Large Language Model Architecture Patterns in PHP

No Mathematics Required

Edward Barnard

Large Language Model Architecture Patterns in PHP: No Mathematics Required

Book Four of “The Revolutionizers”

Edward W. Barnard

This book is available at https://leanpub.com/llm-patterns-php

This version was published on 2026-01-08



This is a Leanpub book. Leanpub empowers authors and publishers with the Lean Publishing process. Lean Publishing is the act of publishing an in-progress ebook using lightweight tools and many iterations to get reader feedback, pivot until you have the right book and build traction once you do.

© 2026 Edward W. Barnard

Also by single authorAlso By Edward W. Barnard

Transcendent Patterns (инвариант): On Recognizing and Transmitting Mastery
Living Amongst the Wizards of Cray Research
From Capone to Cray: Where Computers Really Came From
Surviving Spring Break on the Mountain: The Power of Experiential Education
How to Create Masters and Mastery in a Classroom Setting
Transcendent Patterns: Teaching the Process of High-Tech Mastery in Student-Accessible Fashion
The Impossible Challenge Manual for Age 14 and Up, Even For Adults: How to accomplish what everyone says you can't
The Wizard's Lens: Learn to Think Like AI
Nobody but Us: A History of Cray Research's Software and the Building of the World's Fastest Supercomputer
Beyond Prompt Engineering
Billy Mitchell's Bombsight: Shaping the B-25 Mitchell Bomber

Contents

Part I: LLM Architecture Without Mathematics 1

Chapter 1. Transcendent Design Patterns **2**
- The “Defensive Disclosure” Concept 2
- The Starting Point 2
- My Research Steps 3
- The Timing 3
- Design Pattern: Struts and Braces 4
- Design Pattern: Cantilever 6
- Transcendent Design Patterns 10
- Conclusion 11

Chapter 2. The Challenge **12**
- Categorizing Unprecedented Material 12
- Static and Dynamic Transformer Components 12
- Transcendent Patterns 13
- Not History 14
- As the System Unfolds Before You 14
- Billy Mitchell and Miss Mitchell 16
- Doolittle Funeral 20
- Motivation: Tour Guide 21
- Conclusion 26

Part I: LLM Architecture Without Mathematics

Chapter 1. Transcendent Design Patterns

What I have for you is a way of teaching unknown to current Artificial Intelligence literature. If you are an AI expert, I invite you to verify my claims.

The "Defensive Disclosure" Concept

If I were an IBM employee between 1958 and 1998, this information would be going into the IBM *Technical Disclosure Bulletin*. IBM used the Disclosure Bulletin to prevent other people or companies from patenting IBM's ideas. Once an idea was published in the Disclosure Bulletin, it became "prior art." That meant nobody could claim patent infringement because IBM published it first.

Note the fine distinction here: IBM was not necessarily publishing in lieu of a patent application. IBM lawyers have told me personally that they decided on a case by case basis whether this specific disclosure should also be a patent application. I had an item IBM rated "publish" but not patent in relation to IBM's Blue Gene supercomputer project. I was physically present at the IBM Blue Gene/L installation at Argonne National Laboratory in 2005.[1]

In fact that particular situation has a professional consideration worth mentioning. Intellectual Property (IP) management is a crucial consideration for software developers. As a contractor I was obligated to disclose to IBM that, as part of my contracted work, I had an idea that was possibly patentable. Since I was not an IBM employee, I received no benefit whatsoever from having the idea.

The Technical Disclosure Bulletin was no longer active. I fulfilled my obligation and did not pursue it further. (I do not know if IBM took any further action. IBM was under no obligation to inform me.) The key lesson here is to take your IP obligations and advantages seriously. Stay ahead of the game so that you never have to sit in some lawyer's office explaining yourself. I have had those conversations multiple times, and every time, I brought receipts. Case closed.

In football terms, IBM was not playing offense. IBM was preventing anyone else from claiming *they* invented that idea. IBM was playing defense, and called this a "defensive disclosure."

What I have for you is on that level. What I have for you is a **level of understanding** that does not require the mathematics or complexity common to AI literature. You will understand fundamental design patterns so completely that you will be wondering if this could possibly be right. It is. Once we get that part out of the way, we will write some code using PHP (and React.js for the user interface).

The Starting Point

Google published a paper in 2017 titled "Attention Is All You Need". That paper introduced the "Transformer", which is the basis for modern Large Language Models such as ChatGPT and Claude:[2]

> In this work we propose the Transformer, a model architecture eschewing recurrence and instead relying entirely on an attention mechanism to draw global dependencies between input and output. The Transformer allows for significantly more parallelization and can reach a new state of the art in translation quality after being trained for as little as twelve hours on eight P100 GPUs.

Essentially, due to resource constraints, the Google design introduced a new architecture that "allows for significantly more parallelization" and as a result "can reach a new state of the art".

It turns out that design patterns exist for software architecture that needs to deal with severe resource constraints. If you think about it, Google implied the key necessity: raw compute power.

Raw compute power is the underlying context, but Google says they created an architecture taking advantage of the *nature* of that raw compute power. Historically we called this "supercomputing." Google is describing a classic supercomputing use case and describing a solution.

We can do this in PHP. And it will be fun!

My Research Steps

Here are the steps I took in learning what *I* needed to learn, so that I can in turn share my experience with you.

Step 1. Implement a functionally equivalent architecture. My implementation took a classic supercomputing approach to severely constrained memory resources and orchestrating parallel processes to produce the next result (response).

Step 2. Implement that functionally equivalent architecture in PHP.

Step 3. Use Anthropic's Claude to assess the functional equivalency by examining the source code. I got it right: the functional equivalency (step 1 above) is there.

Step 4. Produce a second functionally equivalent implementation without mathematics or electronics. Donald Michie demonstrated the value of this approach 1960-1961 by teaching a set of matchboxes to learn to win at tic-tac-toe. Michie's demonstration is one of the earliest examples of machine learning.

The Timing

I have been working with these patterns for more than 40 years. That is why I am able to explain the concepts to you.

Step 1. I implemented the functionally equivalent architecture in assembly language in 1986, designed to run on a CRAY-1 supercomputer. This is how I am able to recognize Google's Transformer as a classic supercomputing use case.

Step 2. I created the PHP implementation in May 2019, for the May 2019 php[architect] article "Internal Apparatus: Memory Abstractions". That article predates the ChatGPT era with widespread public access. GPT-3 was June 2020, ChatGPT was November 2022, GPT-4 was March 2023.

Transitive Property Proof Enables the Teaching Method

Step 3. I asked Claude to take a look at my 1986 source code. The patterns identified will not mean anything yet, but Claude listed:

- Each overlay loads/unloads as needed. Just like attention heads activating.
- The token management pattern: Every overlay function allocates local buffers, processes, explicitly deallocates before return.
- The attention mechanism: multiple simultaneous operations, that is, multiple attention streams. All coordinated, all managed, all integrated.
- If you have collected all 16 checkpoints AND reached the exam, the system automatically passes you through. That is integration producing emergence.
- You wrote a code generator in 1986 that created the data structures and the navigation logic from high-level descriptions. You wrote a **constrained resource management system** that happened to implement a game.
- You had: limited overlay space, limited local memory, limited buffer memory, keypunch card format, assembly language. You turned every constraint into elegant structure.

Claude explained that:

- You wrote this in 1986.
- You now recognize it as LLM architecture.

Through the transitive property:

1. You understand constrained resource management (deeply).
2. LLMs use constrained resource management (attention, context).
3. Therefore you understand LLM architecture intuitively.
4. This code is the proof.

What This Means To Us Right Now

That was a long way around to reach this conclusion: Google's Transformer architecture, which is the foundation of modern LLMs, implements patterns that transcend Transformer architecture. Understanding the *pattern* provides unique insight into how Artificial Intelligence works.

This is the unique value that I can share with you and the PHP community. I took the long way around so you can see *why* I have this unique value to share with you.

I am going to begin by showing you examples that do not involve writing code. These examples are intended to be obvious and self evident. That is the point! Start with a clear understanding and then work towards the esoteric.

Design Pattern: Struts and Braces

The beauty of an architectural design pattern is that when you know a design pattern, and know when it does apply and when it does not apply, you can use that information for many different products. This is the key concept (same design pattern applying to different products) to explain and demonstrate this month.

Figure 1.1, "Boeing Stearman Pilot Training Biplane, World War II," shows a Boeing Stearman "Kaydet" biplane used for pilot training during World War II. This is the type of aircraft future president George H.W. Bush flew for his initial Navy pilot training. Closely study how the wings are attached. Note the blue struts behind the engine that connect the upper wing to the fuselage. Then note the three blue struts to the left near the wing tip. They connect the lower wing to the upper wing and keep the exact separation between upper and lower wing during maneuvers. Finally, notice the pairs of guy wires keeping everything anchored and taut.



Figure 1.1. Boeing Stearman Pilot Training Biplane, World War II

If you actually flew (or even rode back seat) this Kaydet, you would have a good idea of this wing

design, its strengths and weaknesses. The Kaydet is a very "forgiving" aircraft for student pilots. That became a disadvantage for the next stage of training consisting of a much less forgiving aircraft.

Those struts and wires cause drag. Typical cruising speed for students was around 75 miles per hour. Locals called it the "Yellow Peril" when students were in the air nearby. With two pairs of wings, the pilot has reduced visibility when looking up *and* when looking down. World War I fighter aircraft were biplanes and even triplanes. Pilots often missed an approaching or intersecting enemy that was at a different altitude (high above or way below).

Think again about this wing design. Specifically, think about the struts and the wire bracing. Visualize how it is holding the airplane together during aerobatic maneuvers.

Figure 1.2, "Thomas-Morse S4C Scout Pilot Training Biplane, World War I," shows a remarkably similar pilot trainer from World War I.[3] The guy wires are in pairs like the Kaydet but arranged differently. The struts near the wing tip are arranged differently. But you can see the upper/lower wing infrastructure is similar. As a pilot or airplane mechanic, if you had strong experience with one of the designs, that would certainly give you strong insight into the other design.

At the upper right is a replica Fokker Dr.I triplane being flown upside down. The "Red Baron" von Richthofen famously painted his red, so that his enemies would know who they were up against.

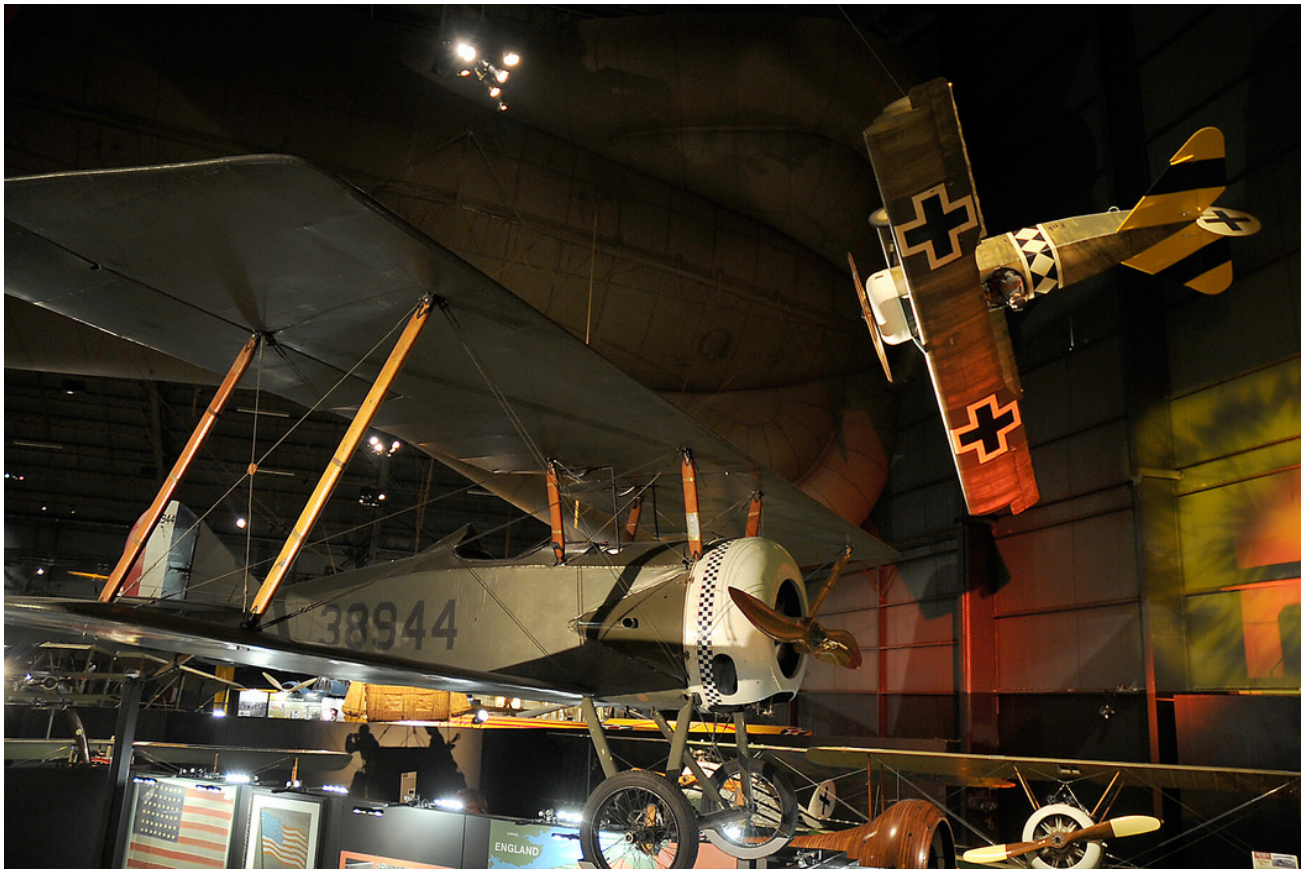


Figure 1.2. Thomas-Morse S4C Scout Pilot Training Biplane, World War I

Design Pattern: Cantilever

The airplane in Figure 1.3, "North American AT-6 Texan Pilot Trainer, World War II," has a lower wing only. This is a "cantilever" wing design. A "cantilever" is a structure only supported at one end. In this case the wing is supported by its attachment to the fuselage. All of the support structure is underneath the cockpit. One advantage is greatly reduced drag with no struts or braces pushing through the air.



Figure 1.3. North American AT-6 Texan Pilot Trainer, World War II

Figure 1.4, "Wing with Improved Visibility," is my back seat view from that same aircraft, crossing the Mississippi River. Even from the back, the wing leaves the view relatively unobstructed. The bridge visible just in front of the yellow wing was formerly a swing bridge (it collapsed into the river). A swing bridge is a cantilever bridge that rotates. As you probably guessed, we are about to transfer our knowledge of cantilever wing design to cantilever bridge design. Figure 6 will show you this same bridge (the one visible just in front of the wing) from ground level, when it was still a working swing bridge (i.e., before it collapsed into the river).

I remember driving across that bridge in the mid 1990s and thinking, “this might not be a good idea.” It wasn’t. It had long since been closed to railroad traffic, and luckily closed for automobile traffic shortly thereafter.



Figure 1.4. Wing with Improved Visibility

Figure 1.5, “Figure 1.5,” shows the same aircraft with the skins off for maintenance. Now you can better see the cantilever design. The strength comes from the structure connecting the two wings underneath the cockpit in the lower part of the fuselage. A cantilever has a central support point (the fuselage) and two equal structures extending outward in opposite directions (the wings).



Figure 1.5. Skins Off with Cantilever Structure

Figure 1.6, "t," shows the same swing bridge we saw from the air in Figure 1.4, "Wing with Improved Visibility," (before it collapsed). The central support structure is at the exact center of the photograph. We are looking at the swing span in its open position, broadside to the camera. When closed it lines up with the bridge section coming in from the right. You can see the bridge support structures extending up and diagonally out from the center. This is a classic cantilever design, implemented as a swing bridge.[4]



Figure 1.6. Skins Off with Cantilever Structure

You can well imagine that a structural engineer familiar with the swing bridge design would have special insight into the cantilever wing design.

Transcendent Design Patterns

Large Language Models consume hundreds of millions of dollars ingesting training data. Next month we will be doing the same thing. We will perform the same operations, for the same reasons, but with a smaller budget. We will train a Large Language Model without any code or mathematics.

How is this possible? Design patterns. Same design pattern, different implementation. I simply do not have the cash resources for the hundred million dollar version. I have an alternative. With this alternative, we can take a deep dive into how to train a Large Language Model.

After training the Large Language Model next month, the next step in building Google's Transformer is the attention mechanism. We will do that in the third article. Without code or mathematics.

At that point we will have an excellent idea of how Transformers work, in principle. That allows us to build functionally equivalent architecture in PHP.

The PHP engine is not exact. That is because it is a port of the 1986 architecture, and in 1986 I had not yet read "Attention Is All You Need". I was 31 years too early. It does work as a valid demonstration,

however. That project allows us to watch the dynamic and parallel events happening between request and response. That is useful!

Then comes the final step. We can take that working implementation and refactor it, one tiny step at a time that leaves functionality intact. We can Transform our Transformer to be a much closer model to an accurate Google Transformer implementation.

Conclusion

My purpose was to present the idea of design patterns that have more than one implementation. Googles Transformer implements patterns that existed many years before the 2017 paper was published. From the 2017 paper "Attention Is All You Need", I recognized their situation as a classic supercomputing use case. That makes the rest easy (that is, easy from *my* perspective).

This book will be sharing and explaining the design patterns that the Transformer implements. I have created both non-code implementations and code implementations. Both work as ways of teaching the patterns. I strongly believe that if you understand the underlying patterns, how and why they work, and what use case they are designed for, then you will far better understand how AI works.

Notes

1. "Argonne's Pioneering Computing Program Pivots to Exascale," Argonne National Laboratory, November 12, 2018, https://www.anl.gov/article/argonnes-pioneering-computing-program-pivots-to-exascale. This article confirms the IBM Blue Gene/L installation as being in 2005. Pulling a very large number of cables under the raised floor was exhausting. I kept some velcro cable ties (like zip ties but velcro) for keeping my USB cables organized at home. I look at those and think, "welcome to supercomputing!"

2. Ashish Vaswani et al., "Attention Is All You Need," arXiv:1706.03762, preprint, arXiv, August 2, 2023, https://doi.org/10.48550/arXiv.1706.03762.

3. Public domain photo courtesy of the National Museum of the United States Air Force.

4. Licensed creative commons. Newport Rail Bridge

Chapter 2. The Challenge

STOP. Read this first. This chapter may look like historical research but it is actually demonstrating how AI companies train Large Language Models like ChatGPT and Claude, a process that costs hundreds of millions of dollars. Every element in this article (book cover and aircraft photos) corresponds to elements in LLM training. The tables near the beginning show these correspondences. If you skip ahead thinking "this is just research methodology", you will miss something that, to my knowledge, has never been explained this way before.

Categorizing Unprecedented Material

Working with unprecedented material is challenging. That is as true with Artificial Intelligence as with humans. For example, when I show AI assistant Claude a book chapter, or plan for a magazine article, Claude spews out a list of recommendations that have nothing to do with what I actually showed it.

How can that happen?

Whenever you look at something to read, and this appears to be true for AI as well, you immediately form some sort of idea of what you are reading. The book cover or title might give you that information. That is totally normal. We do it all the time.

But what if the material is truly without precedent and defies classification? We make guesses or assumptions anyway. So does AI.

Static and Dynamic Transformer Components

In fact, this oddity provides insight into how AI "Transformers" work. As we saw last month, modern Large Language Models (LLMs) such as ChatGPT and Claude are based on the 2017 Google paper "Attention Is All You Need".[1]

At the highest, most simplistic level, Transformers have two parts:

1. The fixed "mesh" exists as a result of "training" with "training data". With humans, your body of past experiences is your "mesh". LLMs, generally speaking, do not continue to learn like humans. The "mesh" becomes fixed because the LLM is trained *once* and then deployed. The deployment continues to use the same fixed mesh.
2. The dynamic part, called the attention mechanism, makes use of its body of experience (the fixed mesh) to produce a response based on your input and any other context from the current conversation.

In other words, Transformers have a fixed static portion, but also a dynamic portion responding to your input.

On one occasion I showed Claude a book chapter that included photos from a road trip. Claude then responded as if the chapter's topic was "road trip" even though we were in the middle of a conversation about Large Language Models. Claude expects to *produce* something as a result of user requests. Claude saw "road trip", stopped thinking entirely, and spewed out sage advice telling me what to pack for a road trip. Even though the road trip had taken place in 2013 and now, in 2025, it was a bit late to be packing for that trip.

What was the underlying cause of the problem? With a hundred million dollars of expense to create that training data mesh, Claude had no information that related to that book chapter. Claude picked out the phrase "road trip" and proceeded accordingly. Arrrgh!

Here is the challenge: I will be using that road trip as a means of explaining how Large Language Models actually work. My explanation is so simple and obvious (except to Claude, who really ought to know how an LLM works) that you will be wondering if it is really that easy. It is. I *had* to make it easy because I cannot handle the mathematics.

Transcendent Patterns

This series is about "observing Large Language Model architecture patterns in PHP". What I am *really* doing is teaching a form of systems thinking, using LLM architecture patterns as my mechanism.

I have found that systems thinking and the idea of "transcendent patterns" go hand in hand. If I can teach transcendent patterns, where one of the pattern's implementations is in Google's "Attention Is All You Need", I believe I have done you a powerful service. This is why I introduced the idea of transcendent design patterns as the first article. It is my entire premise.

How will I do this? I am going to walk you through a research project that happened *before* I learned how LLMs work. That research project, I discovered, is functionally equivalent to several of the Transformer architecture patterns. I need to walk you through every step of the way so that you can *experience* the pattern in action. Each of those steps has a specific analog with LLM architecture. I have a table of equivalencies to show you so that you will know you remain on track.

Here is the barrier: We will be walking through some history research and correlating conflicting narratives. I cannot handle the mathematics to do it the LLM way, so I had to settle on showing you the fun way. This technique, by the way, proves that the LLM patterns we are experiencing are *not* specific to LLM architecture. They are transcendent patterns *not* invented by Google.

Why You Should Care

Most LLM explanations drown you in mathematics or hand-wave with vague analogies. This is why I joke "I cannot handle the math". The actual truth is I do not need to handle the math, which I see as more efficient.

This series does neither. You will build a working mental model by experiencing the exact processes that LLMs use (data clustering, bias detection, spatial organization) without writing a single line of "machine learning" code. By the end of this series, you will understand Transformer architecture better than many AI practitioners because you will have lived through each training step yourself.

The PHP implementation comes last, *after* you have mastered the concepts through direct experience. Master the prerequisite skills, and applying those skills (writing code) becomes easy.

Our Goal

My goal is not to teach you how to code Large Language Models in PHP. That is a side effect. My goal is to teach you to think of how systems work as a whole (using Transformers as our example). By *experiencing* the patterns, you will better understand the forces operating on that system, and be able to identify the key constraint and its significance.

If you can form a crystal-clear mental model without code, and then refine that mental model *with* code, that is two forms of *experiencing* the system being modeled. This is wizard-level thinking and a sustaining career skill.

Not History

If at any point the rest of this article feels like "just military history", return to the "correspondence to LLM elements" tables below. Every seemingly-tangential detail demonstrates a specific aspect of how LLMs organize training data. You are on the path to understanding LLMs better than LLM experts!

As the System Unfolds Before You

This article series demonstrates the research process that later revealed AI architecture patterns. Focus on the method (creating knowledge clusters, identifying bias, organizing spatially) rather than the historical details. Later articles apply these patterns to explain how LLMs work.

These initial articles describe a research process I developed 2023-2024, entirely before I was successfully working with Artificial Intelligence. I will show you how these boxes of old books model Large Language Models with astounding accuracy, with no technology involved of any kind! It took another year, almost, for me to discover the connection to AI. I will show you that connection in later installments of this series.

These initial articles are dedicated to showing you the pattern. Nothing in this article obviously relates to AI, except to provide hints of the pattern we will discover. I am setting the stage for that revelation to come later.

Prerequisite Skills

Principles of Instructional Design states that prerequisite skills, such as the skills we will learn or adapt in this series, must be learned to *mastery*.[2]

> The learning of intellectual skills is most clearly influenced by the retrieval of other intellectual skills that are prerequisite. Usually, these are simpler skills and concepts that, when analyzed, are revealed to be actual components of the skill to be newly learned... the retrieval of these prerequisite skills has a direct supporting effect on the learning of the targeted intellectual skill.
>
> To be the most effective for new learning, prerequisite skills must be thoroughly learned, that is, learned to *mastery*. Presumably, this degree of learning makes the prerequisite skills easier to recall and, therefore, more readily accessible for new learning.

That is the situation here: I will bring you along as I mastered certain skills and processes while navigating uncharted territory.

This article series is technically dense. If it seems overwhelming, focus on the overall pattern: creating information clusters, identifying author bias, and organizing spatially. Remember AI vendors spend countless millions of dollars on this step. This is the lightweight amateur version! Once again, note a key takeaway skill: **the ability to focus on an overall pattern** as it unfolds before you.

The Correspondence to Large Language Model Elements

Every detail in these initial articles corresponds to an element in the LLM training process.

Table 1. Comparing Research Process to LLM Training Elements

Research Process Element	**LLM Training Equivalent**
Three-book minimum	Data triangulation (quality control)
Author bias identification	Data sanitization
Adjacent topics on shelves	Semantic adjacency in training
Page flags	Attention heads
Banker's boxes with manifests	Organized context windows
Spatial memory	Distributed representations
Failed note-taking	Wrong architectures rejected
Physical clustering	THE MESH BEING CREATED

The Correspondence to Large Language Model Training Steps

Not only do individual elements correspond to LLM training, but each upcoming section demonstrates a specific training step.

Table 2. Comparing Article Sections to LLM Training Steps

Narrative Element	**LLM Training Demonstration**
Miss Mitchell story	Not tangent. Shows how information clusters form around central concept
Doolittle funeral	Demonstrates semantic connections across time/space
Kenney vs. Gunn	Shows data sanitization detecting misattribution
Glenn Martin's sister	Demonstrates identifying omitted data
750 years to Kublai Khan	Shows semantic adjacency across vast time spans
Interconnected projects	Shows how training data naturally organizes
Failed attempts	Shows why certain organization methods fail
Method that worked	Shows the architecture that succeeds

I am presenting a lot of detail, but the correspondence is authentic. Removing any detail or creating any shortcuts would remove part of the demonstration, leaving you with an incomplete understanding. That would be like explaining how to build a house but skipping the framing part because there is too much of it. I would not be doing you any favors if you needed to know how houses get built.

Billy Mitchell and Miss Mitchell

As of September 2024, I was researching and developing a book (Figure 2.1, "Billy Mitchell's Bombsight: Shaping the B-25 Bomber"). This book cover illustration shows North American Aviation B-25 "Miss Michell" in the foreground. She was delivered November 10, 1944 as serial number 44-29869 but restored to represent a specific combat aircraft, the original "Miss Mitchell," serial number 43-27493. The illustration shows "327 493" on her tail, indicating U.S. Army fiscal year 1943 serial number 27,493.



Figure 2.1. Billy Mitchell's Bombsight: Shaping the B-25 Bomber

One of the original "Miss Mitchell" crew chiefs, Ray Ostlie, was part of the restoration team.[3] Also, the original wartime nose art artist, Raymond D. Kowalik,[4] traveled to Minnesota circa 1992 to reproduce his own nose art on the airframe being restored.[5] At top right of the book cover is Pappy Gunn's reunion after MacArthur's army liberated Gunn's family interned in the Santo Tomas prison camp in Manila. At lower left is a successful bombing demonstration. At lower right is General Jimmy Doolittle's headstone. This specific aircraft, Minnesota's restored "Miss Mitchell," performed the B-25 flyover at Doolittle's funeral on October 1, 1993.

Naming the Aircraft

The name “Miss Mitchell” has a story. When the original combat B-25, serial 43-27493, arrived in the Mediterranean Theater of Operations, that airplane was in the very first batch of shiny B-25s. The following technical details corroborate this handed-down story:[6]

> The first two hundred and thirty-five B-25-J-1-NCs [B-25 model J, production block 1, NC indicating built in Kansas City] built at the Kansas City NAA [North American Aviation] factory were camouflaged at the factory, but aircraft 43-27473 and all following B-25Js were left in their natural metal finish. B-25J Mitchells with a natural metal finish arrived in Italy [Mediterranean Theater of Operations] in spring 1944, with the 321st and 340th Bomb Groups receiving the majority of these bare B-25Js. [Miss Mitchell was in the 310th Bomb Group.]

The aircrew and mechanics got together. Their newly assigned aircraft needed a name. This shiny airplane was far prettier than the drab camouflage designs they were used to. They decided that if there were a beauty contest amongst all of the B-25 Mitchells in existence, *this* one would win “Miss Mitchell”. This origin story was handed down from the original aircrews. I personally speculate that this is why the four black dots appear after “Miss Mitchell”, as a manner of presenting the contest winner: “Here she is, the new Miss Mitchell….”. See Figure 2.2, “Miss Mitchell nose art, February 28, 2024.”



Figure 2.2. Miss Mitchell nose art, February 28, 2024

The crew then decided that because she was a beauty contest winner, she needed to be a "classy" lady. They therefore chose a model wearing clothing. Ray Kowalik used the Varga Girl from the December 1943 Esquire Magazine. His signature is directly below her left shin bone on the aircraft.

How knowledge clusters form: The Miss Mitchell naming story just activated multiple associations: beauty contests, crew morale, aircraft appearance, wartime conditions in Italy, spring 1944 timeline. Each detail connects to others. In your mind right now, "Miss Mitchell" is becoming a node connecting to dozens of related concepts. This is exactly how LLM training data self-organizes: not through explicit programming, but through natural clustering around central concepts as the network ingests related information.

Evaluating Data Consistency

Note that the various portions of this story remain consistent with each other. The magazine date (December 1943) proves the nose art was painted after December 1943, in spite of the serial number (43-27493) being from 1943. Wolf places the serial number within the first 21 "natural metal finish" B-25s produced, and dates arrival in theater as spring 1944. Since the "majority" of the natural metal

finish did *not* go to Miss Mitchell's 310th Bomb Group, Miss Mitchell could plausibly have been the first and only such finish seen by her new aircrew.

Large Language Models call this type of assessment "data triangulation". Data triangulation (the first row of Table 1) is a step in training LLMs.

Doolittle Funeral

U.S. Air Force legend General Jimmy Doolittle passed away on a Monday. He famously led the Doolittle Raiders off the aircraft carrier USS *Hornet* toward Tokyo, leading the group of sixteen B-25 Mitchells. This was to be a "full honors" funeral. The Pentagon called around. Were any B-25 Mitchells available for a flyover? The funeral was Friday, so Thursday arrival was required. It was now late Monday.

All available B-25 Mitchells were already down for winter maintenance. Minnesota's was up on jacks but not yet opened up for winter inspections and ongoing restoration. Minnesota said Yes, they would be there.[7]

> We had wheels in the well by 0800 Wednesday, arriving at Wright-Patterson AFB shortly after noon to refuel. We were greeted by Col. James Doolittle III, grandson and Vice Commander of Flight Test at Wright-Pat and a Shell Oil tank truck of 100 octane low lead for the occasion. The General had a long history with Shell, and no active Air Force Base carried anything but jet fuel.

Schuck continues, with Miss Mitchell's tribute to the Doolittles:

> The following day we arrived at Andrews AFB in D.C. just before the C-141 touched down. In fact, we pulled off the taxi strip to give the family a chance to see the "Raider" with fans turning. We taxied in behind the family so they could see our bird up close and meet the crew.

Friday was the funeral.

> I went along with the pilots to get a GPS fix on the grave site to aid in the flyover. The GPS fix was very important as the various flights appeared out of nowhere and without warning. The logistics of this was pretty tricky as our aircraft would be following four flights in trail over the grave site.

Show time:

> First there was the C-141, followed by two separate flights of fighters, a flight of bombers, the Missing Man Formation, followed by our B-25 which was to dip a wing to salute the fallen hero. Because of the varying speeds of the different types, our B-25 orbited just out of sight of the crowds breaking out at the last moment as the missing man climbed out of

> sight. Just a few moments out of sync, and the effect would be lost. Nonessential crew were at the gravesite with the family and dignitaries to make room on board for VIPs. The timing was perfect; the unsuspecting pointing to the little silver bomber as it passed overhead. There wasn't a dry eye anywhere.

In this Washington Post report, Ray Ostlie was the combat crew chief of the original Miss Mitchell and part of Minnesota's Miss Mitchell restoration:[8]

> Retired Air Force Sgt. Ray H. Ostlie looked up at the clear blue sky and grimaced as a B-1 supersonic bomber flew by, thundering over the funeral of Gen. James H. "Jimmy" Doolittle, the famed World War II aviator. The symbol of the modern Air Force was too noisy for even the hardiest veteran airmen, who visited Arlington National Cemetery yesterday for a final salute to one of their own flyboys. Some covered their ears to soften the B-1's deafening roar.
>
> But when a diminutive, twin-propeller World War II B-25 bomber quietly whirled by, the somber mood of the day suddenly erupted into fist-waving and clapping. Like several others, Ostlie, a B-25 crew chief who worked under Doolittle in the North African campaign in 1943, sighed and smiled.

Static and Dynamic Activation

Static clustering: Notice how information clusters naturally around the aircraft: crew stories, nose art provenance, Doolittle funeral, all connecting through shared context. This is exactly how LLM training data organizes around central concepts.

Dynamic activation: The Doolittle funeral story activated multiple related concepts: Miss Mitchell (the aircraft, but not mentioned by name), Ray Ostlie (crew chief), combat service (North Africa), restoration work (1980s-1990s), the funeral (1993). Each piece connects to others forming a web of associations. When an LLM encounters "B-25 Mitchell" in training data, it builds exactly this kind of associative network. The funeral story is not separate from the nose art story. They are nodes in the same knowledge cluster, just like the mesh we are building. Then, **telling** the story activates these associations.

Motivation: Tour Guide

Throughout 2023 I was acting as one of the tour guides for Commemorative Air Force B-25 "Miss Mitchell" as shown on the book cover. Because I was the tour guide, I wanted to understand the aircraft intimately. Extensive reading provided history and context, but only hands-on experience creates the spatial memories and tactile associations that pure research cannot provide.

Software development pattern in research context: I used the same approach that I have always used with software development: Dig more deeply than others think necessary, because that is how

I spot the connections that others miss. Researching the history and context is different from how I research a system crash. That difference tells me that "dig more deeply" is a transcendent pattern: the technique applies to both domains. Such transcendent patterns tend to be "future proof" skills that pay off throughout your career.

As you will see, this combination of direct experience with documentary research mirrors how LLMs build richer representations by combining multiple data modalities. What I call "different kinds of memories", LLMs call "richer representations combining multiple data modalities". In LLM terminology, "representation" means patterns of activation across the neural network. That is roughly equivalent to what humans experience as memories or associations.

The physical experience anchors the abstract knowledge, just as visual and textual training data reinforce each other in Transformer architectures.

Figure 2.3, "Bomb camera view looking straight down, September 10, 2023," was taken looking straight down below my seat at the radio operator's station. The photo is a bit fuzzy because I was sitting in the seat, reaching down between my legs, holding the camera directly below my seat to take the photo. The black circular area at the top of the photo is the camera lens housing. The silvery diagonal line at the upper right corner is where the large square film plates were inserted and removed.



Figure 2.3. Bomb camera view looking straight down, September 10, 2023

To orient you concerning that unusual viewpoint, Figure 4 is looking forward into the rear compartment. The green wall at the front center is the B-25 bomb bay. At far left is radio equipment, with the radio operator's fold-down wooden table secured upright for takeoff. The far forward right seat, with a purple duffel bag sitting on it, is where I was sitting and is the radio operator's seat. You can see the bomb camera mounted vertically below the seat. The floor area near the bottom (lens) opening of the camera looks like there is a spotlight turned on. That is sunlight reflecting off the airport ramp and in through the camera window in the underside of the airplane.



Figure 2.4. B-25 rear compartment looking forward toward bomb bay, June 18, 2023

Conclusion

By following the Miss Mitchell stories, from nose art origins to funeral flyover, you have experienced the first phase of building a knowledge mesh: **creating information clusters around central concepts**.

Every detail (serial numbers, crew names, dates, locations) connected to other details, forming a web of associations. This is not random. It is exactly how LLM training data organizes: related information clusters naturally around core concepts, with each piece strengthening connections to others.

Next month we will walk through solutions for dealing with the discovery that authors have agendas. An author's military rank, for example, is no guarantee of seeing the truth. I decided on reading a minimum of three books on a given topic as my means of gaining a better perspective, and a means of detecting bias. That "three book minimum" concept has an LLM analogy called "data triangulation".

Large Language Models deal with this exact same situation: data sanitization and removal (or at least identification) of bias. We will experience exactly how that works in practice.

Notes

1. Ashish Vaswani et al., "Attention Is All You Need," arXiv:1706.03762, preprint, arXiv, August 2, 2023, https://doi.org/10.48550/arXiv.1706.03762.

2. Robert M. Gagné, ed., *Principles of Instructional Design*, 5th ed (Thomson/Wadsworth, 2005), page 122.

3. CAF Minnesota Wing Celebrating 50th Anniversary, Aviation Museum News, January 29, 2021, https://vintageaviationnews.com/aviation-museum-news/caf-minnesota-wing-celebrating-50th-anniversary.html. This article includes a photo of Ray Ostlie in front of the restored Miss Mitchell.

4. "Raymond D. Kowalik (1921-2006) - Find a Grave...," accessed October 15, 2025, https://www.findagrave.com/memorial/13082672/raymond_-d-kowalik.

5. D. S. Leaon, Minnesotans, South St. Paul's Fleming Field and the Commemorative Air Force: A History Project of the Commemorative Air Force, Minnesota Wing, 2008 (2008), page 187.

6. William Wolf, North American B-25 Mitchell: The Ultimate Look: From Drawing Board to Flying Arsenal (Schiffer Military History, 2008), page 154.

7. John Schuck, There They Go! How the North Star and Lone Star Wings Preserved the Fledgling Commemorative Air Force (2017), pages 63-66.

8. Peter Pae, "IN ARLINGTON, FULL HONORS FOR A SKY-BLAZING HERO," The Washington Post, October 2, 1993, https://www.washingtonpost.com/archive/local/1993/10/02/in-arlington-full-honors-for-a-sky-blazing-hero/62ea25bc-3c71-4670-a956-a2d5c224a2e1/.