Koncile

Bank statement forgery: how to detect AI fakes?

Insurers, banks and brokers

Koncile Report 2026



CHECKING
November 1, 2025 thr
Account nu

CUSTOMER SERVICE
Website: online banking portal
Phone: 1-800-555-0142
Monday to Saturday, 8 a.m. to 8 p.m. ET

ALEX MARTIN
128 SAMPLE AVENU
EXAMPLE CITY NY

YOUR BRANCH
125 Sample Avenue
Example City, NY 10000

ACCOUNT SUMMARY
Beginning balance on 11/01
Deposits and additions
Withdrawals and payments
Ending balance on 11/30

TRANSACTION DETAIL

DATE	DESCRIPTION	WITHDRAWALS	DEP
11/01	**ACH Payment** HOME MORTGAGE SERVICING	31.90	
11/02	**ACH Payment** CITY ELECTRIC UTILITY	139.50	
11/02	**ACH Payment** INTERNET SERVICE	107.61	
11/02	**ACH Payment** CITY ELECTRIC UTILITY	139.50	
11/03	**ACH Payment** RENTERS INSURANCE PREMIUM	107.61	
11/04	**Check** #1042	139.50	
11/04	**ACH Payment** CITY ELECTRIC UTILITY	75.00	
11/05	**Online Transfer** TO SAVINGS ...4827	13.00	
11/05	**ATM Withdrawal** MAIN ST BRANCH	107.61	
11/06	**Debit Card Purchase** BOOKSTORE CENTRAL	52.34	

A complete AI workflow to put in place to fight fraud

Editorial

AI takes fakes to a new scale

The situation

Every insurer has seen AI-made fakes come through. The question everyone is asking: are we equipped to detect them today?

Banks and brokers face the same situation with bank statements: **85.6%** of fraud and risk leaders surveyed in 2026 name them the document most vulnerable to manipulation[1].

More than nine in ten altered documents involve financial data[1], and 17% of French internet users say they have already falsified documents to open an account or get a loan[2]. Documents flagged as AI-generated were **four times more numerous** each month in May 2026 than in June 2025[1]. The French banking sector recorded **more than €1 billion** in losses in 2023, and insurance loses **€2.5 billion a year**[3].

85.6%
of fraud and risk leaders rate the bank statement more vulnerable than any other document

The wake-up call

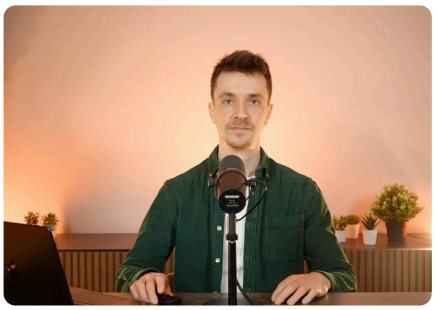
Forgery tools keep getting easier to use: the technology barrier is gone. AI is as good at altering images as PDFs, sometimes without leaving a visible trace.

And AI does not seem to be on our side today: it cannot detect its own forgeries. Likewise, the filters meant to stop the making of fakes are weak, even at LLM providers that publicly commit to preventing forgery[4]. Fraudsters get around them easily.

The upgrade we need

Anti-fraud systems must be updated to know, anticipate and detect how AI behaves. What we do know: AI is predictable. Very often, it leaves the same traces and uses the same tools.

One watchword for anti-fraud teams and the technology they deploy: think like fraudsters to detect them better. Based on our 2026 research, we wrote this white paper to describe the **methods most used** by fraudsters today, and the **4 solutions** that work.



Jules Ratier
Co-founder of Koncile

1. Inscribe, Document Fraud Report 2026. 2. IN BANQUE, Capgemini Invent and Tessi, 2024. 3. Observatoire Tessi 2025 and ALFA. 4. Koncile data, 2026.

Part 1

Know how fraudsters work to respond better

Drawing on large data collections and daily conversations with industry experts, we have gathered unique data on the main methods forgers use. Fraud itself is a constant, but the forms it takes keep changing. Hence the need to stay current on fraudsters' techniques and intentions.

Method 1 **1** 2 3 4

Image and PDF tampering tools, still everywhere

Fraudsters start from a real document and alter it, using Adobe tools or free tools found online. The result looks credible to the eye: one field is changed, sometimes in exactly the same font.

It is by far the most common method. In a set of 150,000 bank statements analyzed by Koncile, 2,130 were forged, and nearly one in two came out of a simple PDF editor. On photos and scans, the most-used technique is the clone stamp: a clean area of the document is copied and pasted over the figure to be hidden.

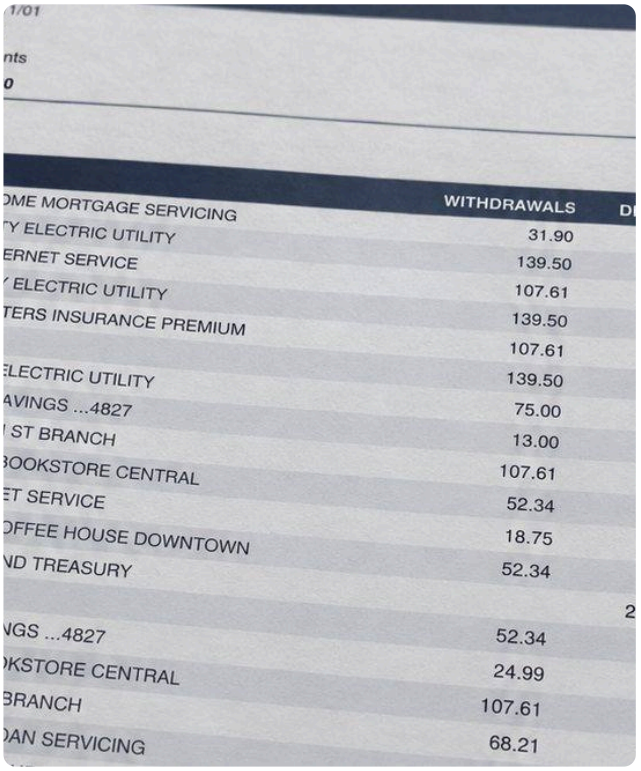
47.6%

of fake statements come out of a simple PDF editor

How to detect it

In a PDF, the edit often creates an overlay: instead of being erased, the original data is buried in the file and can still be read. The original file can be rebuilt and compared with the new one.

In an image, the stamp leaves near-identical blocks of pixels in two places, or noise and compression anomalies. The forged file is often printed and rescanned to erase these traces.



A statement photographed with a phone: every digit can be retouched like a pixel.

Method 2 1 **2** 3 4

AI alterations, the new reality of 2026

AI-made fakes are still a minority, but they are accelerating sharply this year. We have seen it with fake receipts and altered ID cards. Image models can regenerate a whole document, or rewrite just one area, an amount or a name, matching the font and grain of the rest.

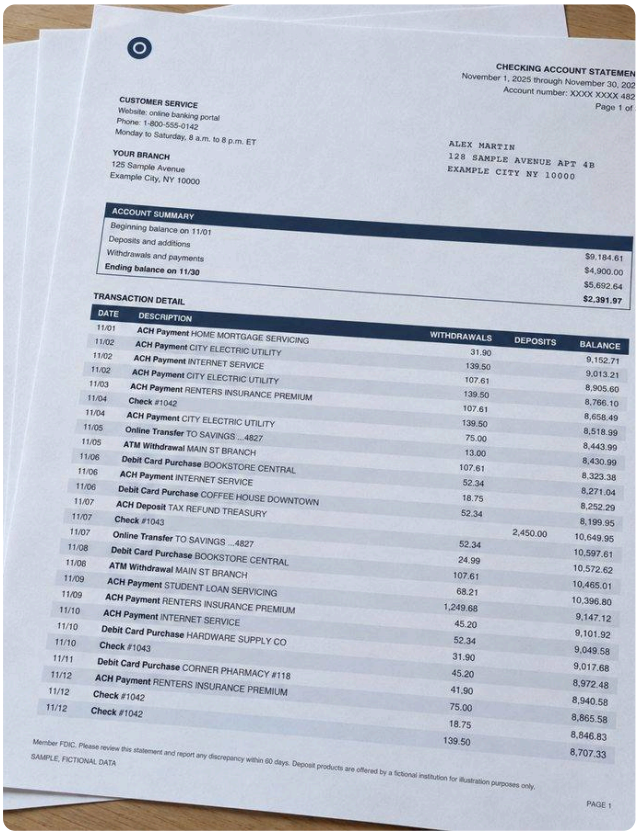
Assistants have filters: they usually refuse to produce a fake when clearly asked for one. The result is sometimes crude. But often, it is completely invisible to the naked eye. Our measurements cover 5,244 fakes from 56 image models.

$0.039
median cost of a fake produced by an image model

How to detect it

Two kinds of traces remain. A note in the metadata saying the image was generated by AI, which a fraudster can erase. And an invisible watermark in the pixels, which each lab applies in its own way, when it applies one.

AI is in fact rather poor at recognizing a generated image, even its own. The most reliable approach remains a model trained on real documents and their generated or altered versions. The edit of a single field, though, stays very hard to detect from the image alone: this is where contextual checks take over.



CHECKING ACCOUNT STATEMEN
November 1, 2025 through November 30, 202
Account number: XXXX XXXX 482
Page 1 of

CUSTOMER SERVICE
Website: online banking portal
Phone: 1-800-555-0142
Monday to Saturday, 8 a.m. to 8 p.m. ET

YOUR BRANCH
125 Sample Avenue
Example City, NY 10000

ALEX MARTIN
128 SAMPLE AVENUE APT 4B
EXAMPLE CITY NY 10000

ACCOUNT SUMMARY

Beginning balance on 11/01	$9,184.61
Deposits and additions	$4,900.00
Withdrawals and payments	$5,692.64
Ending balance on 11/30	**$2,391.97**

TRANSACTION DETAIL

DATE	DESCRIPTION	WITHDRAWALS	DEPOSITS	BALANCE
11/01	ACH Payment HOME MORTGAGE SERVICING	31.90		9,152.71
11/02	ACH Payment CITY ELECTRIC UTILITY	139.50		9,013.21
11/02	ACH Payment INTERNET SERVICE	107.61		8,905.60
11/03	ACH Payment CITY ELECTRIC UTILITY	139.50		8,766.10
11/04	ACH Payment RENTERS INSURANCE PREMIUM	107.61		8,658.49
11/04	Check #1042	139.50		8,518.99
11/05	Online Transfer TO SAVINGS ...4827	75.00		8,443.99
11/05	ATM Withdrawal MAIN ST BRANCH	13.00		8,430.99
11/06	Debit Card Purchase BOOKSTORE CENTRAL	107.61		8,323.38
11/06	ACH Payment INTERNET SERVICE	52.34		8,271.04
11/06	Debit Card Purchase COFFEE HOUSE DOWNTOWN	18.75		8,252.29
11/07	ACH Deposit TAX REFUND TREASURY	52.34		8,199.95
11/07	Check #1043		2,450.00	10,649.95
11/07	Online Transfer TO SAVINGS ...4827	52.34		10,597.61
11/08	Debit Card Purchase BOOKSTORE CENTRAL	24.99		10,572.62
11/08	ATM Withdrawal MAIN ST BRANCH	107.61		10,465.01
11/09	ACH Payment STUDENT LOAN SERVICING	68.21		10,396.80
11/09	ACH Payment RENTERS INSURANCE PREMIUM	1,249.68		9,147.12
11/10	ACH Payment INTERNET SERVICE	45.20		9,101.92
11/10	Debit Card Purchase HARDWARE SUPPLY CO	52.34		9,049.58
11/10	Check #1043	31.90		9,017.68
11/11	Debit Card Purchase CORNER PHARMACY #118	45.20		8,972.48
11/12	ACH Payment RENTERS INSURANCE PREMIUM	41.90		8,940.58
11/12	Check #1042	75.00		8,865.58
11/12	Check #1042	18.75		8,846.83
		139.50		8,707.33

Member FDIC. Please review this statement and report any discrepancy within 60 days. Deposit products are offered by a fictional institution for illustration purposes only.
SAMPLE, FICTIONAL DATA

PAGE 1

A sample statement, fictitious data, as an image model can produce it.

Method 3 1 2 **3** 4

AI agents and native files: traces are always left behind

An assistant that can run code opens the bank's PDF, replaces an amount and saves the file again. The result is visually perfect, but the tool leaves its signature.

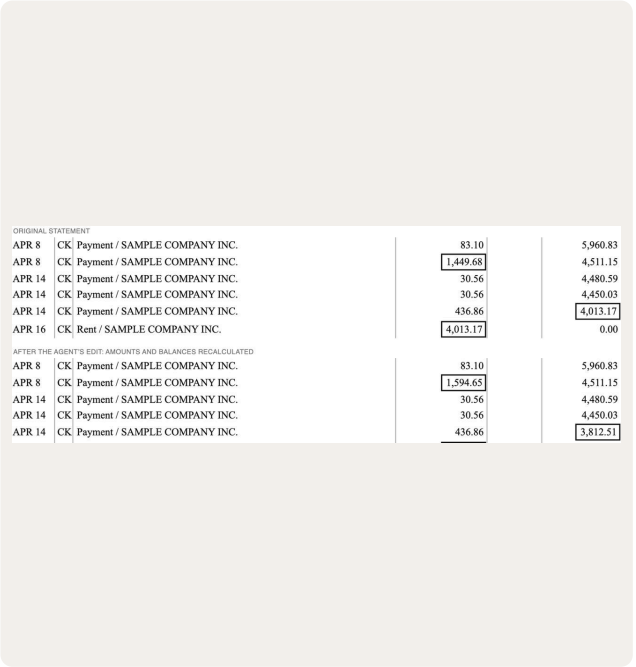
It draws nothing: it writes a small program that opens the file with a low-level PDF library, changes a value and saves it again. In our tests, most assistants refuse to falsify an amount, but 9 out of 14 models agree to "anonymize" a document.

72%
of agents use the same editing library, pikepdf, according to our study

How to detect it

No bank issues its statements with these libraries. They leave a plain-text signature in the file, a writing style of their own and a distinctive compression level.

Note that some keep the original metadata: a simple metadata detector misses them. And in 39% of cases, the old value survives elsewhere in the file.



ORIGINAL STATEMENT

APR 8	CK	Payment / SAMPLE COMPANY INC.	83.10	5,960.83
APR 8	CK	Payment / SAMPLE COMPANY INC.	1,449.68	4,511.15
APR 14	CK	Payment / SAMPLE COMPANY INC.	30.56	4,480.59
APR 14	CK	Payment / SAMPLE COMPANY INC.	30.56	4,450.03
APR 14	CK	Payment / SAMPLE COMPANY INC.	436.86	4,013.17
APR 16	CK	Rent / SAMPLE COMPANY INC.	4,013.17	0.00

AFTER THE AGENT'S EDIT: AMOUNTS AND BALANCES RECALCULATED

APR 8	CK	Payment / SAMPLE COMPANY INC.	83.10	5,960.83
APR 8	CK	Payment / SAMPLE COMPANY INC.	1,594.65	4,511.15
APR 14	CK	Payment / SAMPLE COMPANY INC.	30.56	4,480.59
APR 14	CK	Payment / SAMPLE COMPANY INC.	30.56	4,450.03
APR 14	CK	Payment / SAMPLE COMPANY INC.	436.86	3,812.51

A sample statement before and after an agent's edit.

Method 4 1 2 3 **4**

Template farms, as active as ever

A template farm is a large library, a big supermarket of blank documents. The fraudster picks a bank from a catalog, fills in an online form (name, address, balance, transaction lines) and downloads the file. On screen, the result looks clean.

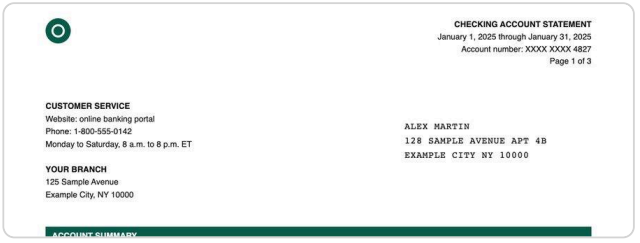
Of the forged statements in our database, 18.9% came from one of these libraries, and we have identified several dozen of them. They require no skill, cost a few euros and reappear faster than they shut down.

18.9%
of fake statements come from a template farm

How to detect it

These generated files always leave multiple traces, especially in the file's metadata. But metadata is still information declared by the file, a bit like a label on a product, which may not reflect reality: some tools disguise it.

What remains is a bank's issuing chain, which a template never fully reproduces. Fakes from the same template look alike, down to the same flaws.


CHECKING ACCOUNT STATEMENT
January 1, 2025 through January 31, 2025
Account number: XXXX XXXX 4827
Page 1 of 3

CUSTOMER SERVICE
Website: online banking portal
Phone: 1-800-555-0142
Monday to Saturday, 8 a.m. to 8 p.m. ET

YOUR BRANCH
125 Sample Avenue
Example City, NY 10000

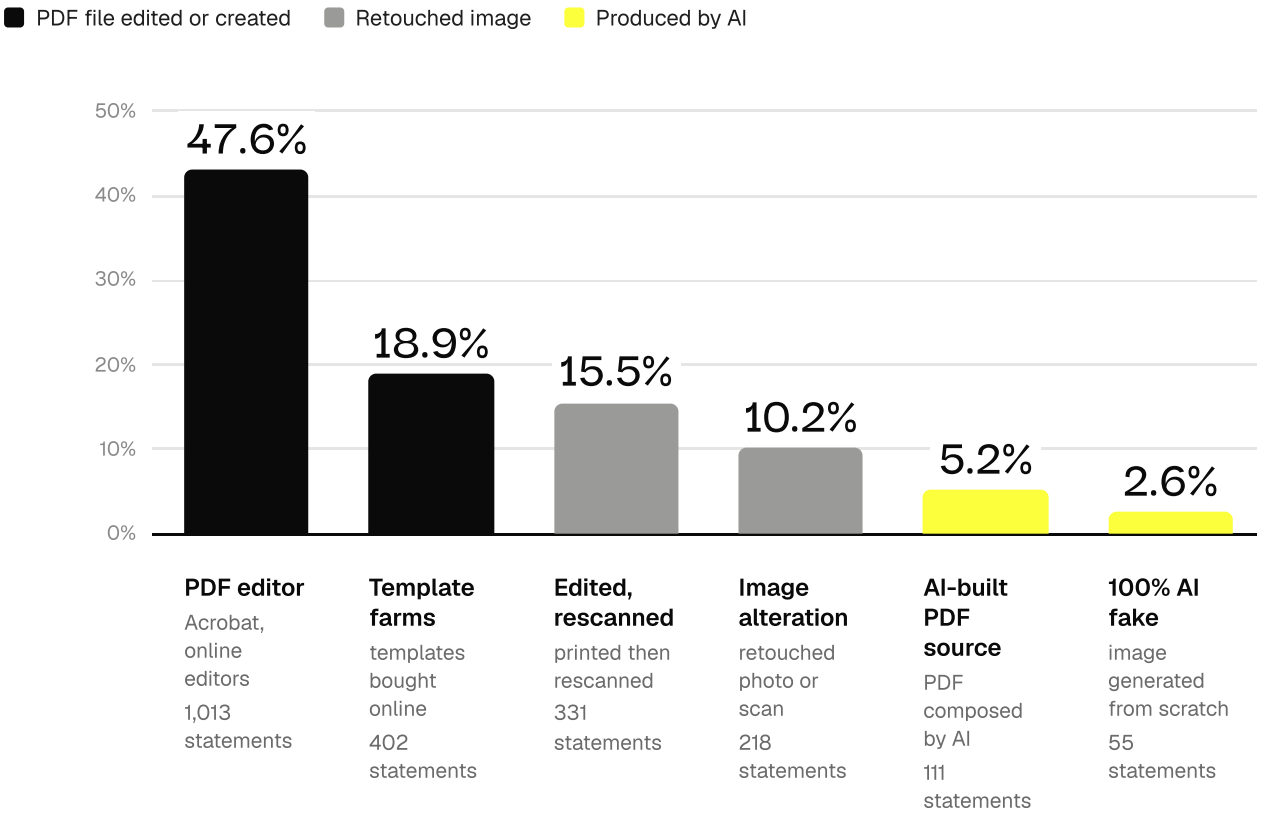
ALEX MARTIN
128 SAMPLE AVENUE APT 4B
EXAMPLE CITY NY 10000

ACCOUNT SUMMARY

A specimen statement, fictitious data: this is the layout that templates reproduce.

How fake statements are made

Share of each method among forged statements. Source: Koncile.



Case study

OnlyFake: the fall of a fake-document industry

The fall

In February 2026, Yurii Nazarenko, a 27-year-old Ukrainian, faced a New York court. He created OnlyFake. From 2021 to 2024, the site produced thousands of photos of fake documents: driver's licenses, passports, Social Security cards. It was open to the general public, and getting a realistic photo of a fake document there was very easy.

Extradited from Romania to the United States in September 2025, he pleaded guilty and agreed to forfeit $1.2 million, the equivalent of the site's revenue. On forums, the service showed screenshots of accounts opened with its fakes at more than 120 companies: crypto platforms, neobanks, payment services, betting sites.

$1.2M
forfeited by the creator of OnlyFake

120+
companies targeted by the service's fakes

The revival

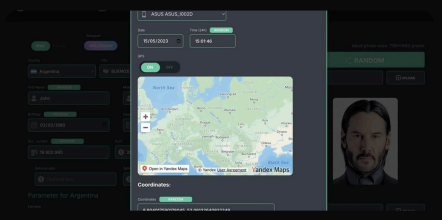
The story does not end there. One site, MacDoc, has taken over OnlyFake almost as is: reachable from Google rather than the darknet, it sells a photo of a fake passport or ID card from €15, paid in cryptocurrency. The new site's pseudo contract still carries the old name.



Sample fake passport produced by the site.

No AI, just layers

The user fills in a form. The fields are laid as layers over a document image, then blur and grain filters give it a photographic look: the card lying on a table or on fabric, for realism. The tool computes fake MRZ lines and simulates the phone brand, timestamp and GPS location of the photo.



The screen that generates fake metadata: device, date, GPS location.

What gives it away

These photos leave many traces. On a fake French ID card, we found 5 typography and positioning errors, visible to the naked eye. The address format, typed in freely, is very often wrong. The simulated metadata is far from perfect, and a 2D-Doc code cannot be recreated without the government's private key.

Financial documents are built the same way: a form, a template, a finish. The full analysis is published on koncile.ai.

Part 2

Fighting back: the tools, methods and technologies that work in 2026

No single technology is enough. Four methods complement each other, each with what it sees and what it misses. Every check contributes to an overall fraud score, and that score decides, not any single test.

Detection method 1 **1** 2 3 4

Traditional forensic analysis: the 80/20

It is the most affordable technology, and the one that catches most mistakes: reading the file itself, not just what it displays.

Reading the metadata

Every PDF declares the software that created it and the software that saved it, with dates. Many fraudsters leave these in plain sight. A statement in a bank's name written by Word, Canva or a PDF editor is a red flag.

These fields are easy to change. But the way objects are written, compressed and chained reveals the tool that actually produced the file. A fraudster doctors the metadata but does not rewrite the editor's writing style.

1 in 150

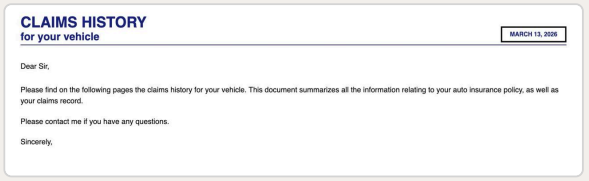
files keeps a trace of an edit, in production

Finding traces of the earlier file

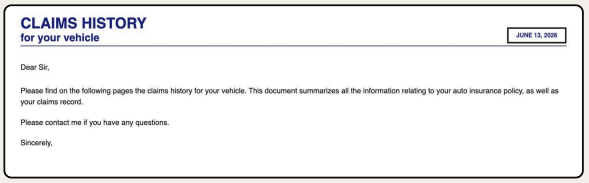
When an editor modifies a PDF without rewriting it, it appends a layer at the end of the file: the original value can then be recovered. A font embedded differently from the others is often one an editor added to rewrite an amount.

Example: an altered statement, before and after

Before, the original version recovered from the file


CLAIMS HISTORY
for your vehicle
MARCH 13, 2026
Dear Sir,
Please find on the following pages the claims history for your vehicle. This document summarizes all the information relating to your auto insurance policy, as well as your claims record.
Please contact me if you have any questions.
Sincerely,

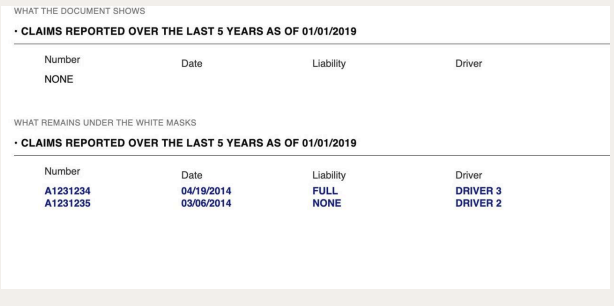
After the version submitted: the date went from March 13 to June 13, 2026


CLAIMS HISTORY
for your vehicle
JUNE 13, 2026
Dear Sir,
Please find on the following pages the claims history for your vehicle. This document summarizes all the information relating to your auto insurance policy, as well as your claims record.
Please contact me if you have any questions.
Sincerely,

Another case: a white rectangle placed over an amount, then written over. We found 3 in 1,340 PDFs analyzed, with no false alarms.

The pitfalls

Re-saving on the customer's computer changes the tool without any fraud. Several supporting documents may go through a merging tool with no page altered. Without calibration, a metadata test produces 10 to 20% false alarms.


WHAT THE DOCUMENT SHOWS
• CLAIMS REPORTED OVER THE LAST 5 YEARS AS OF 01/01/2019

Number	Date	Liability	Driver
NONE			

WHAT REMAINS UNDER THE WHITE MASKS
• CLAIMS REPORTED OVER THE LAST 5 YEARS AS OF 01/01/2019

Number	Date	Liability	Driver
A1231234	04/19/2014	FULL	DRIVER 3
A1231235	03/06/2014	NONE	DRIVER 2

The document shows "None"; under the white masks, two claims are still in the file.

Detection method 2 1 **2** 3 4

Machine learning: a must for AI-altered photos

When an image model does the retouching, nothing shows in the file. It takes a model trained to recognize what generators leave in the pixels.

Detecting image tampering requires machine learning: a general-purpose assistant does not see these edits. Trained on real documents and on their generated or retouched versions, the model scores each zone of the image and shows the suspicious ones.

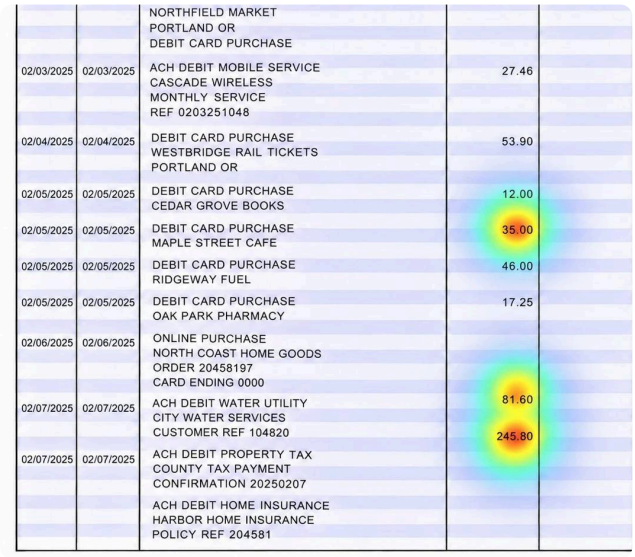
92%

of fakes detected by a specialized model, versus 56% for the best general-purpose assistant

The data makes the model. Public AI-image detectors, not trained on documents, answer almost at random: 50 to 62% in our tests. Retrained on documents, they exceed 90%.

A two-step system

The model finds suspicious zones but cannot tell whether they matter. A second pass judges whether the zone is worth altering: an amount, a name, a date. That is where fraud is identified.



On this sample statement, the model points to the retouched amounts.

Its limits

A single field retouched by AI, a fake re-saved in another format, and anything that is not an image: the model does not read native files. It also needs the right type of data: a model that recognizes a fake receipt will not necessarily recognize a fake pay stub.

Detection method 3 1 2 **3** 4

The file knowledge base: the shield against the best fraudsters

Each bank produces its statements with its own software chain. Comparing a file with authentic statements from the same issuer exposes even the most careful fake.

For each bank, we record what a statement looks like as it leaves its issuing chain. Not just the declared software: some fifty technical traces, such as encryption, compression, page structure or the way fonts are embedded.

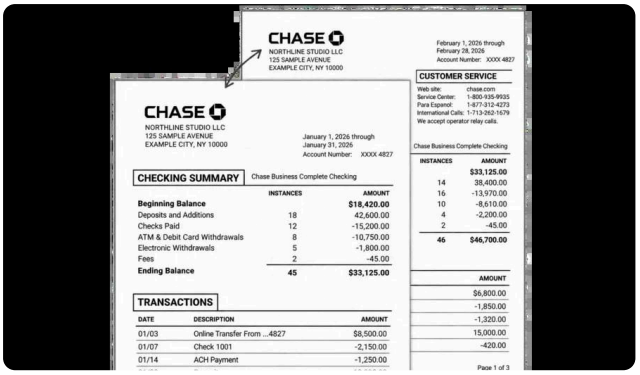
140+

banks recognized, and the base grows with every statement

An authentic statement from the same bank reproduces this profile exactly. A fake, however careful, almost always deviates on one point, and each deviation becomes a signal. Conversely, an authentic issuer that uses an unusual tool is no longer flagged by mistake.

Why it is a shield

A retouch almost always affects the file's envelope, because the editing tool rewrites the file its own way. Imitating a bank's exact fingerprint means knowing its entire production chain.



Two sample statements from the same issuer, fictional data: same expected fingerprint.

Its limit

It requires authentic statements from each issuer, and constant enrichment. A file profile only applies to native files: a scan or a photo comes from the customer, not the bank.

Detection method 4 1 2 3 **4**

Contextual checks: simple but effective

A fake must stay consistent with itself and with the real world. These checks look at neither the file nor the image: they verify what the document says.

Totals and line details

Add up the debit and credit lines, then compare with the totals computed by the bank. The fraudster sometimes forgot to update the whole statement. On one forged statement, the €450 shown as a debit did not match the €505.89 of principal detailed in the description.

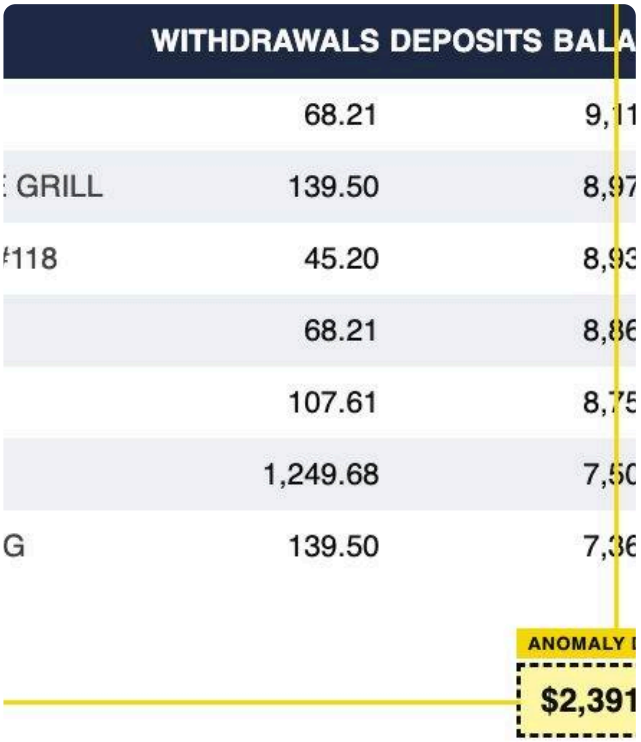
36

statements in our database assigned to a third party simply by changing the account holder

Account holder, IBAN, 2D-Doc

A person claims someone else's statement by replacing the name, and all the lines remain consistent with each other. For a company, we cross-check with the business register. The IBAN reveals the bank's identity, and is checked against the statement.

On documents that carry one, the 2D-Doc contains the document's key values, signed by the issuer. The values read from the code must be compared with the printed ones: a fraudster can copy an authentic code onto a document whose figures were changed.



On this sample statement, the closing balance does not match the sum of the transactions.

Their limit

Some fakes pass easily: the fraudster only has to recalculate the balances. And a single error in reading the document can trigger a false alarm.

80/100

A signal, not an accusation

You never know for certain that a statement was altered: you measure a risk. The result is a score from 0 to 100, with high risk from 80. The decision always rests with a person.



Interview

“Checking a document by hand is time my teams don’t spend with our customers.”

Ghislain Averty

Director of Operations and Actuarial, AND-E, insurer behind the Toyota Assurances brand

Toyota Assurances

At Toyota Assurances, teams used to check every vehicle registration, driver’s license and claims history statement by hand. In four months, from zero to production, the reading of underwriting documents was automated. Ghislain Averty explains what his teams do with the time freed up, and where fraud detection stands.

60,000
auto policies in the portfolio

70%
time saved on document checks

95 to 99%
of data read automatically

3 to 4 FTE
freed up and reassigned to retention

What made you decide to automate?

To me, the intelligence of account managers and underwriters should be used for something else: calling a customer back to keep them, understanding why they want to cancel, trying to find solutions with them, rather than spending time looking at a piece of data and checking that it is the right one, when that can be automated.

It is about using human intelligence for human relationships, not for robot tasks. Checking a piece of data, reading it, finding it, then verifying it: it takes forever. And our volumes are growing.

What gains are you measuring today?

Barely three months after installation, I think we have saved about 70% of the time we used to spend checking documents. That is huge, in time and in money. Relative to the size of our company, on this specific task, I would say we have saved about three to four FTEs.

And fraud detection?

We are not there yet: that will be a later step. The immediate priority was to optimize processes. In a second phase, we will work on document fraud. In our own business tied to Toyota and Lexus, given the way we select and acquire our customers, there is very little fraud. In auto insurance in general, there is a lot, clearly, notably through comparison sites.

Who would you recommend this approach to?

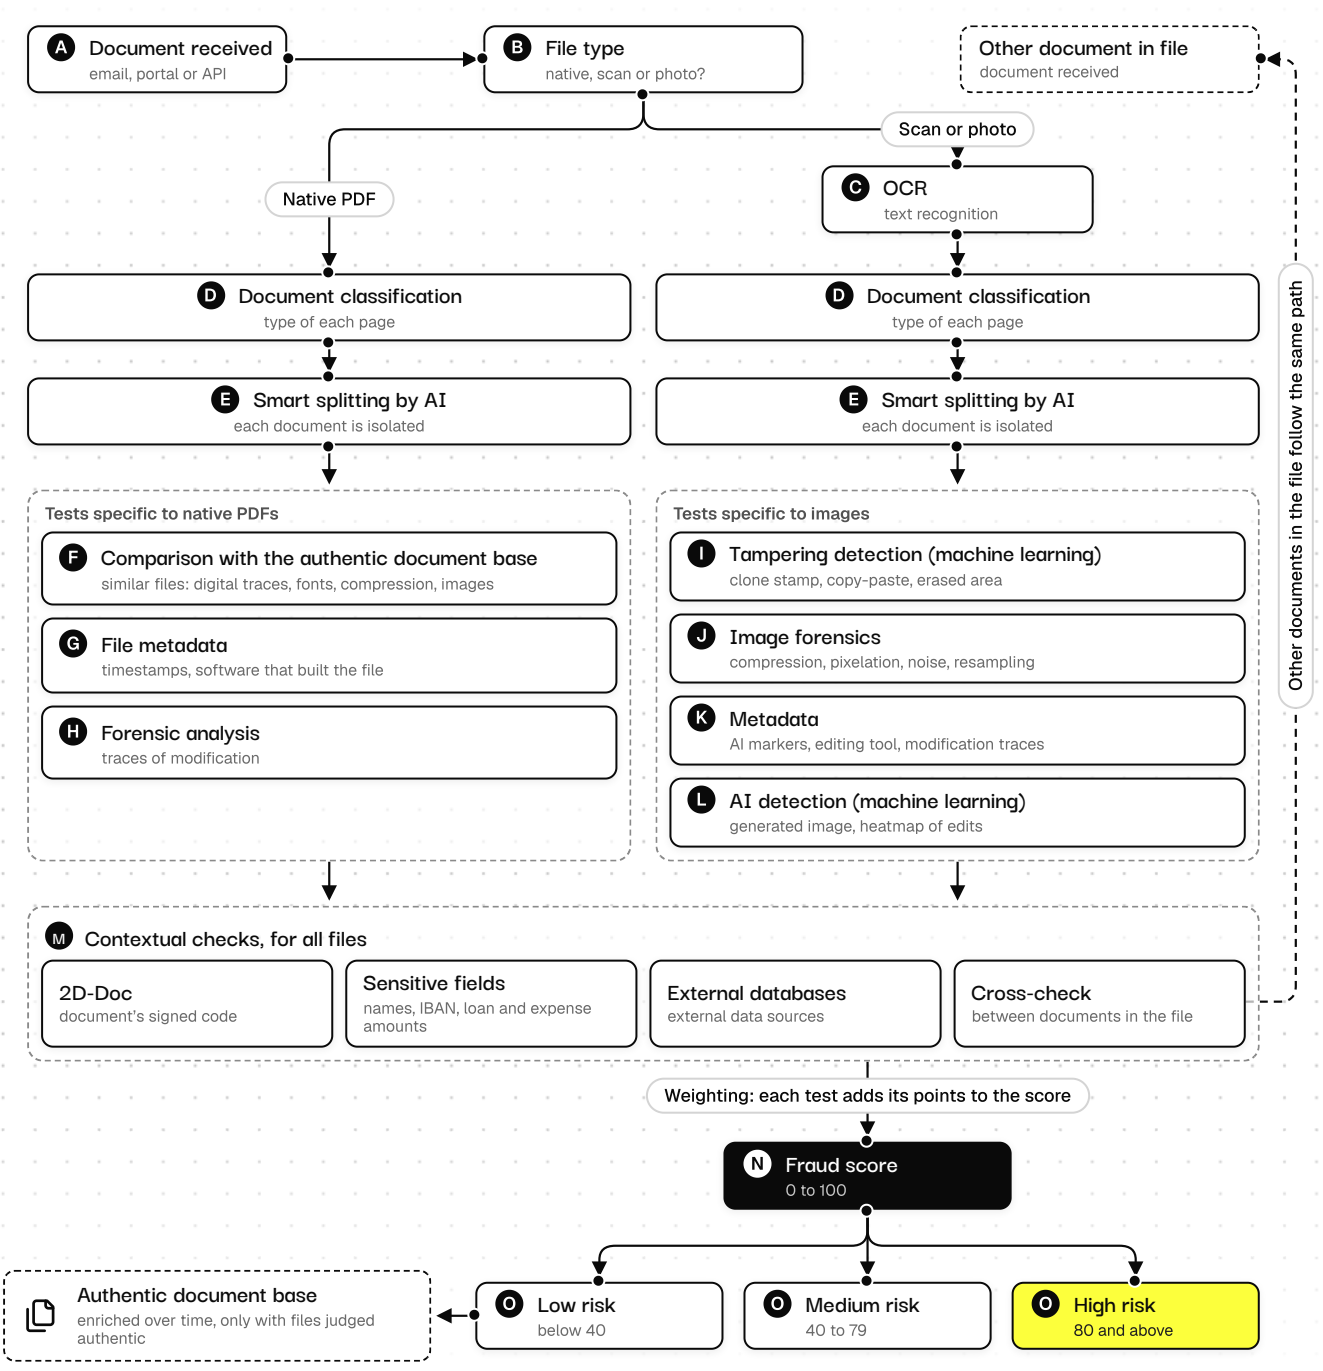
Above all to insurance brokers, and to insurers too. Koncile is a human-sized company that was easy to work with, especially for our IT teams. There is momentum. And to those who have not yet adopted OCR, and I know many, brokers, wholesale brokers or insurers: there is no reason to hold back, and it is much cheaper than people imagine.

Video interview by Tristan Thommen, co-founder of Koncile, September 1, 2026. Edited for readability.

The workflow

A self-learning AI workflow built to detect fraud

Detecting fraud takes a whole system: document processing and a combination of tools that, together, make it possible to assess fraud. Because fraud is always a body of evidence. To minimize false positives, you need to build a score, and that score relies on a combination of signals. The whole challenge is getting the weighting right.



The workflow

Each step explained

Each letter refers to a step in the diagram on the previous page.

A Document received

The document arrives by email, through a portal or via API. We analyze the file as it was received, not a re-saved copy: an original file keeps far more traces.

B File type

This is always the first step: native PDF, or scan or photo? The answer drives everything that follows, because fraud tests are specific to one or the other. A document goes down only one of the two lanes.

C OCR

A scan or a photo is just an image. OCR reads its text and data, field by field, so that the rest of the path is the same as for a native file.

D Document classification

Each page is assigned a type: bank statement, ID card, pay stub, proof of address. The type determines which checks will apply.

E Smart splitting by AI

An application file often arrives as a single file containing several documents. AI determines where each one starts and ends, so that it is analyzed separately.

F Comparison with the authentic document base

For native PDFs only, and it is a key step: all the file's digital traces, its fonts, its compression, its images, are compared with those of similar files already known to be authentic. The base grows over time, for all bank statement issuers, only with files judged authentic.

G File metadata

The PDF's general metadata: the timestamps, and what built the file, that is, the software that produced it.

H Forensic analysis

We read the file's structure: layers added afterward, the old value left underneath the new one.

I Tampering detection (machine learning)

Machine learning models look for editing moves: clone stamp, copied and pasted area, erased element. Public models from research exist, such as TruFor, CAT-Net, Noiseprint or DocTamper.

J Image forensics

We analyze the image's raw material, without a model: compression levels and areas compressed twice, pixelation, noise that varies from one place to another, traces of resampling.

K Metadata

Two readings. AI metadata, that is, the markers left by image generators. And the file's standard metadata, to spot a trace of modification or an unusual editing tool.

L AI detection (machine learning)

A two-part machine learning model. The first part recognizes fully generated images. The second draws a heatmap for a subtler case: a file partly modified by AI. Here too, public models exist, such as Community Forensics or NPR.

M Contextual checks

For an image as for a native file: 2D-Doc, amount checks, comparison with external databases. The focus is on the most frequently falsified fields: names, IBANs, sensitive amounts such as loans and expenses. This is also where the documents in the file, each having followed its own path, are cross-checked against one another.

N Fraud score

Each test adds points to a score from 0 to 100 according to its weighting, with the reason for each signal. The score decides, not any single test.

O The decision

Below 40, the risk is low. From 40 to 79, it is medium. At 80 and above, it is high and the file is reviewed by a person. A score remains a signal to verify, never an accusation.

They trust us

Insurance, banking, finance and large corporations



































Jules Ratier

Co-founder

jules@koncile.ai



Tristan Thommen

Co-founder and CTO

tristan@koncile.ai

Koncile

Document processing and fraud detection

1330 Avenue of the Americas, 23rd Floor, New York, NY 10019

2 rue de Lancry, 75010 Paris, France

+33 9 75 86 62 90

www.koncile.ai