INFOLINE

Volume XI Issue 2 October 2020

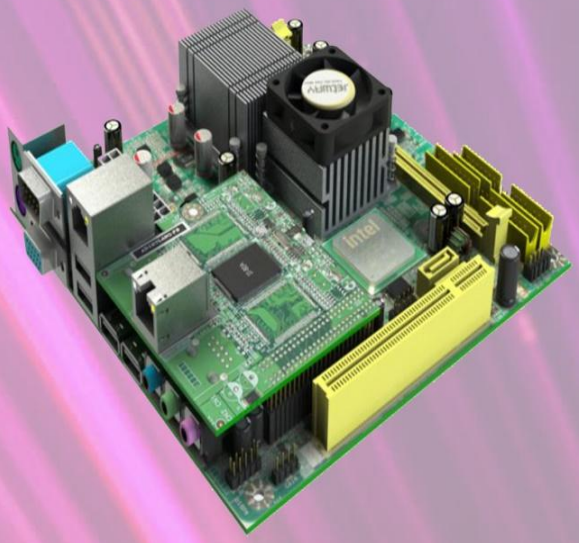

DEPARTMENT OF COMPUTER TECHNOLOGY AND INFORMATION TECHNOLOGY

KONGU ARTS AND SCIENCE COLLEGE
(Autonomous)
Affiliated to Bharathiar University, Coimbatore
Approved by UGC and AICTE, New Delhi
Re-accredited by NAAC, DBT STAR College Scheme
An ISO 9001 : 2015 Certified Institution
Nanjanapuram, Erode - 638 107

INFOLINE

EDITORIAL BOARD

CONTENTS

ROBOTIC PROCESS AUTOMATION (RPA)

Robotic Process Automation (RPA) is a technology used for automating daily tasks, similar to artificial intelligence. Here, the software is used for automating repetitive tasks such as handling and replying to emails, processing transactions, and handling business data. This technology is used for automating tasks for low-level employees to higher-ranking officials. RPA can automate more than 40% of daily tasks. According to McKinsey, more than 60% of all repetitive tasks can be partially automated using RPA. So, this technology is going to threaten a lot of jobs.

When a software program emulates human actions on digital systems to perform business operations, we call it a Robot Process Automation (also known as Robot Process Management). Its abbreviation is RPA. RPA helps organizations in automating repetitive processes which saves their time and resources. The name of RPA is self-explanatory as a robot is something that imitates human actions. A process is a sequence of steps that cause meaningful activity and automation is when a program performs a task without human supervision or intervention.

RPA Architecture Components

The architecture of any RPA tool is very complex and has multiple components. Each of these components has a set of duties to perform. Here is a small list of all elements of an RPA architecture:

- RPA tools
- Platform
- Execution Infrastructure
- Configuration Management

Tools of RPA

RPA tools should have several capabilities. They should be able to automate multiple kinds of application environments, including Citrix, desktop, and web environments. They should be able to develop digital robots and them through programming logic. These robots should understand through configuring and recordings. They should build components you can reuse later and apply them to other robots. This way, you'd have fast deployment speed and modularity. Having such elements also makes the maintenance of them more manageable.

The tools of your RPA software should have the ability to read and write from multiple data sources. They have to build shared application object stores and repositories for keeping locators so they should be capable of that as well.

D.UDHYAKUMAR

I B.Sc. (Computer Technology)

SECURITY SOFTWARE FOR AUTONOMOUS VEHICLES

Before autonomous vehicles participate in road traffic, they must demonstrate conclusively that they do not pose a danger to others. New software developed at the Technical University of Munich (TUM) prevents accidents by predicting different variants of a traffic situation every millisecond. A car approaches an intersection. Another vehicle jets out of the cross street but it is not yet clear whether it will turn right or left. At the same time, a pedestrian steps into the lane directly in front of the car, and there is a cyclist on the other side of the street. People with road traffic experience will in general assess the movements of other traffic participants correctly.

Algorithms that peer into the future

The ultimate goal when developing software for autonomous vehicles is to ensure that they will not cause accidents. Althoff who is a member of the Munich School of Robotics and Machine Intelligence at TUM, and his team have now developed a software module that permanently analyzes and predicts events while driving. Vehicle sensor data are recorded and evaluated every millisecond. The software can calculate all possible movements for every traffic participant provided they adhere to the road traffic regulations allowing the system to look three to six seconds into the future.

Based on these future scenarios, the system determines a variety of movement options for the vehicle. At the same time, the program calculates potential emergency maneuvers in which the vehicle can be moved out of harm's way by accelerating or braking without endangering others. The autonomous vehicle may only follow routes that are free of foreseeable collisions and for which an emergency maneuver option has been identified.

Streamlined models for swift calculations

This kind of detailed traffic situation forecasting was previously considered too time-consuming and thus impractical. But now, the Munich research team has shown not only the theoretical viability of real-time data analysis with simultaneous simulation of future traffic events. They have also demonstrated that it delivers reliable results.

The quick calculations are made possible by simplified dynamic models. So-called reachability analysis is used to calculate

potential future positions a car or a pedestrian might assume. When all characteristics of the road users are taken into account, the calculations become prohibitively time-consuming. That is why Althoff and his team work with simplified models. These are superior to the real ones in terms of their range of motion yet, mathematically easier to handle. This enhanced freedom of movement allows the models to depict a larger number of possible positions but includes the subset of positions expected for actual road users.

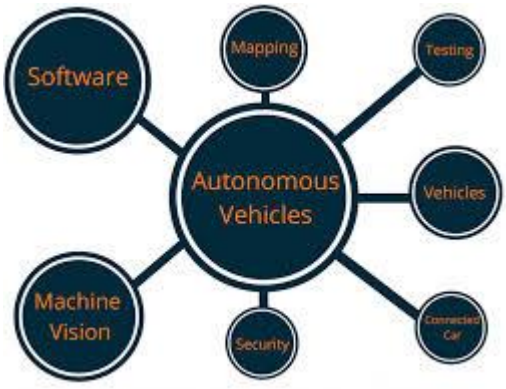

Real traffic data for a virtual test environment

For their evaluation, the computer scientists created a virtual model based on real data they had collected during test drives with an autonomous vehicle in Munich. This allowed them to craft a test environment that closely reflects everyday traffic scenarios. "Using the simulations, we were able to establish that the safety module does not lead to any loss of performance in terms of driving behavior, the predictive calculations are correct, accidents are prevented, and in emergency situations the vehicle is demonstrably brought to a safe stop," Althoff sums up. The computer scientist emphasizes that the new security software could simplify the development of autonomous vehicles because it can be combined with all standard motion control programs.

N.R. SHARMILA

II B.Sc. (Information Technology)

AI-AS-A-SERVICE

Artificial Intelligence (AI) is one of the most transformative tech evolutions of our times. As I highlighted in my book 'Artificial Intelligence in Practice', most companies have started to explore how they can use AI to improve the customer experience and to streamline their business operations. This will continue in 2020 and while people will increasingly become used to working alongside AIs, designing and deploying our own AI-based systems will remain an expensive proposition for most businesses. For this reason, much of the AI applications will continue to be done through providers of as-a-service platforms which allow us to simply feed in our own data and pay for the algorithms or compute resources as we use them.

R.SHOBIKA

II B.Sc. (Computer Technology)

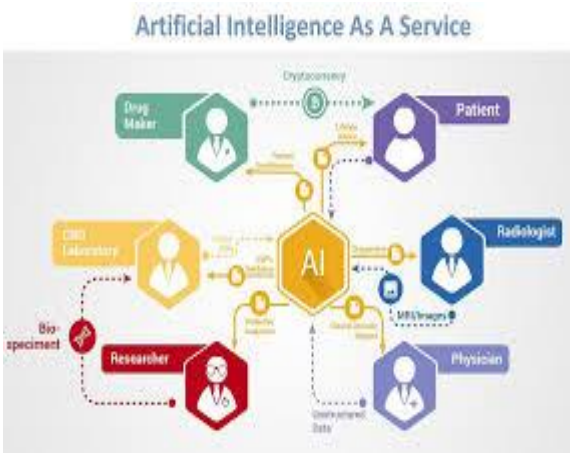

Currently, these platforms, provided by the Amazon, Google and Microsoft tend to be somewhat broad in scope with (often expensive) custom-engineering required to apply them to the specific tasks an organization may require. During 2020, we will see wider adoption and a growing pool of providers that are likely to start offering more tailored applications and services for specific or specialized tasks. This means no company will have any excuses left not to use AI.

Super-fast data networks will not only give us the ability to stream movies and music at higher quality when we're on the move. The greatly increased speeds mean that mobile networks will become more usable even than the wired networks running into our homes and businesses. Companies must consider the business implications of having super-fast and stable internet access anywhere. The increased bandwidth will enable machines, robots, and autonomous vehicles to collect and transfer more data than ever, leading to advances in the area of the Internet of Things (IoT) and smart machinery.

EXTENDED REALITY

Extended Reality (XR) is a catch-all term that covers several new and emerging technologies being used to create more immersive digital experiences. More specifically, it refers to virtual, augmented, and mixed reality. Virtual reality (VR) provides a fully digitally immersive experience where you enter a computer-generated world using headsets that blend out the real world. Augmented reality (AR) overlays digital objects onto the real world via smartphone screens or displays. Mixed reality (MR) is an extension of AR, that means users can interact with digital objects placed in the real world (think playing a holographic piano that you have placed into your room via an AR headset).

These technologies have been around for a few years now but have largely been confined to the world of entertainment with Oculus Rift and Vive headsets providing the current state-of-the-art in videogames and smartphone features such as camera filters and Pokemon Go-style games providing the most visible examples of AR. From 2020 expect all of that to change, as businesses get to grips with the wealth of exciting possibilities offered by both current forms of XR. Virtual and augmented reality will become increasingly

prevalent for training and simulation as well as offering new ways to interact with customers.

D.KRISHNAKUMAR
III B.Sc. (Information Technology)

NEW TRANSISTOR DESIGN DISGUISES KEY COMPUTER CHIP HARDWARE FROM HACKERS

A hacker can reproduce a circuit on a chip by discovering what key transistors are doing in a circuit but not if the transistor "type" is undetectable. Purdue University engineers have demonstrated a way to disguise which transistor is which by building them out of a sheet-like material called black phosphorus. This built-in security measure would prevent hackers from getting enough information about the circuit to reverse engineer it.

Reverse engineering chips is a common practice both for hackers and companies investigating intellectual property infringement. Researchers also are developing x-ray imaging techniques that wouldn't require actually touching a chip to reverse engineer it. The approach that Purdue researchers have demonstrated would increase security on a more fundamental level. How chip manufacturers choose to make this transistor design compatible with their processes would determine the availability of this level of security.

A chip computes using millions of transistors in a circuit. When a voltage is applied, two distinct types of transistors an N type and a P type perform a computation. Replicating the chip would begin with identifying these transistors.

"These two transistor types are key since they do different things in a circuit. They are at the heart of everything that happens on all our chips," said Joerg Appenzeller, Purdue's Barry M. and Patricia L. Epstein Professor of Electrical and Computer Engineering. If these two transistor types appeared identical upon inspection, a hacker wouldn't be able to reproduce a chip by reverse engineering the circuit. Appenzeller's team showed in their study that camouflaging the transistors by fabricating them from a material such as black phosphorus makes it impossible to know which transistor is which. When a voltage toggles the transistors' type, they appear exactly the same to a hacker.

While camouflaging is already a security measure that chip manufacturers use, it is typically done at the circuit level and doesn't attempt to obscure the functionality of individual transistors leaving the chip potentially vulnerable to reverse engineering hacking techniques with the right tools. The camouflaging method that Appenzeller's team demonstrated will be building a security key into the transistors.

"The approach would make N and P type transistors look the same on a fundamental level. You can't really distinguish them without knowing the key," said Peng Wu, a Purdue Ph.D. student of electrical and computer engineering who built and tested a prototype chip with black phosphorus-based transistors in the Birck Nanotechnology Center of Purdue's Discovery Park.

Not even the chip manufacturer would be able to extract this key after the chip is produced. Current camouflaging techniques always require more transistors in order to hide what's going on in the circuit. But hiding the transistor type using a material like black phosphorus a material as thin as an atom requires fewer transistors, taking up less space and power in addition to creating a better disguise, the researchers said.

The idea of obscuring the transistor type to protect chip intellectual property originally came from a theory by University of Notre Dame professor Sharon Hu and her collaborators. Typically, what gives N and P type transistors away is how they carry a current. N type transistors carry a current by transporting electrons while P type transistors use the absence of electrons called holes. Black phosphorus is so thin Appenzeller's team realized, that it would enable electron and hole transport at a similar current level making the two types of transistors appear more fundamentally the same per Hu's proposal.

Appenzeller's team then experimentally demonstrated the camouflaging abilities of black phosphorus-based transistors. These transistors are also known to operate at the low voltages of a computer chip at room temperature due to their smaller dead zone for electron transport described as a small "band gap." But despite the advantages of black phosphorus, the chip manufacturing industry would more likely use a different material to achieve this camouflage effect.

The industry is starting to consider ultrathin, 2D materials because they would allow more transistors to fit on a chip, making them more powerful. Black phosphorus is a little too volatile to be compatible with current processing techniques but showing experimentally how a 2D material could work is a step toward figuring out how to implement this security measure.

B. A.AKSHAYA SHREE
III B.Sc. (Computer Technology)

NEXT-GENERATION COMPUTER CHIP WITH TWO HEADS

EPFL engineers have developed a computer chip that combines two functions logic operations and data storage into a single architecture, paving the way to more efficient devices. Their technology is particularly

promising for applications relying on artificial intelligence.

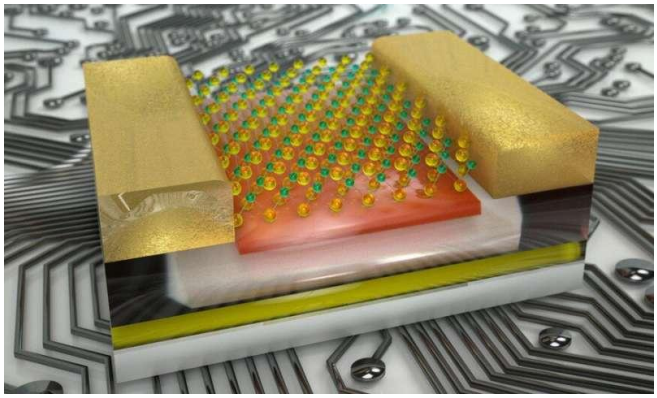

It's a major breakthrough in the field of electronics. Engineers at EPFL's Laboratory of Nanoscale Electronics and Structures (LANES) have developed a next-generation circuit that allows for smaller, faster and more energy-efficient devices which would have major benefits for artificial-intelligence systems. Their revolutionary technology is the first to use a 2-D material for what's called a logic-in-memory architecture, or a single architecture that combines logic operations with a memory function. The research team's findings appear today in Nature.

Until now, the energy efficiency of computer chips has been limited by the von Neumann architecture they currently use, where data processing and data storage take place in two separate units. That means data must constantly be transferred between the two units, using up a considerable amount of time and energy.

By combining the two units into a single structure, engineers can reduce these losses. That's the idea behind the new chip developed at EPFL, although it goes one step beyond existing logic-in-memory devices. The EPFL chip is made from MoS2, which is a 2-D material consisting of a single layer that's only three atoms thick. It's also an excellent semi-conductor. LANES engineers had already studied the specific properties of MoS2 a few years ago, finding that it is particularly well-suited to electronics applications. Now the team has taken that initial research further to create their next-generation technology.

The EPFL chip is based on floating-gate field-effect transistors (FGFETs). The advantage of these transistors is that they can hold electric charges for long periods; they are typically used in flash memory systems for cameras, smartphones and computers. The unique electrical proprieties of MoS2 make it particularly sensitive to charges stored in FGFETs which is what enabled the LANES engineers to develop circuits that work as both memory storage units and programmable transistors. By using MoS2, they were able to incorporate numerous processing functions into a single circuit and then change them as desired.

In-depth expertise

"This ability for circuits to perform two functions is similar to how the human brain works where neurons are involved in both

storing memories and conducting mental calculations," says Andras Kis, the head of LANES. "Our circuit design has several advantages. It can reduce the energy loss associated with transferring data between memory units and processors, cut the amount of time needed for computing operations and shrink the amount of space required. That opens the door to devices that are smaller, more powerful and more energy efficient."

S.KANAGALAKSHMI
I B.Sc. (Information Technology)

LOW-POWER, FLEXIBLE MEMRISTOR CIRCUIT FOR MOBILE AND WEARABLE DEVICES

A KAIST research team succeeded in developing an energy efficient, non-volatile logic-in-memory circuit by using a memristor. This novel technology can be used as an energy efficient computing architecture for battery-powered flexible electronic systems such as mobile and wearable devices.

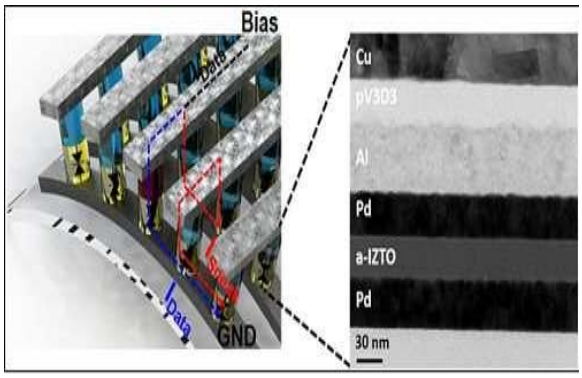

Professor Sung-Yool Choi from the School of Electrical Engineering and Professor Sang-Hee Ko Park from the Department of Materials Science and Engineering developed a memristive non-volatile logic-in-memory circuit.

Transistor-based conventional electronic systems have issues with battery supply and a long standby period due to their volatile computing architecture. The standby power consumption caused by sub threshold leakage current limits their potential applications for mobile electronic devices. Also, their physical separation of memory and processor causes power consumption and time delay during data transfer.

In order to solve this problem, the team developed a logic-in-memory circuit that enables data storage as well as logic operation simultaneously. It can minimize energy consumption and time delay because it does not require data transfer between memory and processor.

The team employed non-volatile, polymer-based memristors and flexible back-to-back Schottky diode selector devices on plastic substrates. Unlike the conventional architecture, this memristive non-volatile logic-in-memory is a novel computing architecture that consumes a minimal amount of standby power. This one-selector-one memristor (1S-

1M) solved the issue of undesirable leakage currents known as 'sneak currents'.

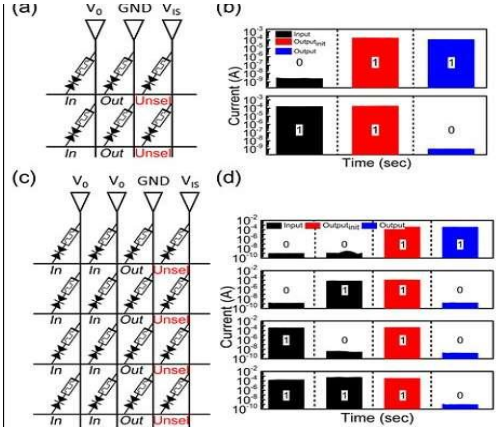

V.LEEPIKA

II B.Sc. (Information Technology)

RESEARCHERS DEVELOP FRAMEWORK THAT IMPROVES FIREFOX SECURITY

Computer scientists develop a technique to protect browsers from buggy third-party libraries

```
// cIndex is primarily used to index into the clusters array which has size
// aLength below. As cIndex is not changing anymore, let's just verify it
// and remove the tainted wrapper.
uint32_t cIndex =
    CopyAndVerifyOrFail(data->cIndex, val < aLength, &failedVerify);
if (failedVerify) {
  return NS_ERROR_ILLEGAL_VALUE;
}
// now put glyphs into the textrun, one cluster at a time
for (uint32_t i = 0; i <= cIndex; ++i) {
  // We makes a local copy of "clusters[i]" which is of type
  // tainted_gr<gr_glyph_to_char_cluster> below. We do this intentionally
  // rather than taking a reference. Taking a reference with the code
  //
  // tainted_volatile_gr<gr_glyph_to_char_cluster>& c = clusters[i];
  //
  // produces a tainted_volatile which means the value can change at any
  // moment allowing for possible time-of-check-time-of-use vuln. We thus
  // make a local copy to simplify the verification.
  tainted_gr<gr_glyph_to_char_cluster> c = clusters[i];

  tainted_gr<float> t_adv;  // total advance of the cluster
  if (rtl) {
    if (i == 0) {
      t_adv = sandbox_invoke(*mSandbox, gr_seg_advance_X, aSegment) -
              xLocs[c.baseGlyph];
    } else {
      t_adv = xLocs[clusters[i - 1].baseGlyph] - xLocs[c.baseGlyph];
    }
  } else {
    if (i == cIndex) {
      t_adv = sandbox_invoke(*mSandbox, gr_seg_advance_X, aSegment) -
              xLocs[c.baseGlyph];
    } else {
      t_adv = xLocs[clusters[i + 1].baseGlyph] - xLocs[c.baseGlyph];
    }
  }

  float adv = t_adv.unverified_safe_because(
      "Per Bug 1569464 - this is the advance width of a glyph or cluster of "
      "glyphs. There are no a-priori limits on what that might be. Incorrect "
      "values will tend to result in bad layout or missing text, or bad "
      "nscoord values. But, these will not result in safety issues.");

  // check unexpected offset - offs used to index into aText
  uint32_t offs =
      CopyAndVerifyOrFail(c.baseChar, val < aLength, &failedVerify);
```

Researchers from the University of California San Diego, University of Texas at Austin, Stanford University and Mozilla have developed a new framework to improve web browser security. The framework, called RLBox, has been integrated into Firefox to complement its other security-hardening efforts. RLBox increases browser security by separating third-party libraries that are vulnerable to attacks from the rest of the browser to contain potential damage a practice called sandboxing. The study will be published in the proceedings of the USENIX Security Symposium in March.

Browsers, like Firefox, rely on third-party libraries to support media decoding (e.g., rendering images or playing audio files) among many other functionalities. These libraries are often written in low-level programming languages, like C and highly optimized for performance.

"Unfortunately, bugs in C code are often security vulnerabilities security vulnerabilities that attackers are really good at exploiting," noted senior author Deian Stefan, an Assistant Professor with UC San Diego's Department of Computer Science and Engineering.

RLBox allows browsers to continue to use off-the-shelf, highly tuned libraries without worrying about the security impact of these libraries. "By isolating libraries we can ensure that attackers can't exploit bugs in these

libraries to compromise the rest of the browser," said the lead PhD student on the project, Shravan Narayan.

A key piece of RLBox is the underlying sandboxing mechanism which keeps a buggy library from interfering with the rest of the browser. The study investigates various sandboxing techniques with different trade-offs. But the team ultimately partnered with the engineering team at San Francisco-based Fastly to adopt a sandboxing technique based on WebAssembly, a new intermediate language designed with sandboxing in mind. The team believes that WebAssembly will be a key part of future secure browsers and secure systems more broadly. The WebAssembly sandboxing effort is detailed in a recent Mozilla Hacks blog post.

"Unfortunately, it's not enough to put a library in a sandbox, you need to carefully check all the data that comes out of the sandbox otherwise a sophisticated attacker can trick the browser into doing the wrong thing and render the sandboxing effort useless," said Stefan. RLBox eliminates these classes of attacks by tagging everything that crosses the boundary and ensuring that all such tagged data are validated before being used.

RLBox has been integrated into Mozilla's Firefox and will be shipping to Linux users in Firefox 74 and Mac users in Firefox 75 with plans to implement in other platforms. "This is a big deal," says Bobby Holley, principal engineer at Mozilla. "Security is a top priority for us, and it's just too easy to make dangerous mistakes in C/C++. We're writing a lot of new code in Rust, but Firefox is a huge codebase with millions of lines of C/C++ that aren't going away any time soon. RLBox makes it quick and easy to isolate existing chunks of code at a granularity that hasn't been possible with the process-level sandboxing used in browsers today."

In the study, the team isolated half a dozen libraries using RLBox. To start, Firefox will ship with their sandboxed Graphite font shaping library. Mozilla plans to apply the sandboxing more broadly in the future, ultimately making millions of users' browsers more secure.

R.JANANI

I B.Sc. (Computer Technology)

TRANSLATING LOST LANGUAGES USING MACHINE LEARNING

Recent research suggests that most languages that have ever existed are no longer spoken. Dozens of these dead languages are also considered to be lost or "undeciphered" that is, we don't know enough about their grammar, vocabulary or syntax to be able to actually understand their texts.

Lost languages are more than a mere academic curiosity; without them, we miss an entire body of knowledge about the people who spoke them. Unfortunately most of them have such minimal records that scientists can't decipher them by using machine-translation algorithms like Google Translate. Some don't have a well-researched "relative" language to be compared to and often lack traditional dividers like white space and punctuation. However researchers at MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) recently made a major development in this area: a new system that has been shown to be able to automatically decipher a lost language, without needing advanced knowledge of its relation to other languages. They also showed that their system can itself determine relationships between languages, and they used it to corroborate recent scholarship suggesting that the language of Iberian is not actually related to Basque.

The team's ultimate goal is for the system to be able to decipher lost languages that have eluded linguists for decades, using just a few thousand words. Spearheaded by MIT Professor Regina Barzilay, the system relies on several principles grounded in insights from historical linguistics, such as the fact that languages generally only evolve in certain predictable ways. For instance, while a given language rarely adds or deletes an entire sound, certain sound substitutions are likely to occur. A word with a "p" in the parent language may change into a "b" in the descendant language, but changing to a "k" is less likely due to the significant pronunciation gap.

By incorporating these and other linguistic constraints, Barzilay and MIT PhD student Jiaming Luo developed a decipherment algorithm that can handle the vast space of possible transformations and the scarcity of a guiding signal in the input. The algorithm learns to embed language sounds into a multidimensional space where differences in pronunciation are reflected in the distance between corresponding vectors. This design enables them to capture pertinent patterns of language change and express them as computational constraints. The resulting model can segment words in an ancient language and map them to counterparts in a related language.

The project builds on a paper Barzilay and Luo wrote last year that deciphered the

dead languages of Ugaritic and Linear B, the latter of which had previously taken decades for humans to decode. However, a key difference with that project was that the team knew that these languages were related to early forms of Hebrew and Greek, respectively.

With the new system, the relationship between languages is inferred by the algorithm. This question is one of the biggest challenges in decipherment. In the case of Linear B, it took several decades to discover the correct known descendant. For Iberian, the scholars still cannot agree on the related language: Some argue for Basque, while others refute this hypothesis and claim that Iberian doesn't relate to any known language.

The proposed algorithm can assess the proximity between two languages in fact when tested on known languages it can even accurately identify language families. The team applied their algorithm to Iberian considering Basque, as well as less-likely candidates from Romance, Germanic, Turkic and Uralic families. While Basque and Latin were closer to Iberian than other languages, they were still too different to be considered related.

In future work, the team hopes to expand their work beyond the act of connecting texts to related words in a known language an approach referred to as "cognate-based decipherment." This paradigm assumes that such a known language exists, but the example of Iberian shows that this is not always the case. The team's new approach would involve identifying semantic meaning of the words, even if they don't know how to read them.

D.KRISHNAKUMAR

III B.Sc. (Information Technology)

SYSTEM BRINGS DEEP LEARNING TO "INTERNET OF THINGS" DEVICES

Advance could enable artificial intelligence on household appliances while enhancing data security and energy efficiency.

Deep learning is everywhere. This branch of artificial intelligence curates your social media and serves your Google search results. Soon, deep learning could also check your vitals or set your thermostat. MIT researchers have developed a system that could bring deep learning neural networks to new and much smaller places like the tiny computer

chips in wearable medical devices, household appliances and the 250 billion other objects that constitute the "internet of things" (IoT).

The system, called MCUNet, designs compact neural networks that deliver unprecedented speed and accuracy for deep learning on IoT devices, despite limited memory and processing power. The technology could facilitate the expansion of the IoT universe while saving energy and improving data security.

The research will be presented at next month's Conference on Neural Information Processing Systems. The lead author is Ji Lin, a PhD student in Song Han's lab in MIT's Department of Electrical Engineering and Computer Science. Co-authors include Han and Yujun Lin of MIT, Wei-Ming Chen of MIT and National University Taiwan, and John Cohn and Chuang Gan of the MIT-IBM Watson AI Lab.

The Internet of Things

The IoT was born in the early 1980s. Grad students at Carnegie Mellon University including Mike Kazar '78, connected a Cola-Cola machine to the internet. The group's motivation was simple: laziness. They wanted to use their computers to confirm the machine was stocked before trekking from their office to make a purchase. It was the world's first internet-connected appliance. "This was pretty much treated as the punchline of a joke," says Kazar, now a Microsoft engineer. "No one expected billions of devices on the internet."

Since that Coke machine, everyday objects have become increasingly networked into the growing IoT. That includes everything from wearable heart monitors to smart fridges that tell you when you're low on milk. IoT devices often run on microcontrollers simple computer chips with no operating system, minimal processing power, and less than one thousandth of the memory of a typical smartphone. So pattern-recognition tasks like deep learning are difficult to run locally on IoT devices. For complex analysis, IoT-collected data is often sent to the cloud, making it vulnerable to hacking.

With MCUNet, Han's group codesigned two components needed for "tiny deep learning" the operation of neural networks on microcontrollers. One component is TinyEngine, an inference engine that directs resource management, akin to an operating system. TinyEngine is optimized to run a particular neural network structure, which is selected by MCUNet's other component: TinyNAS, a neural architecture search algorithm.

System-Algorithm Co-design

Designing a deep network for microcontrollers isn't easy. Existing neural architecture search techniques start with a big pool of possible network structures based on a predefined template, then they gradually find the one with high accuracy and low cost. While the method works, it's not most efficient. "It can work pretty well for GPUs or smartphones," says Lin. "But it's been difficult to directly apply these techniques to tiny microcontrollers, because they are too small."

So Lin developed TinyNAS, a neural architecture search method that creates custom-sized networks. "We have a lot of microcontrollers that come with different power capacities and different memory sizes," says Lin. "So we developed the algorithm [TinyNAS] to optimize the search space for different microcontrollers." The customized nature of TinyNAS means it can generate compact neural networks with the best possible performance for a given microcontroller with no unnecessary parameters.

To run that tiny neural network, a microcontroller also needs a lean inference engine. A typical inference engine carries some dead weight instructions for tasks it may rarely run. The extra code poses no problem for a laptop or smartphone, but it could easily overwhelm a microcontroller. "It doesn't have off-chip memory, and it doesn't have a disk," says Han. "Everything put together is just one megabyte of flash, so we have to really carefully manage such a small resource." Cue TinyEngine.

The researchers developed their inference engine in conjunction with TinyNAS. TinyEngine generates the essential code necessary to run TinyNAS' customized neural network. Any deadweight code is discarded, which cuts down on compile-time. "We keep only what we need," says Han. "And since we designed the neural network, we know exactly what we need. That's the advantage of system-algorithm codesign." In the group's tests of TinyEngine, the size of the compiled binary code was between 1.9 and five times smaller than comparable microcontroller inference engines from Google and ARM. TinyEngine also contains innovations that reduce runtime, including in-place depth-wise convolution, which cuts peak memory usage nearly in half. After codesigning TinyNAS and TinyEngine, Han's team put MCUNet to the test.

MCUNet's first challenge was image classification. The researchers used the ImageNet database to train the system with labeled images, then to test its ability to classify novel ones. On a commercial microcontroller they tested, MCUNet successfully classified 70.7 percent of the novel images the previous state-of-the-art neural

network and inference engine combo was just 54 percent accurate. “Even a 1 percent improvement is considered significant,” says Lin. “So this is a giant leap for microcontroller settings.”

The team found similar results in ImageNet tests of three other microcontrollers. And on both speed and accuracy, MCUNet beat the competition for audio and visual “wake-word” tasks where a user initiates an interaction with a computer using vocal cues (think: “Hey, Siri”) or simply by entering a room. The experiments highlight MCUNet’s adaptability to numerous applications.

B. A.AKSHAYA SHREE

III B.Sc. (Computer Technology)

◆••••••••••••••••••••••••••••••••••◆

NEURAL NETWORKS EVERYWHERE

New chip reduces neural networks power consumption by up to 95 percent, making them practical for battery powered devices. Most recent advances in artificial intelligence systems such as speech or face recognition programs have some courtesy of neural networks, densely interconnected meshes of simple information processors that learn to perform tasks by analysing huge sets of training data.

But neural nets are large, and their computations are energy intensive, so they’re not very practical for handheld devices. Most smartphone apps that rely on neural nets simply upload data to internet servers which process it and send the results back to the phone.

Researchers have developed a special-purpose chip that increases the speed of neural-network computations by three to seven times over its predecessors while reducing power consumption 94 to 95 percent. That could make it practical to run neural networks locally on smartphones or even to embed them in household appliances. The general processor model is that there is a memory in some part of the chip, and there is a processor in another part of the chip, and you move the data back and forth between them when you do these computations

Since these machine-learning algorithms need so many computations, this transferring back and forth of data is the dominant portion of the energy consumption. But the computation these algorithms do can be

simplified to one specific operation, called the dot product. This approach was implemented in dot-product functionality inside the memory so that you don't need to transfer this data back and forth.

Back to Analog

Neural networks are typically arranged into layers. A single processing node in one layer of the network will generally receive data from several nodes in the layer below and pass data to several nodes in the layer above. Each connection between nodes has its own weight which indicates how large a role the output of one node will play in the computation performed by the next. Training the network is a matter of setting those weights.

A node receiving data from multiple nodes in the layer below will multiply each input by the weight of the corresponding connection and sum the results. That operation the summation of multiplications is the definition of a dot product. If the dot product exceeds some threshold value, the node will transmit it to nodes in the next layer over connections with their own weights.

A neural net is an abstraction: The nodes are just weights stored in a computer's memory. Calculating a dot product usually involves fetching a weight from memory, fetching the associated data item, multiplying the two, storing the result somewhere, and then repeating the operation for every input to a node. Given that a neural net will have thousands or even millions of nodes, that's a lot of data to move around.

But that sequence of operations is just a digital approximation of what happens in the brain, where signals travelling along multiple neurons meet at a "synapse," or a gap between bundles of neurons. The neurons' firing rates and the electrochemical signals that cross the synapse correspond to the data values and weights. In the chip, a node's input values are converted into electrical voltages and then multiplied by the appropriate weights. Summing the products is simply a matter of combining the voltages. Only the combined voltages are converted back into a digital representation and stored for further processing.

The chip can thus calculate dot products for multiple nodes 16 at a time, in the prototype in a single step, instead of shuttling between a processor and memory for every computation.

All or Nothing

One of the keys to the system is that all the weights are either 1 or -1. That means that they can be implemented within the memory itself as simple switches that either close a circuit or leave it open. Recent theoretical work

suggests that neural nets trained with only two weights should lose little accuracy somewhere between 1 and 2 percent.

This is a promising real-world demonstration of SRAM based memory analog computing for deep learning applications," says Dario Gil, vice president of artificial intelligence at IBM. The results show impressive specifications for the energy efficient implementation of convolution operations with memory arrays. It certainly will open the possibility to employ more complex convolutional neural networks for image and video classifications in IoT [the Internet Of Things] in the future.

B.THARNIKA

I B.Sc. (Information Technology)

QUANTUM ENGINEERING

Since the 1940s, classical computers have improved at breakneck speed. Today you can buy a wristwatch with more computing power than the state-of-the-art, room-sized computer from half a century ago. These advances have typically come through electrical engineers' ability to fashion ever smaller transistors and circuits and to pack them ever closer together.

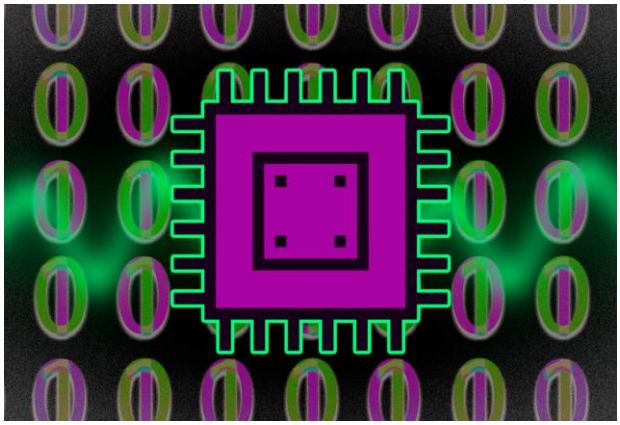

But that downsizing will eventually hit a physical limit as computer electronics approach the atomic level, it will become impossible to control individual components without impacting neighbouring ones. Classical computers cannot keep improving indefinitely using conventional scaling. Quantum computing, an idea spawned in the 1980s, could one day carry the baton into a new era of powerful high-speed computing. The method uses quantum mechanical phenomena to run complex calculations not feasible for classical computers. In theory, quantum computing could solve problems in minutes that would take classical computers millennia. Already, Google has demonstrated quantum computing's ability to outperform the world's best supercomputer for certain tasks.

But it's still early days quantum computing must clear a number of science and engineering hurdles before it can reliably solve practical problems. More than 100 researchers across MIT are helping develop the fundamental technologies necessary scale up

quantum computing and turn its potential into reality.

What is quantum computing?

It helps to first understand the basics of classical computers, like the one you're using to read this story. Classical computers store and process information in binary bits, each of which holds a value of 0 or 1. A typical laptop could contain billions of transistors that use different levels of electrical voltage to represent either of these two values. While the shape, size, and power of classical computers vary widely, they all operate on the same basic system of binary logic.

Quantum computers are fundamentally different. Their quantum bits, called qubits, can each hold a value of 0, 1, or a simultaneous combination of the two states. That's thanks to a quantum mechanical phenomenon called superposition. "A quantum particle can act as if it's in two places at once," explains John Chiaverini, a researcher at the MIT Lincoln Laboratory's Quantum Information and Integrated Nanosystems Group. Particles can also be entangled with each other, as their quantum states become inextricably linked. Superposition and entanglement allow quantum computers to "solve some kinds of problems exponentially faster than classical computers," Chiaverini says.

Chiaverini points to particular applications where quantum computers can shine. For example, they're great at factoring large numbers, a vital tool in cryptography and digital security. They could also simulate complex molecular systems, which could aid drug discovery. In principle, quantum computers could turbocharge many areas of research and industry if only we could build reliable ones.

How do you build a quantum computer?

Quantum systems are not easy to manage, thanks to two related challenges. The first is that a qubit's superposition state is highly sensitive. Minor environmental disturbances or material defects can cause qubits to err and lose their quantum information. This process is called decoherence which limits the useful lifetime of a qubit.

The second challenge lies in controlling the qubit to perform logical functions, often achieved through a finely tuned pulse of electromagnetic radiation. This manipulation process itself can generate enough incidental electromagnetic noise to cause decoherence. To scale up quantum computers, engineers will have to strike a balance between protecting qubits from potential disturbance and still allowing them to be manipulated for calculations. This balance could theoretically be attained by a range of physical systems,

though two technologies currently show the most promise: superconductors and trapped ions.

A superconducting quantum computer uses the flow of paired electrons called "Cooper pairs" through a resistance-free circuit as the qubit. "A superconductor is quite special, because below a certain temperature, its resistance goes away," says William Oliver, who is an associate professor in MIT's Department of Electrical Engineering and Computer Science, a Lincoln Laboratory Fellow, and the director of the MIT Center for Quantum Engineering.

The computers Oliver engineers use qubits composed of superconducting aluminum circuits chilled close to absolute zero. The system acts as an anharmonic oscillator with two energy states, corresponding to 0 and 1, as current flows through the circuit one way or the other. These superconducting qubits are relatively large, about one tenth of a millimeter along each edge that's hundreds of thousands of times larger than a classical transistor. A superconducting qubit's bulk makes it easy to manipulate for calculations.

But it also means Oliver is constantly fighting decoherence, seeking new ways to protect the qubits from environmental noise. His research mission is to iron out these technological kinks that could enable the fabrication of reliable superconducting quantum computers. "I like to do fundamental research, but I like to do it in a way that's practical and scalable," Oliver says. Quantum engineering bridges quantum science and conventional engineering. Both science and engineering will be required to make quantum computing a reality.

Another solution to the challenge of manipulating qubits while protecting them against decoherence is a trapped ion quantum computer, which uses individual atoms and their natural quantum mechanical behaviour as qubits.

R.JANANI

I B.Sc. (Computer Technology)

COMPUTER-GENERATED IMAGERY

Computer-generated imagery (CGI) is the application of the field of computer graphics (or more specifically, 3D computer graphics) to special effects. CGI is used in films, television programs and commercials, and in printed media.

Video games most often use real-time computer graphics (rarely referred to as CGI), but may also include pre-rendered "cut scenes" and intro movies that would be typical CGI applications.

CGI is used for visual effects because the quality is often higher and effects are more controllable than other more physically based processes, such as constructing miniatures for effects shots or hiring extras for crowd scenes, and because it allows the creation of images that would not be feasible using any other technology.

It can also allow a single artist to produce content without the use of actors, expensive set pieces or props. Recent accessibility of CGI software and increased computer speeds has allowed individual artists and small companies to produce professional grade films, games and fine art from their home computers.

M.BHAVAN

II B.Sc. (Computer Technology)

CLOCK PUZZLES

1. What is the time now if 2 hours later it would be half as long until midnight as it would be if it were an hour later?

Answer:9.00 PM

2.Which clock is the odd one out?

Answer: C

3. Which number follows on from these three?

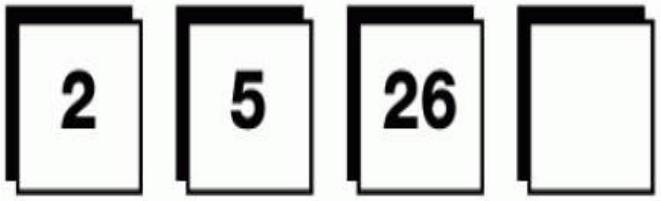

M.BHARANI

II B.Sc. (Computer Technology)

LOGICAL REASONING

The logical reasoning section is an important part of competitive exams. It contains different types of reasoning questions which are intended to judge analytical and logical reasoning skills of the candidate.

The logical reasoning questions can be verbal or non-verbal. In verbal logical reasoning questions, the concepts and problems

are expressed in words. The candidates are required to read and understand the given text or paragraph and according choose the right answer from the given options. In non-verbal logical reasoning questions, the concepts and problems are expressed in the form of figures, images or diagrams and the candidates are required to understand them before choosing the right answer out of the given options.

Logical Reasoning: Verbal Reasoning

Logical reasoning (verbal reasoning) refers to the ability of a candidate to understand and logically work through concepts and problems expressed in words. It checks the ability to extract and work with the meaning, information, and implications from the bulk of the text. The logics are expressed verbally, and you have to understand the logic before solving the questions.

Example

1. Which number should come next in the series, 48, 24, 12, ...?

A. 8
B. 6
C. 4
D. 2
Answer: B

2. RQP, ONM, _, IHG, FED, find the missing letters.

A.CDE
B.LKI
C.LKJ
D.BAC
Answer: C

3. PETAL: FLOWER

A. Pen: Paper
B. Engine: Car
C. Cat: Dog
D. Ball: Game
Answer: B

4. Pointing to a photograph, a man said, "I have no brother, and that man's father is my father's son." Whose photograph was it?

A. His son
B. His own
C. His father
D. His nephew
Answer: C

5. If in a certain language, NOIDA is coded as OPJEB, how is DELHI coded in that language?

A. CDKGH
B. EFMIJ
C. FGNJK
D. IHLED
Answer: **B**

Logical Reasoning: Nov-Verbal Reasoning

Logical reasoning (non-verbal reasoning) refers to the ability of a candidate to understand and logically work through concepts and problems expressed in the form of images, diagrams, etc. It checks the ability to extract and work with the meaning, information, and implications from the given images or diagrams. Here, the logics are expressed non-verbally, and you have to understand the logic before solving them.

Example

1. Select three figures out of the following five figures which when fitted into each other would form a square.

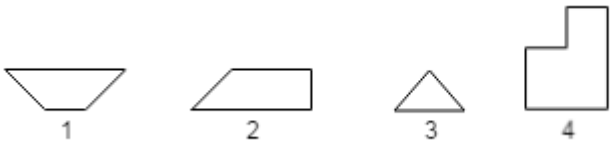

A.1, 2, 4
B.1, 2, 5
C.5, 3, 1
D.2, 3, 4
Answer: B

2. Which of the following figures is different from others?

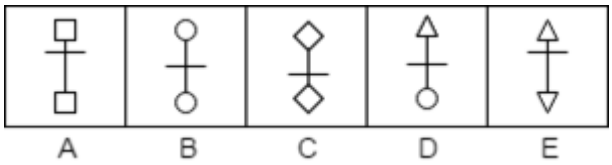

A. figure A
B. figure B
C. figure C
D. figure D
Answer: D

3. Which of the following series of figures follows the given rule?

Rule: The closed figure loses its sides, and the open figure gains its sides as the series proceeds.

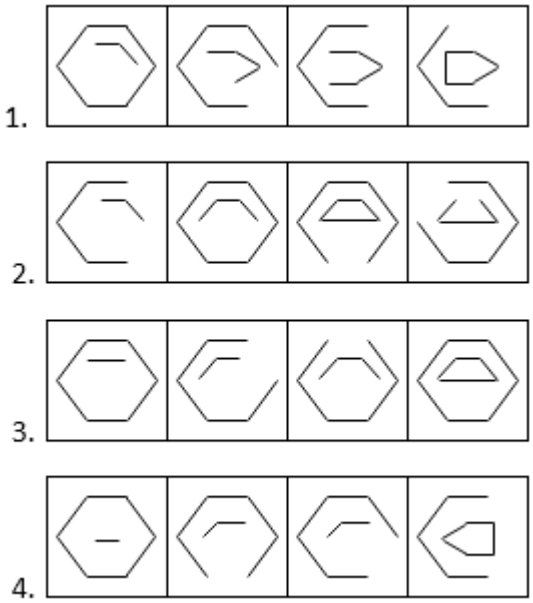

A.1
B.2
C.3
D.4
Answer: A

4. MORTAL

A. Divine
B. Immortal
C. Spiritual
D. Eternal
Answer: B

5. He is too important for tolerating any delay

A. to tolerate
B. to tolerating
C. at tolerating
D. with tolerating

Answer:A

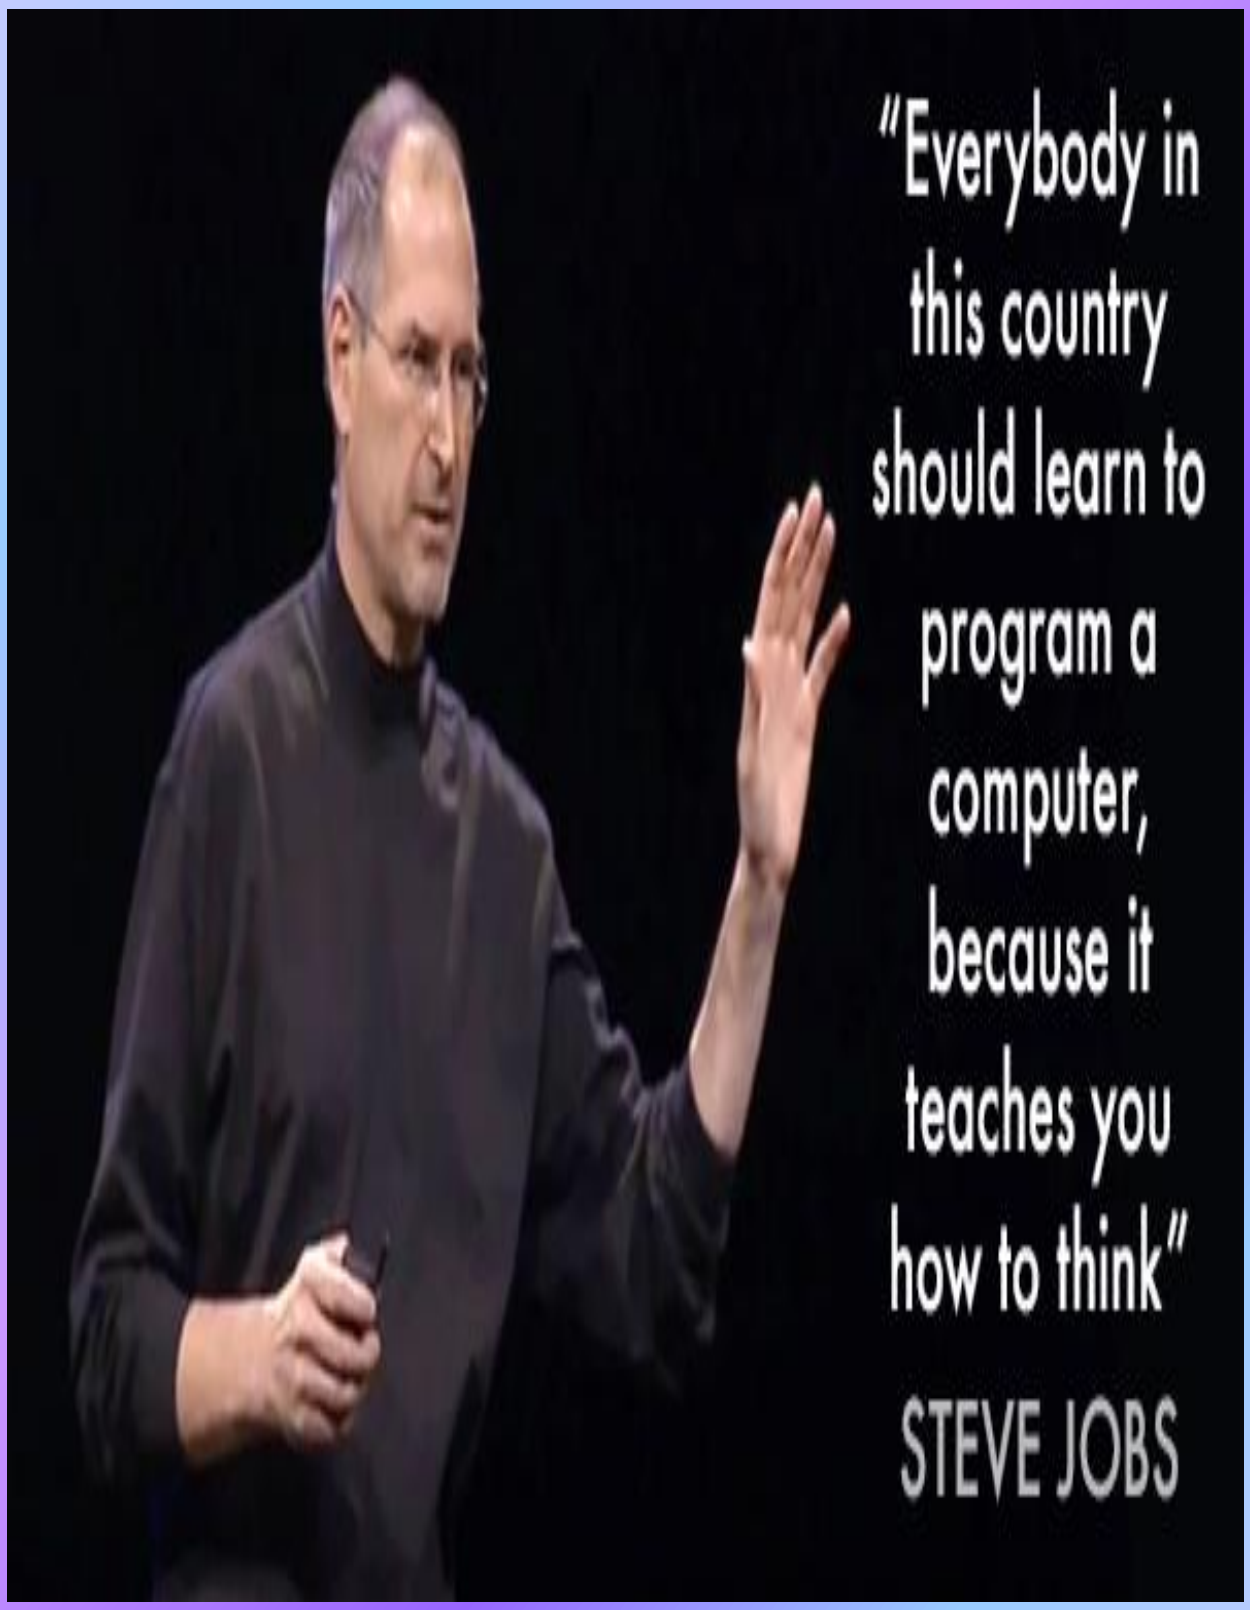
"Everybody in this country should learn to program a computer, because it teaches you how to think"
STEVE JOBS

Steve Jobs