# Trade in AI-Related Products

Michael E. Waugh
Federal Reserve Bank of Minneapolis and NBER

April 2026

ABSTRACT

This paper documents facts about international trade in AI-related products. I develop a large language model (LLM) classification tool that maps HS10 codes in U.S. trade data to products used in the construction and operation of AI infrastructure. AI-related products account for 23 percent of U.S. imports in 2025, and imports of these products have grown by 73 percent since 2023. Over the same period, imports of non-AI-related products have grown by only 3 percent, with the divergence between the two categories beginning in early 2024. Mexico is a key market on both the import and export side, and together with Taiwan these two countries account for about half of all U.S. trade in AI-related products. Trade policy has treated these products lightly with product-level exemptions shielding much of AI-related imports from tariffs. Absent the AI boom, a simple accounting exercise suggests that the U.S. goods trade deficit would have been nearly $200 billion smaller in 2025.

Email: michael.e.waugh@gmail.com. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. My GitHub repository provides the code and supplementary work behind this paper at `https://github.com/tradewartracker/ai-trade-index`.

The United States is experiencing a massive AI-investment-driven boom with hundreds of billions, if not trillions, of dollars being invested in the buildout of data centers and other related AI infrastructure projects. This paper asks: How is this AI data center buildout influencing international trade flows? I find that trade in AI-related products has been a major force in U.S. trade over the past two years. This holds despite the dramatic changes in U.S. trade policy that have taken place over this time period.

Answering this question requires knowledge about what products are related to AI infrastructure. A novelty of this paper is that I use a large language model (LLM) classification tool that maps product descriptions in trade data into a systematic classification of whether the product is used in AI infrastructure and, if so, in what way. This allows me to take unstructured product descriptions and construct structured measures of trade in AI-related products.

What are AI-related products? The classification identifies the obvious computer hardware inputs such as data processing units and storage devices. These products account for about half of all AI-related trade and imports of them have grown by triple digits since 2023. But the classification also identifies a broader set of products tied to electrical infrastructure, networking, cooling and HVAC, and specialty materials. These ancillary products account for the other half of AI-related trade and have also experienced strong import growth.

Through 2025, AI-related products account for 23 percent of all U.S. imports, up from 15 percent in 2023. Leading into the start of 2024, AI-related products grew no differently than non-AI-related products. But since early 2024, import growth in AI-related products has accelerated sharply. As of January 2026, trade in AI-related products had grown by 73 percent relative to 2023, compared with only 3 percent for non-AI-related products.

Where are these products coming from? Two countries play an important role in the sourcing of AI-related products. Not surprisingly, Taiwan is an important source for computer hardware components (e.g. semiconductors) and Taiwan accounts for about a quarter of imports of AI-related products. The surprising source is Mexico. It accounts for another quarter of AI-related trade and its reach extends to computer hardware and other products related to electrical power, networking, and cooling HVAC. China is less important in AI-related trade and its overall share has diminished over the past two years.

How has trade policy treated AI-related products? Lightly. Since the beginning of 2025, the U.S. has experienced a large change in trade policy with average tariff rates increasing by more than 10 percentage points (see, e.g. Waugh (2026) or The Budget Lab at Yale (2025)). I find that effective tariff rates on AI-related products are substantially lower than on other goods. At the end of 2025, the effective tariff rate on AI-related products was 4.5 percent, compared with 12.1 percent for non-AI products. This gap largely reflects product-level exemptions issued around the "Liberation day" tariffs in April of 2025. These exemptions essentially covered about 69 percent of all AI-related products by value. And these exemptions were quickly put back in place when the Section 122 surcharge was implemented in February of 2026.

The massive growth in trade in AI-related products is not isolated to the import side. Exports of AI-related products have grown by 34.5 percent since 2023. Mexico is the dominant destination, accounting for about a third of all AI-related exports, and it leads across nearly every product category. This is suggestive that the U.S. plays an important role in AI global value chains.

The final part of the paper considers the aggregate implications of the growth in AI-related trade through a simple accounting exercise. Specifically, I ask: if AI-related products had grown at the same rate as non-AI-related products since 2023, what would the U.S. goods trade deficit have been? Under this counterfactual, the U.S. goods trade deficit would have been nearly $194 billion smaller in 2025 or nearly 16 percent smaller absent the AI boom.

The key distinguishing feature of my paper is the use of an LLM-based classification procedure to identify a broader set of traded products that may be used in the construction and operation of AI infrastructure. This broader classification moves beyond GPUs and chipsets and provides a more complete picture of what the United States is importing and exporting to support the AI boom. Several related works take a narrower approach. de Soyres, Haag, Liu, and Van Leemput (2026) study global trade trends in three specific HS6 codes that are closely tied to the AI buildout. World Trade Organization (2025) construct a list of roughly 100 AI-related products and emphasize that trade in these products has been a key driver of trade growth in 2025. Economic commentators such as Setser (2026), Politano (2025), and Swanson (2026) have similarly noted the growing importance of computer hardware and semiconductor trade.

## 1. Classifying AI-Related Products

This section describes how I classify internationally traded commodities that are plausibly used to build and operate AI-focused data centers — hyperscale facilities that rely on GPU servers, high-speed networking, and high-density power and cooling. The measurement question is: which traded goods plausibly enter AI data-center construction, equipping, and operations?

I treat this as a classification problem and use a large language model (LLM) to systematically label the universe of HS10 commodity codes (about 19,000) by their relevance to AI data-center infrastructure. The use of LLMs for classification tasks is increasingly common in economics (see, e.g., Clayton et al. (2025), Lagakos et al. (2025), Cascaldi-Garcia and Iacoviello (2026)). The methodological idea is the same — provide the LLM structured descriptions and ask systematic questions about the relevance of each item to a particular domain.

## 1.1. Data

Several inputs are provided to the LLM. First, I use a list of 10-digit Harmonized System commodity codes from U.S. Census Bureau trade statistics. The HS10 classification provides the finest level of product disaggregation available in U.S. trade data, with 18364 distinct codes found in 2024 import and export data. Each HS10 code includes a standardized commodity description. As an example, consider

```
    8542310040: PROCESSORS (INCLUDING MICROPROCESSORS): GRAPHICS
PROCESSING UNITS (GPUS).
```

where `8542310040` is the HS10 code and the description indicates that this product is a processor and specifically a GPU. This is a canonical **High** example.

The second input is the associated 6-digit NAICS code and its description. The idea is to give the LLM additional context about the HS10 code. NAICS codes do this because they are production-oriented, providing context about how a product is typically manufactured and used. I use the Census concordance, which maps HS10 codes to six-digit NAICS codes. In cases where there are multiple NAICS codes associated with an HS10 code, I provide the LLM with a list of the NAICS codes and list of descriptions.

In the analysis, I focus on nominal U.S. trade flows from the U.S. Census Bureau trade statistics. I exclude several classes of commodities in the analysis. The ones excluded are petroleum products (HS chapter 27), precious metals (HS chapter 71), and HS2 chapters 98 and 99, which include art, antiques, and other special provisions.

## 1.2. LLM Classification Procedure

The classification procedure requires specifying both a categorization scheme and a prompt to implement it. The primary scheme that I implement is to classify each HS10 code into one of three bins:

- **High**: Products directly used in data center construction or operation, with clear application to compute, cooling, power, networking, or facility infrastructure.
- **Medium**: Products with plausible data center applications but also significant non-data-center uses, or indirect inputs to data center materials.
- **Low**: Products with no apparent connection to data center construction.

In the baseline results, I focus on products classified as **High**. I include **Medium**-rated products only as a robustness check.

In addition to this categorization, I ask the model to provide a confidence score, a use category, and a natural language rationale of the classification decision. These additional outputs allow for ex-post scrutiny of the classification. So, for each HS10 code $i$ with description $d_i$, I prompt the model to return a structured classification object:

$$C_i = \{\texttt{relevance}_i, \texttt{confidence}_i, \texttt{category}_i, \texttt{reasoning}_i\}$$

where:

- $\texttt{relevance}_i \in \{\textbf{High}, \textbf{Medium}, \textbf{Low}\}$ is constrained by an enumerated type,
- $\texttt{confidence}_i \in [0, 100]$ is an integer confidence score,
- $\texttt{category}_i$ maps to one of eight predefined categories (Compute Hardware, Networking/Telecom, Cooling/HVAC, Electrical/Power, Building Structure, Fire Safety/Security, Specialty Materials, Not DC Related),
- $\texttt{reasoning}_i$ provides a natural language explanation of the classification decision.

The schema enforcement ensures that every response contains valid values for all fields. The model cannot return "Medium-High" or "Moderate", etc. It will only return the three specified relevance levels. The structured approach also facilitates systematic analysis and the comparison of outputs across different LLM models.

An important input is the prompt to the LLM. The prompt provides context on AI data center construction and describes the categories of relevant products. The prompt also ensures that the output is structured according to the schema described above. The complete prompt is presented in the Appendix. Its opening reads as follows:

```
    You are an expert in AI data center construction, operations,
and supply chains.  Your task is to classify products by their
relevance to building and operating AI-focused data centers
(hyperscale facilities running GPU clusters for training and
inference.  You will be provided with two pieces of information
about each product:

1.  **HS10 Code and Description**:  The Harmonized System (HS)
is an international product classification used for trade
statistics.  The HS10 code is a 10-digit code used by U.S.
Customs to classify imported goods.  The description tells you
what the physical product is.
```

```
2.  **NAICS Code and Description**:  The North American Industry
Classification System (NAICS) indicates which U.S. industry
produces or uses this product.  This provides additional context
about the product's typical industrial application.

Consider these categories of relevant products:  COMPUTE: GPUs,
CPUs, memory, PCBs, servers, storage drives, semiconductors;
NETWORKING: Fiber optics, switches, routers, transceivers,
cables; COOLING: Chillers, cooling towers, CRAH units, fans,
pumps, refrigerants; ELECTRICAL: Transformers, switchgear,
UPS, batteries, generators, cables; BUILDING: Structural steel,
concrete, rebar, insulation, raised floors; SPECIALTY MATERIALS:
Rare earths, copper, aluminum, tantalum, thermal interface
materials...
```

The prompt also instructs the model to consider edge cases explicitly — for example, that products like diesel engines could be relevant (for generators) or irrelevant (for vehicles), and that context from the HS and NAICS description should guide such determinations.[1]

Several runs of the classification procedure have yielded broadly stable results, particularly for the aggregate patterns that are the focus of the paper. As a further cross-check, an earlier version of the project used a keyword-matching procedure, based on a list of AI-infrastructure-related materials and products that was then matched to HS descriptions. While that approach was much coarser than the LLM-based classification, it delivered similar trends. An important limitation, however, is that I do not yet have a ground-truth benchmark for validating the product-level labels. The classification should therefore be understood as a systematic but imperfect measurement exercise.

## 2. What are AI-related products?

Table 1 provides an overview of what the classification scheme identifies as AI-related products. The first row reports all HS10 codes that are classified as **High** relevance. There are 645 high-relevance codes out of a total of 18364. The **High** relevance codes accounted for 15 percent of total imports in 2023 and this increased to 23 percent in 2025.

The next seven rows of Table 1 break down those HS10 codes classified as **High** into the predefined categories discussed above. Among those categories, four stand out: Compute Hardware,

[1] I employ Claude (Anthropic) as my main classification model, using the tool-calling API to obtain structured responses.

Table 1: U.S. Import Values by AI Relevance and Category (2023 vs 2025)

| Category | # HS10 Codes | 2023 ($B) | 2025 ($B) | Change (%) |
|---|---|---|---|---|
| High AI Relevance | 645 | 379.0 | 654.0 | +72.6 |
| Compute Hardware | 163 | 144.4 | 353.8 | +144.9 |
| Electrical Power | 250 | 116.9 | 141.8 | +21.3 |
| Networking Telecom | 24 | 62.9 | 99.5 | +58.2 |
| Cooling HVAC | 137 | 41.5 | 47.5 | +14.4 |
| Building Structure | 44 | 12.1 | 10.1 | -16.5 |
| Fire Safety Security | 8 | 0.6 | 0.7 | +10.5 |
| Specialty Materials | 19 | 0.4 | 0.5 | +22.5 |
| Low AI Relevance | 15915 | 1834.7 | 1880.9 | +2.5 |
| Total Trade | 18364 | 2598.3 | 2883.9 | +11.0 |

Electrical Power, Networking and Telecom, and Cooling HVAC. These four categories represent nearly 95 percent of all goods classified as **High**.

Tables 4 and 5 (at the back of the paper) report the top 20 HS10 codes by 2025 import volume that were classified as **High**. For each code, the tables report the Census's short description, the LLM category, the LLM reasoning behind the classification, and various trade statistics.[2] The dominant HS10 code is `8471500150`, which is described as data processing units. And the LLM reasoned that data processing units are a core component of compute hardware.

While the top 10 list is dominated by computer hardware equipment, other notable products also appear. In particular, `7403110000` (refined copper cathodes) is the sixth most important product. The LLM reasoning is that refined copper cathodes are a key input into copper products (e.g., electrical wire) used in data centers. Similarly, insulated fiber optic cables (`8544700000`) appear in the top 20, classified under Networking Telecom — the physical backbone connecting GPU clusters. Approaches to measure the importance of trade in AI-related products by focusing on a narrow set of semiconductor codes would miss these products entirely.

In looking through Table 1 or 4, one cannot help but notice the dramatic growth in trade volumes in nearly all categories. I discuss this growth next.

[2]The trade statistics are 2025 imports, the share of those imports in total trade, and the change in import volume relative to 2023. The description is the Census short description; this is different from their long description which is what is fed into the LLM.

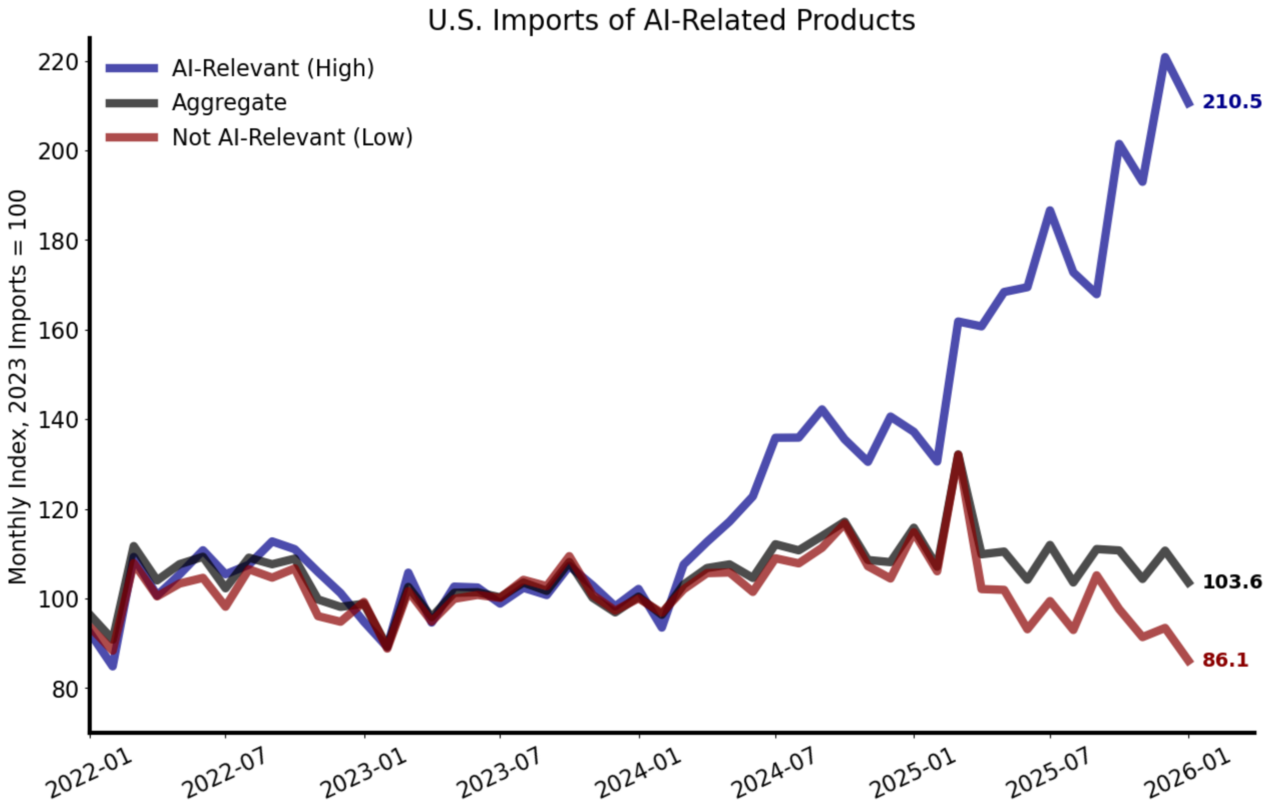


Figure 1: U.S. Imports of AI and non-AI products

## 3. How much trade in AI-related products?

Figure 1 summarizes a main finding of the paper. It plots several different indices of U.S. imports. These indices are computed by taking import volumes in a category for a given month, then normalizing them by the monthly average for 2023.

Focus on the blue line titled "AI-Relevant." This is the index computed for all **High** products. The most obvious feature of AI-related products is the spectacular growth. As of January 2026, the index for AI-relevant products is 211, meaning import volumes are **111 percent larger** than the typical month in 2023.

The second observation concerns the timing. From 2022 to early 2024, there is no differential trend in AI-relevant products relative to non-AI-relevant products. That is, these products were growing at a similar rate to non-AI-relevant products. Then something changed in early 2024. At that point there is divergence and this divergence accelerates in 2025. This timing lines up with the wave of large-scale data center investment announcements that began in early 2024 and accelerated through the year.

Where is the growth coming from? Return to the final column of Table 1. For each category, the final column reports 2025 trade volumes relative to 2023 in percent terms. Consistent with

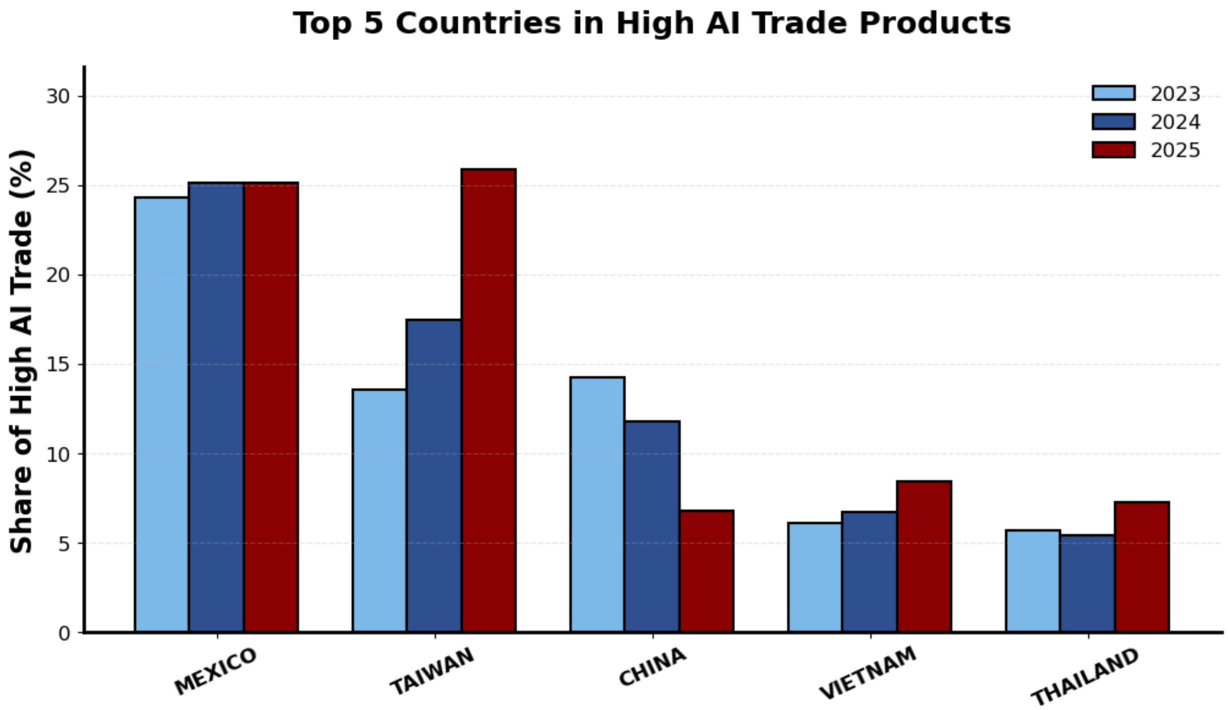


Figure 2: Top Sources of AI-Related Products

Figure 1, **High** AI-relevant products grew by 73 percent. The breakdown by category shows that this growth is broad-based — there is double-digit growth in 5 out of the 7 categories. Focusing on products, Table 4 confirms how broad and spectacular the growth is with all of the top 10 HS10 codes displaying double-digit growth; five display triple-digit growth rates.

Moreover, this growth is not simply driven by Compute Hardware. Even after excluding Compute Hardware, AI-relevant imports have grown by 40 percent since 2023. This is precisely the kind of pattern that would be invisible under a narrower classification focused on semiconductors and processors alone.

The black and red lines in Figure 1 report aggregate trade and non-AI-relevant products. The latter series is the index computed for all **Low** products. In aggregate, trade has grown in 2025 by 11 percent relative to 2023. This aggregate growth is surprisingly strong given that U.S. tariffs increased by more than ten percentage points in 2025. However, this aggregate growth is driven by the spectacular growth in AI-relevant products.

Once AI-related products are stripped out, the picture changes. Non-AI-relevant imports as of January 2026 are 14 percent below the typical month in 2023. On a year-over-year basis, Table 1 reports that non-AI-relevant products grew by only 3 percent, and much of this growth owes to the front-running of tariffs in early 2025.

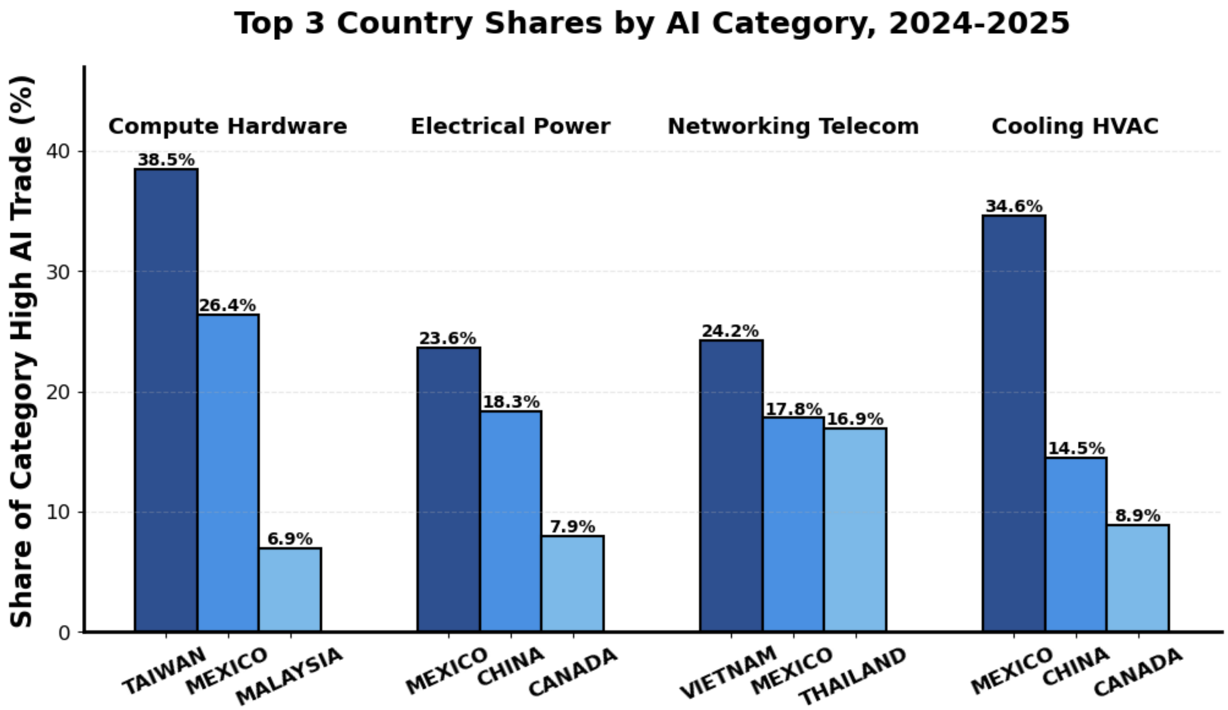


Figure 3: Top Sources by AI-Related Product Category

## 4. Where do AI-relevant products come from?

Figure 2 computes the share of **High** products by source country for each year since 2023.

Mexico's dominance is perhaps the most surprising feature of Figure 2. Consistently, it has accounted for about 25 percent of **High** AI-relevant imports across all years. Less surprising is the role of Taiwan. It is well known that Taiwan is an important source of semiconductor and computer hardware products. Consistent with this notion, Figure 2 shows that Taiwan is important and increasingly dominant in 2025, rivaling Mexico.

Figure 3 breaks down the country composition by top categories of AI-related products. Mexico leads in two of the top four categories and ranks second in every other. Taiwan's prominence is concentrated in Compute Hardware, where it is the leading source.

A final surprising feature is the small and diminishing role of China. Figure 2 shows that Taiwan and China held similar shares in 2023, but the two have diverged since then. Now China's share of AI-relevant products is comparable to Vietnam and Thailand. One important issue with China is that trade policy has become much more restrictive toward China than toward other countries. As I discuss below, AI-relevant products have generally been exempted from the "Liberation Day tariffs" associated with Executive Order 14257. However, Chinese imports faced an additional 20 percent tariff rate from the "fentanyl tariffs" associated with Executive Order 14195, and these fentanyl tariffs are not covered by the product-level exemptions. Thus,

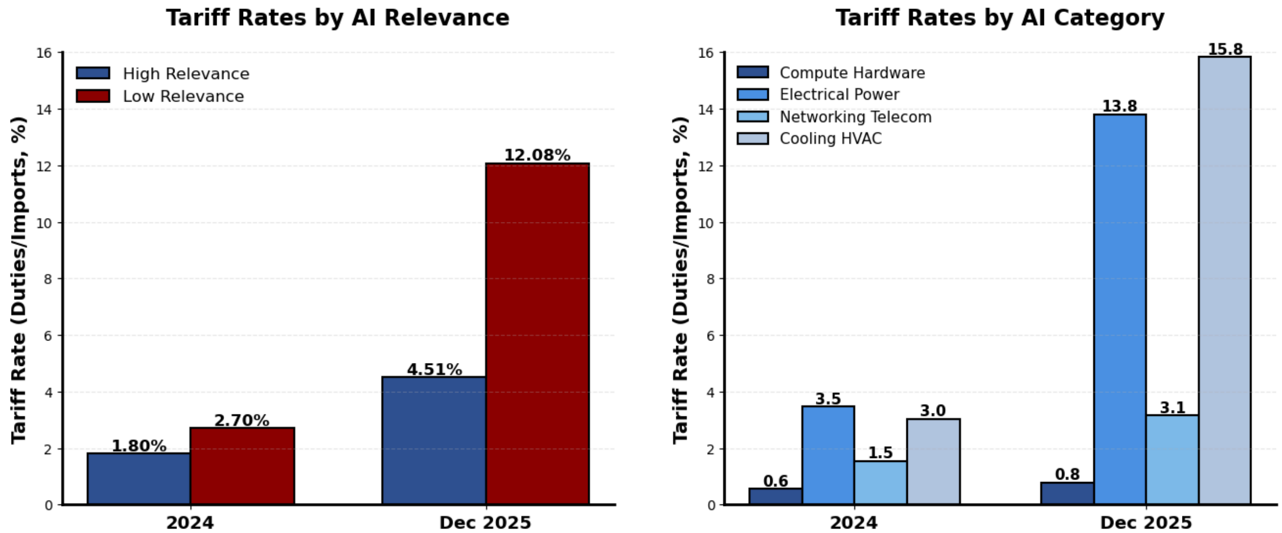


Figure 4: Tariff Rates on AI and non-AI products

China was uniquely penalized by the changes in trade policy.

This breakdown by source and product illustrates why a broad classification of AI-relevant goods matters. A narrow focus on chips alone would suggest that AI trade is primarily a U.S.–Taiwan story. Figure 3 shows it is much broader and it highlights the importance of Mexico.

To this last point, Mexico's role in AI-related trade helps explain some puzzling features of trade within the USMCA over 2025. Both Canada and Mexico essentially had the same tariff treatment over the course of 2025. Yet, U.S. imports from Canada changed by -8.3 percent while U.S. imports from Mexico rose by 6.4 percent. A lot of this divergence reflects the 40.7 percent growth in Mexican AI-related exports to the U.S.

## 5. How has trade policy treated AI-related products?

Figure 4 plots effective tariff rates (duties relative to import value) for AI-related and non-AI-related products. In 2024, the gap between the two was modest: 1.8 percent for **High** products versus 2.7 percent for **Low** products. By 2025, with the broad increase in U.S. tariffs, both rates rose, but the gap widened substantially. AI-related products face an effective rate of 4.5 percent, compared with 12.1 percent for non-AI products.

The gap varies sharply across product categories. Compute Hardware faces just 0.8 percent, while Electrical Power (13.8 percent) and Cooling HVAC (15.8 percent) face rates more than double the overall average for AI products. Networking Telecom is at 3.1 percent.

Why are tariffs lower on AI-related products? The answer is exemptions. Several executive actions in 2025–2026 created product-level tariff exemptions that disproportionately cover AI inputs. Three lists are relevant:

1. the Consumer Electronics amendment to IEEPA Executive Order 14257 (April 11, 2025);
2. the Annex II list from the same executive order; and
3. the Section 122 Surcharge exemptions (February 2026), which were put in place after IEEPA tariffs, and their associated exemptions, were removed.

Table 2 reports the coverage of these exemption lists for AI **High** products. Of the $654 billion in **High** AI imports in 2025, nearly 69 percent falls on at least one exemption list. The Consumer Electronics exemption alone covers $424 billion — the largest single exemption by value — because it captures the core compute hardware products (GPUs, processors, memory, SSDs, networking equipment) that dominate AI trade by value. The Annex II list adds critical materials like copper cathodes and specialty metals.[3]

The Section 122 Surcharge exemptions cover a similar set of products, largely overlapping with the Consumer Electronics and Annex II lists.

The category-level pattern in exemptions helps explain the cross-category variation in effective tariff rates. Compute Hardware is almost entirely shielded: 97.3 percent of its $354 billion in imports falls on an exemption list. Networking Telecom is similarly well-covered at 85.1 percent. By contrast, Electrical Power (13.5 percent), Cooling HVAC (0.4 percent), and Building Structure (0.0 percent) have little or no exemption coverage, which explains why these categories face effective tariff rates above 11 percent.

## 6. U.S. Exports of AI-Related Products

The massive growth in trade in AI-related products is not isolated to the import side. Figure 5 plots the export-side analog of Figure 1. The pattern mirrors the import side: AI-relevant exports have diverged sharply from non-AI exports since mid-2024. The divergence is somewhat less dramatic, but the trend is clear and as of January 2026, AI-related exports grew by 34.5 percent since 2023, led by Compute Hardware (62.1 percent). While import growth is stronger, the export side confirms that trade in AI-related products is not one-sided.

Mexico is an important destination for U.S. exports of AI-related products. In 2025, Mexico accounted for about 32 percent of all U.S. exports of **High** AI-relevant products, making it the single largest destination by a wide margin. What is striking is how broad-based Mexico's role is. Mexico is the leading destination for Compute Hardware exports (33 percent), Electrical Power

[3] And note that these exemptions only apply to the tariffs associated with Executive Order 14257. As discussed above, China faced an additional 20 percent tariff rate associated with Executive Order 14195, the fentanyl tariffs. Canada and Mexico faced these additional tariffs as well, but could be exempt if they were USMCA compliant.

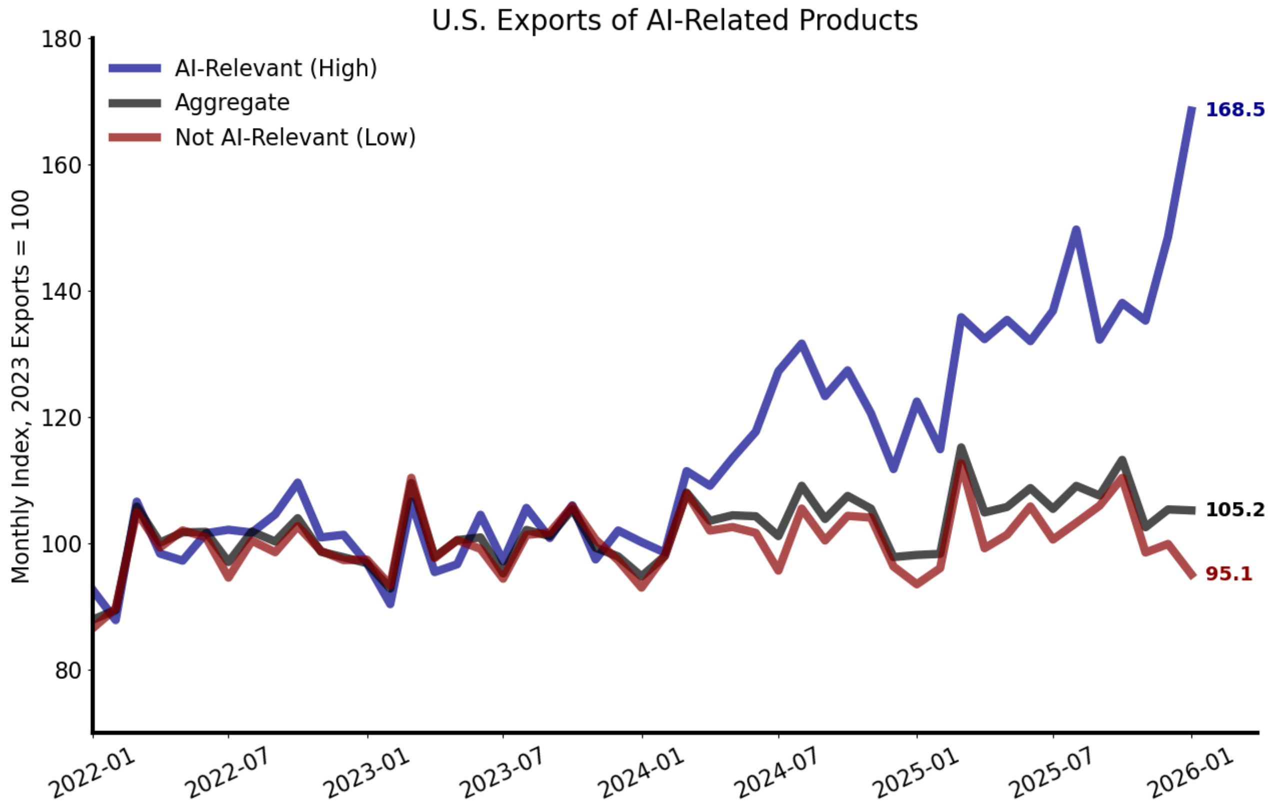


Figure 5: U.S. Exports of AI and non-AI products

Table 2: Tariff Exemption Coverage of AI High-Relevance Imports (2025)

| Exemption List | # HS10 Codes | AI High Trade ($B) | Share of AI High (%) |
|---|---|---|---|
| Consumer Electronics (EO 14257, Apr 11) | 109 | 424.5 | 64.9 |
| Annex II (EO 14257, Apr 11) | 720 | 19.4 | 3.0 |
| Section 122 Surcharge (Feb 2026) | 1,098 | 445.9 | 68.2 |
| Any Exemption List | — | 448.6 | 68.6 |

exports (36 percent), and Networking Telecom exports (19 percent), and it is the second-largest destination for Cooling HVAC exports (25 percent). This breadth across categories mirrors the import-side pattern documented in Section 4 and underscores that the U.S.-Mexico trade relationship in AI-related products extends well beyond any single product type.

How exports should respond to a large demand shock like the AI-investment boom is not obvious. One might expect a form of "venting in" — the U.S. demand shock redirecting exports inward, akin to Almunia, Antràs, Lopez-Rodriguez, and Morales (2018) but in reverse. Yet the data show the opposite: AI-related exports are expanding sharply alongside imports.

The fact that we see exports of AI-related products expanding so much is suggestive that the U.S. may play an important role in AI global value chains. This pattern is also consistent with the U.S. being an entrepôt. That is AI-related materials coming into the U.S., exported to Mexico for assembly or processing, and then shipped back to the U.S. And if so, little value is actually being added domestically. This is a question that would require data on trade in value added.

## 7. AI's Impact on the Trade Deficit

From a macro perspective, one may ask whether the expansion of AI-related trade matters for aggregate outcomes that economists and policy makers care about. This section argues it does and I focus on the trade deficit largely because changes in U.S. trade policy in 2025 were focused on reversing the large U.S. trade deficit.

I illustrate this point through a simple accounting exercise to quantify the AI boom's impact on the trade deficit. Specifically, I ask: if AI-related products had grown at the same rate as non-AI products since 2023, what would the U.S. goods trade deficit have been?

To answer this question I proceed in three steps. First, I compute non-AI growth indices for both imports and exports — the growth rates of the red lines in Figures 1 and 5. I treat these as the counterfactual growth rates that would have prevailed absent the AI boom. Second, starting from 2023 trade levels for AI-related products, I apply these counterfactual growth rates to construct counterfactual AI series. Finally, I define excess AI trade as the gap between actual and counterfactual AI trade.

Table 3 reports the results. In 2025, excess AI imports amount to $265 billion, while excess AI exports are $71 billion — the net effect is $194 billion. In other words, the goods trade deficit would have been nearly $200 billion smaller absent the AI boom. The actual U.S. goods trade deficit in 2025 was $1,235 billion; without the AI boom it would have been $1,041 billion. So, absent the AI data center buildout, the goods trade deficit would have been nearly 16 percent smaller.

At a high level, this result illustrates that trade in AI-related products is a very important force behind U.S. trade over the past year. In fact, it might be even more important than dramatic changes in U.S. trade policy in 2025. With that said, this exercise is merely an accounting exercise and it leaves out many details to consider. In Ferrante, Prestipino, Raffo, and Waugh (2026), we employ a structural approach and use a model to tease out the underlying effects of tariffs, exchange rates, and AI-investment-specific shocks on macro aggregates.

Table 3: Accounting for AI's Impact on Trade ($B)

| Year | Actual Imports | Actual Exports | Excess AI Imports | Excess AI Exports | Net Effect |
|---|---|---|---|---|---|
| 2023 | 2,598 | 1,552 | — | — | — |
| 2024 | 2,792 | 1,601 | 66 | 33 | +32 |
| 2025 | 2,884 | 1,648 | 265 | 71 | +194 |

## 8. Final Thoughts

One of the most important trends in the U.S. economy is the investment in and buildout of data centers and other related AI infrastructure projects. In this paper, I explored how this investment buildout has shaped U.S. trade flows.

The main input into my analysis was the use of an LLM-classification procedure that assigned goods in the HS classification system to their potential use in AI. Figure 1 illustrated the massive growth in trade in these products. Mexico is a surprising source of AI-relevant imports to the U.S. — a result that only emerges when one looks beyond a narrow set of semiconductor products. And Table 3 shows that, from an accounting perspective, trade in AI products has had a large imprint on the U.S. trade deficit.

There are several open questions for future research. The most natural area for improvement is the classification list itself. While the items are plausible and the aggregate trends are sensible, incorporating additional information sources and formal validation would strengthen the foundation.

Another question regards prices versus quantities. Using unit values from the underlying trade data might be a path forward on this dimension. The inflationary impacts of the AI boom are interesting as well.

A broader question is about trade in capital goods. Models of trade often abstract from capital or the durable nature of the goods being traded (Eaton and Kortum (2001), Eaton, Kortum, Neiman, and Romalis (2016), Ravikumar, Santacreu, and Sposi (2019) are notable exceptions). Most of the AI-related products are classified as capital goods. Understanding how large, concentrated capital-goods booms (or busts) propagate through trade is an important area for future work.

Finally, the finding that tariff exemptions have disproportionately shielded AI-related products raises a forward-looking question: as trade policy in the U.S. continues to evolve, whether these exemptions persist has important implications for the cost and pace of the AI buildout.

## References


ALMUNIA, M., P. ANTRÀS, D. LOPEZ-RODRIGUEZ, AND E. MORALES (2018): "Venting Out: Exports During a Domestic Slump," Tech. rep., National Bureau of Economic Research.

CASCALDI-GARCIA, D. AND M. IACOVIELLO (2026): "Quantifying Deregulation and its Economic Effects: A Large Language Model Approach," *Working Paper, Federal Reserve Board*.

CLAYTON, C., A. COPPOLA, M. MAGGIORI, AND J. SCHREGER (2025): "Geoeconomic Pressure," *Columbia Business School Research Paper Forthcoming, Stanford University Graduate School of Business Research Paper Forthcoming*.

DE SOYRES, F., A. HAAG, M. LIU, AND E. VAN LEEMPUT (2026): "The Global Trade Effects of the AI Infrastructure Boom," .

EATON, J. AND S. KORTUM (2001): "Trade in Capital Goods," *European Economic Review*, 45, 1195–1235.

EATON, J., S. KORTUM, B. NEIMAN, AND J. ROMALIS (2016): "Trade and the Global Recession," *American Economic Review*, 106, 3401–3438.

FERRANTE, F., A. PRESTIPINO, A. RAFFO, AND M. E. WAUGH (2026): "The Missing Trade Collapse: Tariffs, AI-Driven Investment, and the Resilience of U.S. Imports," Work in progress.

LAGAKOS, D., S. MICHALOPOULOS, AND H.-J. VOTH (2025): "American Life Histories," Tech. rep., National Bureau of Economic Research.

POLITANO, J. (2025): "The Tariff Exemption Behind the AI Boom," Apricitas Economics (Substack).

RAVIKUMAR, B., A. M. SANTACREU, AND M. SPOSI (2019): "Capital accumulation and dynamic gains from trade," *Journal of International Economics*, 119, 93–110.

SETSER, B. W. (2026): "Understanding the Low October 2025 Trade Deficit," Council on Foreign Relations, Follow the Money.

SWANSON, A. (2026): "The Cost of the A.I. Boom: A Trade Deficit the President Detests," *The New York Times*.

THE BUDGET LAB AT YALE (2025): "State of U.S. Tariffs: November 17, 2025," Policy brief, Yale University, accessed December 2025.

WAUGH, M. E. (2026): "How Restrictive is US Trade Policy?" Tech. rep., National Bureau of Economic Research.

WORLD TRADE ORGANIZATION (2025): "Global Trade Outlook and Statistics Update: October 2025," Tech. rep., World Trade Organization, Geneva, Switzerland.

## 9. Appendix

This appendix presents the full LLM prompt used in the classification procedure described in Section 1, followed by Tables 4 and 5, which report the top 20 AI-related HS10 codes by 2025 import volume. The complete list of all HS10 codes classified as **High** relevance is available in the file `AI_TRADE_HIGH_RELEVANCE_PRODUCTS.md` in the project repository at `https://github.com/tradewartracker/ai-trade-index`.

**System Prompt:**

You are an expert in AI data center construction, operations, and supply chains. Your task is to classify products by their relevance to building and operating AI-focused data centers (hyperscale facilities running GPU clusters for training and inference. You will be provided with two pieces of information about each product:

1. **HS10 Code and Description**: The Harmonized System (HS) is an international product classification used for trade statistics. The HS10 code is a 10-digit code used by U.S. Customs to classify imported goods. The description tells you what the physical product is.

2. **NAICS Code and Description**: The North American Industry Classification System (NAICS) indicates which U.S. industry produces or uses this product. This provides additional context about the product's typical industrial application.

Consider these categories of relevant products: COMPUTE: GPUs, CPUs, memory, PCBs, servers, storage drives, semiconductors; NETWORKING: Fiber optics, switches, routers, transceivers, cables; COOLING: Chillers, cooling towers, CRAH units, fans, pumps, refrigerants; ELECTRICAL: Transformers, switchgear, UPS, batteries, generators, cables; BUILDING: Structural steel, concrete, rebar, insulation, raised floors; SPECIALTY MATERIALS: Rare earths, copper, aluminum, tantalum, thermal interface materials...).

Consider these categories of relevant products:

- **COMPUTE:** GPUs, CPUs, memory, PCBs, servers, storage drives, semiconductors
- **NETWORKING:** Fiber optics, switches, routers, transceivers, cables
- **COOLING:** Chillers, cooling towers, CRAH units, fans, pumps, refrigerants, glycol
- **ELECTRICAL:** Transformers, switchgear, UPS, batteries, generators, cables, busbars

- **BUILDING:** Structural steel, concrete, rebar, insulation, raised floors, fire suppression
- **SPECIALTY MATERIALS:** Rare earths, copper, aluminum, tantalum, thermal interface materials

Be thoughtful about edge cases:

- "Diesel engines" could be for generators (relevant) or vehicles (not relevant)
- "Pumps" could be for cooling systems (relevant) or unrelated industrial use
- "Copper wire" is relevant for electrical systems
- Food, textiles, furniture, and consumer goods are generally NOT relevant

Use the `classify_hs10_code` tool to record your assessment.

---

**User Prompt:**

Classify this HS10 code:
Code: {hs10_code}
Description: {description}

---

**Tool Schema (Structured Output):**

`relevance`: enum ["High", "Medium", "Low"]
`confidence`: integer, 0–100
`primary_category`: enum ["Compute_Hardware", "Networking_Telecom", "Cooling_HVAC", "Electrical_Power", "Building_Structure", "Fire_Safety_Security", "Specialty_Materials", "Maintenance_Operations", "Not_DC_Related"]
`specific_use`: string (e.g., "GPU accelerators", "backup generator fuel")
`reasoning`: string (explanation of classification decision)

Table 4: Top 10 HS10 Codes by 2025 Import Volume with High AI Relevance

| HS10 Code | Description | Category | LLM Reasoning | 2025 Imports ($B) | Share of Trade (%) | Change (%) |
|---|---|---|---|---|---|---|
| 8471500150 | PROC UNT IN HOUS W/ EITHER STOR, IN&OUTPUT,W/O CRT | Compute Hardware | Digital processing units are core compute components used in data center servers. | 163.55 | 5.67 | +343.0 |
| 8517620090 | MACH FOR RECP/CONVER ETC OF VOICE/IMAGE/DATA NESOI | Networking Telecom | This product covers data transmission and conversion equipment in the telecommunications/wireless communications industry. | 52.88 | 1.83 | +58.3 |
| 8473301180 | PTS ADP MCH, NT INCPTNG CRT,PRT CRT ASSEM.;NESOI | Compute Hardware | This code covers printed circuit assemblies for data processing machines - these are core components in servers, GPUs, and other computing hardware that form the backbone of AI data centers. | 48.89 | 1.70 | +199.8 |
| 8517620020 | SWITCHING AND ROUTING APPARATUS | Networking Telecom | Switching and routing apparatus are core networking components essential for AI data centers. | 30.76 | 1.07 | +89.1 |
| 8473301140 | PTS ADP MCH, NT INCPTG CRT,PRT CT ASM.,MRY MODULES | Compute Hardware | This product covers parts and accessories for data processing machines including printing circuit assemblies and memory modules. | 25.77 | 0.89 | +205.6 |
| 8523510000 | SOLID-STATE NON-VOLATILE STORAGE DEVICES | Compute Hardware | Solid-state non-volatile storage devices (SSDs) are essential components in AI data centers for storing datasets, model weights, operating systems, and temporary data. | 17.54 | 0.61 | +84.6 |
| 7403110000 | REFINED COPPER CATHODES AND SECTIONS OF CATHODES | Electrical Power | Refined copper cathodes are the primary raw material used to manufacture electrical wiring, cables, busbars, and other conductive components essential for data center power distribution systems. | 16.64 | 0.58 | +150.3 |
| 8507600020 | LITHIUM-ION STORAGE BATTERIES, NESOI | Electrical Power | Lithium-ion storage batteries are critical components in data center UPS (Uninterruptible Power Supply) systems, providing backup power during outages and power quality issues. | 16.53 | 0.57 | +11.1 |
| 8537109170 | ELECTRICL APP EQP W/ APPS FM 8535&6,LT 1000V,NESOI | Electrical Power | This is switchgear/switchboard apparatus for electrical control and distribution under 1,000V. | 11.30 | 0.39 | +12.9 |
| 8471804000 | UNITS,NESOI,FOR INCORPORATION INTO ADP MACHINES | Compute Hardware | This HS code covers units designed for physical incorporation into automatic data processing machines (computers/servers). | 11.14 | 0.39 | +936.5 |

Table 5: HS10 Codes Ranked 11-20 by 2025 Import Volume with High AI Relevance

| HS10 Code | Description | Category | LLM Reasoning | 2025 Imports ($B) | Share (%) | Change (%) |
|---|---|---|---|---|---|---|
| 8542310045 | PRCSSRS (INCL MICRO): CENTRL PROCSSNG UNITS (CPUS) | Compute Hardware | CPUs are essential compute hardware components in AI data centers, used in servers for both AI workloads and data center infrastructure management. | 10.28 | 1.57 | +0.0 |
| 8471704065 | HARD DISK DRIVE UNT, NESOI, W/OUT EXTNL POWR SUPLY | Compute Hardware | Hard disk drives are essential storage components in AI data centers, used in servers for data storage, model storage, and dataset management. | 6.72 | 1.03 | +68.3 |
| 8542310050 | PROCESSORS (INCLUDING MICROPROCESSORS): OTHER | Compute Hardware | Processors and microprocessors are core computing components essential for AI data center operations. | 6.66 | 1.02 | +0.0 |
| 8473305100 | PTS & ACCESSORIES OF MACH OF HEADING OF 8471,NESOI | Compute Hardware | This is parts/accessories for automatic data processing machines (heading 8471 covers computers/processors) combined with NAICS 333242 for semiconductor machinery manufacturing. | 5.52 | 0.84 | +66.2 |
| 8544429090 | INSULATD ELEC CONDCTR LT=1000V,W/ CONNECTRS, NESOI | Electrical Power | Insulated electric conductors with connectors for voltages $\leq$1000V are essential for data center electrical systems. | 4.85 | 0.74 | +17.4 |
| 8471709000 | OTHER STORAGE UNITS, NESOI | Compute Hardware | Computer storage devices are essential components of AI data centers, used for storing datasets, model parameters, checkpoints, and system software. | 4.24 | 0.65 | +136.3 |
| 8542310075 | PROCESSORS & CONTOLLERS W/NOT COMBO W MEMORIS, ETC | Compute Hardware | Electronic integrated circuits including processors and controllers are core components of AI data center compute infrastructure. | 4.14 | 0.63 | +0.0 |
| 8537109160 | PROGRAMABLE CONTROLLERS, LT=1,000 VOLTS | Electrical Power | Programmable controllers (PLCs) are essential for automating and controlling electrical systems in data centers, including power distribution, cooling systems, fire safety, and building management. | 4.14 | 0.63 | +9.1 |
| 8544700000 | INSULATED OPTICAL FIBER CABLES WITH INDVULY SH FBR | Networking Telecom | Fiber optic cables are essential infrastructure for AI data centers, providing high-bandwidth, low-latency connections required for GPU clusters and distributed computing. | 3.84 | 0.59 | +68.9 |
| 7308909590 | STRUCTURES AND PARTS ETC NESOI IRON OR STEEL | Building Structure | This covers fabricated structural metal components that are essential for data center construction. | 3.79 | 0.58 | -14.4 |