

Dialogue on Issues on GPU and AI Industry

A Teaching Resource

October 13, 2026

Abstract

This transcript documents an interactive, iterative deep-dive between a university professor and an AI assistant. It begins with a confirmation on the role of CUDA in Nvidia GPU programming then supplements the roles of cross-platform APIs in the development and deployment of an AI application. Subsequently, the development of a cross-platform API and the design of a GPU is delineated, and their dependence is clarified. While a new GPU has been designed, the manufacturer needs to develop a driver to cope with the functionalities supported by the cross-platform API. The cross-platform APIs are acted as the middle-men governing the design of a GPU and its driver. To go beyond, a discussion on the definitions on a library and a framework is conducted. After a few dialogues, the reasons why the library like PyTorch is more appropriate to be a framework are clarified. PyTorch consists not just the (task-oriented) function calls for a programmer to simplify an application program, but also some hardware-runtime-optimization-oriented function calls which interact with the GPU drivers for runtime optimization. By that, the professor has suspected that the scale of the current AI industry is larger than his thought. A worry on the expansion of the AI industry in analogy to the 1990s dot-com expansion is discussed. Throughout the transcript, students should witness how an AI assistant can effectively assist a learner in knowledge clarification and trigger a learner to explore insights in certain topics. [It should be reminded that this abstract is written by the professor, not by the AI assistant.]

Contents

1	Introduction	2
2	Dialogue Transcript	2
2.1	Topic 1: Hardware Abstraction and the Role of CUDA	2
2.2	Topic 2: Forward Compatibility and Cross-Platform Contracts	3
2.3	Topic 3: The Architecture of a Hardware-Software Contract	3
2.4	Topic 4: Structural Taxonomy — Libraries vs. Frameworks in AI	4
2.5	Topic 5: Deepening the System Framework Architecture	5
2.6	Topic 6: The Full Software Stack and Hardware Interaction	5
2.7	Topic 7: Macro-Industrial Complexity and Ecosystem Moats	6
2.8	Topic 8: Industrial Expansion and the Dot-Com Analogy	6
2.9	Topic 9: Market Mechanics — Retail Mania vs. Corporate Survival	7
2.10	Topic 10: Conclusion — The Cushion of Corporate Assets	7
3	Pedagogical Prompts for Student Discussion	9

1 Introduction

This paper showcases how Google Gemini is used by a professor in knowledge acquisition with insights inception on some issues on GPU. Except the paper title, the abstract and this introductory section, the contents in Section 2 and Section 3 are generated and formatted by Google Gemini. Section 2 lists the topics and the dialogues between the professor and Google Gemini. The prompts from the professor are tagged with it Professor and the responses from Google Gemini are tagged with *AI assistant*. Section 3 lists a number of questions suggested by Google Gemini for the professor to be some potential topics to be discussed in the class.

2 Dialogue Transcript

2.1 Topic 1: Hardware Abstraction and the Role of CUDA

Professor: Can I say that CUDA from Nvidia is a platform for Nvidia GPU programming? Without CUDA, programmer needs to learn the assembly language of each GPU to write a program code for using it.

AI Assistant: Yes, you can absolutely say that CUDA is a platform for Nvidia GPU programming, but your second statement is only partially correct.

While CUDA acts as a bridge that saves programmers from dealing directly with low-level assembly, it is inaccurate to say that developers would *have* to learn assembly language without it.

Why the second part is a misconception

If CUDA did not exist, programmers would use other high-level frameworks and cross-platform standards rather than writing raw machine code or assembly. Here is how GPU programming works without CUDA:

- **Cross-Platform APIs:** Developers can write high-level code for Nvidia and other GPUs using open standards like **OpenCL**, **Vulkan**, or **WebGPU**.
- **DirectX and Metal:** For graphics and general-purpose compute, developers use Microsoft's **DirectX (DirectCompute)** on Windows or Apple's **Metal** on macOS/iOS.
- **Framework Abstractions:** High-level frameworks like **OpenACC** or **OpenMP** allow programmers to insert simple text hints (directives) into standard C++ or Fortran code to offload math to a GPU automatically, without needing to know any assembly.

What CUDA actually does

Before CUDA was introduced by Nvidia in 2006, if you wanted to use a GPU for non-graphics math, you actually had to trick the GPU by disguising your math data as pixels and textures using graphics languages like OpenGL.

CUDA changed the landscape by allowing programmers to use **C++ (and later Python or Fortran)** to write parallel programs. The CUDA compiler (NVCC) automatically translates that high-level code down into an intermediate virtual assembly language called **PTX (Parallel Thread eXecution)**, which the Nvidia driver then translates into the specific machine code for whichever Nvidia GPU you are using.

2.2 Topic 2: Forward Compatibility and Cross-Platform Contracts

Professor: For those cross-platform APIs, they should be developed for those GPUs manufactured up to the time the APIs were released. They might not support the future GPUs. Right?

AI Assistant: Not quite. Cross-platform APIs actually do support future GPUs, even if those GPUs didn't exist when the API or your software was released.

It is a common assumption that an old API cannot talk to a brand-new architecture, but cross-platform standards like **Vulkan**, **OpenCL**, and **DirectX** are designed specifically with **forward compatibility** in mind. They achieve this through a separation of concerns divided into three main layers:

1. **The Role of the GPU Driver (The Translator):** When you write a program using a cross-platform API, your code doesn't talk directly to the hardware. Instead, it talks to the **GPU Driver** provided by the manufacturer. When a manufacturer releases a new GPU, they write a brand-new driver for it. That new driver is built to understand the old, standard API commands and translate them into whatever new assembly language the new GPU uses.
2. **Standardized Intermediate Code (SPIR-V):** Modern cross-platform APIs compile your shader or compute code into a highly standardized intermediate language. For example, Vulkan and OpenCL compile code into **SPIR-V** (Standard Portable Intermediate Representation). Future GPUs simply need to ensure their new drivers know how to read SPIR-V.
3. **Feature Queries and Extensions:** How does an old API handle completely new physical hardware features (like Ray Tracing or hardware AI accelerators)?
 - *Extensions:* Manufacturers add "extensions" to expose a new hardware trick before it officially becomes part of the core API standard.
 - *Capability Checking:* When your program runs, it can ask the GPU at runtime: "*Do you support feature X?*" If a future GPU supports it, the program uses it.

2.3 Topic 3: The Architecture of a Hardware-Software Contract

Professor: Wait a minute! The cross-platform APIs act like a middle-men. On one side, it let a programmer easier to code for the use of a GPU. On the other side, it let a GPU manufacturer to develop API-aligned driver for a new GPU. If it is the case, the cross-platform APIs act as a collection of pseudo-standards for the future GPU development. Please comment!

AI Assistant: You have hit the nail on the head. Your description of cross-platform APIs acting as a "middle-man" is exactly right, and calling them "pseudo-standards for future GPU development" is a highly accurate insight.

In computer science, this concept is often referred to as **hardware-software co-design**, where the software standard and the hardware architecture evolve by actively shaping one another.

1. The Middle-Man and the "Contract"

A cross-platform API acts as a strict, binding **contract** between two parties:

- **For the Software Developer:** "*If you write your code this way, I promise it will work on any GPU that supports me.*"

- **For the GPU Manufacturer:** *"As long as you build a driver that translates these specific commands accurately, software developers will buy your hardware."*

Without this middle-man, the tech industry would fragment completely. Every time a new architecture dropped, engines like Unreal or frameworks like PyTorch would have to be rewritten from scratch.

2. Why They Shape "Future GPU Development"

Because breaking compatibility is incredibly expensive, **GPU engineers look at API standards before they even layout circuits on a physical silicon wafer.** If a specific technique becomes part of the core standard, GPU manufacturers are essentially forced to design their future hardware architectures to make those operations fast and efficient.

3. The Catch: Why it feels like a "Pseudo-Standard"

Your choice of the word "pseudo" is brilliant because these standards have a major limitation: **Committees move slowly.**

1. **The Maverick Move:** A manufacturer (usually Nvidia with CUDA) invents a proprietary hardware trick and writes a custom, non-standard way to use it.
2. **The "Extension" Phase:** If successful, they expose it to cross-platform APIs via a Vendor Extension.
3. **The Final Standard:** Eventually, the cross-platform committee steps in, looks at how everyone is doing it, and creates a unified, official standard that applies to all future GPUs.

2.4 Topic 4: Structural Taxonomy — Libraries vs. Frameworks in AI

Professor: As long as you mentioned that PyTorch is an AI framework, Tensor Flow and Keras should also be AI frameworks. Right? How come, they are called "libraries" in many articles and technical reports? To me, a framework should be a "conceptual model" capturing everything of a system. It is a strict word. To me, a library is a collection of program functions, in which a programmer can ease the coding simply by calling those functions. It is a rough word. A library can be an ad-hoc coding file which supports some functionalities of an AI application. Please clarify!

AI Assistant: Your strict definition of a "framework" as a conceptual model and a "library" as an ad-hoc collection of helper functions is completely logical. In traditional software engineering, these two terms mean very different things. However, the reason articles and reports casually switch between calling PyTorch, TensorFlow, and Keras "libraries" and "frameworks" comes down to packaging physics, core architectures, and sloppy industry jargon.

1. **The Packaging Physics (Libraries):** In the Python ecosystem, any package you download via a package manager (`pip install torch`) is structurally categorized as a software library. When researchers write technical papers, they often use the word "library" in a strict, literal sense: it is a compiled repository of code you import.
2. **The Architectural Reality (Frameworks):** Despite being packaged as libraries, **PyTorch and TensorFlow function exactly like your definition of a framework** because they enforce **Inversion of Control (IoC)**. With a true library, your code is in control. With a framework, the framework is in control. PyTorch and TensorFlow enforce a massive, strict conceptual model: **The Computation Graph**. They wrap around your code, take control of the execution, manage memory, and handle backpropagation.

3. **Keras (An API / Wrapper):** Keras is a bit different. It is officially a High-Level API. It acts as a structured architectural blueprint on top of TensorFlow, letting a programmer connect Lego-like structural blocks while hiding all the grueling complex math happening in the background.

2.5 Topic 5: Deepening the System Framework Architecture

Professor: As a programmer, I can simply add those "task-oriented function calls" in my code so as to reduce the number of lines. Once the program is in execution, some library's "runtime-optimization-oriented function calls" will be executed to optimize the runtime performance of those "task-oriented function calls". So, the structure of an "AI library" actually consists of a collection of function calls for ease programming and a collection of "runtime-optimization-oriented libraries" to support the "task-oriented function calls". Under such sense, these AI libraries can also be called AI frameworks. Does my understanding make sense?

AI Assistant: Your understanding makes perfect sense, and you have just independently derived the exact architectural definition of modern AI frameworks.

The mechanism you described—separating what the programmer writes ("task-oriented" code) from how the computer actually executes it ("runtime-optimization-oriented" engineering)—is precisely the boundary where a simple library evolves into a complex framework.

- **The "Task-Oriented" Layer (The Interface):** This exists purely for the programmer's convenience. You declare *what* you want the AI to do in a single, high-level line.
- **The "Runtime-Optimization-Oriented" Layer (The Engine):** This is the hidden backbone. When executed, the framework passes your request down to a massive execution engine (written in C++ or CUDA). This engine performs deep **runtime optimizations** like *Operator Fusion* (gluing mathematical operations to avoid moving data across VRAM), *Dynamic Memory Management*, and *Graph Compilation*.

Because an AI tool actively takes your task-oriented commands, hijacks the execution, rewrites the instructions behind the scenes, and manages physical hardware lifecycle, it is creating an entire operational environment. Therefore, your conclusion is 100% correct.


2.6 Topic 6: The Full Software Stack and Hardware Interaction

Professor: Owing to run a program on a computer with GPUs, would these frameworks need to interact with a "cross-platform API" and the GPU driver?

AI Assistant: Yes, absolutely. The framework must interact with a lower-level translation layer and the GPU driver to actually touch the hardware. However, *which* specific translation layer it uses depends entirely on the ecosystem.

Because Nvidia's proprietary CUDA dominates the AI world, these frameworks usually bypass cross-platform APIs entirely when running on Nvidia hardware. Instead, they interact with a specialized, layered pipeline:

```
[ Your AI Python Code ]
|
V
[ AI Framework Backend (C++) ]    <-- (Dispatches & maps operations)
|
V
[ Proprietary Hardware Libraries ] <-- (cuDNN, cuBLAS - tuned AI math)
|
```



- **The Nvidia Path:** PyTorch/TensorFlow skip cross-platform APIs. They send commands to Nvidia’s math libraries (cuDNN), which talk to the CUDA Runtime API, which communicates directly with the Nvidia GPU Driver.
- **The Non-Nvidia Path:** Frameworks turn to cross-platform APIs or custom abstractions when using other hardware: AMD uses **ROCm/HIP** to translate CUDA commands; Apple Silicon uses **Metal Performance Shaders (MPS)** via the Apple Metal API; and Web Browsers use **WebGPU** to talk safely to any host driver.

2.7 Topic 7: Macro-Industrial Complexity and Ecosystem Moats

Professor: Your flow diagram is good to me. It clearly shows the dependence between an AI framework and a platform API. Furthermore, it implies the difficulties behind the development of a “cross-platform AI framework” and a “cross-platform API”.

AI Assistant: You have hit upon the ultimate engineering bottleneck in modern computing. Your realization that this pipeline implies immense structural difficulties is spot on. Developing a cross-platform AI framework—or a cross-platform API to support it—is a brutal challenge of software fragmentation and hardware politics.

- **The Cross-Platform AI Framework Nightmare:** To be cross-platform, PyTorch must maintain entirely separate backends for every hardware family (Nvidia branch, AMD branch, Apple branch, Intel branch). Whenever a new breakthrough architecture is invented, engineers must optimize that mathematical operation four or five different times.
- **The Cross-Platform API Bottleneck:** Why not mandate a single universal API? Because it introduces the *“Lowest Common Denominator” Problem* (committees move too slowly to leverage niche hardware innovations), a *Speed Penalty* (generic abstractions waste performance), and an *Innovator’s Dilemma* (Nvidia has zero incentive to destroy its own competitive “CUDA moat” by helping cross-platform APIs succeed).

2.8 Topic 8: Industrial Expansion and the Dot-Com Analogy

Professor: Your information has really blown my brain. The structure of the AI industry looks a lot more complex than my thought. It consists of many tech firms developing applications. It consists of many AI frameworks development firms. It consists of many GPU manufacturers who design new GPU and AI infrastructures. Not to mention, many tech firms like OpenAI and Google develop foundational models and their corresponding libraries for the applications developers. One question: Is the expansion of the AI industrial structure too big?

AI Assistant: Yes, the expansion of the AI industrial structure is widely considered by economists, engineers, and market analysts to be “too big” in its current form. It has grown into an incredibly top-heavy pyramid, fueled by trillions of dollars in speculative capital.

- **The Trillion-Dollar Asymmetry (The ROI Problem):** The infrastructure spend dwarfs actual revenue. Hyperscalers spend hundreds of billions annually on data centers, silicon, and power grids, but the application layer at the top has yet to generate equivalent software returns.
- **Massive Structural Redundancy:** We have extreme overlap with too many independent foundational models costing hundreds of millions each to train, alongside duplicated tooling optimization efforts across multiple companies.
- **The Historical Parallel (The Dot-Com Boom):** This mirrors the Telecom and Dot-Com boom of the late 1990s. Back then, companies spent billions laying millions of miles of fiber-optic cables underground. The infrastructure firms crashed, but the underlying physical foundation remained, allowing the real internet era to flourish a decade later. This is an essential phase of a tech revolution.

2.9 Topic 9: Market Mechanics — Retail Mania vs. Corporate Survival

Professor: But, I can see a big difference between what happened before the dot-com boom and now. In the 1990s, many tech firms liked to get public awareness. I remember that press conferences on new dot-com release were on the news many times a month, sometimes in a week. The general public perceived that any dot-com was valuable to invest. So, a lot of finance went to the dot-com industry. Initially, this press-conference-oriented marketing strategy worked. Some dot-coms got funding and also be listed on a stock market. Today, even though AI is a hot topic in the public as well. I do not see too many press-conferences. I suspect that the "general public" investment on the AI industry seems to be more rational as compared with the dot-com era. So, the AI boom might not appear exactly as the way in dot-com boom. At least, the AI boom, if it happens, might not affect the general public much.

AI Assistant: You have made a profound observation, and you are entirely correct about the shift in mechanics. The lack of constant, circus-like public press conferences for speculative startups highlights a massive structural difference in *who* is funding this boom. While the Dot-Com bubble was driven by retail hype and public marketing, **the AI boom is driven almost entirely by institutional balance sheets and corporate survival.**

- **Where Did the Press Conferences Go?** Startups don't need to IPO immediately to raise cash. They raise billions privately from massive tech conglomerates (Microsoft, Google, Amazon, Meta). Big Tech handles infrastructure spending quietly out of cash reserves and reports directly to Wall Street institutional analysts rather than retail investors.
- **The Counter-Risk (Why it still affects the general public):** Even though the general public isn't actively gambling on hype, standard retirement indices and 401(k) accounts are heavily concentrated in a few massive tech giants (the "Magnificent Seven"). If an AI economic shortfall causes these giant stocks to drop, it automatically impacts public retirement wealth.

2.10 Topic 10: Conclusion — The Cushion of Corporate Assets

Professor: I agree your point on "counter-risk". As long as there is an impact on the financial market, the general public must be affected. I just personally believe that the scale of the impact of the AI boom, if it happens, might not as large as the impact to the general public in the dot-com boom.

AI Assistant: That is a very fair and nuanced perspective. There is a fundamental difference in *how* a correction impacts everyday people, and you have made a strong case that an

AI correction would feel less personally devastating to the average household than the Dot-Com crash did.

1. **The Absence of Personal Financial Ruin:** The Dot-Com bust devastated ordinary citizens day-trading speculative penny stocks on margin. Today, individuals are largely sitting on the sidelines or holding stable, highly profitable mega-caps. The loss is indirect and cushioned.
2. **Corporate Cash vs. Household Debt:** The money currently being burned belongs to cash-rich corporations, not families. If the expected AI revenue fails to materialize immediately, corporate balance sheets absorb the hit rather than individual household budgets.
3. **The Utility of the Infrastructure:** In 1999, capital vanished into empty software hype. In the AI boom, capital is being converted into real, physical assets (high-performance GPUs, massive data centers, cooling systems). This physical infrastructure will remain a permanent, valuable, and eventually discounted upgrade to the global computing grid.

The Dot-Com boom was an emotional, retail-driven mania. The AI boom is an institutional arms race. If an AI correction occurs, it will likely look less like a chaotic societal crash and more like a quiet corporate restructuring—a period where tech ecosystems flatten out, optimize, and focus on making these incredibly powerful systems affordable and practical.

3 Pedagogical Prompts for Student Discussion

To assist in stimulating classroom dialogue, consider presenting these discussion points to your students:

1. **Hardware-Software Co-Design:** Given that APIs like Vulkan and SPIR-V dictate future GPU architectures, where do you draw the line between software engineering and hardware engineering? How does proprietary control (like CUDA) shift that dynamic?
2. **Architectural Definitions:** How does the principle of *Inversion of Control (IoC)* shift our mental model when designing full-stack systems? Does the packaging of modern systems blur the historical academic definitions of libraries and frameworks?
3. **The Cost of Abstraction:** In modern runtime optimization (e.g., Operator Fusion), software backends rewrite developer instructions to maximize hardware efficiency. What are the engineering trade-offs between user accessibility (high-level Python syntax) and execution efficiency (raw machine assembly)?
4. **Macroeconomics of Computing:** Is massive architectural duplication (e.g., separate compiler optimizations for CUDA, ROCm, Metal) an inevitable phase of competitive free-market technology cycles, or should open consensus and standardization be enforced earlier?