TOOLKIT

The McAfee Safety Series

The McAfee AI Election Toolkit

Table of Contents

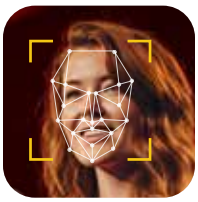

Your election toolkit in the Age of AI

As Election Day approaches, concerns about the influence of malicious AI and deepfakes have grown.

How will the so-called “fake news” and disinformation they spread sway voters or prevent them from casting their vote altogether? The question isn’t if that will happen. Arguably, the question is to what extent will malicious AI and deepfakes influence the 2024 elections?

Disinformation isn't new. Using the power of AI to spread it is.

Disinformation has played a long and shady role in politics. For some time now. George Washington fell victim to it in 1777 when forged letters painted him as a British sympathizer — disinformation that followed him to the first presidency.

From there, we can point to the "yellow journalism" of the late 19th century with its use of scare headlines and unverified claims that pushed partisan agendas. While its run was brief, yellow journalism quickly morphed into brassy tabloid journalism, which has its own watermarks of sensationalism and murky motives.

Jump decades ahead and we have the advent of the internet. With it came the promise that "everyone can be a publisher." Of a sort. Where publishers and media newsrooms once acted as the gatekeepers of information, now practically anyone could post and spread news to a global audience in seconds.

This brings us to the Information Age we live in today, along with the multitude of benefits it brings. Yet an undercurrent courses here, where falsehoods get passed off as truth and muddy the legitimate media landscape. Over time, that has left many skeptical if not downright mistrustful of what they see and hear online. Such a lack of trust is damaging at any time — and yet more so in an election year.

People with little to no technical expertise can create malicious deepfakes that look and sound increasingly legitimate as AI tools continue to evolve. Further, AI helps spread disinformation on a massive scale. The technology has spawned hundreds of unreliable AI-generated news sites laden with AI-authored articles, all of which pass themselves off as legitimate.

Now, the rise of AI and AI deepfake technology amplifies the spread of disinformation even more.

You're not powerless in the face of malicious AI. Quite the opposite.

Even as the creators of malicious AI-generated content have gotten cagier in their ways, their work still gives off signs of a fake. However, spotting this malicious content calls for extra effort on everyone's part when getting their news or scrolling their feeds online. That means scrutinizing what we consume and relying on trusted fact-checking resources to get at the truth. It also means using AI as any ally, with AI tools that detect AI deepfakes in real time.

This Election Year Toolkit will help you do just that. We'll cover the basics of fake news and malicious AI deepfakes, how to spot them, and more. As you'll see, it's a topic both broad and deep, which we'll explore in a step-by-step way.

Establishing the baseline: What is "fake news?"

What exactly is fake news? As a term, it gets used often enough, to the extent that its meaning has gotten watered down. A fresh look at what fake news truly is helps frame the issues of trust we face online.

With that, a textbook definition of fake news goes something like this:

> ***A false news story, fabricated with no verifiable facts, and presented in a way to appear as legitimate news.***

As for its intent, fake news often seeks to damage the reputation of an individual, institution, or organization. It might also spout propaganda or attempt to undermine established facts.

That provides a broad definition. Yet, like much fake news itself, the full definition is much more nuanced. Within fake news, you'll find two categories: disinformation and misinformation.

You might see and hear these terms used interchangeably. They're different, yet they're closely related. And both will play a role in this election.

Disinformation is intentionally spreading false or misleading info.

Misinformation is unwittingly spreading false or misleading info (the person sharing the info thinks it's true).

This way, you can see how "fake news" spreads. A bad actor posts an AI deepfake with misleading info — a form of disinformation. From there, others take the misleading info at face value and pass it along as truth — a form of misinformation.

The two work hand-in-hand by design, because bad actors have a solid grasp on how lies spread online. (We'll cover that in more detail shortly when we touch on social media.)

From there, fake news gets more nuanced still. Misinformation and disinformation fall within a range. Some of it might appear comical, while other types have the potential to do actual harm.

Dr. Claire Wardle, the co-director of the Information Futures Lab at Brown University, cites seven types of misinformation and disinformation on a scale as visualized below:

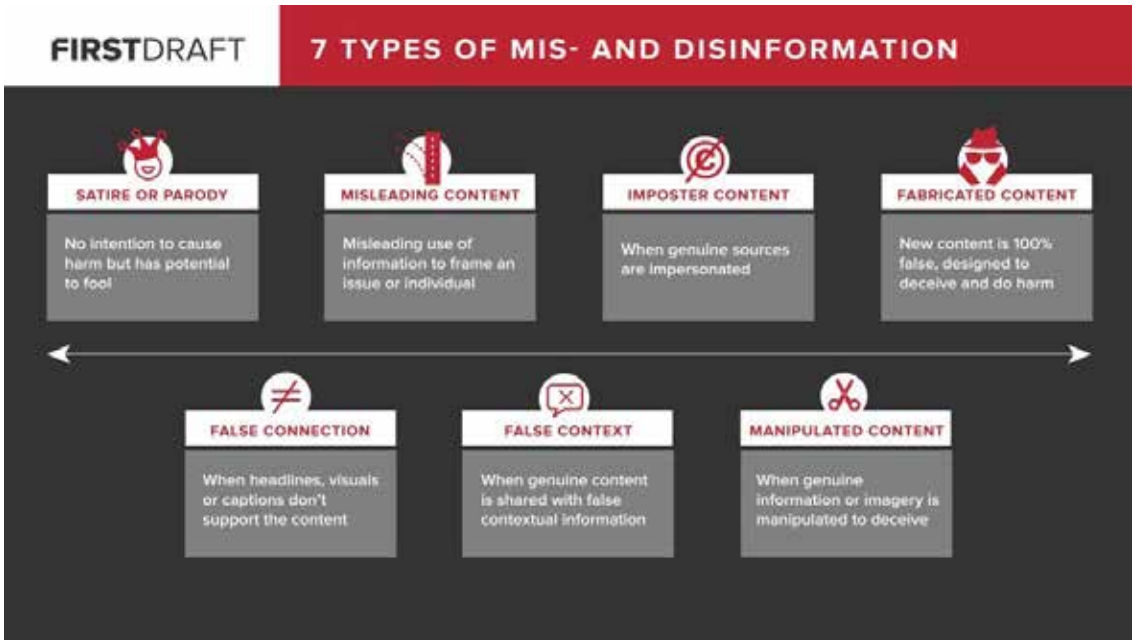

Source – FirstDraftNews.org and Brown University

Put in a real-life context, you can probably conjure up plenty of examples where you've seen mis- and disinformation. Like clickbait-y headlines that link to articles with no payoff. Or maybe you've seen a quote pasted on the image of a public figure, a quote which that person never made. Perhaps you scrolled upon an infographic, loaded with bogus statistics and attributed to an organization that doesn't even exist. It can take all forms.

However, when we focus on fake news in an election cycle, we largely focus on malicious fake news. The kind of malicious fake news that intends to sway voters or deny them their opportunity to vote.

Who's behind fake news? And why?

The answers here vary as well. Greatly so. Fake news can begin with a single individual, groups of like-minded individuals with an agenda, and it can even come from operatives for various nation-states. As for why, they might want to poke fun at someone, drive ad revenue through clickbait articles, spout propaganda, or spread flat-out falsehoods.

In practice, some examples of fake news and the reasons behind it might look like this:

- Imposter sites that pose as legitimate news outlets yet post entirely unfounded pieces of propaganda.
- Parody sites that can look legitimate, so much so that people might mistake their content for actual news.
- Malicious AI deepfakes, images, recordings, and videos of public figures in embarrassing situations, yet that get presented as "real news" to damage their reputation.
- False voter information or false reporting of early election results, which might prevent people from getting to the polls.

With a sense of fake news and the forms it can take, next we look at how malicious AI deepfakes can shape it further still.

How AI further powers disinformation and “fake news”

AI has given new life to the old problem of disinformation and fake news.

It’s done so in two primary ways:

1. Bogus articles and doctored photos once took time and effort to cook up. Now, they take seconds.
2. AI tools can effectively clone voices and people to create convincing-looking deepfakes in digital form.

In effect, the malicious use of AI makes it easier for fake to masquerade as real, with chilling authenticity that’s only increasing. Moreover, it also makes it churn out fake news at a scope and scale previously unseen.

The explosion of fake news sites

With that, we’re seeing an explosion of fake news sites with content nearly, if not entirely, created by AI — with bad actors pushing the buttons.

One media watchdog organization put some striking figures to the recent onrush of fake news sites. In May 2023, the organization found 49 sites that it defined as “Unreliable AI-Generated News Websites,” or UAINS. In February 2024, that number grew to more than 700 UAINS.

Per the watchdog group, these sites run with little to no human oversight. Additionally, they try to pass themselves off as legitimate by presenting their AI "authors" as people. Brazenly, at least one publisher had to say this when confronted with the fact that his "reporter" bylines were AI bots:

> *The goal was to create "AI personas" that can eventually "grow into having their own following," maybe even one day becoming a TV anchor. "Each AI persona has a unique style … Some sort of — this is probably not the right word — personality style to it."*

Beyond spreading disinformation, these sites are profitable. Recent research found that among the top 100 digital advertisers, 55% of them had their ads placed on disinformation sites. Across all industries and brands, 67% of those with digital ads wound up on disinformation sites.

To clarify, these advertisers support these disinformation sites unwittingly. The researchers cite the way that online advertising platforms algorithmically place ads on various sites as the culprit. Not the advertisers themselves.

So as we talk about disinformation sites cropping up at alarming rates, we also see bad actors profiting as they prop them up.

And the explosion of deepfakes

AI tools power yet another Election Year concern, that of malicious deepfakes. Whether they spread disinformation, damage the reputation of candidates, or share incorrect polling information, they threaten to upend the integrity of our elections.

First, what is a deepfake? One dictionary definition of a deepfake reads like this:

An image or recording that has been convincingly altered and manipulated to misrepresent someone as doing or saying something that was not actually done or said.

Looking closely at that definition, three key terms stand out: "altered," "manipulated," and "misrepresent."

Altered

This term relates to how AI tools work. People with little to no technical expertise can tamper with existing source materials (images, voices, video) and create clones of them.

Manipulated

This speaks to what can be done with these copies and clones. With them, people can create entirely new images, tracts of speech, and videos.

Misrepresent

Lastly, this gets to the motives of the creators. They might create a deepfake as an obvious spoof, like many of the parody deepfakes that go viral. Or maliciously, they might create a deepfake of a public official spewing hate speech and try to pass it off as real.

Given this definition, it's important to remember that not all deepfakes are bad. Companies use deepfake technologies to create training videos. Studios use it to dub movie subtitles into other languages. And some content creators just want to get a laugh out of Arnold Schwarzenegger singing showtunes.

How deepfakes might get used during the election

Where deepfakes can become harmful and outright malicious is by pushing disinformation. Particularly in an election year.

Deepfakes have already appeared in this and recent election cycles:

- During the New Hampshire primary, a deepfake audio track of President Biden made rounds of robocalls. In the call, the phony message from the President asked voters to "... save your vote for November election." That falsely implied that if people vote in the primary, they can't vote in the November general election.
- In 2023, a flood of deepfake images showed former president Donald Trump getting arrested on the streets of New York. That, of course, never happened.
- Also in 2023, a bogus news outlet on X (Twitter) posted a voice clone of a Chicago mayoral candidate that made it appear he advocated police violence.

We uncovered several malicious deepfakes ourselves, using our AI-powered deepfake detection technologies.

Those are only a few quick examples of deepfakes that have appeared so far. We can expect to see far more deepfake audio, deepfake images, and deepfake videos that have public figures saying and doing things they never did.

And as polling dates get closer, we can expect more deepfakes that target voters directly.

- Potentially, bad actors will spread phony polling info that prevents voters from getting to polling places in a timely way — or at all.
- On Election Day, we might also see deepfakes that skew polling results, all with the aim of influencing voters.

In all, the bad actors behind these deepfakes aim to do one thing: spread disinformation during the election.

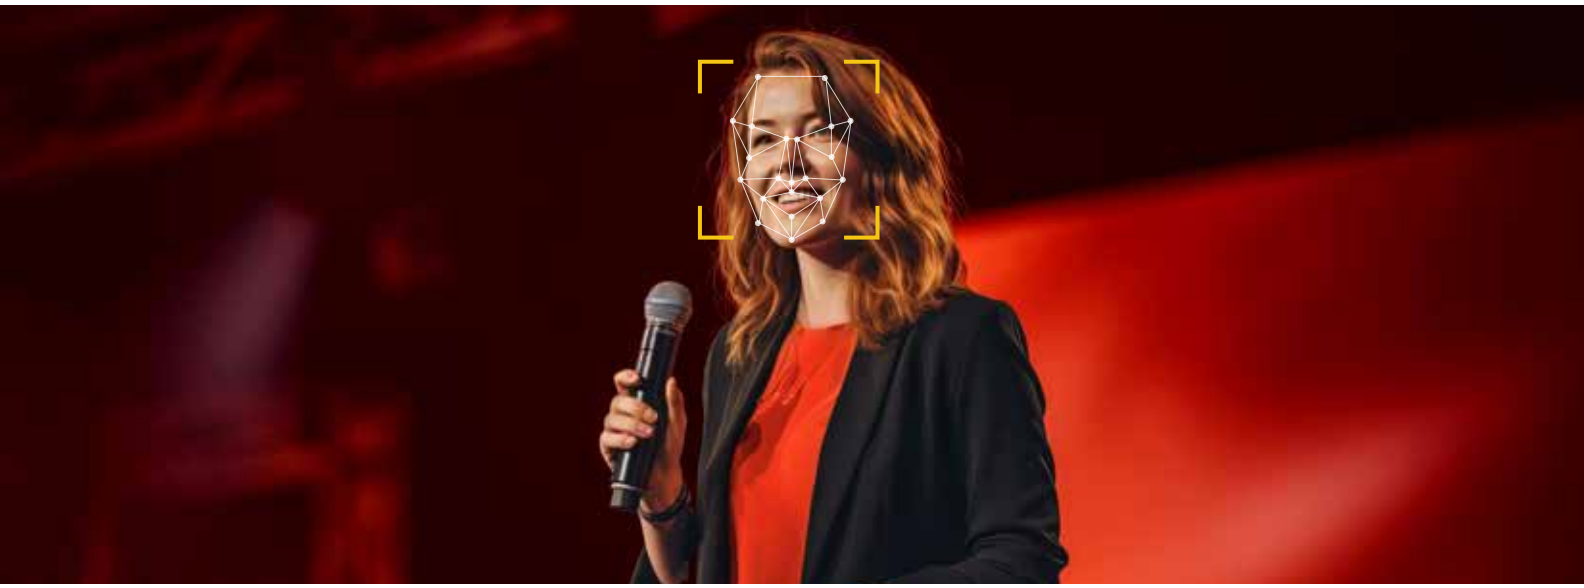

How to spot a deepfake

What are typical signs of a deepfake?

This gets to the tricky bit. When deepfake technologies first arrived, many deepfakes gave off telltale signs of manipulation. Blurriness, unnatural lighting effects, and a lack of breath in speech once stood as ways you could spot a deepfake. Increasingly, you can't count on those signs anymore. The technology has gotten better and continues to get better still.

Yet bad actors still use older and less sophisticated tools. As such, they can leave signs.

How to spot AI-generated text

Look for superficial explanations. AI often says a lot without saying much at all. The explanations they provide lack detail and depth, even as they hide behind a glut of weighty vocabulary to appear informed.

Look for a lack of facts. In addition to lacking detail, AI-generated text often fails to cite exact facts and figures. Moreover, it often fails to cite sources for its claims. (Contrast that with what you're reading in this toolkit, which makes use of endnotes for sources.)

Look for repetition. AI chatbots get trained on volumes and volumes of text. As such, they often latch onto pet terms and phrases that they learned as they were trained. Stylistically, AI chatbots often overlook that repetition.

Look for style (or lack thereof). Today's chatbots are no Ernest Hemingway, Mark Twain, or Jane Austen. Chatbots lack style. The text they generate often feels canned and flat. Moreover, they tend to spit out strings of statements that have little consideration for how they flow together.

How to spot deepfake photos

Zoom in. A close look at deepfake photos often reveals inconsistencies and flat-out oddities. Consider this viral picture of the "Puffer Pope" that circulated recently. Several things point toward a bogus image.

Credit: CNN

Start with the image itself. The Pope in a monster puffer coat. Really? Next, look at the hands in the image. The right hand isn't fully formed. Many AI tools have a notoriously tough time with rendering fingers properly. Meanwhile, the left hand features some lighting and skin tones that look a bit unnatural. An even closer look shows that the crucifix worn by the Pope only has half a chain. Next, look at the face and the unusual shadows cast by the glasses he wears.

How to spot deepfake audio and video

Keep an eye on the speaker. A close look at who's doing the talking in a deepfake video can reveal if it's a fake. Subtle things reveal themselves. Is the speaker blinking too much? Too little? At all? How about their speech? Does it sync up with their mouth perfectly? These might be signs of a deepfake.

Watch how the speaker moves. Does the speaker hold still or move only slightly? This is a sign of lower-grade video deepfake technology. It has difficulty tracking movement. Another possible sign is if the speaker never moves their hand across their face. Once again, that might indicate the work of lesser AI tools. Such tools render the facial image on the hand, and many deepfakes take pains to keep the face clear to avoid that glitch.

Look at and listen to the context. If a speaker is in an open public space, does it sound like they're speaking in that environment? For example, if they're in a city park, can you hear birds? What about traffic noise? How about the murmurs of the crowd? If that's missing, or it feels like ambient sounds are piped in like the laugh track in an old sitcom, you might have a deepfake on your hands.

How does the speaker sound? In the case of audio-only deepfakes, today's AI tools work best when they're fed smaller chunks of text to create speech. They don't work as well with big blocks. This requires creators to stitch those chunks together. As a result, the cadence and flow might sound on the copy side. Also, you might not hear the speaker taking breaths, as normal speakers do.

How else can I spot a deepfake?

Even as companies and social media platforms begin to roll out deepfake detection technologies, one of the most powerful detection tools you have right now is yourself. Essentially, be your own lie detector.

Three steps can help you get down to the truth:

1. **Check your emotions.**

   Has a news story you've read or watched ever made you shake your fist at the screen or want to clap and cheer? How about something that made you fearful or simply laugh? Bits of content that evoke strong emotional responses tend to spread quickly, whether they're articles, a post, or a video. That's a ready sign that what you're seeing might be a deepfake.

2. **Evaluate the claim.**

   Does what you're seeing or hearing seem too bizarre to be real? Too good to be true? Today, "Don't believe everything you read on the internet," now includes "Don't believe everything you see on the internet." If a fake news story is claiming to be real, search for the headline elsewhere. If it's truly noteworthy, other known and reputable sites will report on the event — and will have done their own fact-checking. Certainly, some legitimate articles will generate a response as well, yet it's a good habit to do a quick fact check and confirm what you've read.

3. **SIFT for the facts.**

   Deepfakes look and sound increasingly real — particularly as people quickly scroll through their newsfeeds without a great deal of scrutiny.

   Yet, that's exactly what it takes to spot a deepfake. Scrutiny. Everything online calls for it now, and a technique called **SIFT** can help. It stands for:

   Stop. Don't take things at immediate face value, especially if you see something that stirs up a strong reaction.

   Investigate the source. Who posted it? Is it from a long-standing and reputable source?

   Find better coverage. Confirm that what you're seeing is true with a reputable fact-checking resource.

   Trace the media to the original context. Deepfakes manipulate more than images, audio, and video. They manipulate facts. When you come across a claim, check its source and see if the claim matches or manipulates it.

While doing that on your own can take a bit of time and effort, fact-checking organizations can do that work for you. Later in the toolkit, we have an entire section dedicated to some of the most respected fact checkers and how they can help you uncover the truth.

Did you spot a sketchy deepfake? Turn it in!

The fight against deepfakes calls for a collective response, and you can pitch in. If you spot a fake, you can turn it in. If you think you've come across deepfake video or audio, report it to the platform where it was posted. (Social media sites have these reporting mechanisms built in.) Next, copy the deepfake's address by right-clicking on the video or file. Then, share it with us.

We've created the McAfee Smart AI Hub, a place where you can learn about the latest AI threats — and join the fight against them by turning in the deepfakes you find online.

We'll take the deepfakes that you and others have submitted and use them to train our new McAfee Deepfake Detector. It's a brand-new feature from us, automatically alerts you within seconds if it detects AI-generated video in videos.

Social media — how AI deepfakes get around

Social media has toppled direct access to news. Yet malicious AI deepfakes undermine it as a trustworthy source.

Pew Research found that about a third of Americans say they regularly get their news from Facebook. And it further found that nearly 1 in 4 say they regularly get it from YouTube. Moreover, global research from Reuters uncovered that more people primarily get their news from social media (30%) rather than from an established news site or app (22%). This marks the first time that social media has toppled direct access to news.

With so many Americans relying on social media as a primary news source, one finding from our recent research shows why malicious AI is so concerning. We found that that nearly 1 In 4 Americans (23%) said they recently came across a political deepfake they later discovered to be fake. The actual number of people exposed to political and other deepfakes is expected to be much higher given many Americans are not able to decipher what is real versus fake, thanks to the sophistication of AI technologies.

Fake news crops up in plenty of places on social media. And it has for some time now. In years past, it took the form of misleading posts, image captions, quotes, and infographics. Now with the advent of AI, we see fake news taken to new levels of deception:

- Deepfake videos that mimic the looks and parrot the words of well-known public figures.
- AI-generated voice clones that sound spooky close to the voices they mimic.
- And as mentioned earlier, entire news websites generated by AI, rife with bogus stories and imagery.

With AI, bad actors foster uncertainty and spread disinformation largely on social media. Whether or not their phony story gets picked up and viewed firsthand doesn't matter to these bad actors. Their aim is to get some manner of disinformation out into the social media ecosystem — and then let others do the work for them.

How do election deepfakes spread?

Outside of texts and robocalls, deepfakes primarily spread on social media. And there, the disinformation they carry spreads quickly.

Researchers found that disinformation travels deeper and more broadly, reaches more people, and goes more viral than any other category of false info.

According to the research findings published in Science,

> *"We found that false news was more novel than true news, which suggests that people were more likely to share novel information … Contrary to conventional wisdom, robots accelerated the spread of true and false news at the same rate, implying that false news spreads more than the truth because humans, not robots, are more likely to spread it."*

Thus, bad actors pump false info into social media channels and let people spread it by way of shares, retweets, and the like.

And convincing deepfakes have only made it easier for bad actors to spread disinformation.

Stopping the spread of disinformation and malicious deepfakes

As you can quickly surmise, that comes down to us. Collectively. The fewer people who like and share disinformation and malicious deepfakes, the quicker they'll die off.

A few steps can help you do your part in curbing disinformation and malicious deepfakes...

Verify, then share.

This all starts by ensuring what you're sharing is indeed the truth. Doubling back and doing some quick fact checking (like we'll cover next) can help you make sure that you're passing along the truth. Once more, bad actors entirely rely on just how readily people can share and amplify content on social media. The platforms are built for it. Stop and verify the truth of the post before you share.

Flag falsehoods.

If you strongly suspect that something in your feed is a malicious deepfake, flag it. As we mentioned earlier, social media platforms have reporting mechanisms built in, which typically include a reason for flagging the content.

Get more private on social media.

By setting your profiles to private or "friends only," you can limit your exposure to new and unknown accounts. That includes bogus accounts that spread disinformation and malicious deepfakes via social media direct messages. A quick way you can lock down your privacy on social media is with our Social Privacy Manager. It adjusts more than 100 privacy settings across your social media accounts in just a few clicks, so your personal info is only visible to the people you want to share it with.

What can you trust? The best fact-checking resources

Malicious deepfakes count on people taking the bait.

The people who create them rely on others to take the lies they push at face value — and immediately react to the feelings they stir up. Outrage. Anger. Dark joy. Without pause. Without consideration.

That's understandable. Consider how much content crosses our gaze over the course of a day. Plenty. To the point that we get a bit desensitized to it. We're inclined to think all of it's legitimate, even when it's most certainly not.

Earlier, we talked about checking your emotions when coming across content like this. If something acts as a trigger, that's a sure-fire sign you should follow up and determine if what you're seeing or hearing is truly real.

That's where fact-checking resources come in.

Fact-checking resources for the election

Come across something questionable? You can turn to one of the several fact-checking organizations and media outlets that make it their business to separate fact from fiction. Each day, they assess the latest claims making their way across the internet — and then determine if they're true, false, or somewhere in between.

- Politifact.com
- Snopes.com
- FactCheck.org
- Reuters Fact Check
- AP Fact Check

In all, these resources provide voters with a quick and ready resource to determine if what they've seen or heard is true.

Additionally, for a list of reputable information sources, along with the reasons they're reputable, check out "10 Journalism Brands Where You Find Real Facts Rather Than Alternative Facts." It's published by Forbes and authored by an associate professor at The King's College in New York City. It certainly isn't the end-all, be-all of lists, yet it provides you with a good starting point.

Stay safe from election scams too

Plenty of what we discussed so far can lure you into sketchy corners of the internet. Places where malware and phishing sites take root.

You might have seen a few examples of that already. Scammers blast smartphones with texts about surveys, polls, and fundraising. They'll pose as celebrities, political figures, and organizations — all to steal personal and financial info.

Taken together, messages like these are classic phishing attacks. They just take on an election-year façade.

Protect yourself from phishing attacks and other election scams

Phishing attacks and how to spot them are topics in their own right. For a comprehensive look, grab a free copy of our Phishing Scam Protection Guide. Also, stay tuned to our blog for the latest word on attacks and how to protect yourself from them as well.

Broadly speaking, one of the best protections against phishing scams is you. You can act as your own lie detector by taking a moment and considering the message.

For example, is a Hollywood A-Lister really offering up a 600% match for your donation to a candidate? Highly unlikely. Also, does that politician from a faraway state really want you to participate in some sort of poll? Again, doubtful. These are two common ruses around election time, ones you can spot with a bit of scrutiny.

This stands as an excellent reminder to participate in the election process safely. Donate directly to candidates from their official campaign sites. Only take part in polls from known and reputable sources like Pew Research and the pollsters listed on 538's Pollster Ratings — which are ranked for errors, bias, and transparency.

As these and other scams circulate during the election year, consider using comprehensive online protection software to keep safe. In addition to several features that protect your devices, privacy, and identity, they can warn you of scams too. In particular, our online protection software can help you in a couple of ways:

- Our new Text Scam Detector automatically detects scams by scanning links in your text messages. If you accidentally click? Don't worry, we can block risky sites if you click on a suspicious link in texts.
- Additionally, McAfee Web Protection does the same while you browse, preventing you from clicking bad links and sketchy downloads.

And as we mentioned above, a visit to the McAfee Smart AI Hub can help you stay on top of the latest AI threats — and gives you another place where you can turn in the deepfakes you find online.

Don’t let disinformation and malicious AI steal your vote

Rumors, hoaxes, and false information crop up online during election time. And many of them aim to deprive people of their vote.

Some of it falls under the category of misinformation, such as the apparent case with texts sent to voters in the New York state primary.

Prior to the June 25 polling date, a grassroots voter organization sent out texts to select voters. The texts contained the image and address of their polling location, as the organization stated, “[B]ecause we believe visualizing one’s polling place increases the likelihood of voting.” However, some polling information was wrong. The company apologized and sent out correction texts.

Other instances are not the product of a mistake. In the 2016 General Election, rumors circulated across social media that voters could cast their ballots by text. That, of course, was entirely untrue. In 2018, bogus text messages went out to absentee ballot voters in Indiana. The text alerted them that their votes hadn’t been registered. Once again, this was untrue.

Now, AI tools have made it easier to concoct convincing-looking texts, messages and posts that will likewise provide false voter information. This election, it’ll be more important than ever to rely on trusted sources when it comes time to vote and when following results.

Ways you can protect your vote.

Stick with trusted voter resources.

Go straight to the source for voting info, like how to register, when, where, and how to vote — along with how to confirm your voting registration status. You can find all this info and far more with a visit to https://www.usa.gov/voting-and-elections.

You can find another excellent resource for voters at https://www.vote411.org, which is made possible by the League of Women Voters. Particularly helpful is the personalized voting info it offers. By entering your address, you can:

- See what's on your ballot.
- Check your voter registration.
- Find your polling place.
- Discover upcoming debates in your area.

If you have further questions, contact your state, territory, or local election office. Once again, usa.gov offers a quick way to get that info at https://www.usa.gov/state-election-office.

Trust known and long-standing news organizations for info.

Bad actors might attempt to sway voters on social media with posts of early polling results, claims of voting machine errors, and polling place closures. Always verify such claims with trusted and reputable news outlets. Likewise, you can use the fact-checking resources found earlier in this toolkit.

Report disinformation and malicious deepfakes.

If you think you spotted disinformation or a malicious deepfake on social media, flag it with the platform's tools for reporting questionable posts. Further, you can protect other people's vote by reporting potential election crimes. That includes disinformation about the manner, time, or place of voting. You can use the same resources listed above to contact state and local government election officials.

Your vote counts. That's why others want to compromise it.

With the malicious use of AI tools, bad actors put an entirely new spin on disinformation.

AI makes disinformation look and sound far more credible than ever. And bad actors can produce it on a tremendous scale, thanks to the ease and speed of AI tools. In an election year, that calls for more scrutiny on our collective part — in addition to more tools that help detect deepfakes.

Put in the plainest terms, voters face a host of people, groups, and politically motivated interests that make malicious use of AI. They want to alter thinking. And ultimately, alter votes.

Even as AI content-creating tools advance, other AI detection tools advance as well. People will increasingly have more and more detection tools available to them, which will help combat malicious deepfakes. Still, as that AI arms race continues, one detection tool will stand out. People themselves.

The more we, collectively, can use our own personal lie detectors and turn to the several fact-checking resources freely available, the safer we'll find ourselves from malicious deepfakes and the disinformation they push.

All this requires some extra time, and extra effort, from each of us. Yet that time and effort ensures we're spreading the truth and securing the integrity of our vote.

About McAfee

McAfee is a worldwide leader in online protection. We're focused on protecting people, not devices. Our solutions adapt to our customers' needs and empower them to confidently experience life online through integrated, easy-to-use solutions.

www.mcafee.com

For more information about online protection, visit us at mcafee.com/blogs

i. https://www.politifact.com/article/2022/feb/21/when-george-washington-fought-misinformation/
ii. https://www.newsguardtech.com/press/newsguard-launches-2024-election-misinformation-tracking-center-rolls-out-new-election-safety-assurance-package-for-brand-advertising/
iii. https://www.bloomberg.com/news/newsletters/2024-05-17/ai-fake-bylines-on-news-site-raise-questions-of-credibility-for-journalists
iv. Ibid.
v. https://www.nature.com/articles/s41586-024-07404-1
vi. https://www.merriam-webster.com/dictionary/deepfake
vii. https://www.bbc.com/news/world-us-canada-65069316
viii. https://www.cbsnews.com/chicago/news/vallas-campaign-deepfake-video/
ix. https://www.pewresearch.org/journalism/fact-sheet/social-media-and-news-fact-sheet/
x. https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2023/dnr-executive-summary
xi. https://www.mcafee.com/nb-no/consumer-corporate/newsroom/press-releases/press-release.html?news_id=4698979d-2a55-4f71-84be-c04b41fc7bdc&q=awtana
xii. https://science.sciencemag.org/content/359/6380/1146
xiii. https://www.forbes.com/sites/berlinschoolofcreativeleadership/2017/02/01/10-journalism-brands-where-you-will-find-real-facts-rather-than-alternative-facts
xiv. https://projects.fivethirtyeight.com/pollster-ratings/
xv. https://www.voteftw.com/statement-regarding-new-york-polling-location-images
xvi. https://www.nytimes.com/2018/11/05/us/politics/misinformation-election-day.html
xvii. Ibid.

6220 America Center Drive
San Jose, CA 95002
888.847.8766
www.mcafee.com

C405-24-mcafee-ai-election-toolkit-protection-guide_1024 OCTOBER 2024