二零二五 THE KTI VIBE
VOL 4. Dec 2025

https://www.publicdomainpictures.net/en/view-image.php?image=274467&picture=merry-christmas-banner

TABLE OF CONTENTS

Author: Zhiyu Wen

Bay Area News

https://commons.wikimedia.org/wiki/File:Eric_Swalwell_(48016282941).jpg

Representative Eric Swalwell has announced that he is running for Governor of California in 2026. The East Bay congressman says he wants to focus on three main issues: lowering energy costs, improving public safety and enhancing school quality. Because he already has support in the Bay Area and experience in Congress, his entry makes the upcoming governor's race more interesting.

Waymo was recently approved to expand its driverless operations onto freeways in the Bay Area, including segments of Highway 101. The approval allows its autonomous cars to travel at typical freeway speeds, and the company says this will eventually make trips between San Francisco, the Peninsula, and major travel hubs faster. While the expansion has raised some questions about safety and traffic management, Waymo says it will roll out freeway service gradually and with close monitoring.

Author: Zhiyu Wen

Earthquakes have shaken the Tri-Valley recently, with the strongest reaching the upper-3 magnitude range near San Ramon. According to the USGS, several of the shallow quakes were felt in nearby communities across the San Ramon Valley and parts of the East Bay. No damage has been reported, but the ongoing “swarm” activity has put some residents on edge. Local officials and earthquake experts are reminding the public to secure heavy furniture and stay prepared for possible aftershocks.

https://holdingsforgood.com/

Despite the region’s strong tech sector, a recent report from Joint Venture Silicon Valley highlights growing inequality and high living costs. Housing remains extremely expensive, putting heavy financial pressure on many lower income and low asset households. Some families, including those with full-time workers, are increasingly relying on food banks to make ends meet. Experts predict that if current trends persist, the widening wealth and income gap could pose significant challenges to the region’s social and economic stability.

Author: Zhiyu Wen

Government Shutdown

Imagine working on a group project where nobody can agree on the topic, so the teacher just cancels class and locks the door. Something similar just happened, not in a school, but in the United States government.

A government shutdown occurs when Congress fails to pass a federal budget or a temporary funding measure, resulting in the closure of large portions of the federal government. Although the shutdown is happening 3,000 miles away in Washington, D.C., the consequences are actually huge for us here in the San Francisco Bay Area, where many residents depend on federal services, federal employment, or transportation systems that rely on federal staffing. Understanding why shutdowns happen and how they affect daily life helps clarify the importance of a stable and functional government.

One of the main causes of a shutdown is partisan disagreement in Congress. Lawmakers must approve yearly appropriations bills to fund federal agencies. When disputes arise over issues such as border spending, social programs, environmental regulations, or defense budgets, the process can stall. Recent years have seen deeper political divides, not only between Democrats and Republicans but also within parties themselves. Leadership struggles in the House and Senate can further delay negotiations, sometimes causing the government to run out of money before a deal is reached.

https://commons.wikimedia.org/wiki/File:Government_Shutdown_%2832905547137%29.png

Author: Zhiyu Wen

Shutdowns produce very real consequences in the Bay Area. One major impact involves federal employees, thousands of whom live and work in the region. Employees at the National Park Service, the Environmental Protection Agency, FEMA, NASA Ames in Mountain View, and federal courts in San Francisco may be furloughed or required to work without pay. During the 2018–2019 shutdown, staff at Point Reyes National Seashore were furloughed, resulting in unmaintained trails and suspended visitor services. Many families, especially those living in areas with high housing costs, struggled to cover rent and bills.

Another significant impact is on transportation, particularly airports. Shutdowns often result in staffing shortages among TSA officers and air traffic controllers. At San Francisco International Airport (SFO) and San Jose Mineta International Airport (SJC), reduced staffing can lead to long security lines and flight delays. Some airports around the country are already closing terminals or reducing operations due to staffing shortages, and the Bay Area has begun experiencing delayed departures at SFO and SJC. If the shutdown continues, canceled flights and more restricted airport operations are likely as staffing levels fall further.

https://commons.wikimedia.org/wiki/File:Cancelled_flights.jpg

https://www.flickr.com/photos/usdagov/4622011485

Shutdowns also affect federal assistance programs. While SNAP benefits continue in the short term, a prolonged shutdown can halt new federal funding, causing states to delay or temporarily limit the distribution of benefits. In the Bay Area, where many low-income families rely on SNAP due to high food and housing costs, such funding interruptions can create serious hardship. Local food banks in San Jose, Oakland, and San Francisco are already preparing for increased demand as the shutdown continues and uncertainty grows.

Schools in the Bay Area are less directly affected because K–12 education is primarily funded at the state and local levels. However, some indirect impacts still occur. For example, federal school lunch reimbursements can be delayed, and research funding for universities such as UC Berkeley, Stanford, and UCSF may be paused. Students applying for FAFSA during a shutdown may experience slower processing if federal verification systems are offline. Additionally, children of furloughed federal workers may face financial stress within their families, which can indirectly impact their school performance.

The history of government shutdowns reveals a pattern of recurring political conflict. The modern shutdown era began in 1976, followed by several short shutdowns in the 1980s. Longer and more disruptive shutdowns occurred in the 1990s during disputes between President Clinton and Congress. More recent shutdowns in 2013 and 2018–2019 were caused by disagreements over the Affordable Care Act and funding for the border wall. The increasing frequency of shutdowns reflects the rising polarization and the use of shutdown threats as a form of political leverage.

To prevent future shutdowns, several solutions have been proposed. One approach is to automatically extend previous-year funding if Congress misses a deadline, preventing the government from closing while debates continue. Another idea is strengthening bipartisan budget committees to reduce last-minute negotiations. Some lawmakers propose limiting the ability of small factions to block funding bills, making it easier for majorities to pass essential legislation. Improving budget planning and increasing transparency around funding decisions could also reduce political tension.

Although shutdowns are often framed as political battles in Washington, their real consequences are felt by families, travelers, and workers, especially in high-cost regions like the Bay Area. Whether through delayed SNAP benefits, canceled flights, unpaid federal employees, or slowed student services, shutdowns disrupt the lives of millions. Strengthening the budget process and encouraging cooperation in Congress could help ensure more reliable government operations and reduce the risk of future shutdowns.

Author: Zhiyu Wen

Data Centers in Space

In November 2025, the launch of the Starcloud-1 satellite marked a historic milestone in computing. This 60 kg satellite, approximately the size of a small refrigerator, features an Nvidia H100 GPU for operation in space. It became one of the world's first operational space-based data centers and is expected to operate for about 11 months. According to Starcloud and Nvidia, this mission is a proof-of-concept that demonstrates high-performance data-center-class compute can run in orbit, powered by nearly continuous solar energy and cooled via radiation to the cold vacuum of space.

The primary advantage of computing in space is energy efficiency. On Earth, solar power isn't always available, and some sunlight is blocked by the atmosphere. Solar panels also occupy a significant amount of land, which can be a problem. In space, the Sun shines almost constantly, providing a steady and powerful source of energy. Satellites in certain orbits can get 30–40% more sunlight than solar panels on the ground. Space is huge and free, so we don't have to worry about taking up land. For the AI industry, which requires substantial amounts of power to train large models, space could provide a clean and virtually unlimited energy source—without placing pressure on local electricity grids or relying on fossil fuels.

Cooling computers in space is a big challenge. Space is extremely cold, with a temperature of almost -270°C (2.7 Kelvin), but the vacuum of space acts like a perfect insulator. Because there is no air, heat cannot escape in the usual way, and it must be released as radiation. This means we don't need the millions of gallons of water that Earth's data centers use for cooling, but we do need very large and special radiators to get rid of the heat. Starcloud-1 is an important experiment to determine if powerful computer hardware can operate safely in these extreme temperatures without the usual cooling systems.

Space also changes how computers connect and send data. Light travels about 50% faster in space than in fiber optic cables on Earth. This means that satellites using lasers will send information faster than undersea cables, which is useful for applications that require extremely fast connections, such as stock trading or global data transfer. Even more importantly, unlike traditional methods, transferring huge amounts of data from satellites to computers on Earth, the satellites can do data processing directly in space, then send the important results back to Earth. This saves a lot of bandwidth and makes data transfer much more efficient.

Building a cloud in space, however, is far from simple. Beyond the atmosphere, electronics face a constant barrage of harmful radiation, necessitating heavy, specialized shielding to survive. This is especially risky for complex hardware, such as Nvidia's GPUs, where a single radiation hit can render a sophisticated chip useless. The logistics are equally unforgiving: unlike terrestrial server farms, orbital units cannot be serviced. If a component breaks, it stays broken. And while launch providers like SpaceX have lowered the barrier to entry, the cost of lifting heavy computing infrastructure against gravity remains a massive economic hurdle.

Ultimately, space-based data centers are likely to complement, rather than replace, terrestrial infrastructure. They are ideal for specialized workloads—ultra-low latency computing, sustainable AI training, and edge processing for Earth observation. Starcloud-1 demonstrates that advanced GPUs, such as the H100, can operate in orbit, marking the first steps toward a future where the "cloud" is no longer a metaphor, but a physical reality floating above us.

Author: Zhiyu Wen

Understanding Modern AI

Artificial Intelligence has quickly moved from science-fiction into everyday life. People now mention “training data,” “tokens,” or “neural networks,” as if these terms were obvious, yet the technology behind them often feels mysterious. How can a machine answer questions, write paragraphs, or chat like a real person? To understand this, it helps to look past the idea of a robot brain and instead see AI for what it really is: a powerful mathematical system that predicts language.

The most intuitive way to imagine modern AI is to think of it as the world’s most powerful autocomplete. Your phone tries to guess the next word in a text message by looking at what you usually type. An AI model does something similar, but on a much larger scale. It has been trained on enormous amounts of text—far more than any person could read. To process this text, the AI splits it into smaller pieces called tokens. A token can be a full word, part of a word, or even a single character, depending on the language. By predicting one token after another, the model gradually learns how words and ideas usually connect. It doesn’t memorize the text; it predicts the next token based on the patterns it has seen, effectively guessing the future of a sentence, an argument, or even a line of code.

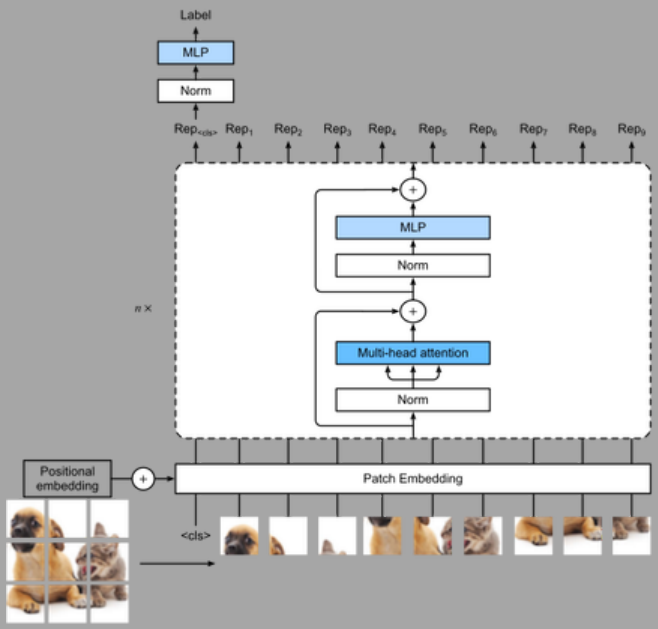

https://commons.wikimedia.org/wiki/File:Vision_Transformer.svg

At the center of this prediction machine is a neural network, a structure filled with billions of tiny adjustable values called parameters. A helpful way to picture this is to imagine a huge sound mixing board with billions of knobs. When a model is new, all the knobs are in random positions, so the output sounds like static—nonsense sentences. During training, the model studies massive amounts of tokenized text. Each time it predicts a token incorrectly, the system measures the error and slightly adjusts the knobs.

After billions of these microscopic corrections, the noise slowly becomes clear, meaningful language. The model hasn't memorized the internet; instead, it has fine-tuned its internal dials to reflect the patterns of human writing.

This is why it's important to distinguish training from inference. Training is like studying in a library for years—absorbing information and slowly adjusting understanding. But when you chat with an AI model, it is no longer learning. It is taking the "final exam." It uses the tuning it learned during training to predict the next tokens, one by one, until it completes a coherent answer. It cannot form new memories or learn during this process; it can only apply what it has already learned.

Does the AI truly understand what it is saying? Not in a human sense. The model has never seen a dog or tasted a sandwich. Instead, it builds a mathematical map of language. On this map, tokens with related meanings sit close together—"dog" near "cat," "pet," and "animal"—while unrelated tokens are far apart. By navigating this map, the model can produce sentences that feel meaningful, even though its "understanding" is really geometry, not lived experience. It behaves as if it understands because the structure of language itself encodes so much about how humans think.

As AI continues to evolve, models are becoming multimodal—able to handle text, images, audio, and even video within one system. Smaller models now run on personal laptops and phones, while larger ones help write software, analyze documents, and support scientific research. Some researchers imagine future systems that help design the next generation of AI, creating a self-improving loop.

But beneath all the advancements, one truth remains constant: AI is a tool built from patterns. It predicts what is most likely to come next, token by token, guided by billions of finely tuned parameters. By viewing AI not as a magical brain but as a massive autocomplete engine, a giant mixing board, and a mathematical map of meaning, the technology becomes far less mysterious—and far more impressive for what it actually is.

Author: Yibin Ma

Persist

In regard to which extent persisting is necessary, persisting in hardship creates innovation and builds resiliency. However, constantly fighting for a goal sometimes leads to failures; nevertheless, persistence creates the opportunity of reaching success. Thus, because persisting leads to innovation, builds personal growth and possibility acquiring successes, to the extent of achieving a goal, persisting is necessary.

https://www.picpedia.org/highway-signs/i/innovation.html

Hardship stimulates the creation of innovations and builds resiliency. Stressed of improving the ineffective AI model, Bard, back in 2024, Google's CEO announced in a blog that Google will merge multiple models and put long term investments to create a better AI model. Stresses from the competitive AI market set Google a goal of creating an effective AI model. Google succeeded in achieving the goal because of its decision of merging different models, which requires consistent effort in building and testing the new model. In addition, when Edison failed for thousands of times of creating a new light bulb, he describes, with optimism, his position as finding ways that does not work, rather than failures. Facing failures crashes hope, while facing thousands of failures crashes the leftover hope of Edison further. However, he persist, which builds his resiliency when facing failures. Therefore, to the extent of achieving a goal, persisting is necessary.

While hardworking sometimes gains failure, however, persisting builds the stepping stone for succeeding. Despite consistently making a huge investment around ten to thirty billions of dollars in the new VR glasses in Apple since 2015, their glass failed when the customers find out that they need to put the battery in their pockets while putting on the glass.

https://www.picpedia.org/highway-signs/h/hard-work.html

It is true that Apple put constant efforts inside the glass. However their lack of technological skills and patience in making innovations before publishing the glass results in the failure of the glass. Nevertheless, Meta's VR glass stands out among the VR market due to their large investment inside the field over 70 billion dollars since late 2020. With massive investments and long persistence, Meta's VR glasses gained favor among consumers. Persistence is a stepping stone for achieving success in the VR field, while other factors like investment of money also plays a crucial role in achieving successes. Thus, persistence is necessary to the extent of achieving a goal.

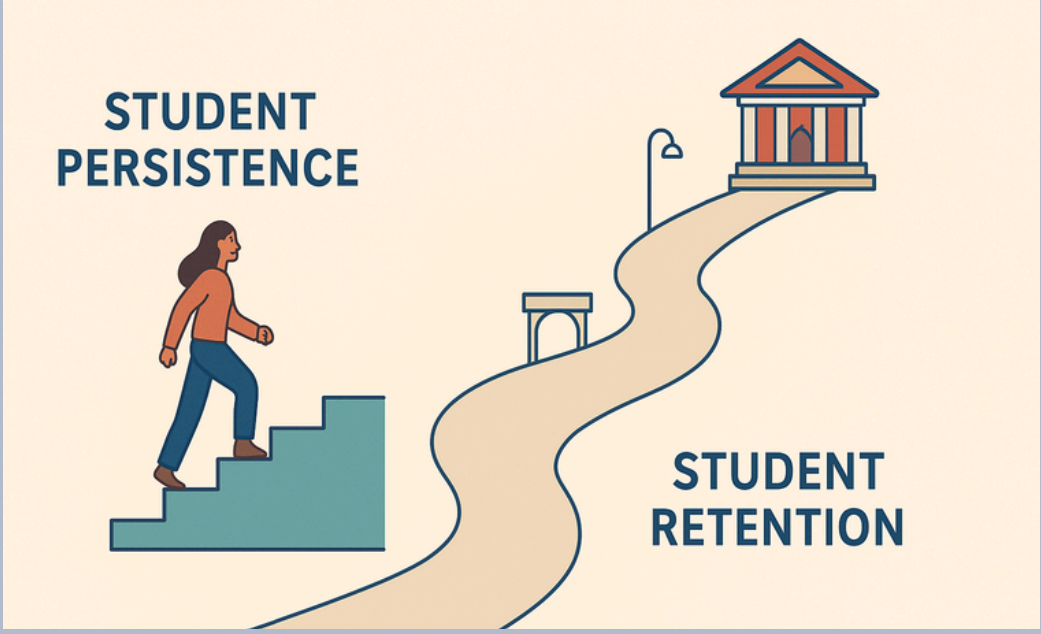

https://moderncampus.com/blog/persistence-vs-retention.html

Constantly working for a goal throughout hardship leads to successes and builds resiliency. However, persistence sometimes results in failures; still, persistence provides the stepping stone of achieving a goal. Therefore, to the extent of achieving a goal, persistence plays a crucial role because it builds successes and resiliency. There are many hardships in life. While the whole world seems to against a goal, persistence is the only choice when making a decision of choosing not be a coward.

https://upload.wikimedia.org/wikipedia/commons/7/7f/THEORY_OF_SUCCESS.jpg

Author: TianHe Si

青春岁月,我与星辰相约

熬得过万丈孤独,藏得下星空大海。

———题记

很多次在夜晚的路上踌躇,心里想的却不是孤单和漫长,而是波澜壮阔的大海和漫天的星光。偶然回首,才发现自己一个人咬着牙走了那么远的路。星常在青春是一个人的孤军奋战,可以路上所遇到的每一个人都像万千繁星中的一颗,虽单个只是影影绰绰,但零零散散地汇总在一起便也成了一片星空,不耀眼但却足以照亮我们前行的路。

一路上,我么会遇到很多人。有些人嬉笑怒骂着人生,却在别人看不到的地方挑灯苦读;有些人号称我用青春赌明天,在生命里消耗着光阴;有些人踌躇满志,步履维艰,却无时无刻的努力前行。在青葱的岁月里,谁在苦苦坚持,谁在肆意挥霍,谁在播撒着无尽的希望。

每个人的人生都是没有彩排的现场直播。在自己的人生里,他们都留下或深或浅的痕迹。唯有对美好人生的渴求,才是我们点亮的夜空的希望。也许随着岁月的消逝,我们会忘记他们坚持的模样,但难以忘怀的是前行路上的点点星光及始终伴我们前行的希望。

夜自明,一直觉得学校是一个残酷的地方,一座一座安静地伫立于荒凉的时间轴上,把青春固定在狭小的空间里。

苦涩的奋战中,还要自欺欺人的说着青春无悔,愿赌服输;明明处在最美好的年华中,却要听信年长者的欺骗,把快乐与希望寄托于毕业和长大。时光从不留恋过往,只是冷眼旁观着大部分人回头寻找记忆,却提供不了一丝余温。就像是无烛无光的夜晚,黑沉沉地掩盖希望,每一朵云都下落不明,每一盏月亮都不知所踪,荒凉了人心,兴起着风浪。

人生不是有了希望才去努力,而是努力了才会有希望。在青春的词奥原理,无尽的希望,成为黎明划破黑夜的那束光。我们在黑夜里寻找着光明,也许努力的唯一意义,便是让我们可以跨越重重的荆棘,不用听从命运的安排,奋力向前,成就更好的自己。别再说为谁而努力,努力只是为了点亮夜空中的自己。别害怕黑暗,黑暗让我们不害怕孤独,相信有了努力的星辰,便无烛无光夜自明。当我们感到孤单,想放弃的时候,别只看到前方路途遥远,要记住身后的一路坚持与头顶的满天星光。

路上会有风,会有浪,会有悲伤,会有孤独,也会有无尽的星辰与希望!

I READ
I TRAVEL
I BECOME

知识树月刊

$5/Month (1 Issue)
$36/Year (12 issues)
订阅或捐款via Zelle 至:
office@knowledgetreeinstitute.com

Knowledge
Tree Institute
Where we grow together

杂志广告请接洽

(510)328-6688

Email:
office@knowledgetreeinstitute.com