

# Why Lab-Perfect Face Matching Fails in Real Stores



SAFR<sup>®</sup>  
from *realnetworks*

## The Wild Faces Reality - How 99.87% accuracy on uncooperative faces separates real prevention from expensive security theater.

Retail theft isn't just rising. It's exploding. Losses hit \$112 billion in 2022, with a 20% jump in just one year. Shoplifting alone is up 93% since the pandemic. But here's what doesn't always make the headlines: about 10% of offenders are responsible for nearly 70% of those losses. That changes the game. This isn't about catching everyone anymore. It's about catching the few who are doing the most damage.

Traditional security tools aren't built for that kind of precision. Surveillance cameras show you what already happened. Security teams are expected to remember faces from watchlists longer than your last grocery receipt. And repeat offenders? They walk right past it all.

So, yes, face matching sounds like the right move. But if your system only works under perfect conditions, what good is it? Real stores don't run on lab lighting and cooperative angles. They run on chaos. The kind of chaos researchers call "wild faces."

### What Makes a Face Wild

Think about someone walking into a store. Hoodie up. Sunglasses on. Maybe a mask. Eyes on their phone. They're not standing still, they're not facing the camera, and they're definitely not smiling for a picture. That's what your system has to recognize. Not a headshot. Not a passport photo. Real life.

The term "wild faces" comes from computer vision research, notably Professor Erik Learned-Miller's Labeled Faces in the Wild (LFW) dataset and testing protocols. It pushed the industry to test systems in the messy, unpredictable conditions where recognition actually matters.

Meanwhile, most algorithms are still trained and tested in ideal settings. Subjects look straight at the camera. Lighting is perfect. No movement. No obstructions. The results look great on paper, with accuracy above 99%. But bring that same system into a store and the performance drops. Sometimes dramatically. It's not uncommon to see real-world accuracy slip into the 70 to 80 percent range. That's the reality gap. And in this space, it's a gap that costs money, time, and trust.

The SAFR software was put through the University of MA's LFW benchmark dataset and under those conditions, SAFR's face matching algorithm scored 99.87% accuracy. Not just on one face, but consistently, across tough-to-capture, real-world inputs.

But high accuracy means nothing if it only works for some people. The facial recognition industry has been plagued by systems that perform well on certain demographics while failing others. A technology that's 99% accurate for one group but 85% accurate for another isn't just biased – it's broken. When accuracy varies by race or gender, you're not just dealing with a technical problem, you're facing a liability that can destroy your reputation and expose your organization to lawsuits.

### Reality Doesn't Grade on a Curve

That's why the National Institute of Standards and Technology (NIST) stepped in. They created tests that actually reflect operational conditions, not just lab setups.

SAFR ranked second globally for demographic consistency out of 103 NDAA compliant algorithms tested. That means it recognizes faces accurately, regardless of race, gender, or skin tone. It doesn't matter who's in the frame - performance stays stable. That kind of fairness isn't just ethical, it's good business. And while we're on the topic of performance, the SAFR model is lightweight. At just 172MB, it can run right on the camera. That eliminates the need for server infrastructure, cuts latency, and avoids compressing images for transmission. You keep your image quality and your edge processing speed.

## Accuracy Isn't Just a Stat

Let's do some quick math. An 85% accurate system sounds pretty solid, until you realize what that means in practice. That's 150 missed identifications for every 1,000 encounters with known offenders.

With SAFR at 99.87% accuracy, you're looking at fewer than two misses out of every thousand. That difference isn't minor. It defines whether your system is quietly failing or actually working. Every missed match gives repeat offenders another opportunity. When those few individuals are causing most of your shrink, every miss matters. Let enough of them through, and the whole investment in face matching starts to look like a bad bet.

Now think about false positives. If your system keeps flagging innocent shoppers, you don't just annoy customers. You break trust with your team. They start ignoring alerts. Eventually, no one believes what the system tells them. This is a common arc. A retailer installs a promising system. At first, the lab numbers generate excitement. Then come the false positives. Settings are adjusted to reduce them, which lowers sensitivity. More offenders get missed. Teams lose faith, and the tech fades into the background.

## What It Looks Like in the Real World

Let's talk about what actually happens on the floor. A shoplifter walks in wearing a mask and hat, eyes on their phone. They've already hit three stores this week. Most systems? They're going to miss the ID.

SAFR is trained for this. It matches based on partial faces, strange angles, and movement blur. It doesn't need a perfect shot. Now picture someone walking through a sunlit doorway. The backlight floods the scene, turning their face into a silhouette. Typical exposure settings can't keep up. SAFR's Face Prioritization Auto-Exposure adjusts for the face, not the entire scene. You get a clear shot where others would get nothing.

Or imagine a coordinated theft team entering all at once. One person blocks the camera. Another distracts the cashier. Others walk by unnoticed. SAFR processes up to twelve faces per second, which gives it enough speed to catch what slower systems miss.

## Why Size and Speed Matter

SAFR has a smaller model size than any other algorithm in its accuracy class. That makes it perfect for edge deployment. No servers. No slowdowns. No compression artifacts.

And speed? It's not just about how fast it runs, it's about how many chances it gets. Someone walking by might only be in frame for a second or two. If your system is slow, it might miss that window. SAFR doesn't. You get multiple frames. Multiple angles. More chances to catch the right person. More signal, less noise.

## You Can't Wing It

Good results require good setup. That means being deliberate.

You need:

- Cameras positioned to capture faces at entrances, not just aisles
- A lighting check so you're not sabotaging your own accuracy
- Real integration with security workflows
- Staff who understand how to respond to alerts
- Clear metrics to track what's working and what's not

This isn't a magic button. But used right, it's a powerful advantage.

## The Payoff

Retailers using SAFR often see a return on investment in under 90 days. That's not an exaggeration. It's because every accurate identification of a known offender interrupts a chain of future thefts.

You're not stopping one incident. You're removing a repeat threat.

This tech also makes your people more effective. Instead of guessing or scanning crowds with gut instinct, they get real-time intelligence. Fewer confrontations. Fewer misses. More prevention. And there's more to gain. Safer employees. Happier customers. Less reliance on bulky deterrents like locked cases that frustrate everyone.

## What's Coming

This is becoming the new standard. As more retailers invest in wild face-accurate systems, the playing field shifts. Criminals will find fewer easy targets. The smarter your systems get, the harder theirs have to work.

We're already seeing retailers link systems together, sharing intelligence on known offenders. When one store flags a repeat thief, others can spot them too. Organized retail crime doesn't stand a chance when stores start acting in sync. And the technology keeps evolving. Better performance in low light. Faster processing. Smarter integrations. It's not slowing down.

## The Bottom Line

Fancy facial recognition might look impressive in a demo. But if it only works when conditions are perfect, it's not prevention. It's theater.

SAFR delivers 99.87% accuracy under real-world conditions. That's NIST-tested, bias-reduced, and battle-ready. So here's the real question: when that same offender walks through your doors again, will your system catch them? Or will it just record another loss?

## Interested in more information?

Contact SAFR to get your pilot started today at [sales@safr.com](mailto:sales@safr.com) or visit us at [www.SAFR.com](http://www.SAFR.com) for more information.

### About SAFR:

Launched in 2018, SAFR is backed by over three decades of technology innovation from RealNetworks—the company that invented streaming and co-developed the Real-Time Streaming Protocol (RTSP), which remains foundational to today's video surveillance systems. That same pioneering spirit now powers a new era of ethical facial recognition, built to protect, perform, and earn trust.

Our mission is to help make the world safer through AI that's designed not just to work—but to work responsibly. Rooted in "AI for Good," SAFR delivers the industry's first unified facial recognition ecosystem—seamlessly spanning access control, surveillance, and mobile solutions—without compromising personal privacy or data integrity. From the team that revolutionized how the world connects, SAFR now redefines how the world protects.

Scan to check out our  
No-Nonsense FAQ on  
Why You Should Know  
SAFR Guard!

