

AI Bias: A Practical Guide for Everyday Users



Photo: Pexels | Edmond Dantès

Developed by the She Leads AI Community of Practice
for AI Governance | September 2026

Table of Contents

Note: This document is optimized for mobile viewing.

For the complete interactive experience, view on a desktop computer or within a dedicated PDF reader app. Interactive performance on mobile browsers or preview apps may vary.

PRE-ASSESSMENT: AI’s Impact on Your Everyday Choices 3

Twelve questions and five minutes reveal where you are already strong and where this guide will best help you.

Why Read the Entire Guide? 7

PART 1: What AI Bias Is and Why It Matters 8

Where Does Algorithmic Bias Come From? 9

Two Mechanisms That Drive Most of the Bias You Encounter 10

How Much Algorithmic Bias Actually Touches Your Life 11

Where Algorithmic Bias Gets In 12

The Myth of Neutral 13

Why “Fair” Is Harder Than It Sounds 14

What About the Law? 15

PART 2: Where AI Bias Shows Up 16

Real Cases 17

Your Résumé 18

Your Face 19

Your Health 20

Your Housing 21

Your Language 22

PART 3: Build Your Own Practice 23

When You Are Using AI 24

What You Check Before You Use AI: 25

When AI Is Being Used On You 26

The Twenty-Minute Personal Audit 27

When Something Goes Wrong 28

Resource Links 30



PRE-ASSESSMENT:

AI’s Impact on Your Everyday Choices

Before You Begin

One truth worth naming upfront: **you are already using AI.**

Not maybe. Already. Just as importantly, AI is being used on you in spaces where you never gave consent.

Examples of How You May Be Using AI in Your Daily Life:	Ways AI May Be Used on You Without Your Explicit Consent:
■ Smartphones: Powering features like voice assistants, predictive text, and battery optimization ■ Navigation Apps: Analyzing real-time traffic to calculate the fastest routes ■ Music Streaming Services: Curating personalized playlists based on your listening habits ■ Camera Systems: Automatically detecting faces and adjusting focus or lighting ■ Fraud Detection: Triggering bank alerts for unusual or suspicious account charges ■ Email Clients: Sorting spam and prioritizing important messages in your inbox	■ Retail and Surveillance Cameras: Using facial recognition or behavioral tracking in public spaces and stores ■ License Plate Readers: Scanning vehicle plates on public roads to map your location and movement ■ Cross-Application Data Processing: Feeding your submitted application details into external AI models for secondary, unrelated evaluations

None of this came with a warning. You never opted in, and no one asked for your permission. Rather than asking whether AI presence exists in your daily routine, this pre-assessment measures your current awareness of it.



PART A: The AI You Use

How carefully you handle it

1. When an AI tool gives you an important answer, do you check it against another source?	☐ Always ☐ Sometimes ☐ Rarely or Never
2. Do you know which of the apps and services you already use have added AI features?	☐ Most of them ☐ Some ☐ I haven't looked
3. When AI output sounds confident and polished, do you question it as closely as you would question a person?	☐ Usually ☐ Sometimes ☐ Rarely
4. Do you know what the AI tools you use do with the information you put into them?	☐ I have checked ☐ Roughly ☐ No

PART A RESULTS

If You Answered	Assessment	Recommendation
“Rarely or Never” / “I haven’t looked” to Questions 1 or 2	You may be relying on automated outputs without realizing where AI is active or checking for accuracy.	Start with Part Three: Use our 10-minute audit to map which AI tools you rely on daily.
“Sometimes” to Question 3	Polished, authoritative AI phrasing often encourages users to drop their guard and accept errors as facts.	See Part Three: Learn why fluency is not accuracy, and the one habit that catches most automated errors.
“No idea” / “Roughly” to question 4	Entering sensitive data into unchecked AI tools risks exposing personal or organizational information.	Read Part Two: Learn how training data and privacy settings handle your inputs.

PART B: The AI Being Used On You

How much of it you can see

1. If you applied for a job tomorrow, would you know if software screened your application first?	☐ Yes ☐ Unsure ☐ No
2. If you were turned down for housing or credit, would you know which report the decision was based on?	☐ Yes ☐ Vaguely ☐ No
3. At a clinic or hospital, would you know whether an automated tool helped decide your care or your eligibility?	☐ Yes ☐ Maybe ☐ No
4. Are you aware of where AI collects data on you in public (cameras, license readers, stores)?	☐ I have a good idea ☐ I have some idea ☐ I haven't thought about it

PART A RESULTS

If You Answered	Assessment	Recommendation
“No” / “Unsure” to Questions 1 or 2	Automated gatekeeping in hiring and housing operates invisibly without explicit disclosure.	Read Part Two: See how screening algorithms and proxy variables affect daily opportunities.
“No” / “Maybe” to Question 3	Healthcare algorithms often use proxy data that can quietly limit access to specialized care.	Review Case 3: Understand how healthcare allocation tools introduce systemic bias.
“Haven’t thought about it” to question 4	Public data collection occurs constantly without opt-in prompts or clear notifications.	Read Part One: Explore the full breakdown of invisible surveillance in everyday life.

PART C: What You Can Do About It

Do you know your options?

1. If an automated decision about you felt wrong, would you know what questions to ask?	☐ Yes ☐ Somewhat ☐ No
2. Do you know if you have the right to request a human review of an automated choice?	☐ Yes ☐ I think so ☐ No
3. Do you know where to file a complaint if an automated decision harmed you?	☐ Yes ☐ 1 or 2 places ☐ No
4. Have you ever requested a copy of a background, credit, or tenant report used on you?	☐ Yes ☐ Once ☐ Never

PART C RESULTS

If You Answered	Assessment	Recommendation
“No” / “Somewhat” to Questions 1 or 2	Most people aren’t aware that automated decisions can often be questioned or formally appealed.	Read Case 4: Learn your legal rights to request human review on automated decisions.
“No” to Question 3	Lacking a reporting path makes it difficult to challenge unfair algorithmic outcomes.	See Part Three: Access our quick-reference list of regulatory oversight bodies and filing options.
“Never” to Question 4	Screening reports frequently contain outdated or inaccurate data that algorithms accept as fact.	See Part Three: Follow our step-by-step checklist to request and audit your personal screening reports.

Why Read the Entire Guide?

High scores still leave blind spots. By mastering every section, you don't just protect yourself, you become a resource for colleagues, family, and community members who scored lower.

Most people only discover how automated decisions work after being harmed by one. Sharing this knowledge is how we build collective leverage.

Where You May Find the Most Value First

Use the table below to find your highest-priority sections.

If You Are	You May Want to START HERE ↓
Job Hunting	▪ Case 1: Your Résumé ▪ then Part Three: Hiring
Looking for Housing or Applying For Credit	▪ Case 4: Your Housing ▪ then Part Three: Housing and Credit
Managing a Health Condition or Caring for Someone Who Is	▪ Case 3: Your Health ▪ then Part Three: Healthcare
Using AI to Write, Apply, or Create	▪ Case 5: Your Language ▪ then Part Three: When You Are Using AI
Concerned About Surveillance Or Policing	▪ Case 2: Your Face
Just Wanting to Understand AI	▪ Case 1: Your Résumé ▪ Part One, then any case that interests you



PART 1:

What AI Bias Is and Why It Matters

You don't need a technical background to understand this toolkit. You don't need to know how AI is built, and nothing in this resource will require that you do.

If you read nothing else, remember this: **AI can be accurate and biased at the same time.** Those are not opposite concepts. A system can work exactly as designed and still produce results that are unfair to you. That single idea is what most people get wrong, and it is the reason this work exists.

Algorithmic Bias IS	Algorithmic Bias IS NOT
When an AI system produces results that are unfair to certain groups of people, usually people who were already being treated unfairly before the software showed up	■ A software bug: A bug is a mistake in execution. Algorithmic bias is a system functioning as designed—faithfully scaling the human bias embedded in its data. ■ Someone being malicious: In nearly every documented case, nobody intended to discriminate against anyone. There is no villain to find and no one to confront. ■ Rare or historical: It is present in tools in use right now, including ones you used this week.

The Main Takeaway:

These systems train on human decisions. Because human decisions carry historical inequity, the software internalizes it—encoding bias into a process that operates faster, reaches wider, and hides behind machine neutrality.



Where Does Algorithmic Bias Come From?

The U.S. government’s standards agency, the National Institute of Standards and Technology, published guidance that sorts the cause of AI bias into three categories.

Knowing the three types is useful as they require different responses, and people often mistake one category for another. (Schwartz et al., 2022)



Kind of Bias	What It Means In Plain Language
Systemic 	Inherited Bias: Algorithmic unfairness isn’t intentional; software simply inherits decades of institutional bias embedded in historical data.
Statistical 	Flawed Data: Skewed training data underrepresents minority groups, making systems fundamentally less reliable for them.
Human 	Unconscious Blind Spots: Designers, developers, and buyers unconsciously code their own assumptions into the technology.

Compounding Impact: These issues aren’t isolated bugs requiring standalone patches—they layer on top of one another. A tool can, and usually does, carry all three issues at once.



TWO Mechanisms That Drive Most of the Bias You Encounter

Every example in this guide traces back to one of two mechanisms:

1

Training Data Bias

AI learns from examples. If the training data lacks representation, the system fails those left out. In a landmark 2018 study, researchers Joy Buolamwini and Timnit Gebru tested commercial facial recognition software:

- Lighter-skinned men: Error rate under **1%**
 - Darker-skinned women: Error rate reached **34.7%**
- (Buolamwini & Gebru, 2018)

The software was **30 times more likely to fail on darker-skinned women** simply because it was trained mostly on photos of lighter-skinned men. While IBM later exited facial recognition in 2020 and other vendors updated their models, the core driver remains unchanged: an AI system only gets good at what it is shown most.

2

Proxy Variables

A proxy is a neutral data point—like zip code, name, or spending history—that quietly correlates with protected traits like race or income. Because historical discrimination is baked into demographic data, an algorithm using these stand-ins produces biased outcomes without ever explicitly referencing race:

- **The Case:** A 2019 study analyzed a nationwide healthcare algorithm that used past spending to predict who needed extra care.
- **The Flaw:** Because less money had historically been spent on Black patients with the same conditions, the tool concluded they were healthier and recommended less care.
- **The Impact:** Correcting this bias would have raised the share of Black patients flagged for extra help from **17.7%** to **46.5%** (Obermeyer et al., 2019).

The Takeaway: When an AI uses one data point to represent another, ask what history is buried inside that variable.



How Much Algorithmic Bias Actually Touches Your Life

The examples in the previous sections prove that bias exists. This section indicates how much it actually impacts you.



Photo: Pexels | Kampus Production

If you have applied for a job recently, an AI very likely touched your application before a person did.

- As of late 2024, **51% of companies** were using AI somewhere in hiring, with **68%** projected by the end of 2025.
- Of the companies using it, **67%** acknowledged the tool produces biased recommendations to some degree. Not *might*. Does. (ResumeBuilder.com, 2024)

That is employers reporting on themselves, against their own interest, which is part of why it is worth taking seriously.

If that makes you uneasy, you are not alone, and it is not just a feeling.

- **71% of Americans** oppose letting AI make a final hiring decision.
- Even earlier in the process, just reviewing applications, **41%** oppose AI being involved at all. Most people's instincts here are already correct. (Pew Research Center, 2023)

The Point Isn't to Panic About AI

The point is to abandon the myth that automated systems are inherently neutral. Now that you understand the mechanics; you realize that default trust is no longer an option.



Where Algorithmic Bias Gets In

Bias doesn’t simply enter at one point that can be easily guarded. Bias can enter at every stage of the process, including after a product is released.



Entry Point	Specifics
When The Data Is Collected	Who got counted, who got missed
When The System Is Designed	What the builders decided to measure, and what they decided to ignore
When It’s Trained	Which patterns it learns to repeat
When It’s Tested	Whether anyone checked how it performs across different groups, or only checked the average
After It’s Deployed	Bias can develop over time in a system that launched clean

That Last Part Surprises People

Even a fair tool drifts into bias as the world evolves, making continuous monitoring essential over one-time certification.



The Myth of Neutral

There's a damaging belief that runs underneath most of the harm here: the computer is objective because it works with numbers.



Data is never neutral; human choices define what is counted, what is ignored, and how success is measured long before an algorithm executes.

There is no way to guarantee any algorithm is free of bias. Bias is inevitable without deliberate, active maintenance; it never just happens by default. (Schwartz et al., 2022)

The greatest danger isn't just that a machine is wrong, but that automation grants its errors an aura of mathematical truth, silencing the scrutiny a human judgment would normally invite.



Why “Fair” Is Harder Than It Sounds

Fairness is not one simple concept with an easy answer.



There are several reasonable definitions for fairness. Mathematicians have proven that in most real situations, you cannot satisfy all of them at once. Making a system fairer by one definition can make it less fair by another.

(Chouldechova, 2017; Kleinberg et al., 2017)

This is not a reason to give up, and it is not a loophole for companies to hide behind. It is motivation to ask sharper questions.

Instead of asking, “Is this tool fair?” as that has no single answer, ask:

- Fair in what sense?
- Who decided that definition?
- Who carries the cost of that choice?

Those are questions that a company has to actually answer.



What About the Law?

Civil rights laws already forbid AI from discriminating, but specific rules for artificial intelligence are still in flux.

Some states are passing stricter protections while federal courts and rollbacks pull the other way, making the legal landscape unpredictable.

Colorado offers the clearest example.

- In 2024 the state passed the first comprehensive AI law in the country, requiring companies using AI in consequential decisions to take reasonable care against algorithmic discrimination. Its start date was pushed back twice.
- In May 2026 it was replaced entirely with a narrower law focused on disclosure rather than a duty of care, taking effect January 1, 2027.

(Colorado General Assembly, 2024, 2025, 2026)

The practical point is not the statute number. It is this: **do not assume someone else has already checked.**

Legal protections against algorithmic bias remain sparse and fragmented across states. While individual awareness can't replace comprehensive regulation, self-vigilance is currently your primary line of defense.



Photo: Pexels | Pavel Danilyuk



PART 2:

Where AI Bias Shows Up

(and What to Do About It)

Common Bias Types You May Encounter

Bias Signal	Where It Operates	Where You Run Into It
Your Name	Résumé screening, tenant checks, background reports	Names carry racial and gender associations, causing trained systems to repeat historical discrimination.
Your Address	Credit, insurance, service eligibility	Zip codes serve as a stand-in for race and class due to decades of discriminatory housing policy.
Your Language	Text screeners, AI tools, customer service bots	Dialect, formality, and non-native English phrasing alter how systems evaluate your competence or intent.
Your Face	Identity verification, surveillance, building access	Accuracy drops significantly based on skin tone and gender, increasing false flags or access denials.
Your Feed	Targeted job ads, housing listings, credit offers	Algorithms filter visible opportunities via behavioral profiling, rendering missed opportunities invisible.
Automation Bias	Human decision-making, workplace tool usage	Polished, confident AI outputs discourage double-checking, turning software errors into final decisions.

Automation Bias is different. While every other row represents something done to you by software, automation bias is a human habit. It is also the single easiest point of failure you can control today.



Real Cases

These are documented cases, linked to their original sources in the Resources section.

Read the ones closest to your life first. Each ends with what could have helped, because knowing what prevents harm matters as much as knowing how it happens.



CASE 1: Your Résumé

Applying for a job through any online system



CASE 2: Your Face

Identity verification, building access, policing



CASE 3: Your Health

Hospital and insurer systems deciding who gets extra care



CASE 4: Your Housing

Applying to rent property



CASE 5: Your Language

Anything that reads your words, including applications, essays, support tickets, and assessments



CASE 1:

Your Résumé

Setting: Applying for jobs through online systems

In 2024, University of Washington researchers conducted a massive AI résumé screening test, running over three million comparisons of **550+ real résumés against 500+ job listings** across three commercial language models. The only variable changed was the applicant’s name at the top.



- **The 2024 Study:** Across 3 million test comparisons, LLMs preferred white names over Black names **85% vs. 9%**, and male names over female names **52% vs. 11%**—never once favoring a Black male candidate over an identical white male candidate. (Wilson & Caliskan, 2024)
- **The 2018 Benchmark:** Amazon abandoned its internal AI recruiter after realizing it penalized the word “women’s” and downgraded graduates of women’s colleges. Trained on a decade of past hiring data, the system simply automated historical bias. (MIT Technology Review, 2018)

The Harm:

- A qualified person is filtered out before any human sees the application. There is no rejection letter to explain why. There is no way to know it happened, and therefore, nothing to appeal.
- The applicant simply never hears back, and concludes they were not good enough.

What Could Have Helped:

- Employers testing screening tools across names, and publishing what they find
- Disclosure at the point of application that automated screening is in use
- A stated route to request human review of an automated rejection
- Human review of the rejection pile, not only the shortlist. The discarded pile is where this damage lives.

What This Means For You:

Assume software screens your résumé first. Use exact keywords from the posting (matching is often literal) and keep your formatting simple. Most importantly, reach out directly to a recruiter, connection, or hiring manager whenever possible. Bypassing the automated filter entirely isn’t a shortcut; it’s your single best advantage.

Photo: Unsplash/Bram Naus



CASE 2:

Your Face

Setting: Identity verification, building access, policing

In 2020, Detroit police arrested Robert Williams on his front lawn in front of his family for a theft he didn’t commit.



- A facial recognition system incorrectly matched his driver’s license to blurry security footage, holding him for nearly **30 hours** before anyone manually verified the match.
- While his was the first publicly reported wrongful arrest caused by facial recognition, the ACLU has documented **fourteen known cases** in total—nearly all involving Black individuals. (American Civil Liberties Union, 2026)

The Harm:

- Beyond the trauma and arrest record, Williams’ case exposed a broader police pattern: treating algorithmic matches as automatic proof of guilt rather than an unverified investigative lead, while skipping basic checks like alibis or location verification.

What Could Have Helped:

- Treating facial recognition output strictly as a lead requiring independent evidence; never as sole grounds for arrest
- Checking physical descriptions, alibis, and location data before taking action
- Formally disclosing the use of facial recognition to defense teams

The 2024 Settlement: Detroit police established the nation’s strictest department limits: arrests or photo lineups based solely on facial recognition matches are now banned

What This Means For You:

If told you were identified by facial recognition, explicitly dispute the match and ask what independent evidence exists. If under arrest, request a lawyer immediately to challenge the identification.

Photo: Pexels/Matias Mango



CASE 3:

Your Health

Setting: Hospital and insurer systems deciding who gets extra care

An algorithm deployed across the U.S. healthcare system was tasked with identifying patients who needed specialized care.



- It used past healthcare spending as its stand-in measure for illness severity.
- Because less money was historically spent on Black patients with the same conditions, the algorithm interpreted that lower spending as better health, cutting the number of Black patients flagged for extra support by more than half. (Obermeyer et al., 2019)

The Harm:	What Could Have Helped:
■ Qualifying patients were denied critical care management and never knew an automated determination had even occurred. ■ They experienced it as routine medical care, illustrating what makes proxy bias so dangerous: from the inside, it is completely invisible.	■ Testing outcomes across racial and demographic groups prior to rollout, rather than relying on overall aggregate accuracy. ■ Measuring true clinical need (symptom severity, chronic conditions) rather than financial proxies like past spending. ■ Maintaining continuous internal monitoring. This flaw was only uncovered because independent external researchers investigated it.

What This Means For You:

Ask your healthcare provider if any automated tool was used to assess your risk, determine program eligibility, or prioritize your care. If a system’s conclusion contradicts your lived symptoms, state the discrepancy clearly. Request that your provider explicitly note your dispute in your chart. A formal entry in your medical record carries legal and clinical weight that verbal comments lack.

Photo: Pexels/kaboompics



CASE 4:

Your Housing

Setting: Applying to rent property

In 2022, the Consumer Financial Protection Bureau (CFPB) analyzed over **24,000 renter complaints**, revealing widespread failures in the tenant background check industry. More than **16,000 complaints** involved inaccurate reports, and **4,500** involved difficulties getting errors fixed.



- The primary driver? Screening companies frequently matched records using name alone—omitting Social Security numbers or dates of birth. (Consumer Financial Protection Bureau, 2022).
- Applicants were routinely denied housing for criminal records, evictions, or debts belonging to strangers with the same name.

The Harm:

- Prioritizing speed for landlords over accuracy for renters leads to lost application fees, wasted time, and missed housing
- Because common naming patterns track along racial and ethnic lines, these algorithmic mismatches disproportionately penalize communities of color. The CFPB has affirmed that name-only matching violates the Fair Credit Reporting Act (FCRA).

What Could Have Helped:

- Requiring full identifying criteria (SSN, DOB) rather than surface-level name matching
- Automatically purging sealed or legally outdated records
- Immediately identifying the screening agency to applicants upon denial

What This Means For You:

You are legally entitled to a free copy of your report and the name of the screening company if denied.

- Request your report and check for mismatched records or old entries.
- Send formal dispute letters to both the screening company and landlord; keep copies.
- If you have a common name, check your reports before applying. Fixing errors takes weeks.
- Submit a report at consumerfinance.gov; public complaints drive regulatory enforcement.

Photo: Pexels/Erik Mclean



CASE 5:

Your Language

Setting: Anything that reads your words, including applications, essays, support tickets, and assessments



A landmark 2024 Nature study tested how large language models respond to African American English (AAE).

- When given identical content varied only by dialect, models generated consistently negative covert judgments—assigning AAE speakers lower-prestige jobs and, in hypothetical criminal scenarios, recommending the death penalty more frequently. (Hofmann et al., 2024)
- The stereotypes expressed were more severe than any human racial stereotypes recorded in experimental literature.

The Harm:

- Because this bias is covert, asking a model direct questions about race yields polished, egalitarian responses.
- The prejudice operates silently underneath standard alignment guardrails, judging individuals not on what they say, but how they say it.

What Could Have Helped:

- Testing models against real-world language patterns rather than exclusively standard formal English
- Classifying linguistic and dialectal variation as an explicit equity vector alongside race and gender
- Prohibiting automated model recommendations from operating without human review in high-stakes decisions

What This Means For You:

The data your organization collects about the people you serve does not always stay with you. Who can access that data is part of your duty of care to your community.

Photo: Pexels/cottonbro studio



PART 3:

Build Your Own Practice

Reclaim Your Agency

These are two fundamentally different situations, and each requires a different response:

- **When you are using AI:** You control the inputs you enter and the outputs you accept.
- **When AI is being used on you:** You can't change the system, but you can control how you respond to it.





When You Are Using AI

What You Put In: the CLEAR Framework

Vague prompts produce biased defaults, because a system with no context fills the gap from its training data.

Adding specifics is the single most effective thing you can do. Below is an effective structure adapted from the CLEAR framework in the Nonprofit Toolkit.

Letter	Stands For	To Answer	Example
C	Context	Who this is for and what the situation is	“I need an image for a flyer for my kid’s school fundraiser. It is a pancake breakfast in the school gym.”
L	Lived Experience Awareness	Who is affected, and how they should be represented	“Show families from our school. A range of ages, skin tones, body types, and abilities. People eating and talking, not posed.”
E	Ethical Guardrails	What to avoid, including stereotypes, pity framing, and assumptions about income, family, or ability	“Avoid showing anyone as needy or as a helper. No assumptions about family structure. Not everyone is a two-parent family.”
A	Audience Sensitivity	Who will read or see this	“Parents and staff at the school. It will be printed and posted in the hallways.”
R	Response Style	Tone, length, format	“Warm, natural light, candid rather than staged. Horizontal, so it works as a header.”

CLEAR framework adapted from AI Colab Studio, 2026



For everyday content, you do not need the full framework each time.

The minimum that works: say who it is for, and say what to avoid.

- “Write a social media caption for a food pantry in [city] serving [community]. Use dignified language. No pity framing.”

That is enough to move the output away from defaults.

For Images Specifically:

- For images, be specific about who appears and how.
- “Diverse” produces vague results. Skin tone, age, body type, ability, and setting are all worth naming, and it is worth looking hard at what you get.



What You Check Before You Use AI:

- Does this describe real people accurately, or a default version of them?
- What assumptions did it make that I did not give it, about income, family, language, or ability?
- Would the person described recognize themselves in this?
- What is missing? Who is not here?

If something is off, re-prompt with more direction. And notice the pattern. A tool you keep correcting in the same direction is telling you something about its baseline.

The habit that matters most: when output sounds confident and polished, that is precisely when to slow down.

Fluency Is NOT Accuracy.

A biased answer from a person invites argument. The same answer from a machine tends to end the conversation. Do not let it.



When AI Is Being Used On You

You usually cannot change the system. You can change what happens around it, and these are questions you have standing to ask.



Hiring

- **Ask:** “Is automated screening used in this process, and at what stage?”
- **Ask:** “Can I request that a person review my application?”
- Reach a human where you can. Referral routes skip the filter.
- Some jurisdictions require employers to disclose automated employment decision tools, and in some cases publish bias audits. Check your state and city, since rules vary and frequently change.



Housing and Credit

- **Ask:** “Which company produced the report this decision was based on?” You have a legal right to know, and to a free copy.
- Request your reports before you need them, especially if your name is common.
- Dispute errors in writing, and keep records.
- File a CFPB complaint. It is free, and it is tracked.



Healthcare

- **Ask:** “Was any automated tool used to assess my risk, eligibility, or priority?”
- “What information was that based on?”
- If a system’s conclusion contradicts your experience of your own body, say so and ask for it in your chart.



Anywhere a Decision Affects You

- **Ask:** “Was this decision made or influenced by an automated system?”
- **Ask:** “What can I do to have it reviewed by a person?”
- Get the answer in writing where you can.

You will sometimes be told nobody knows. That answer is itself information, and worth documenting.



The Twenty-Minute Personal Audit

Once, then once a year.

You'll need a notepad or another way to record your answers.



Photo: Pexels | RDNE Stock Project

MINUTES 0:01–0:05 List the Decisions

- What consequential decisions were made about you in the past year, across jobs, housing, credit, insurance, healthcare, benefits, or schooling?
- Write them down.

MINUTES 0:06–0:10 Mark the Unexplained Decisions

- Which came back without a reason, or a reason that did not fit the facts?
- Those are your candidates.

MINUTES 0:11–0:15 Pull One Report

- Take the highest-stakes one and request the underlying record, such as a credit report, tenant screening report, or background check.
- Read it. Errors are common.

MINUTES 0:16–0:20 Fix Your Defaults

- Turn off ad personalization where offered.
- Review app permissions.
- Decide which AI tools you actually trust and for what purpose.
- Write down the one thing you will do differently.



When Something Goes Wrong



1. **Write it down now.** Date, organization, what was decided, what you were told, who you spoke to. Memory fades and details matter later.
2. **Ask whether an automated system** was involved, in writing if possible.
3. **Request human review explicitly.** Many organizations have a process they do not advertise.
4. **Get the underlying record** and dispute errors formally.
5. **Escalate.** The CFPB for credit, housing, and background reports. Your state attorney general or civil rights agency for discrimination. The EEOC for employment. Legal aid organizations handle these cases and are free.
6. **Report it.** The Algorithmic Justice League collects accounts of AI harm. Individual reports are how patterns become visible.



Keep It Current

Revisit this **once a year, or whenever your circumstances change**, such as a job search, a move, a new diagnosis, or a child starting school.

Ask Yourself:

- What new tools have I started using, and what do they do with what I give them?
- What consequential decisions are coming up, and where might automation sit inside them?
- Have the rules changed where I live?



A Final Note

You are not powerless here, and you are also not responsible for fixing this. Both are true.

The burden should not fall on individuals to detect and route around biased systems. It should fall on the people building and buying them. Until that changes, knowing what to ask is real protection, and asking is itself a form of pressure. Systems get audited because people notice. Laws get written because people complain. Every case in this guide became known because somebody refused to accept an unexplained answer.

You do not need to become an expert. You need one habit: when a decision about your life does not add up, ask whether software was involved, and ask who will review it.

That question is available to everyone. It requires no technical background. And it is the thing these systems are least prepared for.

***We must keep the information accessible,
or WE will be the bias.***



Resource Links

For those who want
to do further research



All sources were opened and checked against
the claims they support in August 2026.

American Civil Liberties Union. (n.d.).

Williams v. City of Detroit.

<https://www.aclu.org/cases/williams-v-city-of-detroit-face-recognition-false-arrest>

American Civil Liberties Union. (2026, April).

More than a dozen wrongful arrests due to police reliance on facial recognition technology.

<https://www.aclu.org/news/privacy-technology/more-than-a-dozen-wrongful-arrests-due-to-police-reliance-on-facial-recognition-technology>

Bertrand, M., & Mullainathan, S. (2004).

Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. American Economic Review, 94(4), 991–1013.

<https://www.aeaweb.org/articles?id=10.1257/0002828042002561>

Buolamwini, J., & Gebru, T. (2018).

Gender shades: Intersectional accuracy disparities in commercial gender classification. Proceedings of Machine Learning Research, 81, 1–15.

<https://proceedings.mlr.press/v81/buolamwini18a.html>

Chouldechova, A. (2017).

Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. Big Data, 5(2), 153–163.

<https://arxiv.org/pdf/1703.00056>

Colorado General Assembly. (2024). SB 24-205:

Consumer protections for artificial intelligence.

<https://leg.colorado.gov/bills/sb24-205>

Colorado General Assembly. (2025).

<https://leg.colorado.gov/bills/sb25b-004>

Colorado General Assembly. (2026).

SB 26-189: Automated Decision-Making Technology Act.

<https://leg.colorado.gov/bills/sb26-189>

**Consumer Financial Protection Bureau. (2022, November 15).**

CFPB reports highlight problems with tenant background checks.

<https://www.consumerfinance.gov/about-us/newsroom/cfpb-reports-highlight-problems-with-tenant-background-checks/>

Consumer Financial Protection Bureau. (n.d.).

Tenant background checks.

<https://www.consumerfinance.gov/rules-policy/tenant-background-checks/>

Hofmann, V., Kalluri, P. R., Jurafsky, D., & King, S. (2024).

AI generates covertly racist decisions about people based on their dialect.

Nature, 633(8028), 147–154.

<https://ora.ox.ac.uk/objects/uuid:4e0fe173-52f2-433d-8042-0ff941919401>

Kleinberg, J., Mullainathan, S., & Raghavan, M. (2017).

Inherent trade-offs in the fair determination of risk scores. Proceedings

of the 8th Innovations in Theoretical Computer Science Conference.

<https://drops.dagstuhl.de/storage/00lipics/lipics-vol067-itcs2017/LIPIcs.ITCS.2017.43/LIPIcs.ITCS.2017.43.pdf>

MIT Technology Review. (2018, October 10).

Amazon ditched AI recruitment software because it was biased against women.

<https://www.technologyreview.com/2018/10/10/139858/amazon-ditched-ai-recruitment-software-because-it-was-biased-against-women/>

Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019).

Dissecting racial bias in an algorithm used to manage the health of populations.

Science, 366(6464), 447–453.

<https://pubmed.ncbi.nlm.nih.gov/31649194/>

Pew Research Center. (2023, April 20).

Americans' views on use of AI in hiring.

<https://www.pewresearch.org/internet/2023/04/20/americans-views-on-use-of-ai-in-hiring/>

ResumeBuilder.com. (2024, October).

7 in 10 companies will use AI in the hiring process in 2025, despite most saying it's biased. Survey of 948 business leaders conducted via Pollfish.

<https://www.resumebuilder.com/7-in-10-companies-will-use-ai-in-the-hiring-process-in-2025-despite-most-saying-its-biased/>



Schwartz, R., Vassilev, A., Greene, K., Perine, L., Burt, A., & Hall, P. (2022).

Towards a standard for identifying and managing bias in artificial intelligence (NIST Special Publication 1270). National Institute of Standards and Technology.

<https://www.nist.gov/publications/towards-standard-identifying-and-managing-bias-artificial-intelligence>

Wilson, K., & Caliskan, A. (2024, October).

AI tools show biases in ranking job applicants’ names according to perceived race and gender. University of Washington. Presented at the AAAI/ACM Conference on AI, Ethics, and Society.

<https://www.washington.edu/news/2024/10/31/ai-bias-résumé-screening-race-gender/>

Organizations Doing This Work

Algorithmic Justice League. (n.d.).

<https://www.ajl.org>

AI Now Institute. (n.d.).

<https://ainowinstitute.org>

Developed by Jasmine Schwarz, Schwarz Strategic Services, for the She Leads AI AI Governance Community of Practice, Pod 2.

Format and structure adapted from AI Bias: A Practical Guide for Nonprofits by Kim Olpin, Executive Director, Jericho Road Pasadena, with Tomar Beh (AI Colab Studio). The pre-assessment model, the Think, Feel, Do framework, and the case format of ‘Harms, What Could Have Helped,’ and ‘What This Means for You’ are hers. The CLEAR prompting framework in Part Three is adapted from the same guide, which credits AI Colab Studio (2026).

Foundational research synthesized by Audrey Schaub, COO, Vativ.ai, for Pod 2.

Pod 2: Audrey Schaub (Pod Lead) · Carmela Pereira (Communications and Creative Direction) · Jasmine Schwarz (Artifact Lead) · Tomar Beh (Tech Steward) · Marcia Narine Weldon · Kim Olpin.

All sources verified August 2026. Last updated: September 2026.