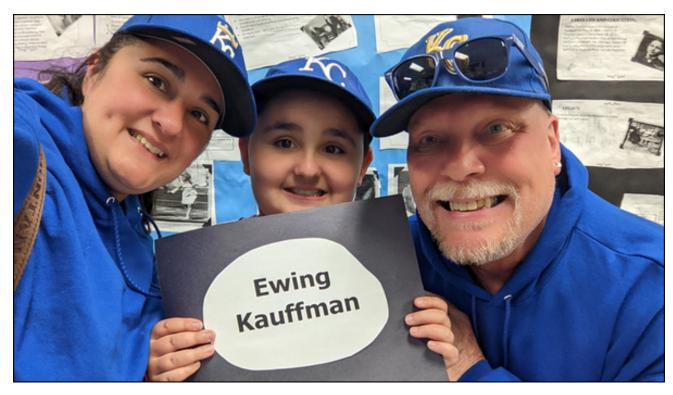# In This Issue...

From The Chief Editor's Desk

My wife and I work very hard at being GOOD parents. We strive to make sure our kids know right from wrong. We strive to make sure they behave, and have a good "moral compass." We are involved in their school life, much more so than many of the parents of their classmates. We hear it all of the time from their teachers at school.

After all, our children are the longest lasting part of "our legacy," as children are for the vast majority of people across the planet. By teaching them and guiding them, they will (hopefully) grow into responsible, productive, well-balanced adults that contribute to that proverbial "greater good."

It's no secret that there are DEEP divisions across many sectors of society. While this column is not political in nature (and if you're thinking that I'm taking a political stance, then you are the one reading more into what I've written than is here), the biggest divide is probably across political boundaries and beliefs.

I've discussed my son, Ryan, before in this magazine. In my article in November 2020, I laid out the challenges he faces. He continues to improve in all aspects of his "challenges." Just get him talking about Transformers, Minecraft, or Sonic the Hedgehog, and he'll talk your ears off. You'll probably learn WAY more than you ever thought possible and more than you probably wanted to know about any of those

topics. His reading is coming along fantastically, although we are now focused on improving reading comprehension. Math remains his favorite school subject. His new favorite reading material involves the graphic novel adventures of Dog Man and Cat Kid.

It's easy to fool yourself into thinking that he isn't paying attention to other "things" going on in his life. He typically doesn't acknowledge them. But every once in a while, he surprises you. Just the other night, he told us, "Idiots are ruining our world. I'm going to stop them."

In just that brief, passing statement, he let us know that he really is paying attention to those other "things" going on in his life. He senses and feels the tension in his world brought on by the massive division in society. At three months shy of his tenth birthday, he has shown a depth of thought, insight and understanding that many adults lack. He may not understand all the intricacies of the things that separate us, but he knows it isn't right.

The innocence of a child.

His mom and I really don't talk politics much at home. And, on those relatively rare times when we do, it usually isn't in front of the kids. Nope. These are things that he has noticed all on his own.

Proud? You bet! His mom and I couldn't possibly be prouder. And, it's validation that our efforts to teach and guide our children about right and wrong are working.

****

This month's cover photo comes from my wife, Laura. One of her patients at the hospital where she works gave her a flower. She thought it was pretty, so she took a picture of it. That picture (which you see on the cover) came out exceptionally good. Being Spring (in the Northern Hemisphere, at least), it made what I felt was a perfect image for this issue of the magazine.

****

Until next month, I bid you peace, happiness, serenity, prosperity, and continued good health!

The PCLinuxOS Magazine

Created with Scribus

Screenshot Showcase

Posted by tuxlink, on April 12, 2023, running KDE.

Altered Reality: AI Generated Images

by Paul Arnote (parnote)

"A picture is worth a thousand words."

"Pictures, or it didn't happen."

"A photo doesn't lie."

We've all most likely heard at least one of the above common sayings in describing images or pictures, if not all of them. And, until just very recently, there was a high level of "truth" in those sayings. But now, that line between what's real and what's not real has been blurred beyond distinction, blurring our perception of what is real and what is "Memorex."

The entire computer world, as well as mainstream media outlets, has been abuzz within the past year with the emergence of GPT-4 AI content creation. You have probably already heard of ChatGPT or DALL-E. Reports about its abilities have permeated recent news cycles.

The tech industry has certainly heard about it, as virtually everyone and their brother is jumping on the AI airship. Google has announced their upcoming use of AI to help power their search engine. Microsoft has announced what appears to be a full embrace of new AI capabilities. Amazon is preparing to launch their answer to the rush to AI power. Adobe has announced Firefly, their answer to the image creator DALL-E. Every week, it seems that one or more tech giants are embracing AI in one form or another. In fact, even the April 2023 cover image was generated by me (not that I count myself among "tech giants"), using simple text to describe the image I was wanting to create on the DALL-E AI image creator website (shown below).

Creating the image was as simple as describing what I wanted in the image, using plain text descriptors. For the above image, I entered "Cartoonish Linux mascot Tux dressed as the Easter Bunny, in a large green meadow with mounds of Easter candy and Easter eggs strewn all about."

That generated four images for me to choose from, shown below. And boy, it created those images in like exceptionally little time. As in, just a matter of seconds.

As you can see, only ONE of the four images is actually Linux mascot Tux dressed as the Easter Bunny (second from the left). The other three depict an actual bunny as the Easter Bunny.

For fun coupled with a bit of "what if," I repeated my search by slightly altering my image creation criteria. For this second round, my creation text read "Cartoonish Linux mascot Tux dressed as the Easter Bunny, in a large green meadow with mounds of Easter candy and Easter eggs strewn all about holding an Easter Basket."

That generated another four images for me to choose from, shown below. Again, the images were generated and presented to me in a matter of seconds.

This time, all of the images were of a bunny posing as the Easter Bunny. It seems that the "Linux mascot Tux dressed as the Easter Bunny" part of my image creation text was completely ignored. The generated images were nice images, but the lack of Linux mascot Tux rendered the generated images unusable for my intended purpose.

By this point, I was "stuck" on penguins in my image creation criteria. So, I repeated my creation request with different scenarios. In my first two image creation terms, I was deliberate about specifying a "cartoonish" image. So, how would DALL-E do producing photorealistic images? I was about to find out, in pretty short order.

So, for my next "experiment" with DALL-E, I entered the following text as my image creation criteria: "Photo with penguin in the foreground, walking through Times Square." Below are the four images that DALL-E presented, based on that specific image creation criteria.

Has there ever been a penguin walking in Times Square? I don't know. Maybe. But I wanted to "test" to see how realistic of an image DALL-E would present. To be perfectly honest, I was blown away with how genuine and realistic the images appeared. Again, the images were created and presented to me in a matter of seconds.

But my curiosity wasn't satiated yet. I decided to do another test, with "Photo with penguin in the foreground, walking on the Titanic" as my image creation criteria. Below are the results.
Once again, these were nice images generated by DALL-E. But I realized that the Titanic is depicted in the background of most of the images, so the penguin couldn't possibly be walking ON the Titanic.

So, I repeated my image creation text, with a slight modification. I used ""Photo with penguin in the foreground, walking on the deck of the Titanic" instead.

Now we're talking! The image results are very good. They look realistic and genuine, with appropriately positioned shadows under the penguin even. And for the most part, it gave me exactly what I asked for.

By this time, I was informed by the OpenAI website that I had just 10 "credits" remaining, so I stopped creating images. I wanted to "save" my credits for that "just in case" moment that I'd need to "create" other images. Apparently, users are "given" 15 credits per month, and each creation of the four-image sets costs one credit. At least, that's how I understand how the "credits" work. Users can also purchase additional credits. You can get 115 credits for $15 (U.S.). Monthly credits (for users who signed up prior to April 6, 2023) expire one month after they are granted, and do not carry over from month-to-month. Paid credits expire one year after they are purchased. You can read all about how DALL-E credits are used here. And yes, you do have to sign up for an account. I have yet to receive an email from OpenAI (who runs the DALL-E website) beyond the initial account setup.

More Than Just DALL-E

While DALL-E is just the first to go public with its AI image creation, it's not necessarily the only game in town … or at least, soon won't be. Microsoft's Bing search engine also allows users to create AI images, but it currently uses the OpenAI.org DALL-E to create them. So, everything you've learned about DALL-E applies to those created via Bing. You have to "sign up" on the Bing image creation site, which (according to a disclaimer on the site) also signs you up for emails from Microsoft Rewards. Uhm … no thank you! Unless I tune up my spam filter first. At least they tell you up front that signing up for the image creator also signs you up for a flood of Microsoft spam.

Adobe's Firefly is currently in beta. Firefly is touted as being easier to use, although it's difficult to see how much easier AI image creation can be

than it is with DALL-E. Interested users can request access here, by filling out the form at the link. Access to the beta will be expanded over time, so signing up will most likely mean that you'll have to "wait in line" for access. I've not been able to uncover any costs associated with the final product. But, knowing Adobe from their track record, access won't come for cheap.

Google, believe it or not, is playing catch-up, after falling asleep at the proverbial AI wheel. They have announced that their AI solution is "coming soon," but no date has been revealed at the time of the writing of this article.

Amazon has also announced that they have an AI competitor "coming soon." But, like with Google, no date has been revealed at the time of the writing of this article.

You can rest assured that you're likely to see other AI products come along as it gains steam.

The Danger Zone: Altered Reality

The images created by DALL-E are good. Damn good. In fact, maybe they're TOO good. And therein lies the real problem.

GIMP does display EXIF metadata for the image downloaded from DALL-E. To be honest, I've always associated EXIF data with JPG files, so I was quite surprised to find it embedded in the PNG file I downloaded from DALL-E.

But how many people are going to even think of looking at or for the metadata included in an image? Plus, all actions I performed on the downloaded image to either resize it or to convert it between formats resulted in the EXIF data being deleted or overwritten. In my experiences, the EXIF data is ONLY available in the image that was originally downloaded from DALL-E.

So, with the reality being that most people will never look at (or think of looking at) the metadata embedded in an image, the risk for abuse is literally off the charts. With the photorealism of the images created by AI, it's difficult to ascertain whether an image is real or an AI creation. Perform a couple of actions on the image (other than renaming it), and all you have is someone's word about whether the image is genuine or not.

There is a Google Chrome plugin, Fake Profile Detector, that can be used to spot "fake AI generated images." It comes from V7 Labs, and is typically used to help spot fake profiles and profile images. It's not known if it is capable of spotting AI generated images that are NOT profile images of people. It's a start, but I do wonder by what methodology it detects fake images. If it's just the EXIF data we found earlier, that's easy enough to bypass, as we've shown.

You can bet your life savings that users with malicious intent have already taken this information and put it in their hat of tricks. It's only a matter of time before bad actors leverage AI created images for their own private gain, to advance their agenda, or to inflict harm on those who think differently.

Today's political environment has separated us more than ever before. Disinformation is pervasive on both sides of the political divide. Without some safeguards (provided via legislative measures and AI providers themselves), the proliferation of AI generated content will make it even more difficult to weed out the disinformation from the genuine information.

As with most other areas, I doubt that the AI content providers will do an adequate job of self-regulation and self-policing. In just about every instance of self-regulating and self-policing that has been tried, it has failed miserably. We should learn from our past mistakes, but never seem to. And to be perfectly honest, I doubt the success of legislative measures because legislators in most jurisdictions (the ones who write the laws) don't understand the technology that they are attempting to regulate, in most cases. Do you have doubts? Just look at the current harvesting of your private and personal information from the web that continues unchecked by many government and corporate entities, mostly due to legislative inaction. That legislative inaction is fueled by two main factors: ignorance of the technology they are trying to regulate, and graft/payoffs of legislators from corporate lobbyists.

What a Catch-22!

Then there's the whole concept of AI generated content. When you stop and think about it, AI generated content is scary and unsettling. It's the new great unknown. I've heard it said something to the effect that whoever controls the content, controls the world and controls the minds of the people. Will life imitate art (in particular, the movies)? Like at the end of the Jodie Foster movie of Carl Sagan's Contact, will we be the long lost creators of technology that continues to function long after we are dead and gone, and the civilizations that remain continue to use the technology

because it just works? Or will the AI become self-aware, and morph into something resembling Skynet from the Terminator movie series, where there are actual battles between humans and machines? Like I said, it's scary stuff.

Remember those three common sayings I began this article with? The AI image creators throw all of those out the proverbial window – baby, bathwater and all.

Screenshot Showcase

Posted by Archie, on April 15, 2023, running KDE.

How We Think About Copyright & AI Art

by **Kit Walsh**
Electronic Frontier Foundation
Reprinted under Creative Commons License

Artists are understandably concerned about the possibility that automatic image generators like Stable Diffusion will undercut the market for their work. We live in a society that does not support people who are automated out of a job, and being a visual artist is an already precarious career.

In this context, it's natural to look to copyright law, because copyright is supposed to help ensure that artists get paid for their work. Unfortunately, one copyright theory advanced in a class-action lawsuit by some artists against Stable Diffusion is extremely dangerous for human creators. Other theories—both in that lawsuit and another suit by Getty Images—propose to alter and expand copyright restrictions in ways that would interfere with research, search engines, and the ability to make new technology interoperate with old.

This legal analysis is a companion piece to our post describing AI image-generating technology and how we see its potential risks and benefits. We suggest that you read that post first for context, then come back to this one for our view on how the copyright questions play out under U.S. law.

Copyright law is supposed to embody a balance between giving artists a sufficient incentive to create, by granting them control of some of the ways their art can be used, and giving the public the right to build on and/or use that art in new and interesting ways. Here, the question is whether those who own the copyright in the images used to train the AI generator model have a right to prohibit this kind of use.

To answer that question, let's start with a few basic principles.

First, copyright law doesn't prevent you from making factual observations about a work or copying the facts embodied in a work (this is called the "idea/expression distinction"). Rather, copyright forbids you from copying the work's creative expression in a way that could substitute for the original, and from making "derivative works" when those works copy too much creative expression from the original.

Second, even if a person makes a copy or a derivative work, the use is not infringing if it is a "fair use." Whether a use is fair depends on a number of factors, including the purpose of the use, the nature of the original work, how much is used, and potential harm to the market for the original work.

Copyright and Training Sets

Here's how fair use would apply to AI art generation:

Step 1: Scraping the Images from the Web

Like copying to create search engines or other analytical uses, downloading images to analyze and index them in service of creating new, noninfringing images is very likely to be fair use. When an act potentially implicates copyright but is a necessary step in enabling noninfringing uses, it frequently qualifies as a fair use itself. After all, the right to make a noninfringing use of a work is only meaningful if you are also permitted to perform the steps that lead up to that use. Thus, as both an intermediate use and an analytical use, scraping is not likely to violate copyright law.

Step 2: Storing Information About the Images

In this step, the system analyzes the images and stores information about how the pixel arrangements correlate with words in the text annotations.

The Stable Diffusion model makes four gigabytes of observations regarding more than five billion images. That means that its model contains less than one byte of information per image analyzed (a byte is just eight bits—a zero or a one).

The complaint against Stable Diffusion characterizes this as "compressing" (and thus storing) the training images, but that's just wrong. With few exceptions, there is no way to recreate the images used in the model based on the facts about them that are stored. Even the tiniest image file contains many thousands of bytes; most will include millions. Mathematically speaking, Stable Diffusion cannot be storing copies of all of its training images (for now, let's put a pin in the question of whether it stores a copy of any of them).

So the model isn't storing copies. But is it generating and storing infringing derivative works of all of the images in the training data?

Probably not, for at least three reasons:

First, a derivative work still has to be "substantially similar" to the original in order to be infringing. If the original is transformed or abridged or adapted to such an extent that this is no longer true, then it's not a derivative work. A 10-line summary of a 15,000-line epic isn't a derivative work, and neither are most summaries of books that people make in order to describe those copyrighted works to others.

Second, copyright doesn't grant a monopoly on a genre or subject's tropes and motifs, including expressive elements like wavy lines to denote shaking, giving animals more human facial expressions, and similar common—even if creative—choices. What's more, copyright does not apply at all to non-creative choices—like representing a cat as having four legs and a tail. Much of the information stored by and produced by an AI art generator falls into these categories.

Third, the amount of copyrightable expression taken from each original image in the training set could be considered "de minimis," a legal term that means "too minimal to qualify as infringing."

Even if a court concludes that a model is a derivative work under copyright law, creating the model is likely a lawful fair use. Fair use protects reverse engineering, indexing for search engines, and other forms of analysis that create new knowledge about works or bodies of works. Here, the fact that the model is used to create new works weighs in favor of fair use as does the fact that the model consists of original analysis of the training images in comparison with one another.

The class-action lawsuit against Stable Diffusion doesn't focus on the "outputs," (the actual images that the model produces in response to text input). Instead, the artists allege that the system itself is a derivative work. But, as discussed, it's no more illegal for the model to learn a style from existing work than for human artists to do the same in a class, or

individually, to make some of the same creative choices as artists they admire.

Moreover, AI systems learn to imitate a style not just from a single artist's work, but from other human creations that are tagged as being "in the style of" another artist. Much of the information contributing to the AI's imitation of style originates with images by other artists who are enjoying the freedom copyright law affords them to imitate a style without being considered a derivative work.

Step 3: Creating Output Images

Unlike the Stable Diffusion case, Getty Images's suit focuses on outputs, claiming that outputs are sometimes "substantially similar" to training data. Getty doesn't provide an example of this, apart from the presence of its watermark in some Stable Diffusion outputs.

It's not surprising that the complaints don't include examples of substantially similar images. Research regarding privacy concerns suggests it is unlikely it is that a diffusion-based model will produce outputs that closely resemble one of the inputs.

According to this research, there is a small chance that a diffusion model will store information that makes it possible to recreate something close to an image in its training data, provided that the image in question is duplicated many times during training. But the chances of an image in the training data set being duplicated in output, even from a prompt

specifically designed to do just that, is literally less than one in a million.

This means that there exists, at most, a handful of rightsholders out there that might have a copyright claim. Thus far, neither lawsuit suggests the plaintiffs are in that category.

Of course, the statistical standards used in this research aren't the same as the legal standards used in copyright law, but we can nonetheless take them as informative, and they are consistent with the tiny amount of data per image the diffusion model stores.

To sum up: a diffusion model can, in rare circumstances, generate images that resemble elements of the training data. De-duplication can substantially reduce the risk of this occurring. But the strongest copyright suit against a diffusion-based AI art generator would likely be one brought by the holder of the copyright in an image that subsequently was actually reproduced this way.

As with most creative tools, it is possible that a user could be the one who causes the system to output a new infringing work by giving it a series of prompts that steer it towards reproducing another work. In this instance, the user, not the tool's maker or provider, would be liable for infringement.
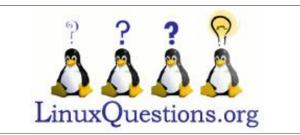

What Would it Mean for Art if the Court Finds that Stable Diffusion Infringes Copyright?

The theory of the class-action suit is extremely dangerous for artists. If the plaintiffs convince the court that you've created a derivative work if you incorporate any aspect of someone else's art in your own work, even if the end result isn't substantially similar, then something as common as copying the way your favorite artist draws eyes could put you in legal jeopardy.

Currently, copyright law protects artists who are influenced by colleagues and mentors and the media they admire by permitting them to mimic elements of others' work as long as their art isn't "substantially similar" and/or is a fair use. Thus, the same legal doctrines that give artists the breathing room to find inspiration in others' works also protect diffusion models. Rewriting those doctrines could cause harm far beyond any damage Stable Diffusion is causing.

In our companion blog post, we explore some of the other consequences. In particular, we discuss who would likely benefit from such a regime (spoiler: it's not individual creators). We also discuss some alternative approaches that might actually help creators.

Done right, copyright law is supposed to encourage new creativity. Stretching it to outlaw tools like AI image generators—or to effectively put them in the exclusive hands of powerful economic actors who already use that economic muscle to squeeze creators—would have the opposite effect.

AI Art Generators & The Online Image Market

by **Katharine Trendacosta** & **Cory Doctorow**
Electronic Frontier Foundation
Reprinted under Creative Commons License

Now that computer-generated imaging is accessible to anyone with a weird idea and an internet connection, the creation of "AI art" is raising questions—and lawsuits. The key questions seem to be 1) how does it actually work, 2) what work can it replace, and 3) how can the labor of artists be respected through this change?

The lawsuits over AI turn, in large part, on copyright. These copyright issues are so complex that we've devoted a whole, separate post to them. Here, we focus on thornier non-legal issues.

How Do AI Art Generators Work?

There are two different parts of the life of an AI art generator. First are the data that teaches it what a "dog" is or, more abstractly, what "anger" looks like. Second are the outputs that the machine gives in response to prompts. Early, when the generator has not had enough training, those outputs only loosely reflect the prompts. But eventually, the generator will have seen enough images to figure out how to properly respond to a prompt (this is just how people do it, too). AI-generated creative content can run the gamut from "prompt based on an image I saw in a fever dream" to "very poorly written blog post."

How Does an AI Art Generator "Learn"?

AI art generators depend on "machine learning." In a machine learning process, a training algorithm takes in an enormous set of data and analyzes the relationships between its different aspects. An AI art generator is trained on images and on the text that describes those images.

Once it has analyzed the relationships between the words and features of the image data, the generator can use this set of associations to produce new images. This is how it is able to take text input—a "prompt"—like "dog" and generate (that is, "output") arrangements of pixels that it associates with the word, based on its training data.

The nature of these "outputs" depends on the system's training data, its training model, and the choices its human creators make.

For instance: a model trained by feeding it images labeled with text that appeared close to those images on public web-pages will not be as good at matching "prompts" as it would be if it had been trained with images that had been manually annotated with explicit, human-generated labels.

This process is not too different from how babies learn things. For example, a lot of kids basically think all animals are "doggies" until they have enough exposure and correction by adults to distinguish "doggie" from "horsie." Machine learning can make similar mistakes, finding connections that, to humans, are obscure. For example, a cancer classifier can "learn" that an image shows a tumor if that image contains a ruler. The AI learned a shortcut: images of structures that a radiologist has identified as cancerous tumors have pictures with rulers for scale and to track size. The training images of benign growths were from a different set, and they didn't have rulers.

Beyond the effect of training data quality, there is also the effect of different training "models." These models have names like "diffusion" or "generative adversarial networks" (GANs). Each of these models has different strengths and

weaknesses (as of this writing, diffusion models are generally considered the state of the art).

During training, programmers introduce variables that determine the similarity of the model's output to the images in its training data. Other variables determine whether the system prioritizes creating close matches for the prompt, or being more experimental by outputting images for which the model has less "confidence" (a mathematical term describing a kind of statistical certainty) as a match for the users' prompts. Some models allow users to adjust such variables when they issue prompts to the model.

Where Does the Training Data Come From?

In general, the training data comes from scraping the web: finding available images that have text associated with them (in some cases, annotations are added afterwards). This means the creators of the images or people depicted in them likely do not know or specifically consent to being included in the analysis. For the "Stable Diffusion" system that is the subject of two recent lawsuits—a class action complaint on behalf of several visual artists and another filed by Getty Images—the dataset is five billion images indexed by a nonprofit called LAION.

For an analysis of the copyright concerns related to those training sets, see our other blog.

Work Replacement and AI

Many artists are concerned that the availability of AI art will mean less of a market for their work. That's a valid concern: there are some services provided by artists that could likely be replaced by AI generators. This happened previously, with transcription: machine learning systems replaced some human transcription. However, these automated systems produce output that is of generally low quality, as anyone who has seen auto-generated closed captions can attest.

In fact, the issues that come with automating labor go back centuries: automated replacements that can be owned outright by employers or are simply cheaper than paying a worker can result in fewer people with jobs. In a perfect world, automation would be used to free people to pursue matters they care about, but that's not the world we live in (yet), so it's natural and valid for workers to worry about automation driving down wages or pushing them out of their industry altogether.

The debate over AI art isn't limited to general concerns about automation and the lack of support for people automated out of a job; it's also about whether that AI art generation is especially unfair because much of its training data consists of copyrighted images used without permission. We discuss this in the other post.

Beyond labor market and fairness concerns, there's a real risk that AI art will give a few corporations even more control over future creativity. Most access to art is already controlled by a few major gatekeepers, which have no interest in the livelihood of artists and no appetite for risk.

For example, Getty Images, a plaintiff in one of the lawsuits against AI art generation, has cornered the market on stock and event images. Most news organizations use Getty because it's a near-certainty that Getty will have an image of the subject of a given article standing on a red carpet. Without Getty, media companies would have had to either send a photographer to every event, or figure out which freelance photographers were present at it and try to license their images. As a monopoly, Getty both undercuts independent photographers and gouges news organizations.

In its lawsuit, Getty cites an AI-generated image that produced a distorted version of its watermark. Getty claims that this is proof that its copyrighted materials are found in the output of an art generator, but what's really happening is that the image generator has "learned" that any image of a red carpet contains a Getty watermark, so it draws the watermark into images that seem "Getty-like." In other words, Getty has such a lock on a certain kind of news photography that a statistical analysis of all newsworthy photos of celebrities will conclude that Getty is inseparable from that kind of photography. A Getty watermark is to a celebrity image as a ruler is to a tumor.

Letting Corporations Control AI Will Flatten Our Creative World

At the moment, there are freely available, open-source models for AI art generators, and anyone can tweak them in innovative ways and innovate with them. But if the legal environment or the technology changed so that only a few large companies could make or use AI art models, it would make our creative world even more homogenous and sanitized.

For example, large commercial deployments of diffusion models already refuse queries that might lead to nude images, which of course are not intrinsically harmful, illegal, or immoral, and have a long history in artistic expression. Heavy-handed restrictions on "adult" subject matter are especially hard on people whose identities are wrongly labeled obscene, intrinsically sexual, or "adult only" (including queer people), erasing them from the world generated by these tools.

AI art generators' bias needn't be the result of explicit, active censorship; it can also come from bias in their training data. For example, an AI art tool may generate images of white people as the default, reinforcing racial inequality, or tend towards lighter skin in response to requests for "beautiful" people. Images of women are more likely to be coded as sexual in nature than images of men in similar states of dress and activity, because of widespread cultural objectification of women in both images and its accompanying text. An AI art generator can "learn" to embody injustice and the biases of the

era and culture of the training data on which it is trained. AI art generators sometimes produce surprising novelty, but they predominantly favor past values and aesthetics. Models tend to recreate what they see over and over, making their output tend toward the average and typical, at the expense of minority aesthetics and identity.

Another thing to watch for: AI art generators may depend on and reveal private information. Imagine asking an AI art generator to generate images related to a medical condition and seeing a recognizable person in the output (this could happen if the model wasn't trained on many images related to that condition).

Finally, as has been the case with "deep fakes," it is possible to use machine learning to generate deceptive images depicting real people doing things they never did. Those images can shame or defame the person or otherwise harm their social and economic lives.

However, such images can also be used for important social commentary, or simply as art, when they are not passed off as true occurrences. We understand when we see an image of a politician lighting the Constitution on fire that they did not literally burn the document—rather, that the creator of that image is commenting on the politician's policies. This is a situation where each should be evaluated on its own merits, rather than banning technology that has both positive and negative uses. As with photomanipulation, it's important that we learn how to determine what is real. For

example, norms around parody photomanipulation exaggerate the edited feel both as part of the parody and to make the joke clear.

What The World Looks Like If Ai Creators Need Permission From Rightsholders

See our other blog for our thoughts on copyright and why we don't think AI art generators are likely to infringe. For purposes of this discussion, however, imagine that you can't train an AI model on copyrighted information without permission.

Requiring a person using an AI generator to get a license from everyone who has rights in an image in the training data set is unlikely to eliminate this kind of technology. Rather, it will have the perverse effect of limiting this technology development to the very largest companies, who can assemble a data set by compelling their workers to assign the "training right" as a condition of employment or content creation.

This would be a pyrrhic victory for opponents of the very idea of AI art: in the short term, AI tools would cease to exist or would produce lower-quality outputs, reducing the potential to drive down creators' wages.

But in the medium to long term, this effect is likely to be the reverse. Creative labor markets are intensely concentrated: a small number of companies—including Getty—commission

millions of works every year from working creators. These companies already enjoy tremendous bargaining power, which means they can subject artists to standard, non-negotiable terms that give the firms too much control, for too little compensation.

If the right to train a model is contingent on a copyright holder's permission, then these very large firms could simply amend their boilerplate contracts to require creators to sign away their model-training rights as a condition of doing business. That's what game companies that employ legions of voice-actors are doing, requiring voice actors to begin every session by recording themselves waiving any right to control whether a model can be trained from their voices.

If large firms like Getty win the right to control model training, they could simply acquire the training rights to any creative worker hoping to do business with them. And since Getty's largest single expense is the fees it pays to creative workers—fees that it wouldn't owe in the event that it could use a model to substitute for its workers' images—it has a powerful incentive to produce a high-quality model to replace those workers.

This would result in the worst of all worlds: the companies that today have cornered the market for creative labor could use AI models to replace their workers, while the individuals who rarely—or never—have cause to commission a creative work would be barred from using AI tools to express themselves.

This would let the handful of firms that pay creative workers for illustration—like the duopoly that controls nearly all comic book creation, or the monopoly that controls the majority of role-playing games—require illustrators to sign away their model-training rights, and replace their paid illustrators with models. Giant corporations wouldn't have to pay creators—and the GM at your weekly gaming session couldn't use an AI model to make a visual aid for a key encounter, nor could a kid make their own comic book using text prompts.

Approaches To AI That Respect Artists

The Writer's Guild of America-West is in the middle of renegotiating its minimum basic agreement. That agreement creates the floor for how to credit and pay writers across a number of creative industries, including film and television. The Guild's AI proposal has some technical problems reflecting an incomplete understanding of how the technology works, but from a labor perspective, it demonstrates an excellent proposal for AI-generated output that, while not perfectly understanding how the technology works, does grasp the central concern at issue very well.

The Guild's core proposal is this: AI-generated material can't replace a human writer. AI-generated material cannot qualify as source material for adaptation in any way. AI-generated work can be used as research material, just as a Wikipedia article could, but because of the unclear nature of the sources that go into its output and how the output is generated, it has no place as an "author" in the world of copyright. AI outputs are not, in the Guild's opinion, copyrightable.

That means that if a studio wants to use an AI-generated script, there can be no credited author and no copyright. In a world where studios jealously guard the rights to their work, that's a major poison pill. Under this proposal, studios must choose between the upfront cost of paying a writer what they are worth, and the backend cost of not having control over the copyright to the product.

This is a smart strategy that zeroes in on the Guild's domain: protecting its members. That said, the Guild's conception of the technology is a little off: the Guild claims that AI creates a mosaic from its training data. This is less true than the Guild claims, and the output from AI doesn't infringe as often as they imply. But despite these technical misapprehensions, the Guild's way of thinking about it as a tool is very smart (again, here is our analysis of the copyright status of AI art).

For the Guild, AI-generated writing has no place in Guild-covered works. If a studio makes something that is covered by this agreement, you have to hire a human writer and pay the Guild's negotiated rate (or more). AI material cannot be used to undercut that labor. AI is a tool to help writers, not a replacement for writers.

That is the way all technology should be seen in relation to artistic work: as an artistic tool, not as a replacement for artists. A broad ban on AI will not fix the inequities of a highly concentrated market - but it could cost us the exciting uses of this technology for creative expression.

Exciting Things About AI Art Generation

Any development that gives more people the ability to express themselves in a new way is an exciting one. For every image that displaces a potential low-dollar commission for a working artist, there are countless more that don't displace anyone's living—images created by people expressing themselves or adding art to projects that would simply not have been illustrated. Remember: the major impact of automated translation technology wasn't displacing translators—it was the creation of free and simple ways to read tweets and webpages in other languages when a person would otherwise just not know what was being said.

When people use AI tools, it leads to a different kind of "creativity" than human artists produce on their own, as the tool finds associations and imagery that unassisted people hadn't previously made. AI art generators can also help working artists in several ways, for example, by producing a rough first pass or automating time-consuming tasks like shading a flat image. This would be the art equivalent of the research material argument made by the WGA.

There is a lot to like about art generators. The problem going forward is keeping the good things—open-source technology that researchers can audit, cutting down on the tedious parts of making things—without letting the concerns give power to the same companies that disempower artists every day.

The PCLinuxOS Magazine

Created with Scribus

Screenshot Showcase

Posted by astronaut, on April 3, 2023, running Openbox.

PCLinuxOS Recipe Corner

from the kitchen of youcantoo

Hashbrown Breakfast Casserole

Serves: 6-8

INGREDIENTS:

20 ounces shredded hash browns thawed
1 pound sausage cooked, crumbled, and drained
¼ cup onion finely diced
½ red bell pepper diced
½ green bell pepper diced
2 cups cheddar cheese shredded, divided use
8 eggs
1 can evaporated milk 12 ounces,
    or 1 ⅓ cups milk or light cream
½ teaspoon Italian seasoning or your favorite
    herbs/spices (optional)
½ teaspoon Kosher salt or to taste
¼ teaspoon black pepper

**Linux Docs
Linux Man Pages**

die.net

DIRECTIONS:

Preheat the oven to 350F (if baking immediately). Grease a 9×13 pan or a 3qt baking dish or spray with cooking spray.

Brown sausage in a large skillet until no pink remains. Drain fat.

Place the hash browns, cooked sausage, onions, bell peppers, and 1 ½ cups of cheese in the prepared dish. Gently mix the ingredients together and spread evenly in the pan.

Add eggs, evaporated milk, salt & pepper, and Italian seasoning to a large bowl. Whisk until well combined.

Pour the egg mixture over the hash brown mixture and sprinkle the remaining cheese on top. Cover with foil and refrigerate overnight if desired (or bake immediately).

Bake uncovered for 55-65 minutes or until cooked through.

TIPS:

Swap out the Sausage with Turkey, Ham, Bacon or last night's Taco meat. Better yet, add some of each to the breakfast casserole. If you don't like bell peppers, substitute with sauteed mushrooms. Substitute the Sharp Cheddar

cheese with pepper jack cheese or any other bold cheese flavor works well.

*To Bake Late*r: Cover the mixture tightly and refrigerate for up to 48 hours (or 3 days if using pasteurized egg product). Remove from the fridge at least 30 minutes before baking (while preheating the oven).

Fridge: They'll keep in the fridge for 3-4 days. Reheat in the microwave or the oven.

Freezer: Can you freeze breakfast casserole? Yes, you sure can! This will last in the freezer for up to four months. Wrap individually and store in freezer bags. Defrost in the fridge overnight and reheat in the microwave or the oven.

NUTRITION:

Calories: 474   Carbs: 19g   Fiber: 1g
Sodium: 816mg   Protein: 25g

Screenshot Showcase

Posted by bones113, on April 22, 2023, running KDE.

GIMP Tutorial: Colorizing A Black & White Photo

by Meemaw

On YouTube, I have found several tutorials for colorizing a photo. Michael Davies Design has one, as well as Logos By Nick and others. I also found one by JBColourisation. I also did an article on colorizing a photo back in 2013, but there are several methods, so we'll do a different method this time.

We'll be using layer masks for this, and I just heard an easier description of a layer mask. It's like "a window that's been blacked out" so you can't see what's inside, and when you paint the layer mask, you're "making a hole" that shows what you want seen.

Depending on your photo, you can choose whatever color you want, but if you are doing a face, you might find a reference photo or two to help you get the skin tone right. So, I've opened the image I want to colorize (my mother was so pretty!)

Since it's a portrait, and I want accurate skin tones, I'm going to look on Wikimedia Commons for a color photo that has skin tones I can use.
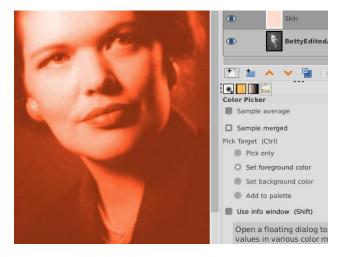
The first step, however, after you open your images, is to make sure the photo you want to colorize is desaturated. Some photos that look black & white still have some color in them. Click on **Colors** > **Desaturate** > **Desaturate**. Then change to the other photo and use the color picker to choose the skin color. Now go back to your photo and add a transparent layer to it with the settings shown at top right:

Notice it says Mode: HSL Color, and Fill with: Foreground Color. Since you've chosen the skin tone, it should fill with your chosen skin tone (right).
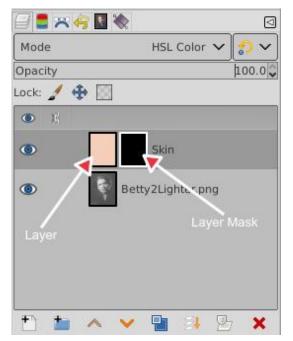
Yikes! That's really dark! From the color I chose, I wouldn't have thought it would be that dark, but it's fixable. Go ahead and add a Layer Mask to your layer. Make it Black (Full

Transparency). When you do that, your photo will change to black & white again, and your layer will have an extra rectangle in the layers dialog. See the difference below.
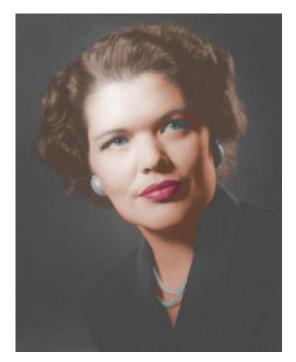

Click on the layer mask, and using a soft brush (mine's about 50%) and the color white, paint on the skin. Don't worry about overlapping things, because you can always change back to black and fill in the places that shouldn't be skin-colored.

I know it's too intense still, so click on the layer itself rather than the layer mask, click on **Colors > Hue and Saturation**, and decrease the saturation until it looks like you want it. You'll have to do that with each layer.

When you have the skin the way you want it, add a new layer for the hair, and more for the lips, and anything else that is in the photo (I need to do the eyes, the jewelry, her suit and the background, too). I have the skin, hair & lips done here:

Feel free to go back to different layers and zoom in to touch up mistakes as you go. Then, pick the colors, add the layer and the layer mask, and use your brush to paint the appropriate sections. You can also choose a color (I think my mother's suit was blue) without having a reference photo, and use that in your new layer. You'll make a separate layer (and layer mask) for each item that needs to be a different color.

Here's the end result:

I'm sure yours will turn out just as good, or better!

ICYMI: Musk, Tech Leaders Call For AI "Pause"

by Paul Arnote (parnote)

While "dumb phones" may be falling out of fashion on a global scale, it's a different story **in the U.S., where "dumb phones" are on the rise in the U.S. as Gen Z looks to limit screen time**, according to an article on CNBC.

Amazon just announced it has opened its Sidewalk long-range network to developer testing, and provided a map showing the low-bandwidth coverage area within reach of 90% of the US population, according to an article on Fierce Electronics. Yes! Let's hand over even MORE of our data over to Amazon! [/sarcasm]

Image by vectorpouch on Freepik

Elon Musk and dozens of other technology leaders have called on AI labs to pause the development of systems that can compete with human-level intelligence, according to an article on CNBC. In an open letter from the Future of Life Institute, signed by Musk, Apple co-founder Steve Wozniak and 2020 presidential candidate Andrew Yang, AI labs were urged to cease training models more powerful than GPT-4, the latest version of the large language model software developed by U.S. startup OpenAI.

A new class of compounds could offer the benefits of gastric bypass surgery without surgery (reduce weight, lower blood glucose, and boost calorie burn), while avoiding the side effects of nausea and vomiting commonly associated with current weight loss and diabetes drugs, according to an article on SciTechDaily. (Move along … no potential for abuse here!)

A portion of the Twitter site's source code was posted on GitHub shortly after widespread layoffs at Twitter, according to an article on TechRepublic. Twitter has taken down a GitHub listing in which a significant amount of the social media site's source code was leaked, according to a legal finding on March 24, 2023 acquired by The New York Times. The leaked code appeared to have been available on GitHub for several months before Twitter sent a copyright infringement takedown. It included "proprietary source code for Twitter's platform and internal tools," according to the filing. The code is no longer available on GitHub since that time. Meanwhile, one week earlier, on March 17, 2023, Musk tweeted that all code used to recommend tweets would be open sourced on March 31.

Image by upklyak on Freepik

NordPass has compiled a list of passwords used by the world's largest companies across 20 industries and 31 countries to "secure" their business accounts, according to a report from the company. Spoiler alert: they're terrible.

A new Android trojan is targeting all your mobile bank accounts, according to an article on TechRadar. The trojan has been around since June 2022, and was offered up as a MaaS (malware-as-a-service) "rental" for $3,000 (U.S.) per month.

ChatGPT AI banned in Italy over privacy concerns, according to an article from Irish Examiner. The Italian data protection authority is now investigating whether ChatGPT has complied with European GDPR rules.

A rock crystal sat in a museum for nearly two centuries. Then they found it wasn't a rock, according to an article in the Miami Herald. It turned out, instead, to be possibly the oldest known dinosaur egg in existence.

British billionaire Richard Branson's rocket company, Virgin Orbit, is ceasing its operations, according to an article on The Street. The company is ceasing operations "for the foreseeable future" after failing to secure a funding lifeline, only five years after it was created, CEO Dan Hart told employees on

March 30, according to CNBC. Subsequent widespread reports cite that Virgin Orbit has filed for bankruptcy, and is seeking a buyer.

Hackers exploit a WordPress plugin flaw that gives full control of millions of sites, according to an article on ArsTechnica. Elementor Pro fixed the vulnerability, but not everyone has installed the patch.

Image by Leopictures from Pixabay

From the crazy "you can't make this stuff up" department, overlapped with the crazier "I have so much money I don't know what to do with it all" department, **the Bill & Melinda Gates Foundation this month awarded a $4.8 million grant to ZELP, which claims its artificial intelligence mask technology for livestock will reduce methane emissions and curb climate change**, according to an article on the Children's Health Defense website. The mask goes around the cow's head and captures the methane gas exhaled by the animal (created from their digestion of food), oxidizing it, and then releasing it into the air as carbon dioxide and water vapor, according to ZELP.

An article from Lifehacker covers **"How to Spot a Fake-A$$ Website."** Scammers make websites that look similar to the real deal—but there are ways to tell if they're fake.

This one ought to chafe your chaps, if you live in the U.S. According to an article on MedPage Today, **third-party tracking is used on almost all U.S. hospital websites, endangering patient privacy, a cross-sectional observational study found**. Of 3,747 hospitals included in the 2019 American Hospital Association (AHA) annual survey, 98.6% of their website home pages had at least one third-party data transfer, and 94.3% had at least one third-party cookie. The tracking data most commonly went to Google's parent company Alphabet (98.5% of homepages), followed by Meta (formerly Facebook), which was used in 55.6% of hospital homepages. Adobe Systems and AT&T collected data from 31.4% and 24.6% of hospital pages, respectively. You can read the complete research article from the report/study from Health Affairs.

Two Russia-associated groups hit Israel and Finland with DDoS attacks, in early April 2023. Cybersecurity experts say the actions represent a marked increase in exploits and a harbinger of cyberattacks to come, according to an article on TechRepublic. The pro-Russia hacker group NoName057(16) reportedly claimed it was behind Denial of Service (DoS) attacks against the Finnish parliament's website on the same day the country joined NATO. Russia-aligned hacktivists also attacked one of the biggest names in security, Check Point,

along with universities and medical centers in Israel, the Jerusalem Post reported.

Oops! **Samsung workers have unwittingly leaked top secret data whilst using ChatGPT to help them with tasks**, according to an article on TechRadar. The company allowed engineers at its semiconductor arm to use the AI writer to help fix problems with their source code. But in doing so, the workers input confidential data, such as the source code itself for a new program, internal meeting notes data relating to their hardware. Since ChatGPT retains user input data to further train itself, these trade secrets from Samsung are now effectively in the hands of OpenAI, the company behind the AI service.
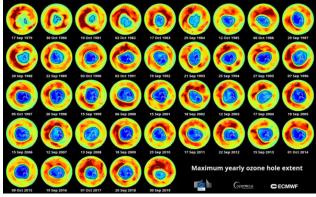

Image from Copernicus, E.U.

According to an article on Vox, **the ozone hole is on track to be healed by mid-century**. New UN data shows that 99 percent of ozone-destroying chemicals have been phased out, underscoring a hopeful environmental story.

Hide your children! Circle the wagons! Go to your bunkers! **Google is planning to release**

their natural language chatbot for its search platform to compete with ChatGPT, according to an article from TechRepublic. Yes, from the company that has shown ZERO restraint in monetizing every picogram of your personal and private information it can get its greedy little paws on (and built a tech empire from it, in the process), now comes their answer to the AI empowered natural language search tool. Oh, yeah! Sign me up! NOT!

Surprise! Surprise! Surprise! Another month, another Google Chrome zero-day vulnerability. On April 14, 2023, **Google released an emergency update for Chrome to address a zero-day security flaw**, according to an article on The Register. The vulnerability, tracked as CVE-2023-2033, can be exploited by a malicious webpage to run arbitrary code in the browser. Thus, surfing to a bad website with a vulnerable browser could lead to your device being hijacked. Exploit code for this hole is said to be circulating, and may well be in use already.

Amazon has added generative artificial intelligence models into Amazon Web Services. Generative AI is the type of artificial intelligence that can create text or images, similar to ChatGPT and DALL-E, according to an article on TechRepublic. Bedrock, now in limited preview, opens the door to two new large language models hosted on AWS.

Image Source: Clipart-library.com

In the U.S., the annual income tax filing ~~fiasco~~ period just ended. **Despite promoting filing your taxes for FREE** (as happens in many, many countries), an article on ProPublica profiles how **a consortium of tax filing software vendors have tricked users into paying to file their taxes**, and otherwise doing everything in their power to lower the number of free tax return filers. All in order to maximize their profits. Yes, this article is about four years old, but the exact same tactics are still employed today. Buyer beware! There are a LOT of articles and videos on the very same topic, here. They aren't difficult to find on your own, but this list should make a good start on the topic. Watch/read if you dare … and if your blood pressure can handle it!

Chromebooks are popular with students but are not built to last and have created massive amounts of electronic trash, according to a recent report featured on the Breitbart website. The April report (PDF) from the U.S. Public Interest Research Group (U.S. PIRG) Education

Fund highlights the fact the devices have an expiration date. The 31 million Chromebooks sold globally in the first year of the pandemic represent approximately 9 million tons of CO2e emissions. Doubling the life of just Chromebooks sold in 2020 could cut emissions equivalent to taking 900 thousand cars off the road for a year, more than the number of cars registered in Mississippi. Assuming no additional maintenance costs, in the U.S., longer lasting Chromebooks could save taxpayers $1.8 billion dollars across all K-12 students.

A team of explorers announced it found a sunken Japanese ship that was transporting Allied prisoners of war when it was torpedoed off the coast of the Philippines in 1942, resulting in Australia's largest maritime wartime loss with a total of 1,080 lives, according to an article on CBS News. The wreck of the Montevideo Maru was located after a 12-day search at a depth of over 4000 meters (13,120 feet) — deeper than the Titanic — off Luzon island in the South China Sea, using an autonomous underwater vehicle with in-built sonar.

Image Source: Clipart-library.com

Since 1906, Miller High Life has used the "Champagne of Beers" slogan. This week, that appropriation became impossible to swallow, according to an article on the NPR website, and widely reported elsewhere in the media. At the request of the trade body defending the interests of houses and growers of the northeastern French sparkling wine, **Belgian customs crushed more than 2,000 cans of Miller High Life** advertised as such. The Comité Champagne asked for the destruction of a shipment of 2,352 cans on the grounds that the century-old motto used by the American brewery infringes the protected designation of origin "Champagne." Like someone is going to buy BEER and be disappointed it's not CHAMPAGNE? Wow! There must be nothing more important to worry about.

This ought to be of interest to all the MLUs reading this. **Scientists believe they have discovered the mechanism for hair turning gray**, which could help develop treatment to alter cells in order to reverse or halt the process, according to an article on The Guardian, and widely reported across most media outlets. A new study suggests stem cells may get stuck as hair ages and lose their ability to mature and maintain hair color.

A bicyclist in Cleveland took matters into his own hands after being hit by a car that ran a red light and hit the rear wheel of the cyclist last October and throwing the cyclist from his bike, according to an article on Bicycling. Tiring of the inaction by "authorities," he put up his OWN traffic signs, many expressing the sentiment "Slow The F(iretru)ck Down" (or something similar).

In another article from Bicycling, **The Terminator** himself – former California governor and box-office top tier actor, Arnold Schwarzenegger – ~~filled~~ **terminated a pothole himself**, after waiting three weeks for city workers to perform the task (they never showed up). The pothole had been plaguing motorists and bicyclists. At least one passerby motorist stopped and thanked the 75 year old Scwarzenegger for repairing the pothole.

Be Skeptical Of FBI Warnings About Phone Chargers

by **Jacob Hoffman-Andrews**
Electronic Frontier Foundation
Reprinted under Creative Commons License

Every few years, an unsourced report circulates that "the FBI says plugging into public charging kiosks is dangerous." Here's why you should ignore the freakout and install software updates regularly.

Your phone is designed to communicate safely with lots of things – chargers , web sites, Bluetooth devices such as earbuds or speakers, Wi-Fi, and even other phones, for instance when sending and receiving text messages. If doing any of these normal phone things can give your phone malware, that is a security vulnerability (which is a type of bug).

Security vulnerabilities happen with some frequency. That is why your phone prompts you to update your software so often – the makers of its software find out about bugs and fix them.

So, when you hear a report that public chargers are giving people malware, you should ask "what is the vulnerability being used, and when will it be fixed?" as well as "how widespread is the problem? How many people are affected?" Unfortunately, the periodic reports of "juice jacking" never have such details, usually because they are recycled from earlier reports which themselves lack details.

The most recent news reports reference a tweet from the FBI Denver field office. According to reporter Dan Goodin's conversation with an FBI spokesperson, the field office relied on an article the FCC published in 2019 warning about USB charging stations. The only source for that article was a warning from the Los Angeles County District Attorney's Office that did not itself allege any specific bug or specific instances of charging stations being used for attacks. The FCC later quietly removed the sourcing from its article, allowing itself to be incorrectly treated as a primary source for juice jacking claims.

While the video from the LA County D.A. doesn't mention it, the ultimate source for the term "juice jacking" is a Brian Krebs article from 2011 reporting on a vulnerability demonstrated at DEFCON that year. As you can imagine, phone security has changed dramatically since 2011. And so far there have been no reports of widespread exploitation of USB vulnerabilities in the wild.

As a complex protocol, USB does present a large attack surface– and there are some built-in risks, like the ability for a USB device to pretend to be a keyboard (so lock your phone while charging). You may also want to bring your own charger or battery for electrical reasons. Phone manufacturers often recommend charging only with approved chargers, to avoid charging too slow or (worse) too fast, and potentially damaging your phone or battery. But realistic security is about risk management, and for most people the risk of a public USB charger is very low.

Undoubtedly there will continue to be bugs in phones' USB stacks in the future, just as there will be bugs in web browsers and chat apps. Some of those bugs will have the potential to infect your phone with malware, particularly if large numbers of people forget to update their software. But with a little skepticism and common sense, we can stop zombie scaremongering about charging stations from making the rounds again.

GIMP Tip: Join Text

by tuxlink

When I first registered for PCLinuxOS and started getting more involved in the community, I noticed how in the artwork of the distro logo, two letters in the middle are cleverly 'joined' to look like only one letter or character. The U and the N look like they are merged or magically melted together. Being a fan of graphics, in time when I needed to use the logo, I wasn't sure how to go about joining the two letters. I asked in the forum, but the answers I got didn't yield a good enough result. There was always something sticking out of place, or something didn't look quite right. In time, I eventually figured out how to join the two letters seamlessly. Here, in this Gimp tip, we'll see how I made it work for me. I am sure there are many other different ways to accomplish the same task, but this is how I do it. (See samples A1, A2, A3 below)

When you open up Gimp, choose a **New** project, and you'll see the default size will be the size of your monitor resolution. If you prefer a size that's smaller, by all means choose that. Just make sure it's wide enough to hold the wor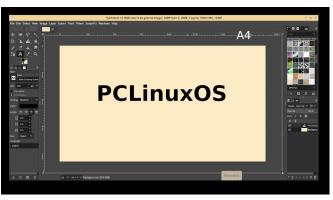d PCLinuxOS in a fairly large sized font. I will leave you to choose the background color, it can be black or white, or any other color. Just make sure you use a contrasting color for your font and characters, ie, white background, black font. The larger the font size you use, the easier it will be to move characters around, and the better the final result will be.
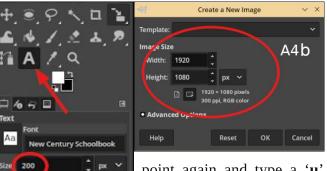
I am going to purposely pick a font that is standard in all PCLinuxOS systems, **New Century Schoolbook Bold**. In the size box, insert **200 pixels**. I also chose this font because it has clean edges. There are no serifs or other pieces that might stick out and ruin our 'joining' project. Using the **Text tool**, and aforementioned font, type out the word PCLinuxOS in caps and lowercase. This step is to make sure yours looks like the font used in the screenshot.

Next, over on the bottom right corner is the layers tab. The text layer should be sitting on top of the background layer. Right click that text layer, and choose, **Delete Layer**.
(See samples A4, A4b, A5 at top right.)

The next step is to choose your **Text tool**, on the left side close to the edge of the page, type in **PCLi** only! Take your mouse and point it unto a different part of the page, then, just type the character '**n**' in lowercase only. Then again on a different part of the page, lift your mouse and point again and type a '**u**' by itself. The rest of the word will be by itself also, type **xOS**. By now you should have 4 individual pieces of text on the page. If you look in the Layers tab there should be 4 layers of text on top of the background layer.
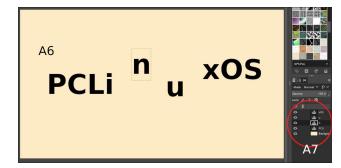(See samples A6_A7 on the next page.)

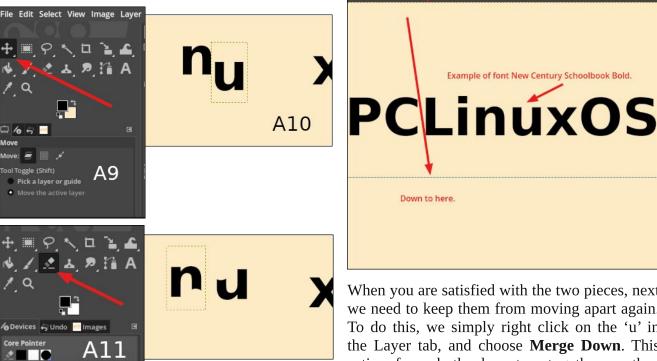horizontal guide line will help with accuracy. (See A9, A10, A11, A12)

Next, we're going to be using the **Move tool**. It's in the top left corner of the **Tools tab**. To move anything in Gimp, you must first select it in the Layers tab. After selecting the **Move tool**, click on the text layer containing the lowercase '**u**'. This allows you to move just that item on the page. Move it close to the lowercase '**n**'. Because I'm using a black font on a cream bright background, I can easily drop the **u** on one side of the **n** and it looks like it could fit together. But when we line them up with the other pieces, the corners will protrude and look ugly.

To fix this, I use the **Eraser tool**. Choose a hard round brush set to 70 pixels and 100 hardness. Make sure the letter '**u**' is selected in the layers tab, and then with the Eraser, remove half of the left side of the letter '**u**'. Then do the exact same with the right side of the letter '**n**'. It should now look like this. (A12) Now it's simply a matter of selecting the '**u**', and with the **Move tool**, move it over the '**n**' section. Now comes the accuracy part. You have the choice of using your eye for alignment, (Bless your heart and eyesight, I wish mine was still as good!) but if you need assistance with alignment, a

A guideline will allow you to line up the various parts that we want to join together accurately.

Reach up above the numbered grid at the top of the page, and click your mouse and drag a line down to about 3/4 of the page and release it. Now you have a straight edge to place all characters on a line for alignment. (See A5b at top right)
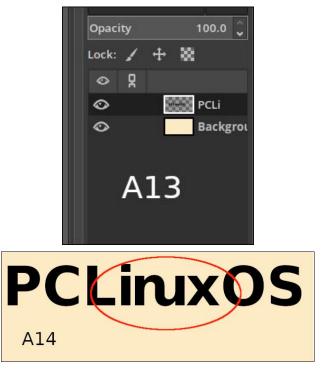
When you are satisfied with the two pieces, next we need to keep them from moving apart again. To do this, we simply right click on the 'u' in the Layer tab, and choose **Merge Down**. This action fuses both characters together so they now move as one item. They also become one layer. From here, we simply use the **Move tool** to line up the other pieces to spell out the complete word. Use the guideline. If you need to make fine small movements, use the up and down and left and right arrows on the keyboard. When you are satisfied with how it all fits, making sure that all spaces between each character are equal, you can then right click on the **xOS Layer**, and once again choose **Merge Down**. One more time and you should have the complete word as one item, and just 2 layers left in the **Layers tab**. The text layer, and the background.

With your **Move tool**, you can now click on the text layer and move the finished word anywhere you need it.

Lastly, you need to save the completed text word as a PNG file. I usually delete the background layer, and save the singular text layer as a PNG. That way, I can use it as an individual piece on many other projects. (See A13, A14.)

A13

A14

Screenshot Showcase

Posted by francesco_bat, on April 10, 2023, running icewm.

PCLinuxOS Recipe Corner Bonus

Try substituting Peppermint with Wintergreen, Spearmint or Cinnamon for a different flavor.

Ritz thin mints will keep for 3-4 months or even longer in the freezer. Place in zippered baggies on the door of your freezer so that yummy chocolate fix will always be in easy reach! No thawing required!

from the kitchen of youcantoo

Ritz Cracker Thin Mints

Servings: 40

INGREDIENTS:

40 salted Ritz crackers
16 ounces chocolate
   (1 pound or approx 2 ½ cups)
½ teaspoon peppermint extract
½ teaspoon vegetable oil or coconut oil

DIRECTIONS:

Place chocolate and vegetable oil in a bowl. Place the bowl over a pot of gently simmering water (ensuring the bowl doesn't touch the water). Stir gently until melted. Once melted, stir in the extract.

Dip crackers into the chocolate mixture and allow excess to drip off.

Place on a parchment-lined pan and refrigerate until cooled.

NUTRITION:

Calories: 73   Carbs: 9g   Fiber: 1g
Sodium: 28mg   Protein: 1g

The PCLinuxOS Magazine

Created with Scribus

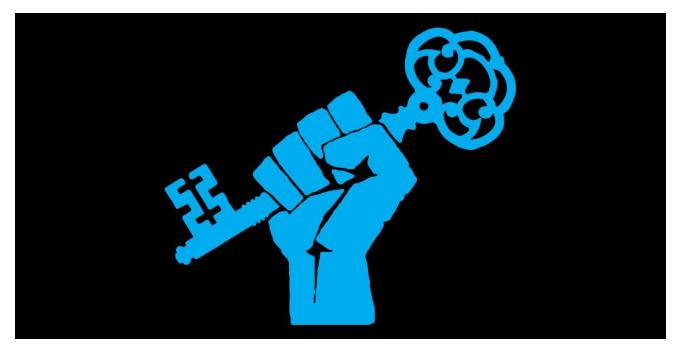# The EARN IT Bill Is Back, Seeking To Scan Our Messages & Photos

by **Joe Mullin**
Electronic Frontier Foundation
Reprinted under Creative Commons License

In a free society, people should not have their private correspondence constantly examined. U.S. lawmakers, we would hope, understand that individuals have the right to a private conversation without the government looking over their shoulder.

So it's dismaying to see a group of U.S. Senators attempting for a third time to pass the EARN IT Act (S. 1207)—a law that could lead to suspicionless scans of every online message, photo, and hosted file. In the name of fighting crime, the EARN IT Act treats all internet users like we should be in a permanent criminal lineup, under suspicion for child abuse.

What The New "EARN IT" Does

The EARN IT Act creates an unelected government commission, stacks it with law enforcement personnel, and then tasks it with creating "best practices" for running an internet website or app. The act then removes nearly 30-year-old legal protections for users and website owners, allowing state legislatures to encourage civil lawsuits and prosecutions against those who don't follow the government's "best practices."

As long as they somehow tie changes in law to child sexual abuse, state lawmakers will be able to avoid longstanding legal protections, and pass new rules that allow for criminal prosecutions and civil lawsuits against websites that don't give police special access to user messages and photos. Websites and apps that use end-to-end encryption to protect user privacy will be pressured to remove or compromise the security of their services, or they'll face prosecutions and lawsuits.

If EARN IT passes, we're likely to see state lawmakers step in and mandate scanning of messages and other files similar to the plan that Apple wisely walked away from last year.

There's no doubt the sponsors intend this bill to scan user messages, photos, and files, and they wrote it with that goal in mind. They even suggested specific scanning software that could be used on users in a document published last year. The bill also makes specific allowances to allow the use of encryption to constitute evidence in court against service providers.

Linux Docs
Linux Man Pages

Bill Language Purporting To Protect Encryption Doesn't Do The Job

Under pressure, the bill sponsors did add language that purports to protect encryption. But once you take a closer look, it's a shell game. The bill clearly leaves room to impose forms of "client-side scanning," which is a method of violating user privacy by sending data to law enforcement straight from user devices, before a message is encrypted. EFF has long held that client-side scanning violates the privacy promise of end-to-end encryption, even though it allows the encryption process to proceed in a narrow, limited sense. A 2021 paper by 10 leading technologists held that client-side scanners are a danger to democracy, amounting to "bugs in our pockets."

The Chat-Scanning Software Being Pushed By This Bill Doesn't Work

But the available evidence shows that scanning software that looks for Child Sexual Abuse Material, or CSAM, is far from perfect. Creators of scanning software say they can't be fully audited, for legal and ethical reasons. But here's the evidence so far:

* Last year, a New York Times story showed how Google's CSAM scanners falsely accused two fathers of sending child pornography. Even after the dads were explicitly cleared by police, Google kept their accounts shut down.

* Data being sent to cops by the U.S. National Center for Missing and Exploited Children (NCMEC)—the government agency that will be tasked with analyzing vastly more user data if EARN IT passes—is far from accurate. In 2020, the Irish police received 4,192 reports from NCMEC. Of those, only 852 (20.3%) were confirmed as actual CSAM. Only 9.7% of the reports were deemed to be "actionable."

* A Facebook study found that 75% of the messages flagged by its scanning system to detect child abuse material were not "malicious," and included messages like bad jokes and memes.

* LinkedIn reported 75 cases of suspected CSAM to EU authorities in 2021. After manual review, only 31 of those cases—about 41%—involved confirmed CSAM.

The idea of subjecting millions of people to false accusations of child abuse is horrific. NCMEC will export those false accusations to vulnerable communities around the world, where they can be wielded by police forces that have even less accountability than law enforcement in the United States. False accusations are a price that EARN IT supporters seem willing to pay.

We need your support to stop the EARN IT Act one more time. Digital rights supporters sent more than 200,000 messages to Congress to kill earlier versions of this bill. We've beaten it twice before, and we can do it again.

There are currently dangerous proposals that could mandate client-side scanning schemes in the U.K. and European Union, as well. But we don't need to resign ourselves to a world of constant surveillance. In democratic nations, supporters of a free, secure, and private internet can win—if we speak up now.

Disclaimer

1. All the contents of The PCLinuxOS Magazine are only for general information and/or use. Such contents do not constitute advice and should not be relied upon in making (or refraining from making) any decision. Any specific advice or replies to queries in any part of the magazine is/are the person opinion of such experts/consultants/persons and are not subscribed to by The PCLinuxOS Magazine.

2. The information in The PCLinuxOS Magazine is provided on an "AS IS" basis, and all warranties, expressed or implied of any kind, regarding any matter pertaining to any information, advice or replies are disclaimed and excluded.

3. The PCLinuxOS Magazine and its associates shall not be liable, at any time, for damages (including, but not limited to, without limitation, damages of any kind) arising in contract, rot or otherwise, from the use of or inability to use the magazine, or any of its contents, or from any action taken (or refrained from being taken) as a result of using the magazine or any such contents or for any failure of performance, error, omission, interruption, deletion, defect, delay in operation or transmission, computer virus, communications line failure, theft or destruction or unauthorized access to, alteration of, or use of information contained on the magazine.

4. No representations, warranties or guarantees whatsoever are made as to the accuracy, adequacy, reliability, completeness, suitability, or applicability of the information to a particular situation. All trademarks are the property of their respective owners.

5. Certain links on the magazine lead to resources located on servers maintained by third parties over whom The PCLinuxOS Magazine has no control or connection, business or otherwise. These sites are external to The PCLinuxOS Magazine and by visiting these, you are doing so of your own accord and assume all responsibility and liability for such action.

Material Submitted by Users
A majority of sections in the magazine contain materials submitted by users. The PCLinuxOS Magazine accepts no responsibility for the content, accuracy, conformity to applicable laws of such material.

Entire Agreement
These terms constitute the entire agreement between the parties with respect to the subject matter hereof and supersedes and replaces all prior or contemporaneous understandings or agreements, written or oral, regarding such subject matter.

Like Us On Facebook!
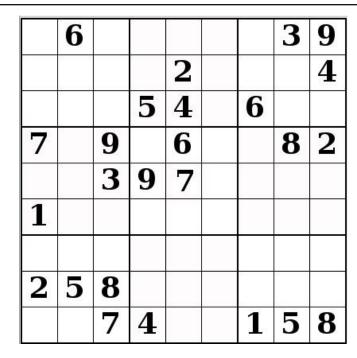*The PCLinuxOS Magazine*
PCLinuxOS Fan Club

Screenshot Showcase

Posted by parnote, on April 12, 2023, running Xfce.

PCLinuxOS Puzzled Partitions

SUDOKU RULES: There is only one valid solution to each Sudoku puzzle. The only way the puzzle can be considered solved correctly is when all 81 boxes contain numbers and the other Sudoku rules have been followed.

When you start a game of Sudoku, some blocks will be prefilled for you. You cannot change these numbers in the course of the game.

Each column must contain all of the numbers 1 through 9 and no two numbers in the same column of a Sudoku puzzle can be the same. Each row must contain all of the numbers 1 through 9 and no two numbers in the same row of a Sudoku puzzle can be the same.

Each block must contain all of the numbers 1 through 9 and no two numbers in the same block of a Sudoku puzzle can be the same.

SCRAPPLER RULES:
1. Follow the rules of Scrabble®. You can view them here. You have seven (7) letter tiles with which to make as long of a word as you possibly can. Words are based on the English language. Non-English language words are NOT allowed.
2. Red letters are scored double points. Green letters are scored triple points.
3. Add up the score of all the letters that you used. Unused letters are not scored. For red or green letters, apply the multiplier when tallying up your score. Next, apply any additional scoring multipliers, such as double or triple word score.
4. An additional 50 points is added for using all seven (7) of your tiles in a set to make your word. You will not necessarily be able to use all seven (7) of the letters in your set to form a "legal" word.
5. In case you are having difficulty seeing the point value on the letter tiles, here is a list of how they are scored:
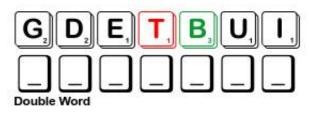 0 points: 2 blank tiles
  1 point: E, A, I, O, N, R, T, L, S, U
  2 points: D, G
  3 points: B, C, M, P
  4 points: F, H, V, W, Y
  5 points: K
  8 points: J, X
  10 points: Q, Z
6. Optionally, a time limit of 60 minutes should apply to the game, averaging to 12 minutes per letter tile set.
7. Have fun! It's only a game!

Download Puzzle Solutions Here

Possible score 212, average score 148.
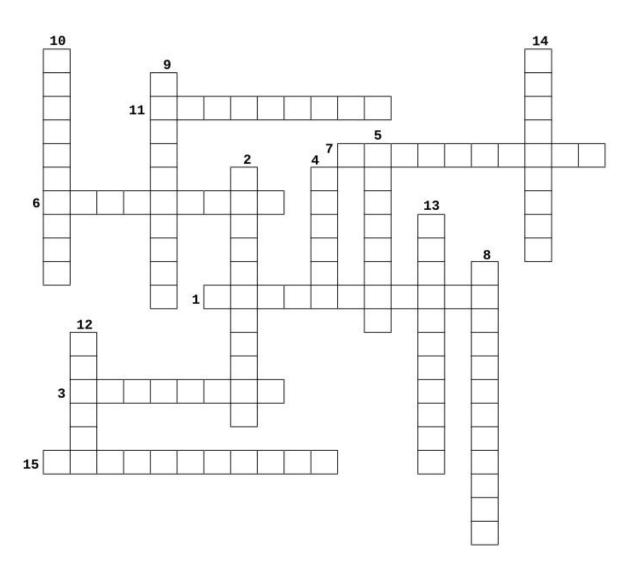
May 2023 Word Find
Office

```
G K O N W J R B J G K R U E N E R P E R T N E H Y F J J J R P
C L J N D C S U P E R V I S O R C T P I E C E R V S L T E Y
E G D K V N Y L K E A G E N D A A Q E E X I L V G I Y J D G
C M P W P F B Z P D A C L D D M T D Q G O H C T X S C W L P
U E T A R T S E R E T N I J R M W I Q L N M O R K N N W O U
N N E G O T I A T I O N M B X B C V W L L N K G M L E D H R
E M N D H A J V S Y O A H V L V D S I G N A T U R E R G K C
M D H L D S Z C N B N D E F L A T I O N A Z C A X H R X C H
P O C P L B S A L A R Y U J E R D Y M P Y R T E G D U B O A
L K O I W S O S G Q Z Y A M S B L K F S C I M O N O C E T S
O I R G B J V E S L D A L L F J E N L A R E T A L L O C S I
Y B P S R S M V Q E A B I P J E A D A S H S B U S M U J J N
M F O R J E R H T H N A R A W F R N C O M P E T I T I O N G
E M R O N Z A E M S B I T E P E S E I V T S Z L S N A S W X
N M A T U G R D O I Z V S R M Q W D M S P S T I R I H W Y D
T Y T C U W B W L V Q M O U Y I U I N B D F L S W Y O X L Z
R N I E P M I I I X E F Y Y B W E V B E Z A H N Y Z L M B E
E V O R R T T I J E I S O V N G Q I W V T P J Z P D S H K J
D P N I O Y R K K T N D I O Y N W D N I I T O Y K B Y Y F T
L M B D D X A W C A F D I T J I N F P T D D M R D K V E D N
O C S F U D T M D X L T I V R P H A D U Y J X U K E A P U E
H E Z O C R I V D S A S C D V E C Y L C L E J S W C F L X I
E J O D T G O Q U P T R D F G E V D H E F C H A I R M A N L
R H E R I L N N U O I P A T H K W D A X A C I E D E G I C C
A R K A O Z I C T Z O E L J R K S P A E L O D R G M L C N L
H D Z O N M C H K Z N C B T A O A P U W H F S T G M H N B O
S K O B K O V M J E B A P O B O J X M Y Z R S R U O L A I Q
W B T T Q E B L I F X P R M K B N F B N J V E J O C O N B K
C O P Y R I G H T V X M Y E T Y D E E D D D R V E C J I X M
H K V D C V U J W Q Q J C H T P U L N O W H Z K O M N F M O
```

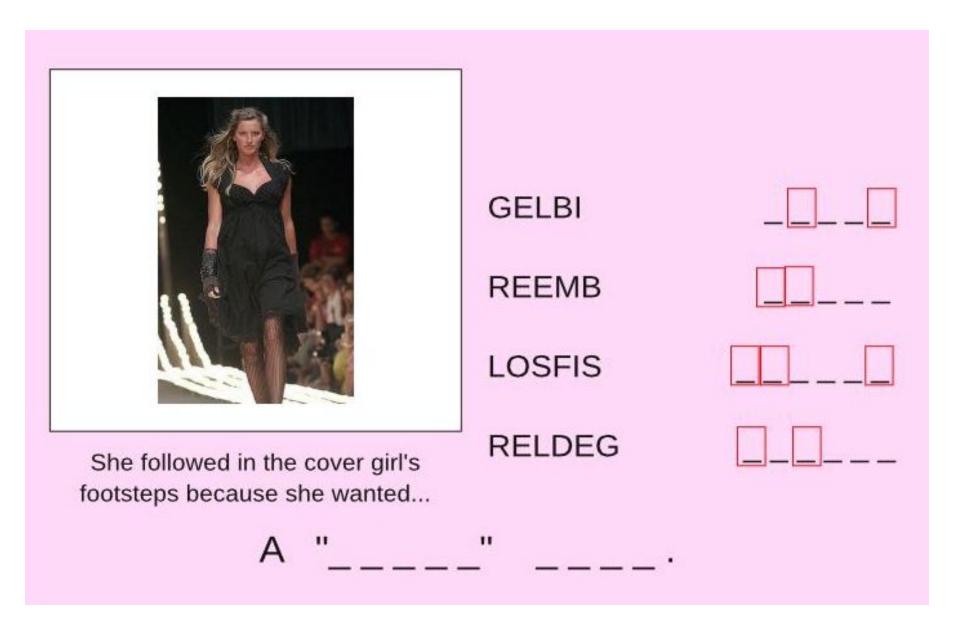|	
ADVERTISE	AGENDA
ARBITRATION	BOARD OF DIRECTORS
BOOKKEEPING	BUDGET
BUSINESS	CAPITALIST
CHAIRMAN	CLIENT
COLLATERAL	COMMERCE
COMPETITION	COPYRIGHT
CORPORATION	CURRENCY
DEFLATION	DIVIDEND
ECONOMICS	ENTREPRENEUR
EXECUTIVE	FINANCIAL
INFLATION	INTEREST RATE
LIABILITY	MANAGEMENT
NEGOTIATION	OCCUPATION
OVERHEAD	PRODUCTION
PROFIT	PURCHASING
RECEIPT	SALARY
SHAREHOLDER	SIGNATURE
STOCKHOLDER	SUPERVISOR
TREASURY	UNEMPLOYMENT

Download Puzzle Solutions Here

May 2023 Crossword Office

1. The process of discussing something with someone in order to reach an agreement.
2. The practice or profession of recording the accounts and transactions of a business.
3. A reward paid to the shareholders for their investment in a company's equity.
4. Financial gain from a transaction or from a period of investment or business activity.
5. Business expenses (such as rent, insurance, or heating) not chargeable to a particular part of the work.
6. A persistent increase in the level of consumer prices.
7. Security pledged for the repayment of a loan.
8. One who takes the initiative to create a product or establish a business for profit.
9. An investor of capital in business or person of great wealth.
10. One who is in charge of a particular department or unit
11. To make known; call attention to.
12. A systematic plan for the expenditure of a usually fixed resource, such as money or time, during a given period.
13. The process of resolving a dispute or grievance by presenting it to an impartial third party.
14. Exclusive right to multiply and to dispose of copies of an intellectual production.
15. One who owns a share or shares of stock in a company.

Download Puzzle Solutions Here

Mixed-Up-Meme Scrambler

GELBI

REEMB

LOSFIS

RELDEG

She followed in the cover girl's footsteps because she wanted...

A "_ _ _ _ _" _ _ _ _.

More Screenshot Showcase

Posted by jogurtmen, on April 4, 2023, running KDE.

Posted by luikki, on April 7, 2023, running KDE.

Posted by Meemaw, on April 3, 2023, running Xfce

Posted by mutse, on April 20, 2023, running Mate.