



DISTINGUISHED LECTURER PROGRAM

This ASHRAE Distinguished Lecturer is brought to you by the ASHRAE Society Chapter Technology Transfer Committee

Lecturer presentations and/or opinions do not necessarily reflect the policies or position of ASHRAE or the chapter.

Please!

Silence Phones

Distinguished Lecturer Evaluation Forms are very important. Please complete at the end of the presentation and return to the CTTC or Program Chair.

More information on the DL program available at:

www.ashrae.org/distinguishedlecturers



Distinguished Lecture
2025-26



AI in HVAC Industry: Past, Present and Future

Krishnan Gowri, *Ph.D., BEMP, LEED AP, Fellow ASHRAE*
Senior Engineering Fellow, Energy Solutions

Acknowledgements

- Ben Bartling, Smart Building Engineer, Intellimation LLC
- Bilal Sher, CEO, Building Diagnostic Robotics
- Kevin Lawson, Ripple Automated Building Design
- Nathaniel Jones, ARUP
- Liang Zhang, University of Arizona
- Tianzhen Hong, Lawrence Berkely National Laboratory
- Gang Jiang, University of Utah
- Weili Xu, Pacific Northwest National Laboratory
- Jayson Bursil, Delta Controls

- Members of GenAI.MTG

Outline/Agenda

- From calculators/computers to AI - How did we get here?
- Why is AI now different from the 80's hype?
- What are some example applications of genAI?
- How can ASHRAE members benefit from the recent trends in AI?

What is special about this number?

1729

Hardy–Ramanujan number

"I remember once going to see him when he was ill at Putney. I had ridden in taxi cab number 1729 and remarked that the number seemed to me rather a dull one, and that I hoped it was not an unfavorable omen. "No," he replied, "it is a very interesting number; it is the smallest number expressible as the sum of two cubes in two different ways."

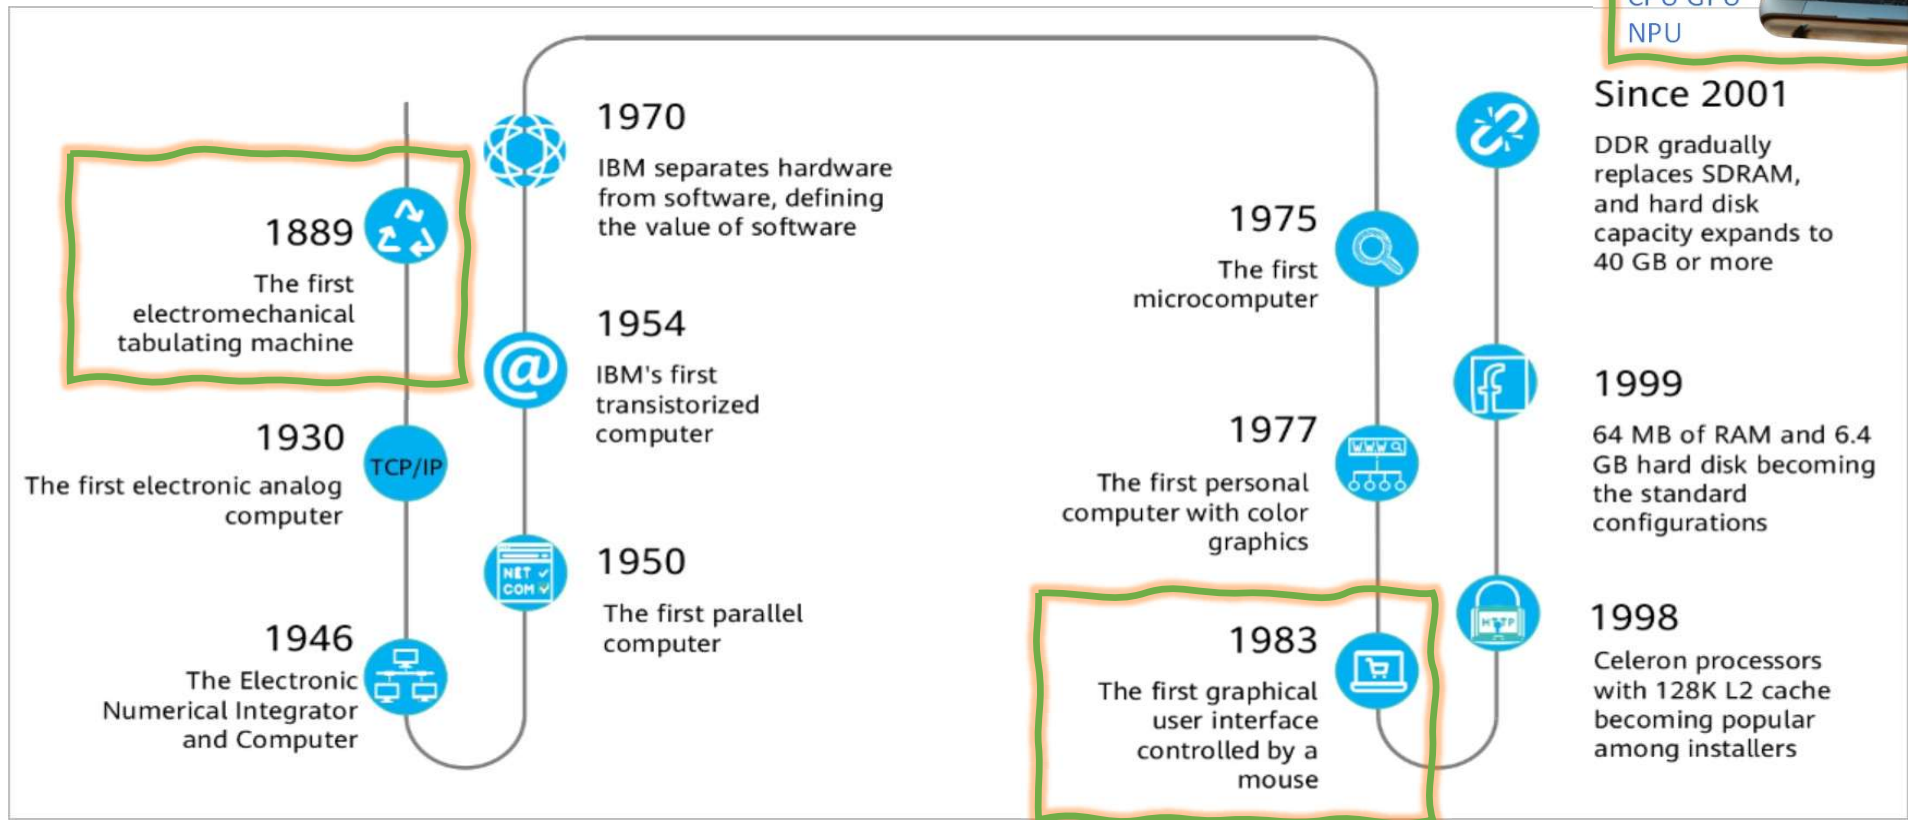
-Hardy

$$1729 = 1^3 + 12^3 = 9^3 + 10^3$$

- Just a four digit number – nothing special
- Represents sum of two cubes
- Represents sum of two cubes in two different ways
- Represents smallest number expressible as sum of two cubes in two different ways
- Number divisible by the sum of its digits

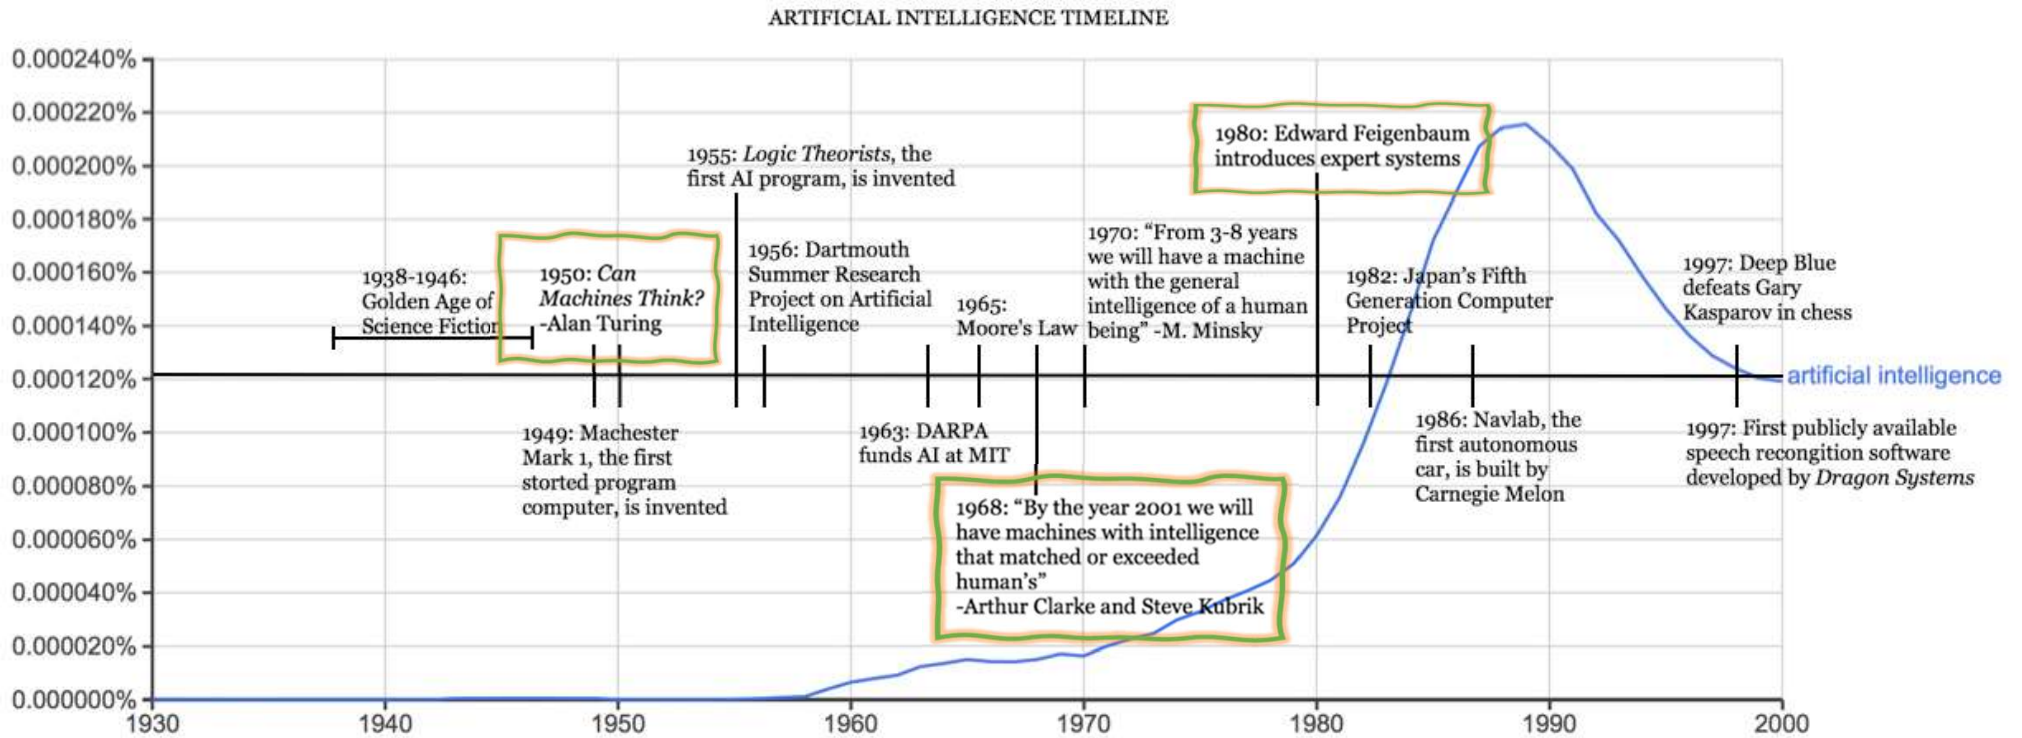
Computer Technology Evolution

2024
CPU GPU
NPU



(Source: <https://fastvoicemedia.de/image-collection/what-is-the-timeline-of-computer-history.aspx>)

Artificial Intelligence - Timeline



(Source: The history of Artificial Intelligence, Harvard student blog available at: <https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/>)

Artificial Intelligence in the 1980's

Goal of AI:

“Replication of the **functionality** of human mind”

Technologies:

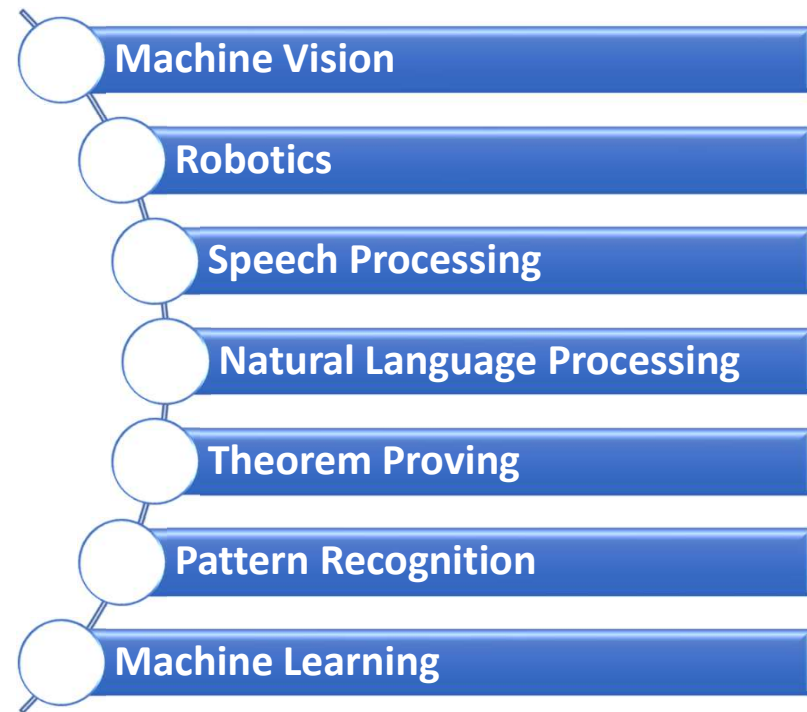
Expert Systems or Knowledge-Based Systems

Neural Networks

Fuzzy Logic

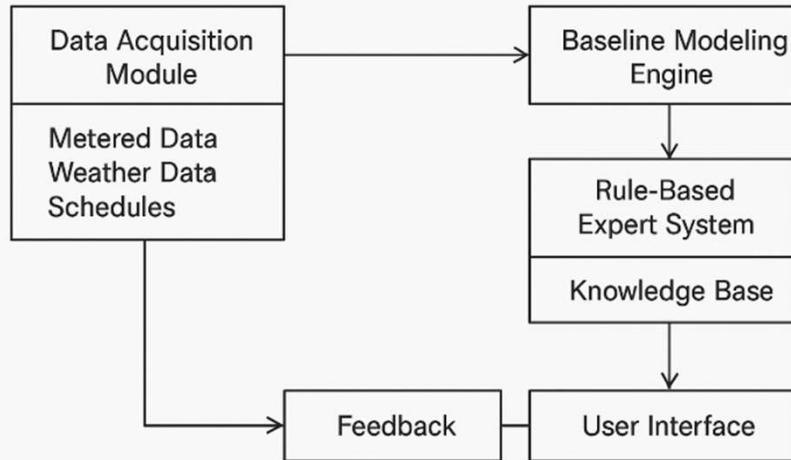
New Career Title e.g.:

Knowledge Engineer



Rule-Based Systems – Example: BEACON

Architecture of BEACON



Component	Function
Data Acquisition Module	Collects hourly metered data (electricity, thermal energy, weather)
Baseline Modeling Engine	Builds regression models to predict expected energy use
Rule-Based Expert System	Applies diagnostic rules to detect anomalies and inefficiencies
Knowledge Base	Stores operational rules, fault patterns, and corrective strategies
Inference Engine	Matches observed data against rules to generate diagnoses
User Interface	Displays alerts, recommendations, and performance feedback
Feedback Loop Module	Tracks post-intervention performance to validate and refine actions

Haberl, J. S., Smith, L., & Kreider, J. F. (1988). *Metered Data Analysis and Knowledge Based Methods that Reduce HVAC Operation and Maintenance Problems*. Proceedings of the Symposium on Improving Building Systems in Hot and Humid Climates. Energy Systems Laboratory, Texas A&M University.

Knowledge-Based Systems – Example: BEADS

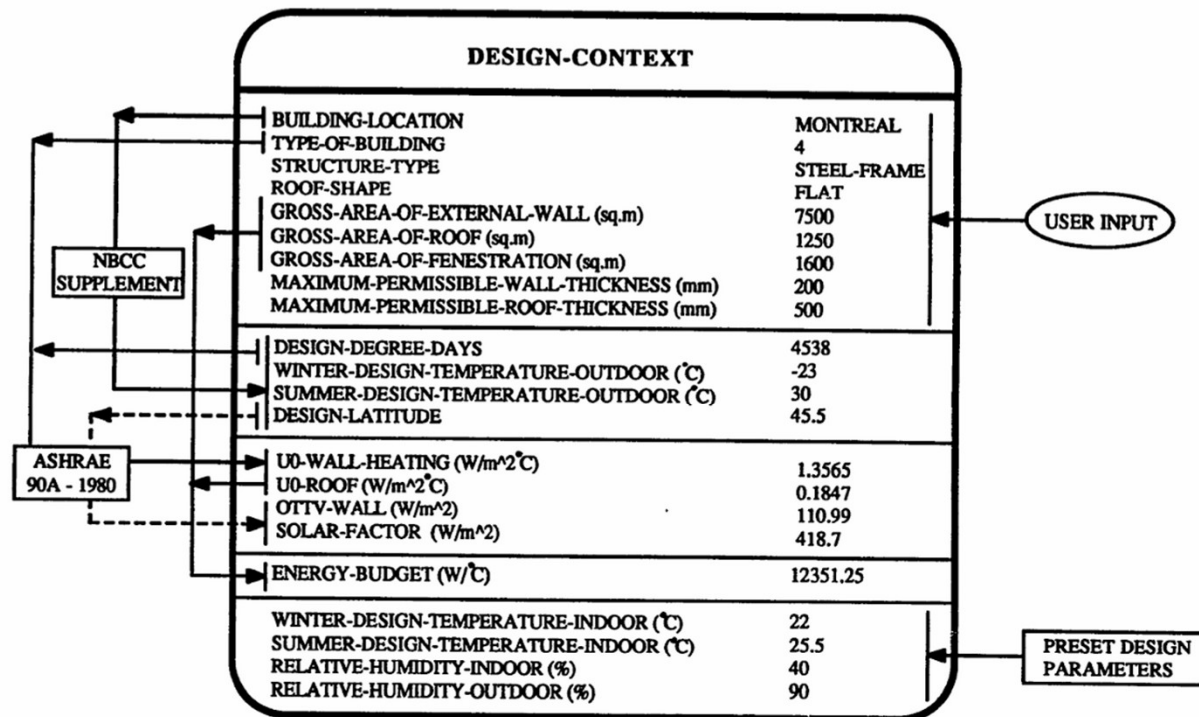
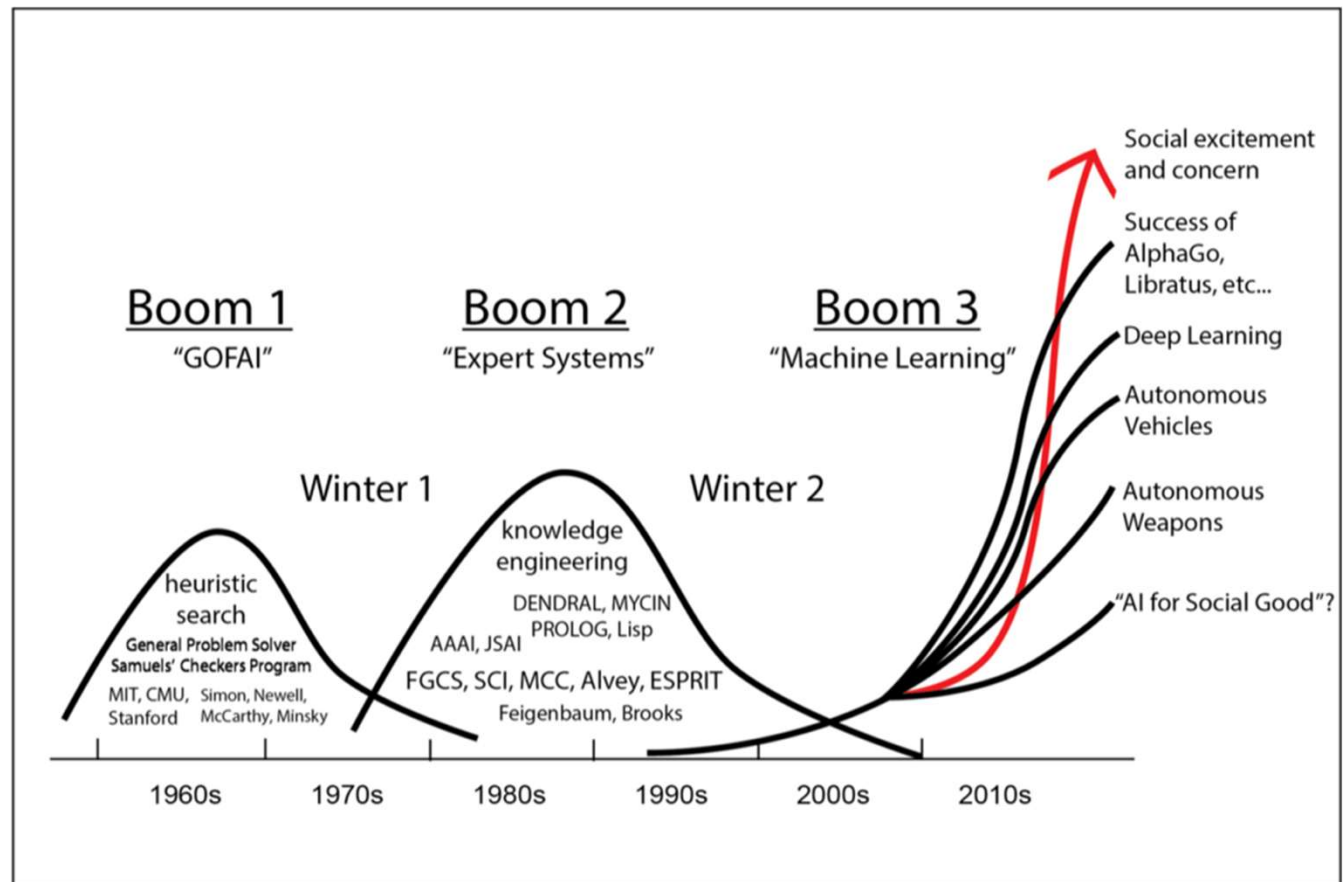


Figure 4.3: Example of a Design Context and Information Sources

Expert Systems and AI Limitations of 1980s

- ❖ Lack of computing power for text processing and collapse of the LISP machine market
- ❖ Slowdown in development and deployment of Commercial Expert systems - they could solve only well-defined problems in narrow domains



Technology Advancements Triggering AI Revolution (2022)

The AI innovations on the Hype Cycle reflect priorities across four main categories:

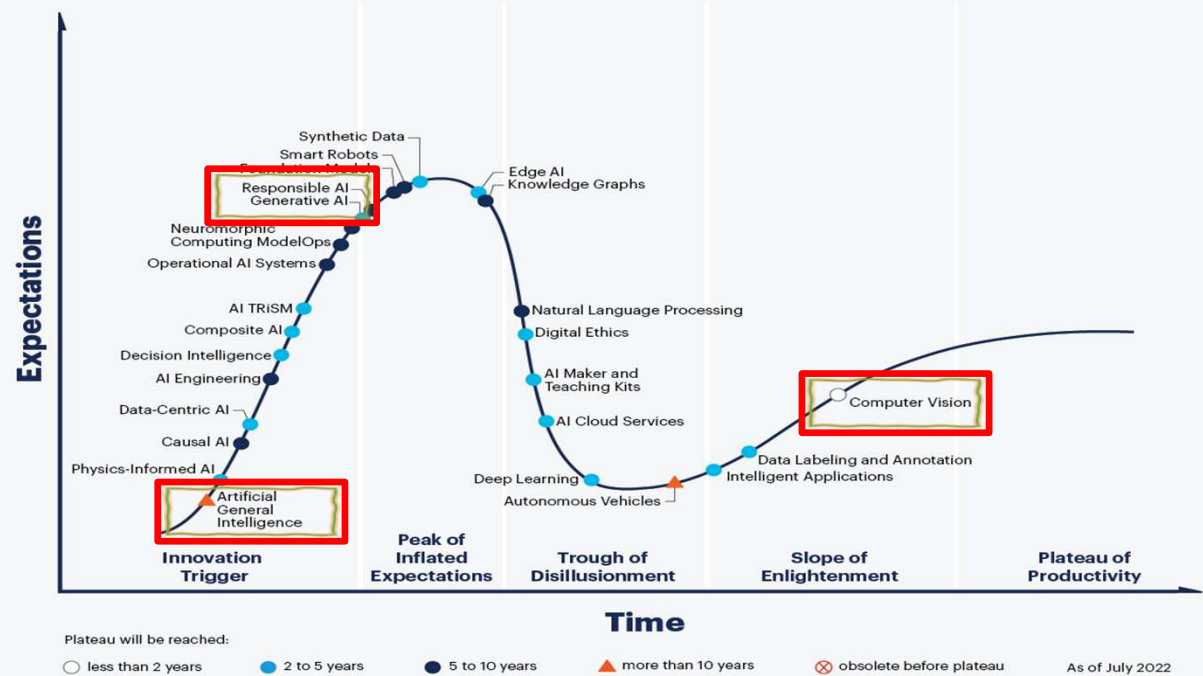
1. Data-centric AI
2. Model-centric AI
3. Applications-centric AI
4. Human-centric AI

Technologies Triggering the Hype:

1. Big Data
2. Cloud Computing
3. Data Science
4. Machine Learning
5. Large Language Models

(Source: <https://www.gartner.com/en/articles/what-s-new-in-artificial-intelligence-from-the-2022-gartner-hype-cycle>)

Hype Cycle for Artificial Intelligence, 2022



gartner.com

Source: Gartner
© 2022 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner and Hype Cycle are registered trademarks of Gartner, Inc. and its affiliates in the U.S. 1957302

Gartner

HVAC Applications

- **Generative AI**
LLMs and diffusion models for design generation, documentation, and automated content creation
- **Machine Learning**
Predictive analytics for system performance, energy optimization, and predictive maintenance
- **Computer Vision**
Thermal imaging analysis, equipment condition monitoring, and on-site inspections
- **Agentic AI**
Autonomous systems for workflow automation, decision-making, and complex project management

Generative AI

Artificial Intelligence (AI)

Machine Learning (ML)

Deep Learning

Generative AI



Learn from relationships

Learn from data

Emulate human intelligence

Large Language Models

- Deep learning models trained on vast amounts of linguistic data, so they understand and respond in natural language text



Capabilities

Natural Language Understanding (NLU)
Natural Language Generation (NLG)
Context Awareness



Components

Transformers
Large Datasets
Learning Models



Current AI Models in Use

Proven Technologies Deployed Today

Large Language Models (LLMs)

- GPT-4, Claude for design specification
- BIM documentation generation
- Code compliance analysis & interpretation
- RFP and proposal automation
- Stakeholder communication drafts

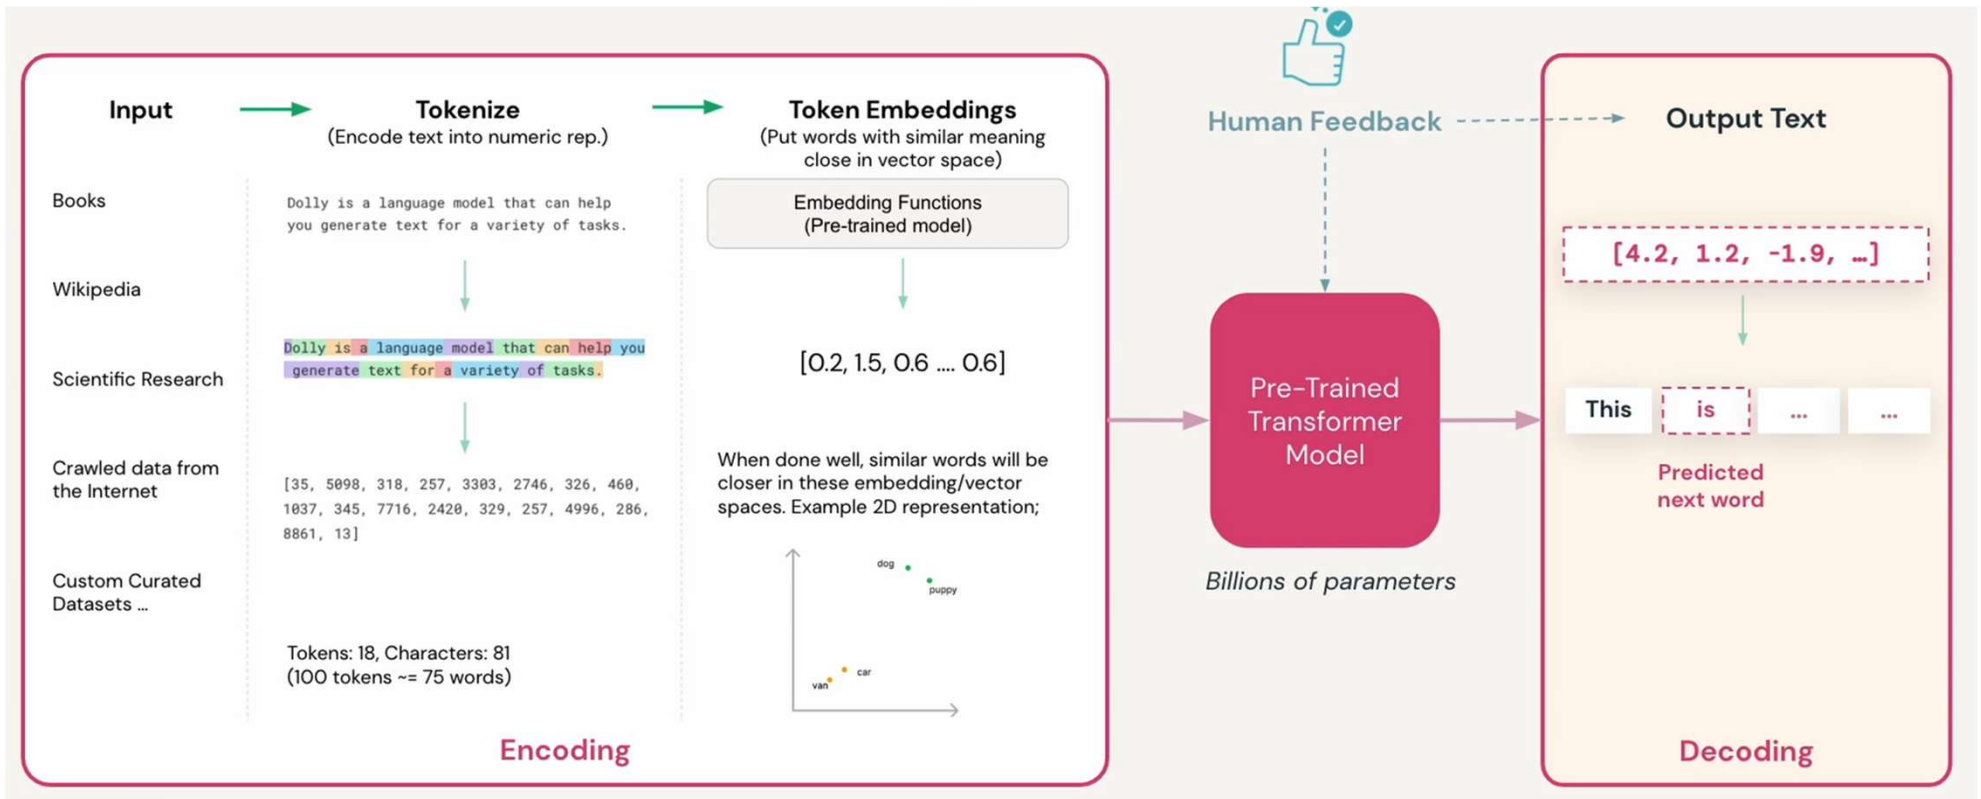
Specialized ML Models

- Energy optimization algorithms
- Load forecasting models
- Equipment failure prediction
- Sensor data anomaly detection
- System performance analytics

Computer Vision & Image Analysis

- Thermal imaging interpretation
- Equipment damage assessment
- Installation quality verification
- 3D model generation from photographs

LLM – How do they work?



How to use LLMs in HVAC Consulting?

Helpful prompts for HVAC design:

Rephrase this technical report for a non-technical audience.

Write an abstract for this report.

Explain this architectural concept to me.

Make this email sound more polite/professional.

“What the large language models are good at is **saying what an answer should sound like**, which is different from **what an answer should be**.”

—Rodney Brooks, *Robust.AI*

Generative AI Chatbot - How does it work?

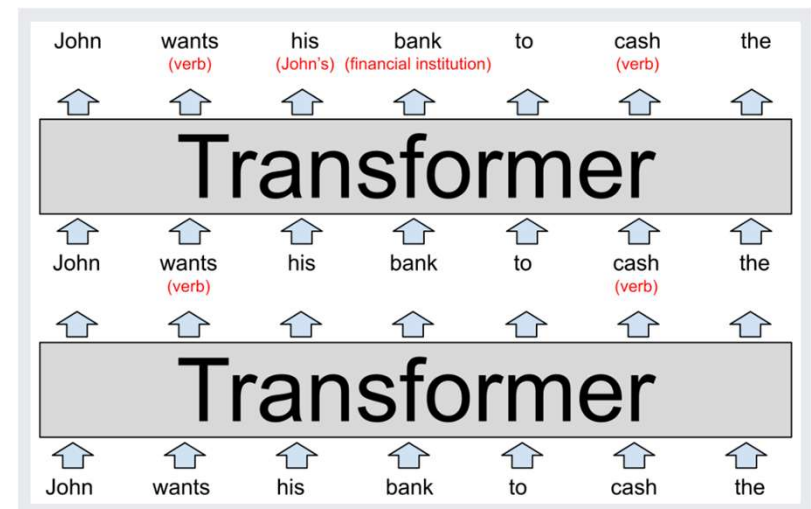
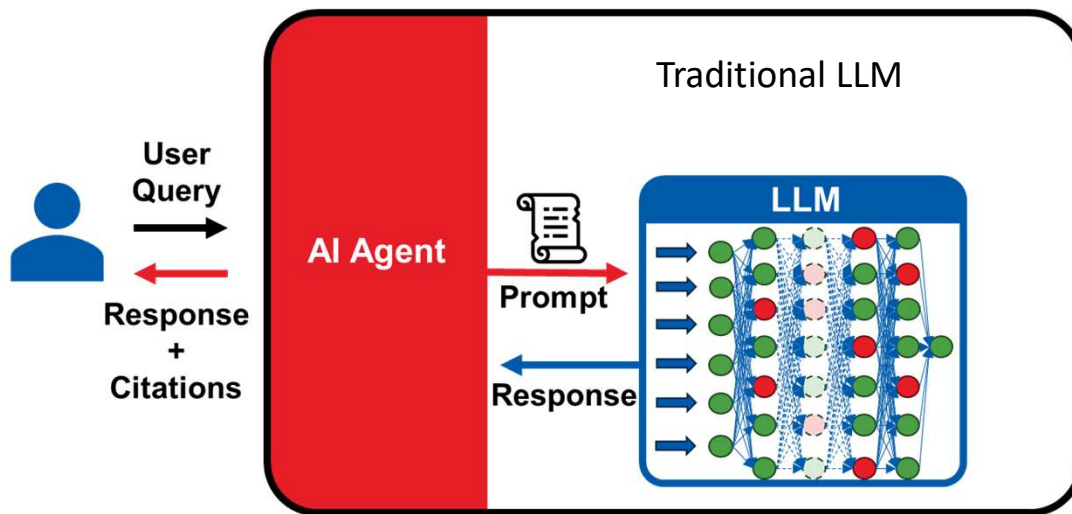
Train the Chatbot with data



Give a prompt



Get output, review and refine



Prompt Engineering – What is it?

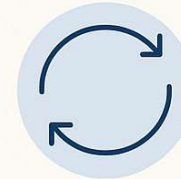
- Understanding the model's capabilities (e.g., text generation, summarization, reasoning)
- Formulating clear, specific, and contextual prompts to achieve desired outputs
- Iterating and refining prompts based on the quality of the output



**Model's
Capabilities**

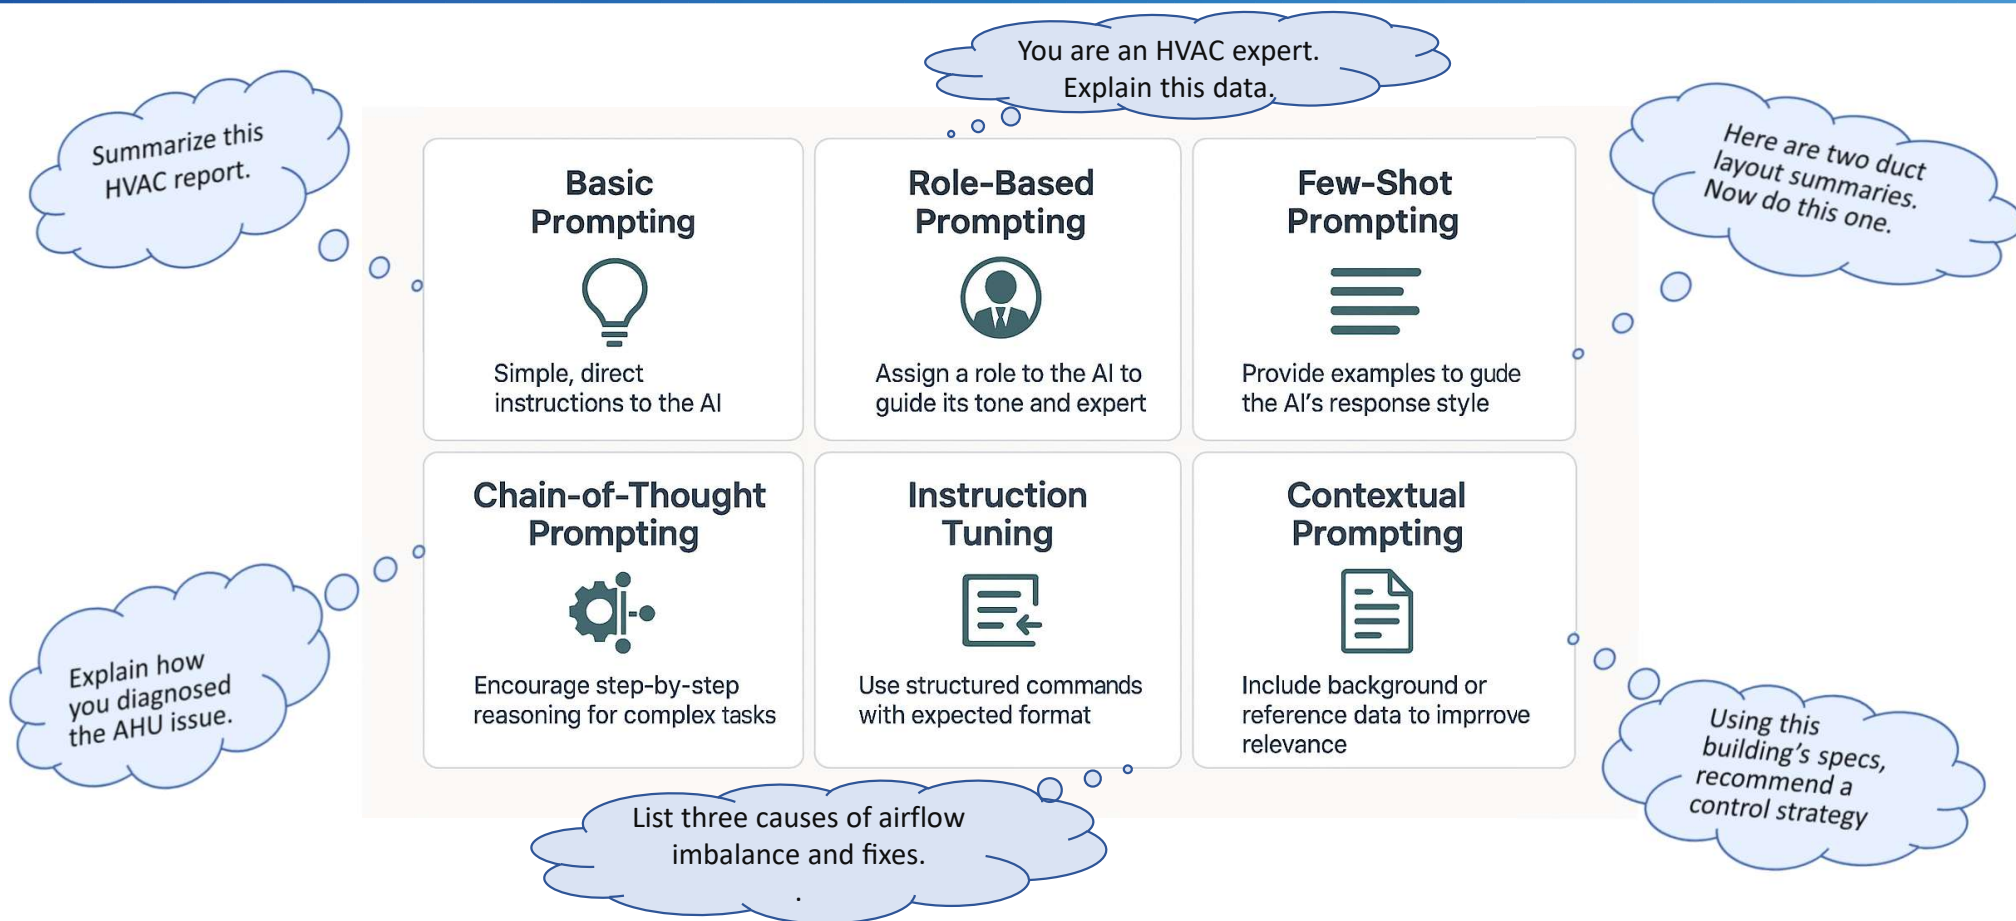


**Formulating
Prompts**



**Iterating &
Refining**

Prompting Techniques



Anatomy of a Good Prompt

Context	Who you are, what you are working on
Task	What you need or want from the LLM
Formal	The structure of the output
Constraints	Any limitations or boundaries
Examples	Optional – like required result

Prompt Engineering for Energy Modeling - *Examples*

Technique	Prompt Content	Response Summary
Zer-shot/Basic	what is the annual EUI for a typical 2000 sf residential building in phoenix, AZ?	6.78 kWh/ft ² -yr (~23 kBtu/ft ² -yr), which aligns to ~13,500 kWh/year for 2,000 ft ² .
Role-Based	You are an energy analyst setting EUI target for a typical 2000 sf residential building in Phoenix, AZ.	Set an annual site EUI target of ~27 kBtu/ft ² -yr for a typical, code-compliant 2,000 sf residential building in Phoenix, AZ.
Few-shot	Use the EnergyStar portfolio manager data on residential building stock and the PNNL residential prototype results as examples	≤ 28 kBtu/ft ² -yr Represents a realistic yet efficient benchmark for new construction
Chain-of-thoughtAssume air tight construction, WWR of 15%, south facing glazing. calculate the total heat/cooling load to select efficient space conditioning system, calculate the ERI and compare with DOE zero energy ready home target.	Annual site EUI: ≈26 kBtu/ft ² -yr ERI (no PV): ~55 ZERH compliant: Yes
Contextual	...and meeting the prescriptive requirements for envelope and mechanical systems as per DOE ZERH. Create an energy model to size the equipment, run annua simulation and calculate the target	modeled EUI target is ~26 kBtu/ft ² -yr (site), supported by realistic load sizing, high-efficiency systems, and annual simulation.

Prompt Techniques - Comparison

Context – Calculate EUI for a residential building model

Technique	Token Count	Accuracy vs Zero-Shot	Key Additions	Best Use Cases	Limitations
Zero-Shot	~30	Baseline	Direct task only	Quick estimates, well-known problems, rapid prototyping	Vague outputs, inconsistent format, broad ranges
Role-Based	~60	+25%	Expert persona, credentials	Technical responses, domain-specific vocabulary, professional tone	Still lacks format guidance, may be overconfident
Few-Shot	~320	+30-45%	2-5 demonstration examples	Consistent formatting, standardized outputs, quality templates	Higher token cost, need good examples, may overfit to examples
Chain-of-Thought	~420	+50-65%	Step-by-step reasoning instructions	Complex calculations, multi-step problems, transparency needs	Longer responses, verbose, may show incorrect reasoning
Contextual	~750	+60-75%	35+ parameters, standards, specs	Engineering-grade accuracy, compliance work, detailed analysis	Very high token cost, requires complete specifications

Depends on LLM and Context

AI Throughout Design & Construction

Building Life-Cycle Applications

DESIGN STAGE

- System sizing
- Duct layout generation
- Energy modeling
- Load calculations
- Spec writing

CONSTRUCTION STAGE

- Progress monitoring
- Quality assurance checks
- Installation verification
- Change order analysis
- Schedule optimization

OPERATIONS

- Predictive maintenance
- Performance tuning
- Fault diagnostics
- Energy optimization
- Occupant comfort

Expected Impact

Reduction in design time • Improvement in energy efficiency • Faster construction quality checks
Reduction in maintenance costs • 24/7 system monitoring • Predictive insights

AI in Design: Detailed Applications

Automated System Sizing

AI analyzes building geometry, climate data, and occupancy patterns to recommend optimal HVAC equipment sizes, reducing undersizing and oversizing errors by 40%.

Intelligent Duct Routing

Computer vision and optimization algorithms design duct layouts that minimize pressure drop and noise while fitting building constraints, saving 20+ hours per design.

Energy Modeling & Simulation

Automated hourly energy simulations with hourly weather data and occupancy patterns, providing detailed equipment selection and code compliance verification.

Automated Specifications

Generate comprehensive, code-compliant specifications and technical documentation in minutes, customized to local codes and standards.

Code Compliance Analysis

AI cross-references designs against ASHRAE, building codes, and local regulations, identifying violations before submission and reducing rejections.

BIM Integration

Automatically generate and update BIM models with HVAC systems, coordinating with architectural and structural elements.

AI in Construction & Operations

Site Progress Monitoring

Computer vision analyzes daily site photos to track equipment installation, ductwork progress, and compliance with construction schedules automatically.

Quality Assurance & Testing

AI systems verify ductwork sealing, airflow measurements, and vibration levels against specifications, documenting compliance automatically.

Automated Change Orders

NLP analysis of change requests estimates cost and schedule impacts, routes approvals, and updates project documentation in real-time.

Predictive Maintenance

ML models analyze sensor data (vibration, temperature, pressure) to predict equipment failures 4-6 weeks in advance, reducing downtime.

Energy Optimization

Real-time AI adjusts setpoints, ventilation rates, and equipment runtime based on occupancy and weather, reducing energy use.

Fault Diagnosis & Alerts

AI detects refrigerant leaks, control failures, and airflow issues from sensor patterns, alerting technicians before occupant complaints arise.

Generative AI for Building Energy Modeling



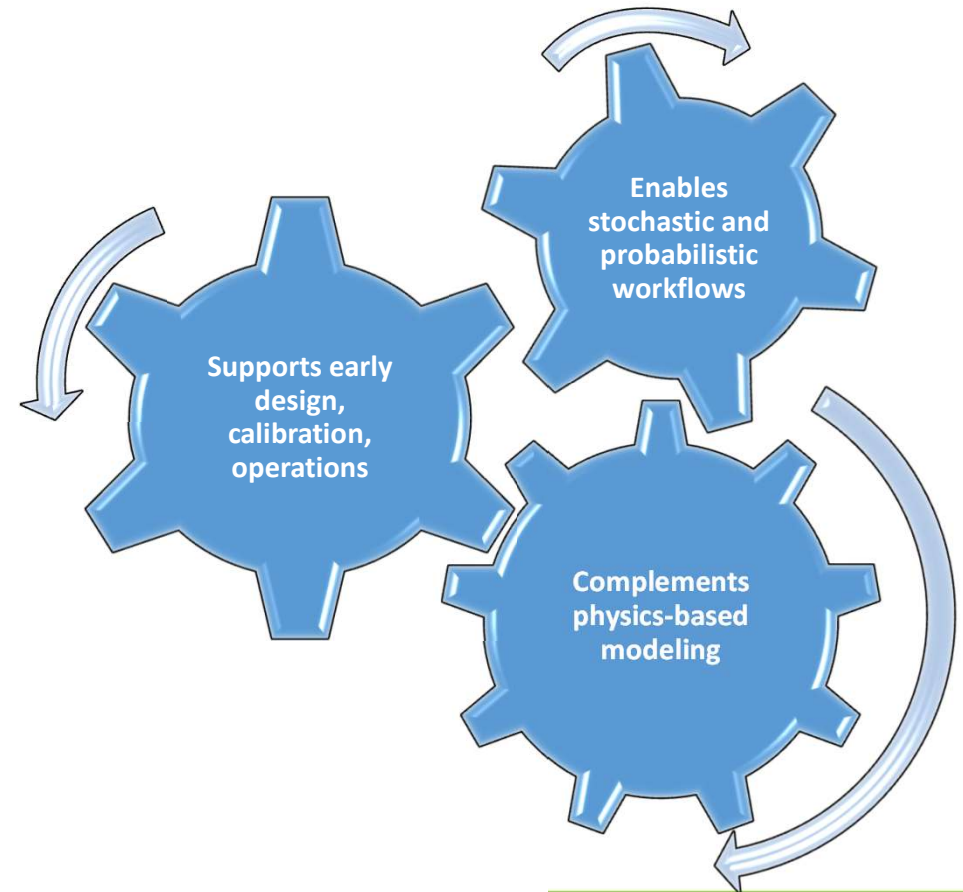
Traditional simulations are slow and expert-intensive, high latency, deterministic inputs



Need for scalable, automated modeling



GenAI enables rapid exploration and probabilistic modeling



Generative AI Techniques for Building Energy Modeling



Surrogate Models



Diffusion Models



Transformer-Based Architectures



Hybrid Physics-GenAI

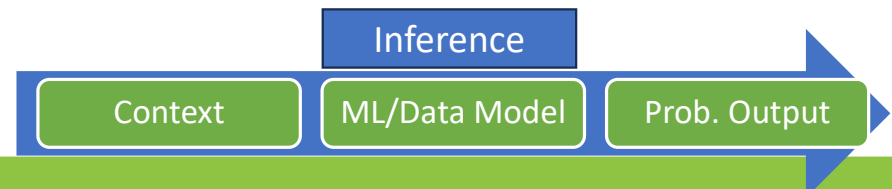
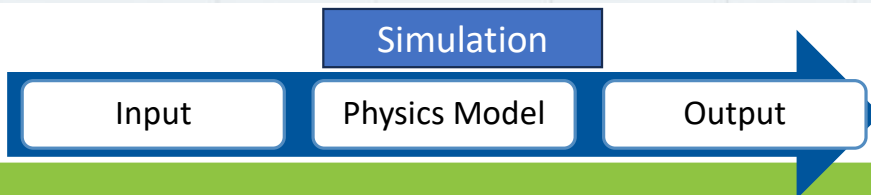
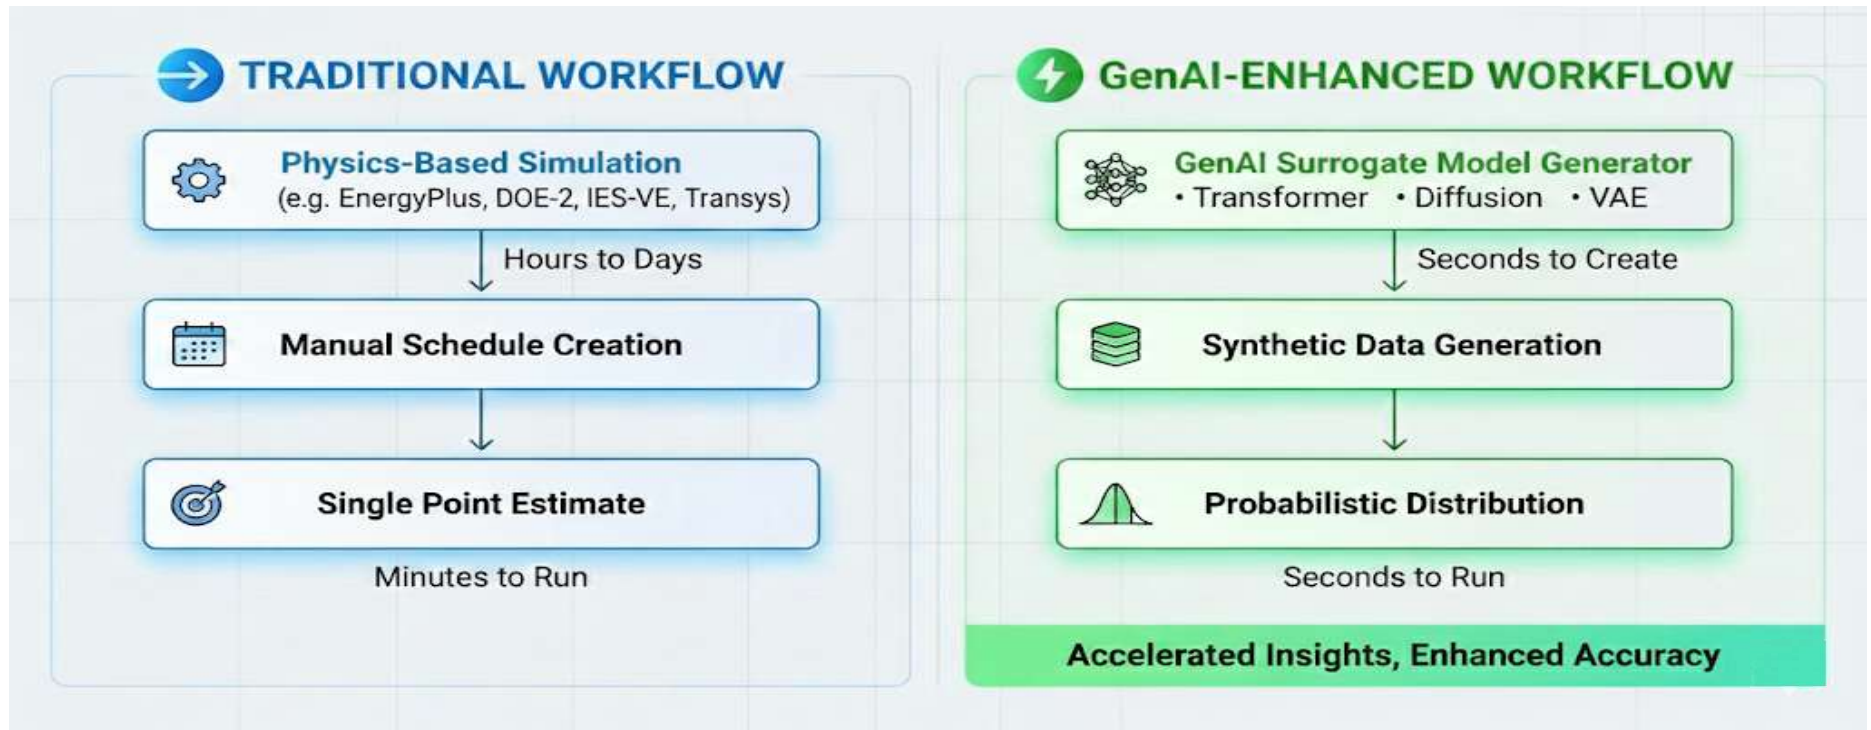


Prompt Engineering



Agentic AI - Model Control Protocol (MCP)

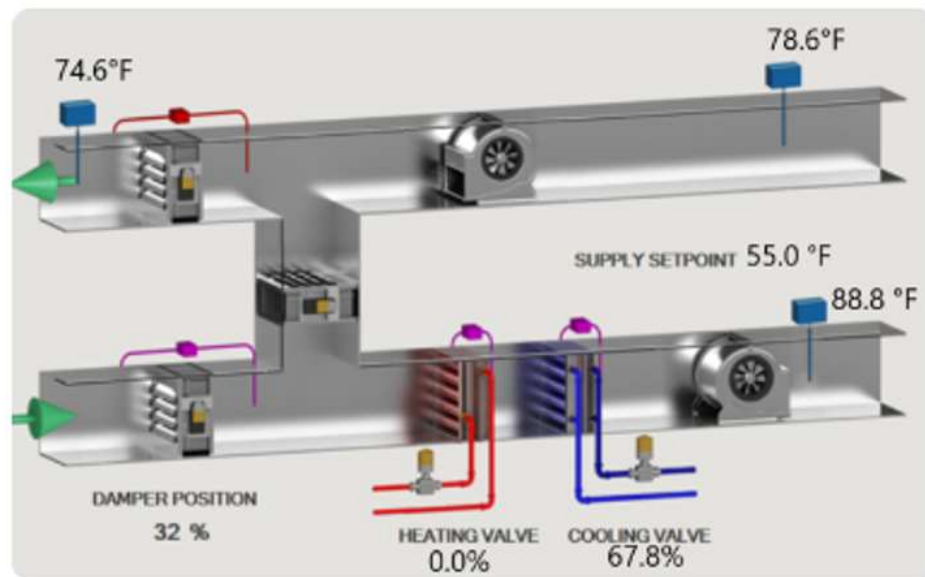
Traditional vs. GenAI Energy Modeling



HVAC Operational Data Analysis – Experiments

- Provide data to Chatbot in small batches.
 - *i.e., a weeks worth of electrical power meter data on 15-minute intervals.*
- Can be useful to the smart building IoT sectors where logic and data processing can be automated to provide a “virtual analyst” in describing results to the end client.
- Typical commissioning efforts or building retuning require an “expert” to explain results to the client where AI can potentially provide a “virtual expert.”

Broken AHU Heating Coil BAS Graphics



can you tell me what is wrong with this air handling unit and please layout some steps in troubleshooting it?

Results on Broken AHU Heating Coil

(LLM response – sample 1)

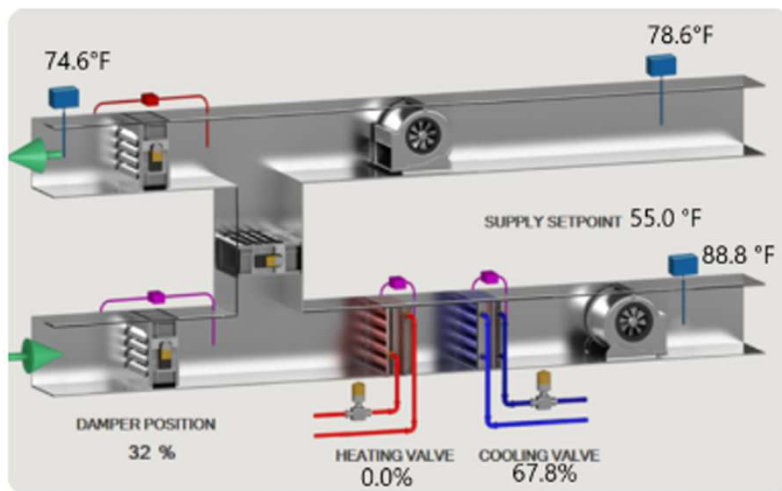
- **Cooling Valve Position:** The cooling valve is 67.8% open, which suggests the system is trying to cool the air but is not reaching the setpoint.
- **Heating Valve Position:** The heating valve is 0% open, indicating no heating is being applied, which is correct since cooling is needed.
- **Damper Position:** The damper position is 32%, which might be insufficient for proper airflow. The damper position should be checked to ensure it is providing the right amount of outside air and return air for effective cooling.

Troubleshooting Steps

1. Check the Cooling Coil
2. Inspect the Cooling Valve
3. Assess Airflow
4. Evaluate Temperature Sensors
5. System Controls
6. Check for Air Leaks
7. Review Historical Data

Results on Broken AHU Heating Coil

(LLM response – sample 2)



Operational Issues in Dual-Duct VAV Unit

- **Cooling valve overused**
→ 67.8% open, but zones still too warm
- **Damper too restricted**
→ Only 32% open, limiting airflow to zones
- **Heating valve inactive**
→ 0.0% despite high zone temperatures
- **Uneven zone temperatures**
→ Ranges from 74.6°F to 88.8°F – poor distribution
- **Energy inefficiency**
→ High cooling effort with low comfort delivery
- **Possible control mismatch**
→ Sensors, actuators, or logic may be misaligned

Future AI Models & Emerging Tech

1 **Multimodal AI Systems**

Combining text, image, and sensor data in a single model for comprehensive design review, installation verification, and operational diagnostics.

2 **Graph Neural Networks**

Understanding complex interdependencies in building systems—how changes to one zone affect pressure, temperature, and energy in adjacent spaces.

3 **Digital Twins with Real-Time Sync**

Live virtual replicas of HVAC systems that predict performance under various conditions and optimize ahead of actual building changes.

4 **Autonomous Design Agents**

AI agents that iteratively refine designs based on performance goals, cost targets, and constraints without human intervention between iterations.

Where are we at?

Three Years at a Glance

Tracking AI's journey from hype to production across three pivotal years

2024

The GenAI Gold Rush

- GenAI at **Peak of Inflated Expectations**
- Composite AI & Autonomous AI at **Innovation Trigger**
- Knowledge Graphs climbing **Slope of Enlightenment**

THEME Massive VC inflows, proof-of-concept proliferation

2025

The Reality Check

- GenAI slides into **Trough of Disillusionment** — avg \$1.9M spend, under 30% CEO satisfaction
- AI Agents & Multimodal AI hit **Peak**
- AI TRiSM & ModelOps climb toward **Plateau of Productivity**

THEME Governance and ROI pressure intensify

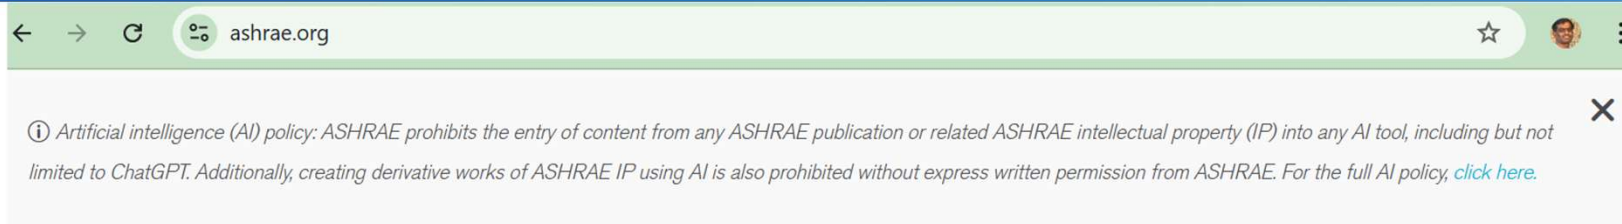
2026

The Agentic Era

- \$2.52T global AI spend projected (**44% YoY growth**)
- Agentic AI earns its own dedicated Hype Cycle
- GenAI recovers onto **Slope of Enlightenment**
- 71% of orgs have deployed AI agents — but only 11% in full production

THEME Orchestration and production scaling become the challenge

ASHRAE Policy for Use of AI (Feb 9, 2025)



ASHRAE Policy for Use of Artificial Intelligence (AI)

Board of Directors Approved February 9, 2025

Purpose:

The purpose of this Artificial Intelligence (AI) policy is to guide the ethical and effective use of AI technologies to serve ASHRAE members. As a publisher of technical guidelines and standards, organizer of conferences, provider of educational programs, and a research institution, ASHRAE strives to ensure that AI is used in ways that align with our mission, upholds intellectual property rights, and promotes fairness, transparency, and accountability.

GenAI.MTG

scope

*Initiated as part of
presidential theme by Dennis
M. Knight in 2023-24*

Support the adoption of generative AI by compiling technology adoption resources and best practice development guidelines

Undertake research projects to benefit ASHRAE members and for society operations and members services.

GenAI.MTG

status

*First meeting held on Feb 10, 2025
– Orlando, FL*

Three Work Groups:

1. genAI.Applications
2. genAI.ASHRAE Products
3. genAI.Tools

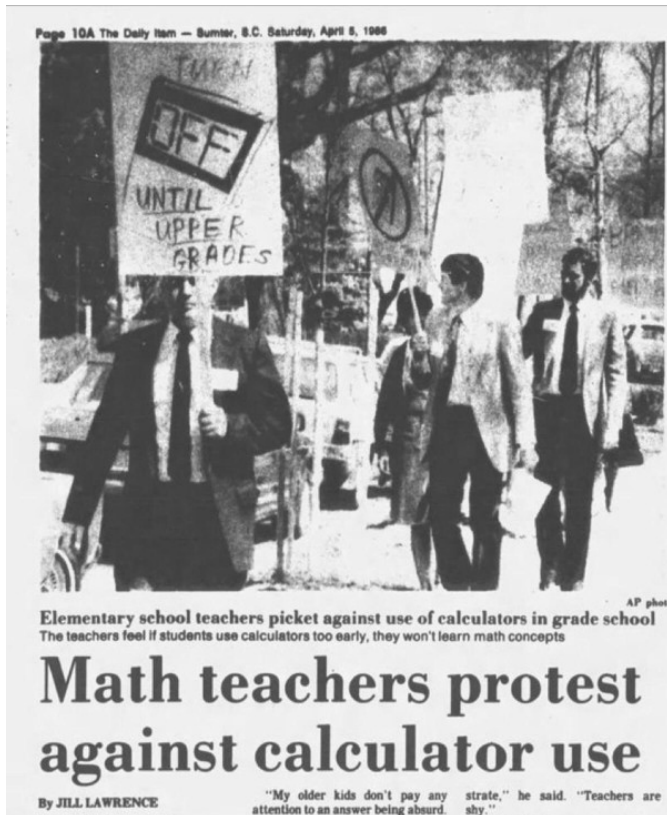
Programs:

Three programs submitted and scheduled for presentation at Phoenix, more initiatives and programs planned for Las Vegas, 2026

In Closing...

Key Takeaways

- **AI is a tool, not magic** - Uses math you can understand
- **Data quality matters most** - Good data + simple AI > bad data + complex AI
- **Start with the problem** - Don't use AI because it's cool
- **Correlation \neq Causation** - AI finds patterns, you determine meaning
 - The future is Physics+AI, not just AI



- Generative AI applications will transform all aspects of the HVAC engineering workflow.
- Human in the loop is critical, and AI will help improve productivity but not replace the human.
- Be thoughtful, cautious and drill down the sources before relying on AI tools.
- Adapt, learn and get ahead with AI

Questions

Contact Information:

Krishnan Gowri, Ph.D., BEMP, LEED AP, Fellow ASHRAE

kgowri@energy-solution.com

Krishnan.gowri@outlook.com

LEADERSHIP WANTED!

www.ashrae.org/volunteer

**BECOME A FUTURE LEADER IN ASHRAE –
WRITE THE NEXT CHAPTER IN YOUR CAREER**

ASHRAE Members who are active at their chapter and society become leaders and bring information and technology back to their job.

YOU ARE NEEDED FOR:

- ❖ Society Technical Committees
- ❖ Society Standard Committees
- ❖ Young Engineers in ASHRAE
- ❖ Chapter Membership Promotion
- ❖ Chapter Research Promotion
- ❖ Chapter Student Activities
- ❖ Chapter Technology Transfer



Find your Place in ASHRAE and volunteer