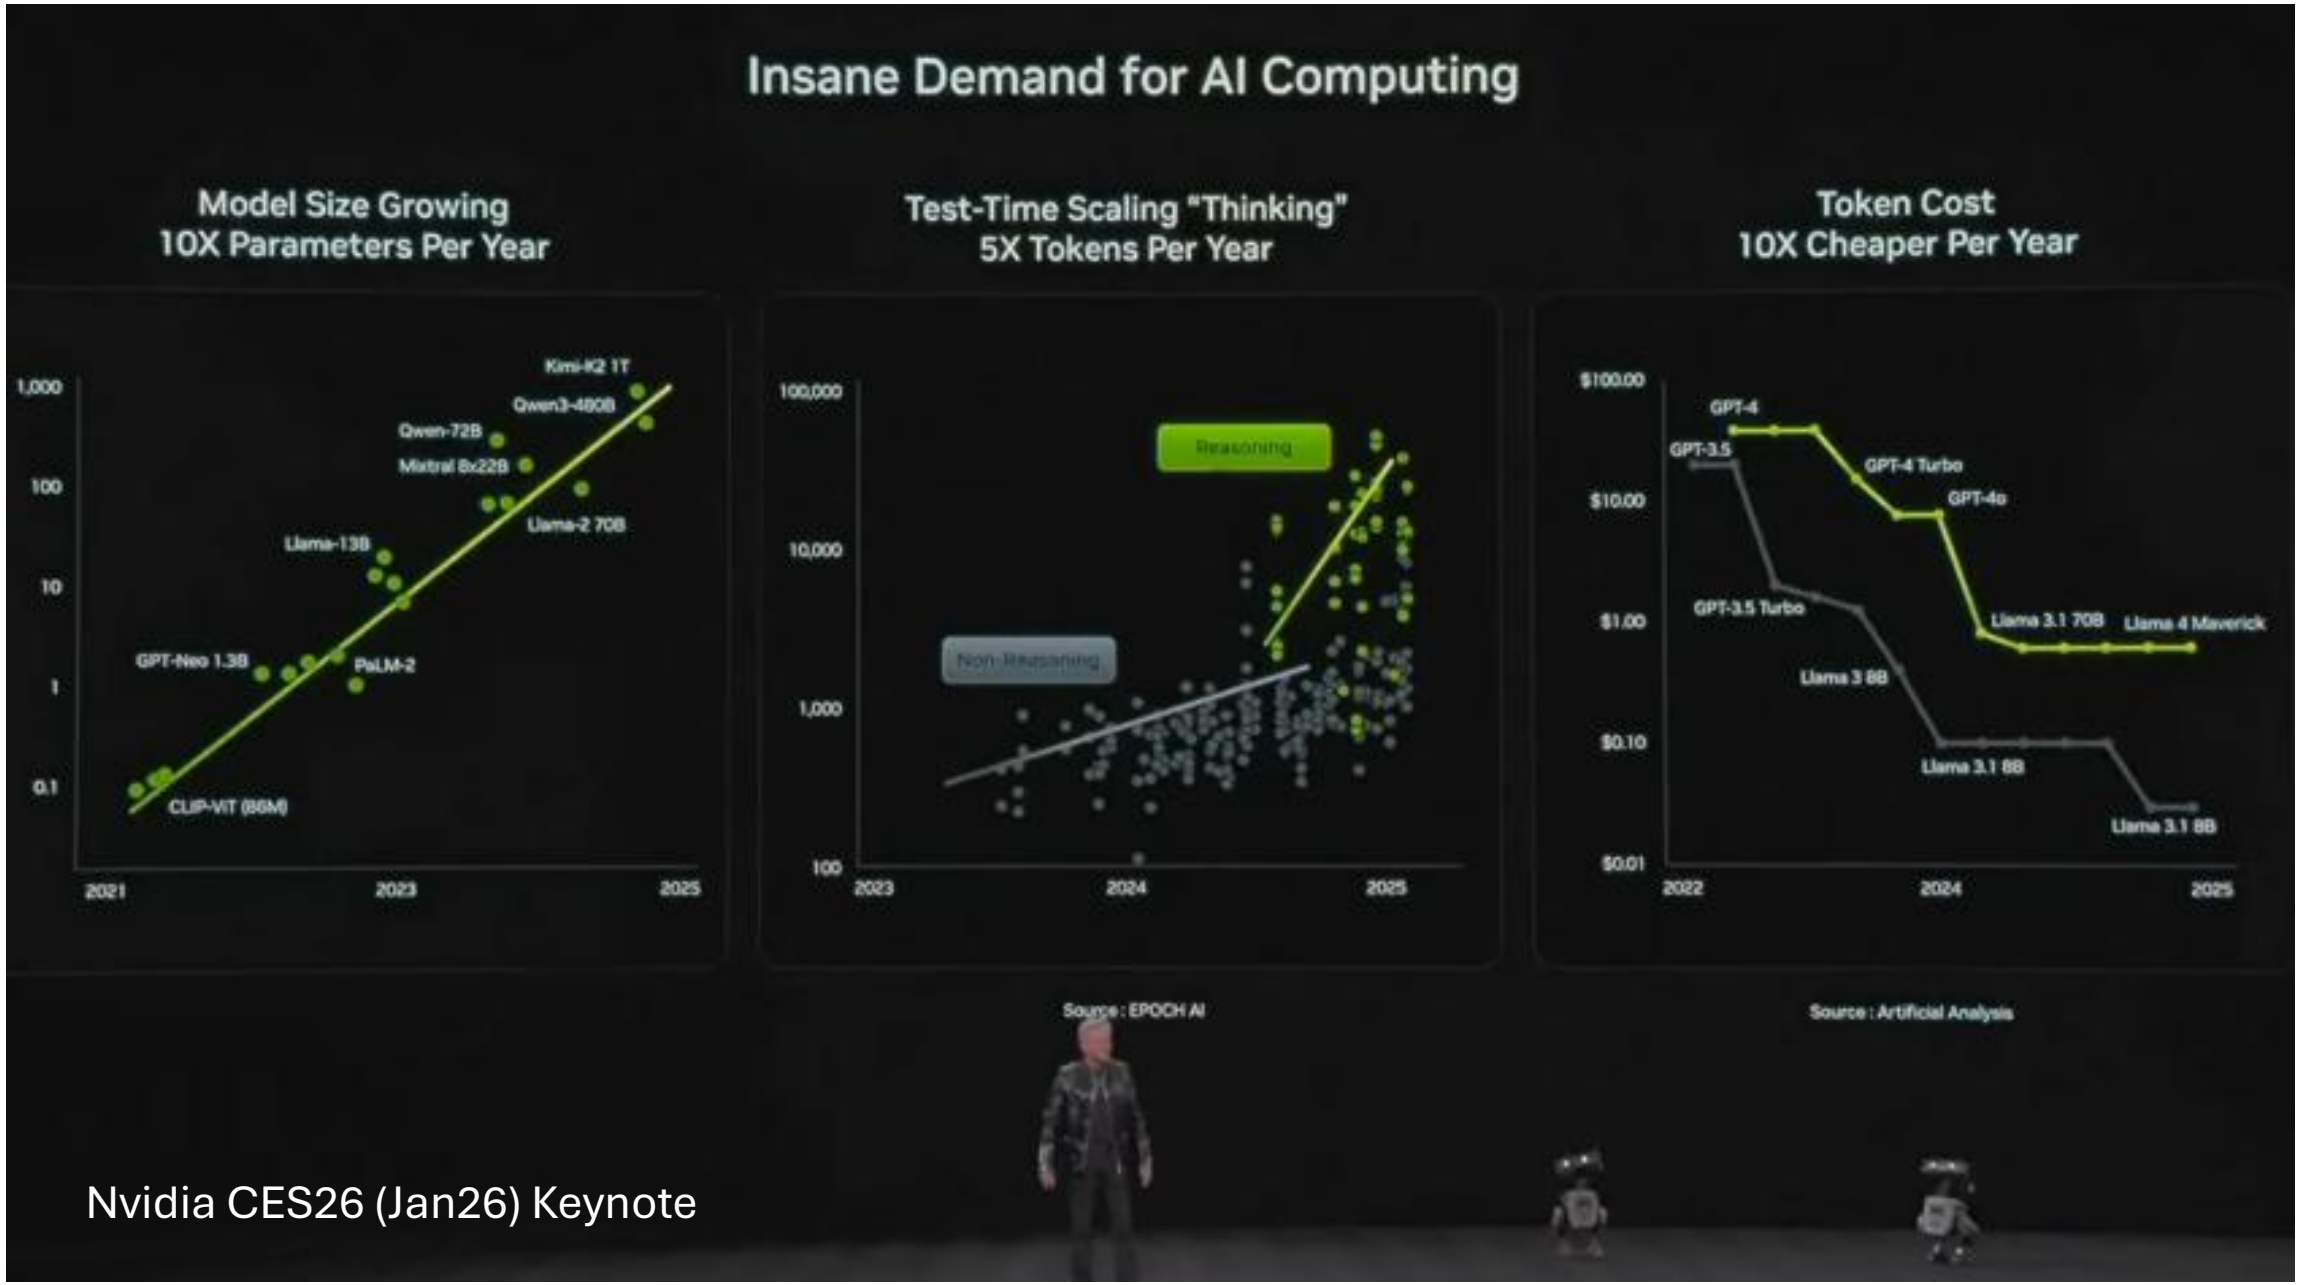

Nvidia CES26 (Jan26) Keynote

Note: these are all log scales...

University of Southampton

A current (recurring?) question…
Will AI replace us?

- “AI won’t replace humans, but humans who use AI will replace humans that don’t”
    - before 2020, by either Prof. Fei Fei Li, Stanford or Prof. Karim Lakhani, Harvard

- Augmentation, Not Replacement
    - The core message is that AI is a powerful tool, not a substitute for human ingenuity, but adept **users will gain a significant advantage**
- Evolution of Work
    - Instead of eliminating jobs, AI reallocates human effort to **higher-value tasks** requiring judgement, empathy, and creativity, making uniquely human skills more valuable
- Skill Development
    - It's a call to action for professionals and organizations to **learn and integrate AI**, much like past technological shifts

Agenda

- AI?
- Areas of EDA that could benefit from AI
- Status (Dec 25) of EDA products (market positioning) & research
  - From main 3, From startups, From research, From AI providers
- RISC-V processor verification – and AI…
- Some of the challenges going forward
  - Data for EDA
- Futures…

Where AI can help:
Predict, Optimize, Generate, Assist

- Predict (explore alternatives)
  - AI: Estimate QoR, risks
  - Method: supervised ML, GNNs (DL: Graph NN), CNNs (DL: Convolutional NN)
  - Example: **testcase failure prediction, bug probability, stimulus coverage accuracy**
- Optimize (physical design & verification have huge search spaces)
  - AI: search large design space under constraints, agentic
  - Method: RL (Reinforcement Learning), BO (Bayesian Optimization), evolutionary search
  - Example: PPA, macro placement, **test case redundancy, guide test generators**
- Generate ((take care…) need to confirm correctness)
  - AI: produce candidate artifacts under constraints, agentic
  - Method: LLMs, diffusion, program synthesis, copilots
  - Example : **script stubs, SystemVerilog/UVM scaffolding, constraint templates, SVA**
- Assist (lowest risk entry as outputs can be advisory)
  - AI: augment human workflows, decisions, agentic
  - Method: LLMs+ retrieval classifiers, copilots
  - Example: **triage, summarization, documentation**, Q&A

Challenges for AI with EDA

- Hard constraints
  - Many outputs must satisfy non-negotiable constraints, (e.g. ISA rules)
- Expensive labels
  - ‘ground truth’ often comes from long tool runs or measurement on silicon
- Constant changes
  - New blocks, new architectures, new methodologies change data distributions
- High cost of error
  - Confident but wrong recommendations can create sign-off risk
- Data availability for training
  - <discussed later>

Agenda

- AI?
- Areas of EDA that could benefit from AI
- Status (Jan 26) of EDA products (market positioning) & research
  - From main 3, From startups
  - From research, From AI providers
- RISC-V processor verification – and AI…
- Some of the challenges going forward
  - Data for EDA
- Futures…

What is the status of AI in EDA…

- Have been papers on ML in EDA for over 20 years
- Big 3 EDA started including it in products for optimization 5-8 yrs ago
- 2022 ChatGPT got everybody excited… and it evolved into ‘copilots’
- 2025 – DAC – ‘the Verilog Moment’ for AI
  - with 49 sessions, 32% of papers, and most companies talking about AI
- 2025+ many have moved from adoption of co-pilots to starting use of agents*…

  * Agents perceive environment, reasoning over complex goals, taking multi-step actions

Nvidia CES26 on Cadence

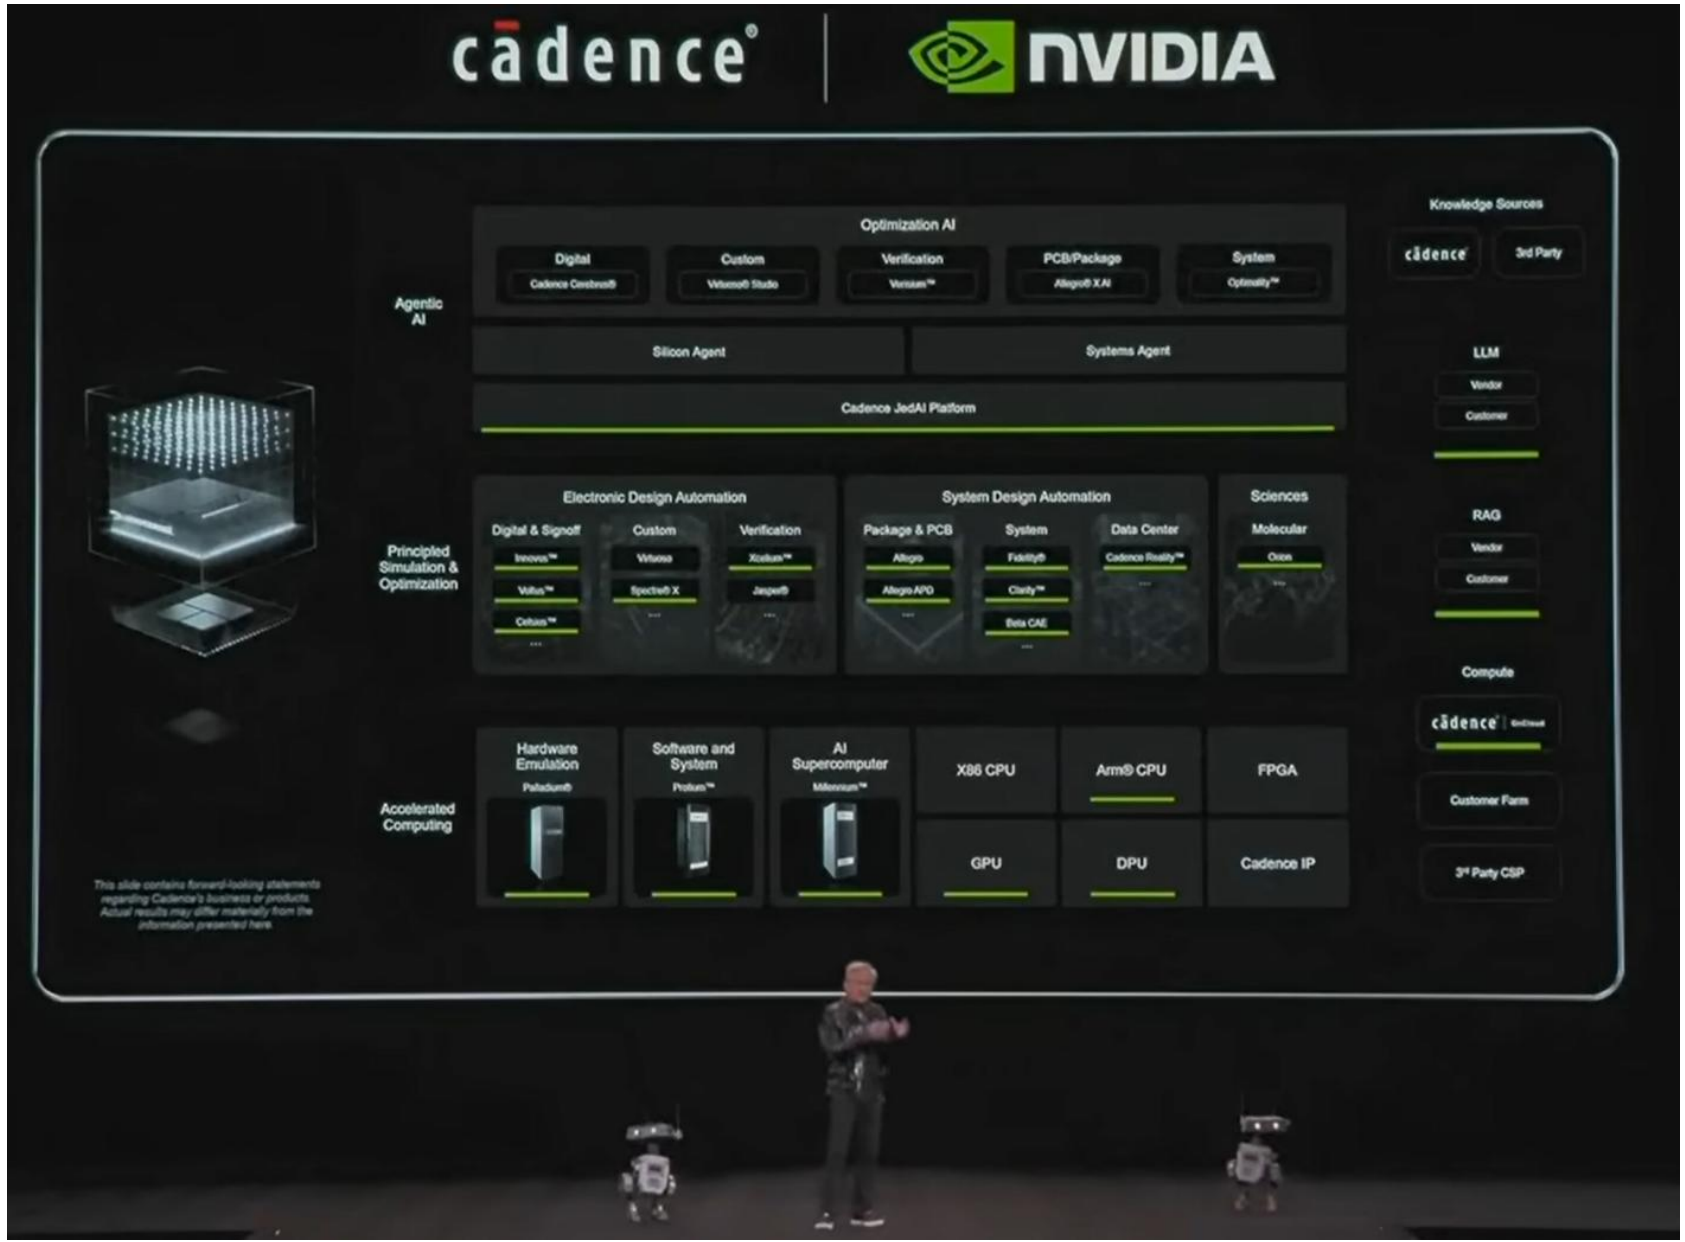

Cadence:

- JedAI (was ChipGPT)
- Verisium
  - CodeMiner
  - PinDown
  - WaveMiner
  - AutoTriage
    AutoFocus
  - SmartRun
  - Copilot
  - SpecMiner
  - …
- DebugAI
- SimAI / XceliumML
- FormalAI / JasperGold
- ChipStack

…

University of Southampton

Nvidia CES26 on Synopsys

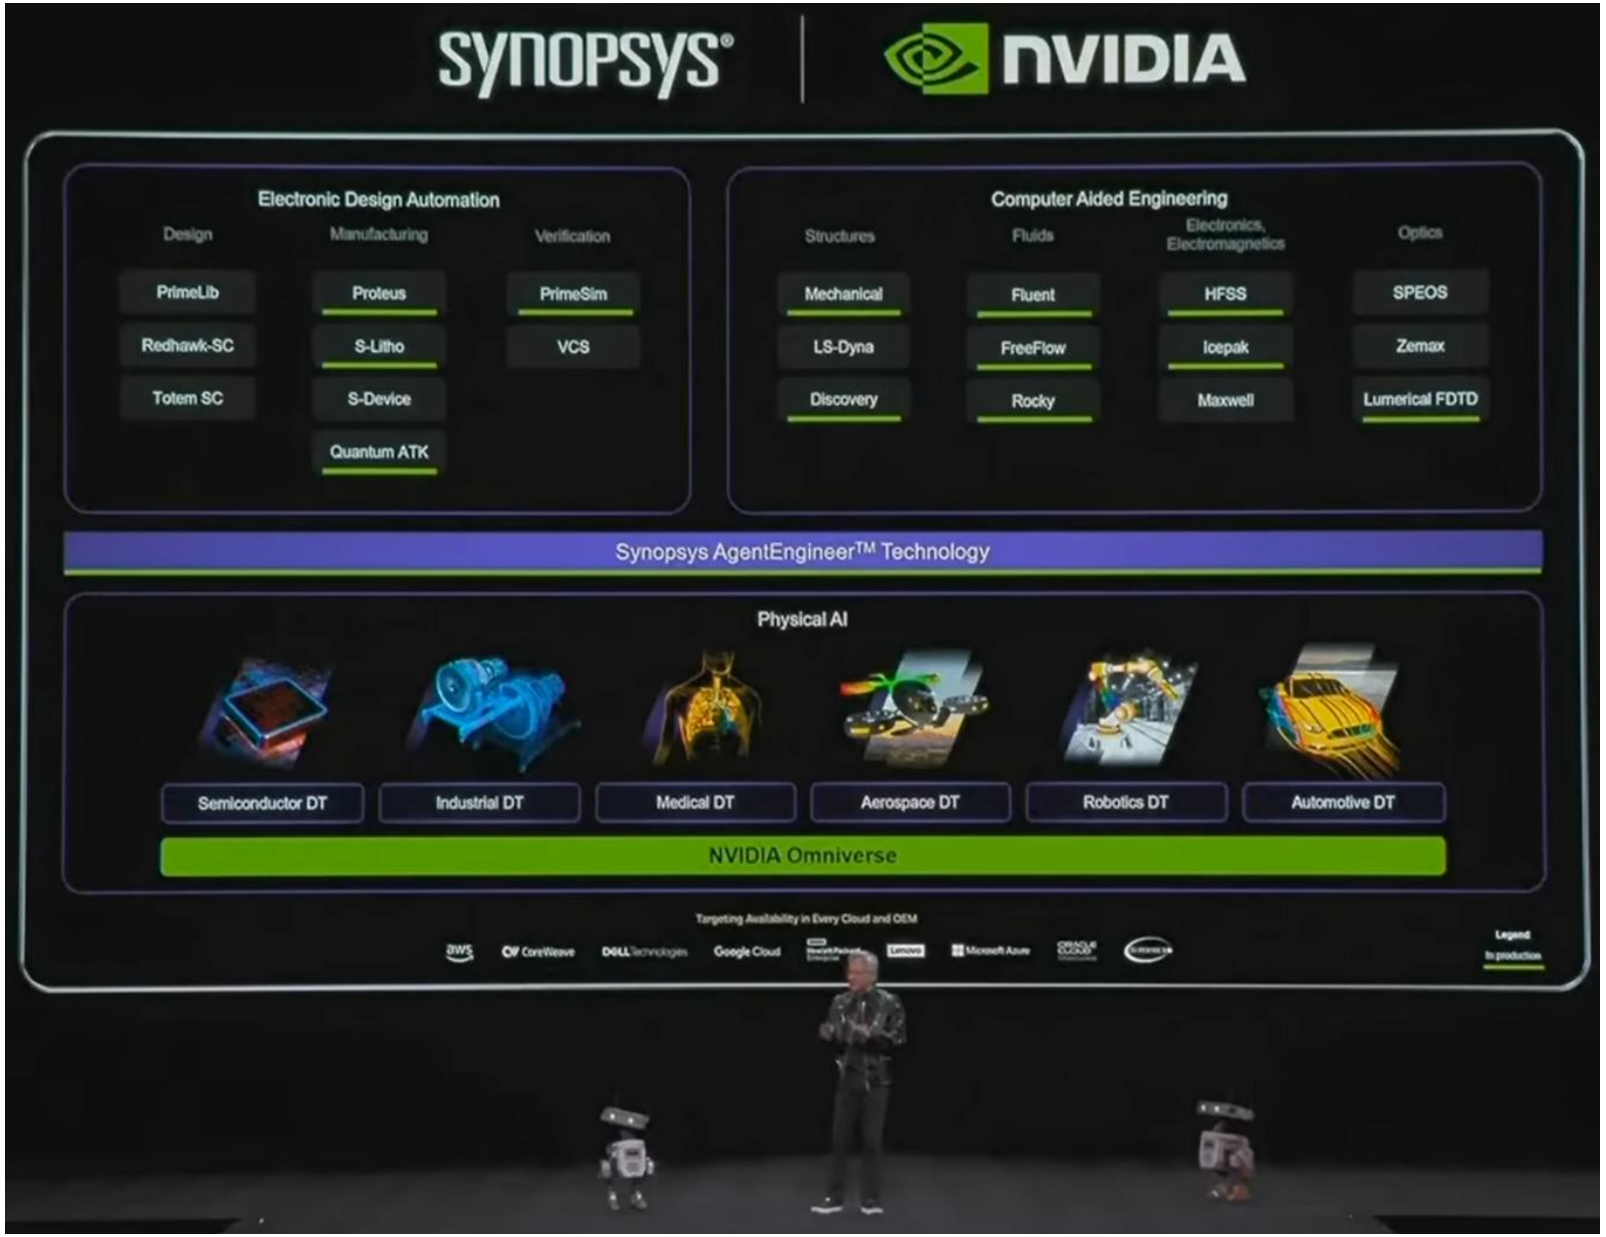

Synopsys.ai:

- Copilot
- VSO.ai
- VCS
- Verdi
- DSO.ai
- TSO.ai
- 3DSO.ai
- ASO.ai

…

University of Southampton

Nvidia CES26 on Siemens EDA

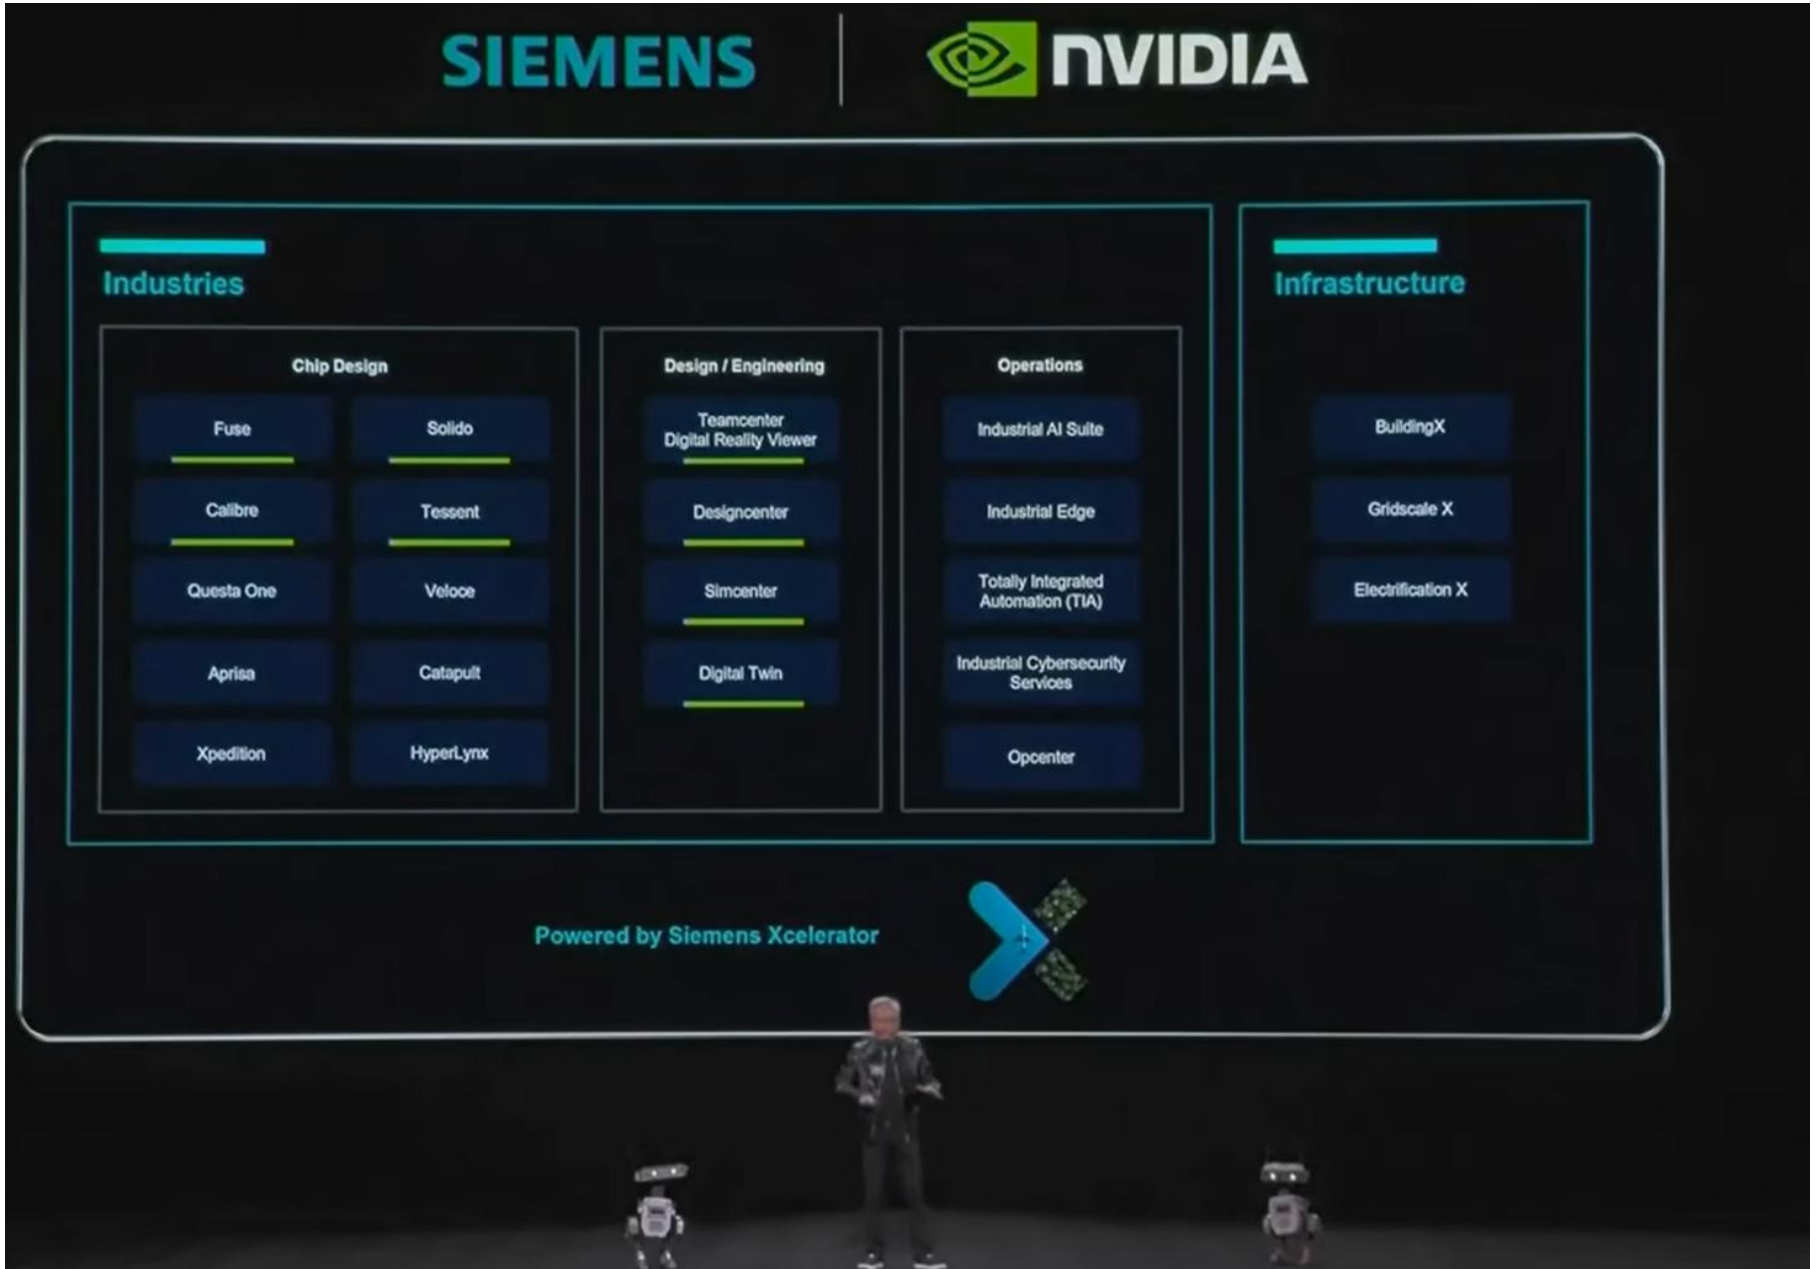

Siemens EDA:

- Questa One
- Stim Free Verif
- Regress Navigator
- Coverage Analyzer
- Avery Verif VIP
- Questa OneSpin
- Veloce CS
- Nvidia NIM

...

University of Southampton

EDA tools with AI

- Big 3: getting very familiar and comfortable with adopting / promoting AI
  - Synopsys.ai Copilot, VSO.ai, VCS, Verdi, DSO.ai, TSO.ai, 3DSO.ai, ASO.ai, …
  - Cadence.ai JedAI, Verisium (AutoTriage, SemanticDiff, WaveMiner, …), XceliumML, JasperGold, ChipStack, …
  - Siemens EDA QuestaOne VerificationIQ, SmartCreation, …

- Established (smaller) Verification companies:
  - Breker VerificationOS (AI Planner, …)
  - Lubis EDA (smart prompting, LLM orchestration, formal tools)

Several EDA+AI startups & examples of their focus

- PrimisAI – **genAI agents** for copilot to create RTL for FPGA and ASIC
- ChipAgents – **agentic** root-cause analysis from waveforms, spec to testbench
- ChipStack (Cadence) – **agentic** test bench verification
- BroncoAI – AI based **regression analysis** and simulation debug
- MooresLabAI – **agentic** AI to automate test plans with tool orchestration
- Silimate – **co-pilots** for bug finding and PPA estimation
- Rise Design Automation – raise abstraction with C++ and hardware focused **agents**
- VerifAI – RL with multiple LLMs – an **AI test engineer**
- …
- CogniChip - AI-enabled chip design, 10 levels of AI-driven automation
  - **creating a new segment**… between fabless companies and EDA companies
- …

Note – they are pretty much all '**agentic AI**'

Agenda

- AI?
- Areas of EDA that could benefit from AI
- Status (Jan 26) of EDA products (market positioning) & research
  - From main 3, From startups
    - What are Agents and what is “Agentic AI”
  - From research, From AI providers
- RISC-V processor verification – and AI…
- Some of the challenges going forward
  - Data for EDA
- Futures…

Change in 2025: Agentic AI and Agentic Automation

- **New class of autonomous AI systems** designed to act and react independently to achieve complex, high-level goals
  - genAI -> creates content in response to prompts
  - agenticAI -> uses LLMs as a central "brain" to plan, **use tools**, and execute multi-step workflows
- Characteristics
  - Autonomy
    - Operate independently without user oversight, **making decisions, pursuing defined objectives**
  - Proactivity
    - Anticipates needs and **initiates tasks on its own** rather than waiting for direct user commands
  - Goal Orientation
    - Breaks down a **high-level outcome** (e.g., "increase coverage") into smaller, manageable sub-tasks
  - Adaptability
    - Continually learns from its environment and past outcomes, refining its strategy through a cycle of **perception, reasoning, action, and reflection**

=> They don’t just react or follow preset rules - act with **autonomy, initiative, and adaptability**

AI Agents in EDA

Benefits

- Conversational Control
  - "Simulate the current design and tell me why the timing fails on the reset path,“ - run the tools and provide a technical summary
- Autonomous Scripting
  - Replace rigid template scripts with dynamic sequences based on the design's unique requirements
- Extensibility
  - Add support for other tools by defining new "tools" within the framework

Impact

- Transition from execution-focused roles to architectural-oversight roles
- Instead of writing test bench / code sequences use agents to create and execute actionable steps to achieve goals
- Use agents to iterate through simulation failures, triaging, suggesting or fixing until goals met
- Use agents to optimize regressions (test optimization, test sequencing)

Recap new EDA startups using AI – most were at DAC 2025
- They are pretty much all ‘agentic AI’

- PrimisAI – genAI agents for copilot to **create RTL** for FPGA and ASIC
- ChipAgents – agentic **root-cause analysis** from waveforms, spec to testbench
- ChipStack (-> Cadence) – agentic **test bench verification**
- BroncoAI – AI based **regression analysis** and simulation debug
- MooresLabAI – agentic AI to **automate test plans** with **tool orchestration**
- Silimate – co-pilots for **bug finding** and PPA estimation
- Rise Design Automation – raise abstraction with C++ and **hardware focused agents**
- VerifAI – RL with multiple LLMs – an AI **test engineer**
- …
- CogniChip - **AI-enabled chip design**, 10 levels of AI-driven automation
  - creating a new segment… between fabless companies and EDA companies
- …

=> “Are we in the ‘**agenticEDA**’ era?”

Examples of agentic flows in use in 2025

- Synopsys + Microsoft demonstrated AgentEngineer
  - Verification: testbench creation, setup, and routine design analysis
  - “reduce iteration cycles and manage the growing workforce gap by automating implementation details”

- Cadence + Renesas adoption of Versium AI + JedAI
  - Generative AI agents used to ensure alignment between high-level functional specifications and final RTL design
  - “significant gains in design team productivity and a reduction in the ‘cost of verification’”

- Siemens EDA - ‘Agentic AI will become foundational’
  - Harry Foster article – “Reflections from inside an industry undergoing its biggest transformation in decades”

- Good example video
  - “Accelerating RISC-V Design and Verification with AI Agents”, RISCV Summit May 2025
  - https://www.youtube.com/watch?v=_Xj0KlcYMCs

2025 funding of D&V EDA + AI startups…

- Ricursive Intelligence - $35M seed
  - AI models to automate all stages of **chip design and verification**
- Cognichip - $33M seed
  - Artificial Chip Intelligence (ACI®) - AI-enabled **chip design**
- ChipAgents - $21M series A
  - Agentic AI **chip design** environment (-> Cadence)
- Chipmind - $2.4M pre-seed
  - AI agents to automate and optimize complex **chip design and verification** tasks
- Maieutic Semiconductors - $1.8M seed+
  - Generative AI copilot for analogue **chip design**

- Vinci - $46M series A
  - Physics-based AI for hardware design and simulation, particularly advanced packaging and 2.5D/3D IC
- Quilter - $25M series B
  - Automates PCB design using physics-driven AI

Agenda

- AI?
- Areas of EDA that could benefit from AI
- Status (Jan 26) of EDA products (market positioning) & research
  - From main 3, From startups
  - From research, From AI providers
- RISC-V processor verification – and AI…
- Some of the challenges going forward
  - Data for EDA
- Futures…

APRIL (AI for Productive Research and Innovation in eLectronics) AI Hub

UK consortium of academic and industrial partners
Focus on world-leading AI research relevant to the electronics and semiconductor industry

- £80M funded by UKRI, EPSRC
- All about AI
  - Accelerating Discovery, Enhancing Efficiency, Commercial Adoption, Talent Pipeline
- Five ‘pillars’ of research
  - Materials discovery, Device design, **Circuit & System Design**, **Testing & Verification**, Modelling
- ~20 UK Universities
- ~30 industrial members including Google DeepMind (internships)

APRIL AI Hub (2)

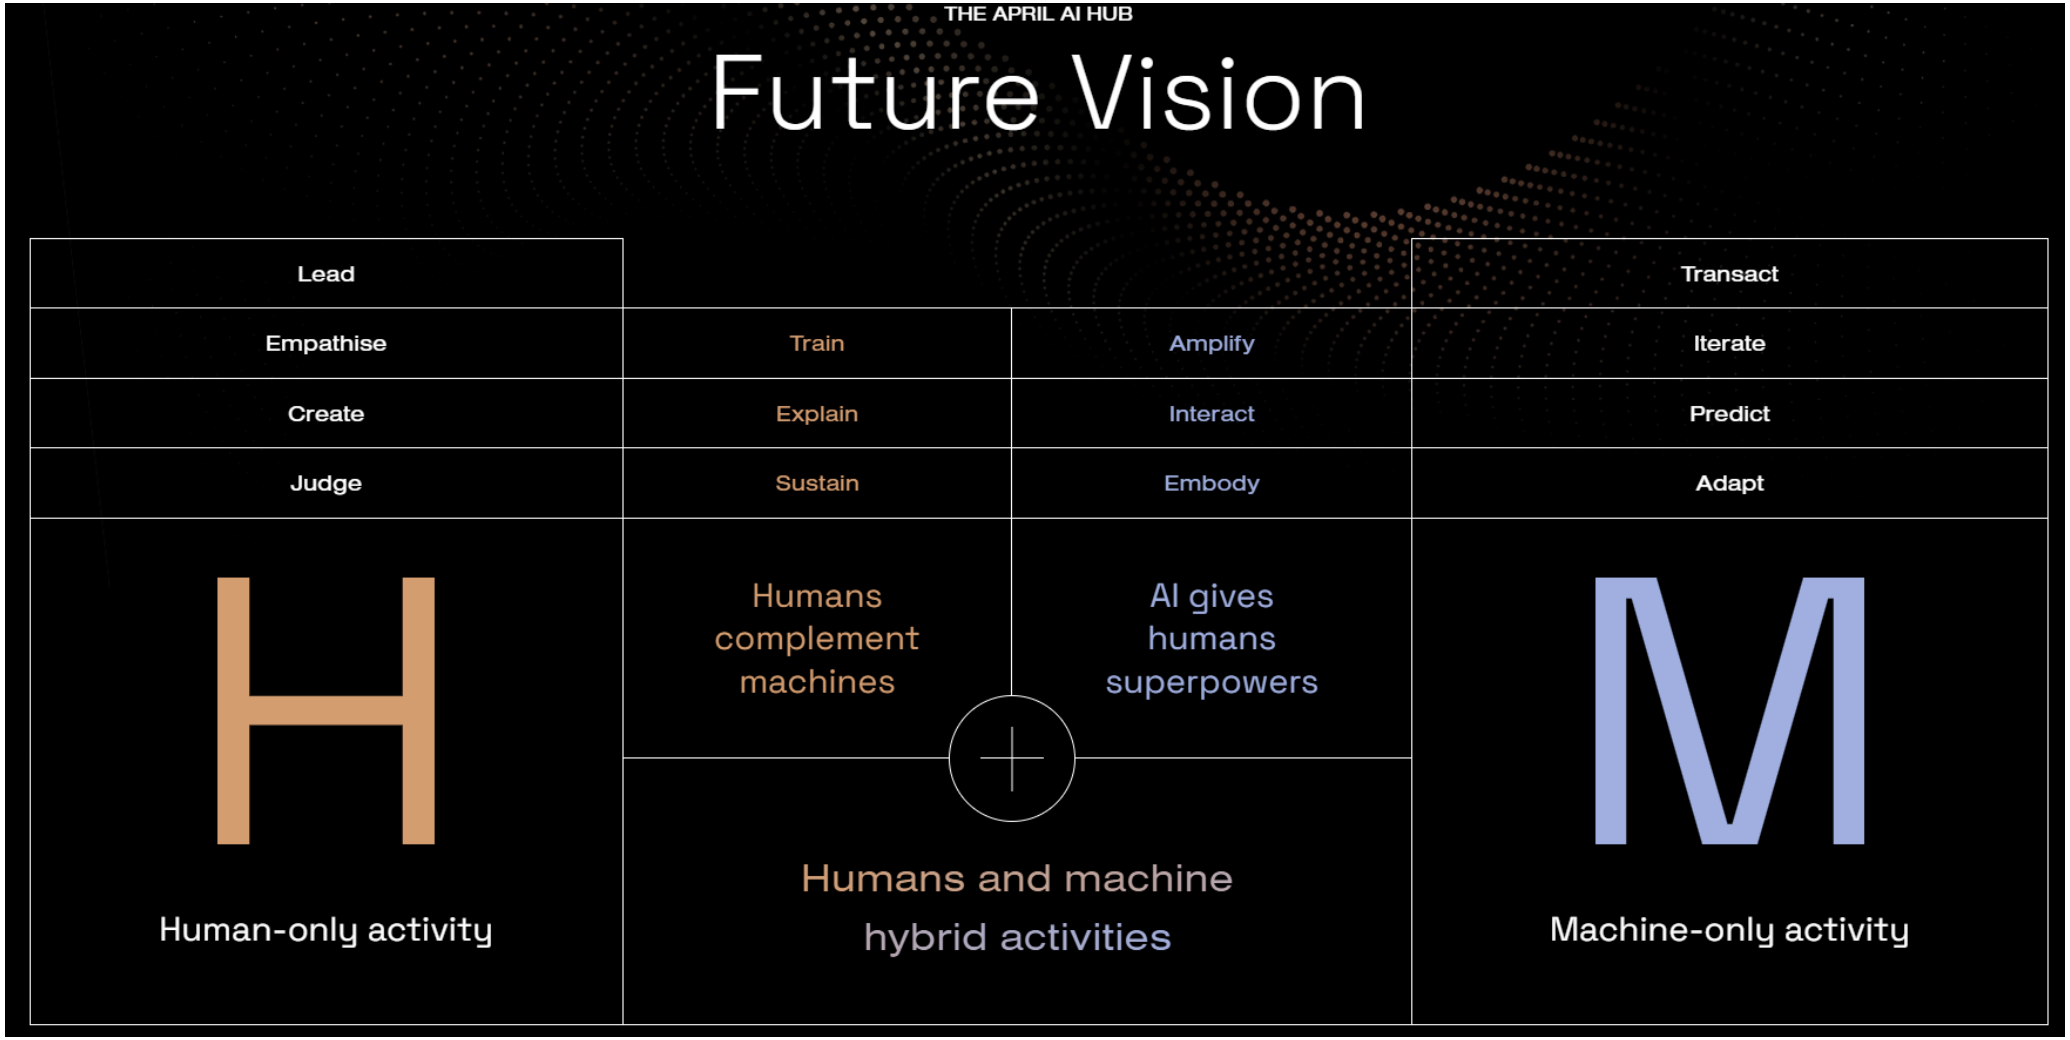

I like the superpowers bit… as I believe domain knowledge enables the superpower…

EDA AI tools / research in Design & Verification

Example research from APRIL AI Hub (https://www.april.ac.uk/)

LLM-Driven Circuit Creation, RF Design (KCL, Imperial)

Dynamic Verification Loops, Beyond-CMOS (York, Bristol)

IC Building Block Automation (Glasgow)

Other interesting AI+EDA technologies (from Univs. Around the world)

- AssertLLM - automatic assertion generation framework (HK Univ)
- AutoBench – evaluates test benches using LLMs (3 German Univs)
- VerilogReader – LLM coverage directed/aided hardware test gen (2 China Univs)
- LLM4DV – benchmarking framework to evaluate LLMs for hardware test gen (LowRISC, Univs: Cambridge, Imperial, Edinburgh)
- EvoVerilog – LLM & evolutionary algorithms (EAs) to generate RTL (Univs: Maryland, Hong Kong)
- …

Tech developed in AI hardware developers… (not very public)

- AlphaEvolve, AlphaOpen… (Google DeepMind)
  - LLMs to autonomously discover and optimize complex algorithms
  - Refines entire codebases using a continuous, automated feedback loop
  - Suggested Verilog **RTL optimizations** for Google's TPU (Tensor Processing Unit) circuit designs, which were integrated into hardware

- NVBugs/ChipNeMo (internal Nvidia) (Neural Modules)
  - Design, verification, bug tracking / management using ML / AI
  - Does not use off-shelf LLMs, instead use domain adaptation techniques
  - Target EDA usage:
    - **Code/script generation**: design code, test benches, assertions, tool scripts
    - **Engineering assistant chatbot**, Q&A about designs, tools, infrastructures
    - **Bug triage, summarization and analysis**

  => "Domain-specific LLMs can be substantially (5x) smaller and run effectively on smaller compute platforms"

Agenda

- AI?
- Areas of EDA that could benefit from AI
- Status (Jan 26) of EDA products (market positioning) & research
  - From main 3, From startups
  - From research, From AI providers
- RISC-V processor verification – and AI…
- Some of the challenges going forward
  - Data for EDA
- Futures…

RISC-V market (Jan 2026)

- Two different “RISC-V economies”
  - Internal “**utility core**” inside big SoCs
    - Huge activity in shipped cores, but mostly “invisible”
    - Massive adoption as internal infrastructure (e.g. Nvidia)
      - not “RISC-V replacing x86 / Arm as the main CPU”

    -> Moving from controller / management cores first then gradually expanding upward (e.g. Qualcomm)

    -> Moving to certified / robust processors + ecosystem (e.g. NXP, Quintauris, automotive)
  - **Licensable CPU-IP** product (SiFive / Andes / MIPS / Codasip / ARC-V, …)
    - Lots of design activity, smaller dollars / volumes not the same magnitude as shipped cores
    - Optimize time-to-market + risk reduction: its ready, validated, got software toolchains

- And there is the business changes…
  - Qualcomm buys Ventana
  - Meta buys Rivos
  - Global Foundries buys Synopsys ARC-V (adds it to its recently acquired MIPS RISC-V cores)

RISC-V Verification
Controller cores -> Application cores

- Move from validating tiny embedded control cores to Linux-class
  - Controller class
    - m-mode only, single hart, basic interrupts, extension ‘bingo’, no MMU, no hypervisor, no vector context-switch chaos
    - “ISA compliance level testing will do”
  - Application class
    - coherent multi-hart CPU complexes, cache-management, virtualization, profile-aligned, vectors, interrupt virtualization (AIA), security + safety
    - Require full industry standardization, certification

New areas of verification required for ‘Application’ class processors

Coherency-era of multi-hart

- Challenges of RVWMO (weak memory ordering model)
- Correctness becomes about interleaving with real parallel stimulus, timing-sensitive races
- No deadlocks under interrupt storms
- Concurrency infrastructure: memory ordering, LR/SC failure modes, AMO atomicity
- Atomic sequences + fences + cache-block ops, cache evictions, coherence invalidations, interrupts / exceptions, misaligned / boundary cases, …

AIA interrupt architecture

- PLIC-ish to AIA is a major leap
- Interrupts need distributed state correctness
- Per-hart interrupt IMSIC
- APLIC routing, domain partitions
- Virtualization paths, guest direct vs hypervisor mediation
- Interrupt storms, nested traps, priority inversions
- Interrupt injection vs pipeline, e.g. vectors

MMU + TLB for real OS workloads

- Jump from “no MMU” to “real OS workloads” causes a huge bug class explosion
- Page faults: instruction/data, read/write/exec, privilege violations
- ASID/VMID behaviours
- TLB invalidation correctness (global vs ASID vs address-range)
- Permission changes while cores run

Virtualization (H extension) and guest OS

- HS / VS contexts, trap delegation, guest interrupts, …
- HS <-> VS trap flows
- Nested translations (guest page tables, second stages)
- Chaos of interrupts during page walks
- CSR access in HS vs VS
- Hypervisor exception, interrupt delegation

Where can AI / ML help…

- Coverage generation
  - Generative agents **understand specs and generate functional coverage** to ensure sequencing, racing, coherency, hazards, … are targets to be hit in tests
  - Complex, e.g. driven by cross: profile x privilege x exception x concurrency x power state
- Property generation
  - Generative agents **create libraries of assertions**, e.g. protocol transitions – for formal analysis
- Test generation
  - Reinforcement Learning agents **focus instruction generation** across harts focusing on architectural corner cases and coverage goals
    - Reward for reaching new protocol states, snoops etc., …**force into 'high pressure' states**, …
  - Genetic (evolutionary) algorithms **evolve 'winning' tests** to increase hitting states in coherency protocols
- Anomaly detection
  - Graph Neural Networks capture how data moves – train on normal behaviour, monitor simulations **identify coherence breaches**
- Bug Hunting / Root Cause Analysis
  - AI driven triage analyses RTL/waveforms and clusters similar failures (**reduces debug**)
  - RCA copilots analyse diagnostics (logs, waveforms, …) to '**reason' about failures**
  - Similarity mapping comparing failure signatures with history incidents and **predict root cause categories** & explanations…
  - Planning algorithms do **scenario synthesis of test sequences** needed to reproduce

The new DV problem - why AI becomes necessary

As RISC-V processors become ‘real’, DV changes dramatically…

- Stops being: ‘generate more tests’
- Becomes:
  - Search – in exploding state space
  - Optimize – expensive regression across farms
  - Explain failures quickly as debug gets harder/bottleneck
  - Prove hard properties (ordering, isolation, liveness)
  - And … need to learn from all previous run data to steer next runs

DV direction…

Verification cost

- For RISC-V out-of-order + multi-hart + coherent + virtualisation + vectors + crypto, …
  -> **verification effort becomes ***the*** cost centre**

And for general DV - highest ROI in adopting AI / ML in DV - today

- Regression optimization – select right tests, engines, configs
- Coverage guided stimulus generation
- Failure triage at scale – clustering, predict culprits, summarizes ‘why it failed’
- AI assisted formal property creation
- Automated specification -> verification plan -> coverage models
- Waveform intelligence building cause-effect summaries

- See Cadence VF 2023 Matt Graham “Application of AI in IC DV”
- See Synopsys VF 2024 Bradley Geden “Auto Verification – are we there yet?”

Agenda

- AI?
- Areas of EDA that could benefit from AI
- Status (Jan 26) of EDA products (market positioning) & research
- RISC-V processor verification – and AI…
- Some of the challenges going forward
  - Data for EDA
- Futures…

Issues with training AI models for EDA

- Process of training an AI model is:
  - Clean data -> train model -> validate model -> test model -> then deploy
  - Bad results if:
    - **Too much data** or long training then model learns data too well and **fails to generalize**
    - **Too little data** and model **fails to learn** input -> output relationship == lots of errors
- Care is needed when evaluating to see if a model is any good…
  - What if you have trained that model on the test you are evaluating it on
    - It will do just great on that test… often un-intentional … but not so well on other tests?
- Big challenge for RTL design & verification & AI is getting good quality data…

Training data issues – is data used legally?

- EU AI Act (2025) Compliance - providers of general-purpose AI models used in EDA **must publish** a "sufficiently detailed **summary" of training content**

- Hidden Vulnerabilities and Licensing Risks in LLM Pre-Training Datasets
  - LLM4Code 2024 study found LLM outputs can **trigger copyleft license violations** (e.g., GPL) if the training data is not strictly audited

- DAC 2025 paper/study on the ethical and legal industrialization of Generative AI for hardware verification found over 70% of GitHub RTL did **not have a license** that meant it was **useable for training**…

- GTC2025 Nvidia Multi-AI Agents paper emphasizes an approach where one agent generates code while a second ‘checker’ **agent verifies compliance** with design rules and licensing headers

Example data sets for training of AI models

- VerilogEval v2 — 156 problems, functional correctness
- RTLLM 2.0 — 50 handcrafted designs for synthesis and function
- RealBench — IP level real designs, cores, blocks
- CVDP (Nvidia) — 783 detailed problems to test agentic workflows
- Pluto — 114 problems, focus on synthesis efficiency
- RTLBench — 160 designs from blocks to complex state machines
- ResBench — for FPGA specific resource constraint evaluation
- ProtocolLLM — examples of e.g. SPI, I2C, Uart, AXI, …
- HDLEval — 100 problems, for formal, supports Verilog, VHDL, …
- RTLCoder-27K — prompts & RTL code
- AssertEval — 18 designs for SVA assertion generators
- AssertionBench — for SVA assertion generators
- …

And for RISC-V designs

- Use the RTL source of processor models from: OpenHW, ChipsAlliance, LowRISC, …

And then there are EDA pre-trained models

This is where an AI model is pre-trained with EDA related data – like RTL design, or Verification data sets

- VeriGen for RTL code gen
- DeepRTL for RTL code gen, and understanding existing RTL code, created own training data sets
- BetterV for RTL code gen, trained on augmented data, creates own data sets
- RTLCoder …
- …

And then there are frameworks to create Verilog datasets

They address "data scarcity" in the hardware design domain by creating correct-by-construction synthetic data to train more reliably

- CraftRTL (Nvidia)
  - Uses information derived from state-transition diagrams, Karnaugh maps, and waveforms
  - Strategy is to create massive paired datasets of "buggy code" and their corresponding "fixed code“
    - injects mistakes into code
- …

So there is a requirement – for EDA data sets

- There is much talk and work on AI creating RTL
  - And there are more and more datasets and models focused on the genAI RTL designs task

=>There needs to be more work for RTL verification to use AI

- **This needs addressing**

=> and for RISC-V verification using AI there seems very little so far…

- **This needs addressing**

Agenda

- AI?
- Areas of EDA that could benefit from AI
- Status (Jan 26) of EDA products (market positioning) & research
- RISC-V processor verification – and AI…
- Some of the challenges going forward
  - Data for EDA
- Futures…

RTL Verification… Predictions – what's next

- Much more adoption as verification is where AI delivers its highest measurable return on investment (ROI)
- Use Agents – replace static scripts with context aware auto scripts
- AI-driven verification will automatically close coverage on standard IP
- Debug process will be dramatically accelerated with use of integrated AI co-pilots
  - e.g. those in Synopsys Verdi or Cadence SimVision using assisted waveform visualization and failure debug

RISC-V Processor Verification - Predictions

- AI models will be tailored for hardware processor verification
- Tools will provide more insights into what goes wrong
- AI will explore designs to focus the verification efforts
- Shift from AI-as-an-assistant to autonomous agentic AI that orchestrates workflows
- AI will change processor verification from ‘generate more tests’ to ‘search, identify, understand, and explain’

Things we need to do

- Collect and preserve data to train on
- Treat AI / ML models as replaceable (expect upgrades / refresh)
- Record user prompts and use them as part of continuous improvement

- Develop clear process boundaries (what AI is allowed to generate vs what must be verified by engineers)
- Maintain traceability (inputs → prompts → outputs → validation)

- Include AI agents as integrated "team members" alongside human engineers

Summary

- Old style EDA will not cope with current & future DV requirements
- AI is finding its way into most of the EDA DV flows
- Agentic AI adds AI + tools together to provide autonomous workflows
- Use of AI augments human engineers (gives superpowers)
- AI is going to transition from experimental pilots to "mission-critical infrastructure" in EDA
- Requirements changing rapidly so don’t lock in specific models

And…

- Industry / researchers / users – we need to create more public training data

AI+EDA for Design Verification Engineers

(with examples for RISC-V processors)

Simon Davidmann

AI+EDA Researcher, School of Electronics and Computer Science

Southampton University, UK

DVClub, Bristol, UK

Wed 28 Jan 2026

Note: some of this data may be incorrect as I used AI for some of the research… (and it kept making things up)

What is AI? – convergence of data, algorithms, computation (the fuel, logic, engines, and now the model, flows,…)

- Artificial Intelligence (AI) – ‘stuff’ that used to only be done by humans (and our intelligence): **learn, infer, reason**
- Machine Learning (ML) – machines using large amounts of data, **finds patterns, predicts, does useful stuff**
- Deep Learning (DL) – use models/layers of Neural Networks to **apply weights into models to guide decisions**
- Large Language Models (LLMs) – one type of foundation DL model, **token predictors**
  - Transformers (“Attention is all you need paper”) (c.2017)
  - Provides natural language interface to data
  - More than language, e.g. now images, etc.
- Generative AI (GenAI) – use models (eg LLM), see words, **predict next word** (like…auto-complete…), generate
  - ChatGPT – Generative Pre-trained Transformer
- Agentic AI – use the models, genAI, ML, but **add tools, objectives, collaboration, checking**
  - Workflows with multi-agent systems reasoning over data, and constraints and using tools, with iteration, and self reflection (checking)
- Orchestration (new) - coordination is everything – **conductor of agents**
- …

=> **Can accelerate much of the EDA design & verification flows…**

DVClub Bristol, Jan 2026 - Processor Verification

Where leading experts and practitioners will explore the workflows, tools, and collaboration models driving the next generation of verification excellence. This session will highlight practical strategies for improving efficiency, consistency, and quality in processor verification projects

Abstract

- Machine learning and assistance has been used in EDA tooling for many years and now there are more and more intersection points between Artificial Intelligence and Machine Learning (AI/ML) and EDA.
- At DAC62 in 2025 there was much discussion about "Are we living through the 'Verilog Moment' for AI?". The Verilog Moment was when the chip design industry accepted that its world had fundamentally changed by moving from schematics to HDLs. At the DAC it seemed that over half of the papers, presentations, talks, demos included ML/AI. There was also a buzz around many of the new 'AI startups'.
- This talk will introduce several of the interesting ML/AI technologies being developed/used and the areas of the EDA Design and Verification space that they are affecting. It will also survey some of the existing available commercial DV products using these ML/AI techniques and includes information about ongoing AI/ML related research in UK Universities.
- It also includes discussion on how AI/ML can be used in chip DV workflows and covers benefits and concerns in their adoption. The issue of training data sets and their licensing is also introduced.

3 key points;

- AI is here now and in use in EDA flows
- Agentic tools are becoming more and more effective for DV
- A big challenge is how we address the shortage of appropriately licensed training data