Eidgenössische Finanzmarktaufsicht FINMA
Autorité fédérale de surveillance des marchés financiers FINMA
Autorità federale di vigilanza sui mercati finanziari FINMA
Swiss Financial Market Supervisory Authority FINMA

SBA Digital Finance Day
9 november 2026
Marlene Amstad, FINMA Chair

AI & Digital Transformation – New Opportunities for Private and Public Organisations

Ladies and Gentlemen,

Thank you to the Swiss Bankers Association for the invitation, and to August Benz for the warm welcome. It is a pleasure to open this afternoon right after a panel that asked exactly the right question: how do we balance innovation, trust and competitiveness?

Looking at today's programme, I notice that much of the afternoon is devoted to artificial intelligence. Only five years ago, hardly anyone would have predicted that. Back then, decentralised finance and NFTs were at their peak. Anyone who spoke of "tokens" meant digital units of value – not the building blocks of large language models. Same word, entirely different world. That alone tells you how fast this field is moving.

For more than 30 years, I have worked where financial markets meet technology – in Asia, the US and Europe, at central banks, in asset management, and now at FINMA. And I have never been more convinced: AI is not just another tool. It is a fundamental shift – in how we interact with data, and in how decisions are made.

The title of my talk speaks of new opportunities for private and public organisations. I will take that title literally. Let me cover three topics. First: why finance is more exposed to AI than most other sectors. Second: where Swiss financial institutions stand today; and what that means for you. And third: what FINMA, as a public organisation, is doing with AI itself. Here I am particularly pleased to present, for the first time, a new pilot that sits exactly where the private and the public side meet.

1. AI in the financial market

AI was already used in business and research well before the public launch of ChatGPT at the end of 2022. Since then, progress and investment have only accelerated. Business models and entire fields of science are being turned upside down or rethought overnight. The momentum is still too great for a final assessment of the opportunities and risks of this technological leap. And yet one thing is already clear: AI is likely to have a profound impact on our economy, and on the financial centre and the banking sector in particular.

Laupenstrasse 27
3003 Bern
Tel. +41 (0)31 327 91 00
www.finma.ch

Finance is particularly exposed to AI. The reason is simple: finance runs on data, speed and trust.

Data. Processing information is the core of banking, and the wider you cast the net, the truer this is. Data can be a video recording with which a client identifies themselves when opening an account, as well as a building plan or the sales documentation in a mortgage file. Together, this data is the basis for decisions such as granting a loan or taking on a new client. Most information is found not in tables but in unstructured data, and text is precisely where the new generation of AI is strongest.

Earlier forms of data processing already had a major impact on financial business. AI goes further: it can recognise patterns in vast datasets, process new types of data, understand text and make recommendations – at a speed and breadth we never knew from conventional software. It also opens great potential for automating processes. AI systems can – if we let them – take decisions themselves, or provide the basis for human decisions. They can also draw insights from structured data or process it further automatically.

Speed. AI compresses time. Credit decisions that once took days can be taken in seconds. But this also narrows the window in which an error can be caught before it spreads.

And trust. The entire financial system rests on it. Every efficiency gain is worthless if it erodes the confidence of clients, investors and the public.

Data, speed, trust: that is why AI matters more for finance than for most sectors. This brings me to my second point.

2. Where Swiss financial institutions stand in AI – and what it means for you

What many people do not realise is that Switzerland is at the forefront of this development. The reason is the diversity of our financial centre: a wide range of products and services, and providers from global players to small regional banks and young fintechs. That diversity positions Switzerland as a natural laboratory for AI in finance.

As an integrated supervisor, FINMA sees AI arriving at banks, insurers and asset managers alike. FINMA has a front-row seat to how AI is transforming finance. We started surveying the industry on its use of AI as early as 2021, well before ChatGPT was released. Our most recent survey, in 2025, covered around 400 institutions – banks, securities firms, insurers, fund management companies, asset managers and financial market infrastructures. Three findings stand out:

First, momentum is high. For every application already in use, two more are in development. On average, the institutions surveyed run around five AI applications and are working on nine more – larger firms considerably more intensively than smaller ones.

Second, look at *where* exactly AI is being used: the rollout of applications goes broad before it goes deep. Today, the emphasis is clearly on an early phase and on less critical use cases. Chatbots for client advice and communication are at the forefront, followed by text generation and general process optimisation. That explains one of the most striking figures in our survey: over 90 per cent of the institutions using AI already rely on the newest, generative methods. These applications are comparatively quick to implement, easy to delimit and, if something goes wrong, manageable in their impact.

Core processes however are a different matter: credit decisions, the calculation of regulatory metrics, processes with a direct external impact. AI is moving into these fields as well, as we see internationally, but much more cautiously. From a supervisory perspective, that is a sensible sequence.

Only about half of the institutions using AI have embedded it in an explicit AI strategy. Few develop their applications in-house. Most rely on external providers – and increasingly on the same small number of large technology groups. This places high demands on governance and risk management: an institution remains responsible for the applications it uses, including those it has bought in. Suppose a large part of the industry comes to depend on the same few models, for example for risk scoring or fraud detection. Then a single outage, or a single flawed model, could affect financial stability.

What does this mean for you? The question is no longer whether you use AI, but whether you are in control of it. This is all the more true of agentic AI, which several of you will discuss in this afternoon's breakout session. AI systems no longer just answer questions; increasingly, they plan, execute and chain tasks on their own. This sharpens a principle I consider non-negotiable: accountability cannot be delegated to an algorithm. An agent can act on your behalf. It cannot be responsible on your behalf. For every AI application, someone in your organisation has to own its purpose, its limits and its outcomes. Clear responsibilities are not bureaucracy; they are simply good management.

The third finding is a very topical one: a considerable share of the AI risks facing the financial centre does not arise from the institutions' own use of AI at all, but from malicious actors using the technology to do harm. AI is a classic dual-use technology.

Take cyber attacks as an example. The most capable frontier models have shown that they can find previously unknown vulnerabilities in widely used software that had been scrutinised for decades – and turn them into working attacks. The same capability that helps defenders find weaknesses helps attackers exploit them. For every institution in this room, the consequence is simple: the time you have to close a known vulnerability is shrinking fast. Patch faster.

Take digital fraud as a second example. A figure from Interpol's latest global threat assessment on financial fraud shows how much difference AI makes: AI-enabled fraud operations are around four and a half times more profitable than traditional methods. Interpol speaks of a veritable industrialisation of fraud, made possible by cheap digital tools and globally networked criminal structures. Fake voices on payment orders, identity fraud in digital onboarding, real-time phishing – these attacks hit institutions regardless of how well they govern their own AI.

At the same time in our supervisory guidance on digital fraud risks at Swiss financial institutions this spring, we noted that the number of such cases has risen steadily since 2022 – driven by AI, automation and the advancing digitalisation of financial services. Yet our survey shows gaps in defences: a considerable share of the institutions surveyed have no dedicated digital fraud policy, and some do not even have a body responsible for it.

So my message to the industry is twofold. Use AI – and stay in control of it. And use AI to defend yourselves – because your attackers already do.

3. AI in FINMA's supervision

This brings me to my third point: the enormous opportunities AI offers – from the perspective of a supervisory authority.

For FINMA, the use of technology, and of AI in particular, is a strategic priority of the past few years. We pursue two goals. FINMA wants to become more *efficient*: reaching the same or better results faster and with less effort. And we want to be more *effective*: raising the quality of supervision through deeper insights.

What excites me most is something classic data-based supervision could never offer: insights from unstructured data. Data-based supervision has so far relied on structured data: key figures, reports, balance sheet items. But a large part of what matters for supervision is found not in tables but in running text: in minutes, regulations and e-mails, as well as in publicly available sources, from the prospectus to the social media post. AI makes it possible to process and analyse text systematically. This allows classic data-based supervision to be extended to new areas, such as governance or anti-money laundering.

Several applications are already in use at FINMA. I presented three in some detail at the Point Zero Forum in June so let me keep it short today. The first is a pipeline for AI-assisted document analysis: a modular process that can handle both structured and unstructured documents. The second is an AI-supported tool that, ahead of on-site inspections, automatically identifies possible anomalies in documents on the basis of the inspection programme, before our staff begin the actual review. Both serve the same purpose: they prepare data for internal use. The AI reads, sorts, flags and always provides a source. But it does not decide; its role remains a supporting one. Responsibility for any decisions that follow stays with our supervisors.

This brings me to something I have been very much looking forward to. Today, I am pleased to present for the first time, an application that takes us one step further. We call it Subsuma. Until now, we have used AI for internal analysis. Here, we are entering new territory: using AI to help prepare a product that goes to the outside world.

Every year, close to 100 fintech start-ups ask FINMA: does my business model require a licence, or some other regulatory set-up? Every case is individual, but the basic questions are often very similar. That makes it a clearly defined, tech-savvy field – with plenty of past cases to learn from.

And this is exactly where the private and the public meet. On one side of the desk, a private start-up – often AI-native itself – describes its business model. On the other side, a public authority assesses whether a licence or other regulatory requirements are needed. One request, two sides of the desk – and AI can help on both.

How does it work? The system analyses the request and summarises the business model so that all elements relevant for the legal assessment are captured. On that basis, the case is assessed against the various statutory provisions from which a licensing requirement could arise. The system must be able to back up every statement with sources – references to the case file – and make it fully traceable.

The first aim of the project was to deliver a proof of concept with small models and to evaluate the possibilities and limits of the different models. For productive use, we plan to move to more powerful variants later.

A key feature of this prototype is that we do not rely on just one LLM. Instead, we are using three open-source models, all running on FINMA's own hardware. This set-up gives us a variety of capabilities, flexibility and independence from large providers. AI is rapidly becoming a strategic resource. For government use, but also for the Swiss economy as a whole, it may be advantageous for Switzerland to have models of its own that it can rely on. We need to develop our own AI capabilities and strengthen our AI sovereignty by investing in this key technology.

That's why I am particularly pleased that one of the models in this pilot is Apertus – a fully open, transparent and multilingual language model developed in Switzerland by ETH Zurich, EPFL and the Swiss National Supercomputing Centre. ETH also advised us on the design of the evaluation, and we would be glad if FINMA's experts' feedback could, in turn, contribute to the further development of Apertus.

To evaluate the three models, we organised what you might call a blind test. First, the models were trained on past requests. Then, new requests were fed into the system and our subject-matter experts rated the outputs of the different models without knowing which model had produced which answer. At the same time, AI models were used to assess one another. This is an iterative process in which our lawyers, IT specialists and AI experts work side by side. That interdisciplinary cooperation has proven to be the single most important success factor in developing AI use cases.

The honest interim result: high-quality output requires larger models than the small ones we started with. But the methodological foundations and the front end are in place for the next phase. It is too early for figures. But one thing is already clear: this will speed up our answers. For a start-up, a faster, more consistent answer to the question "Do I need a licence?" or any other regulatory requirement, is not a detail. It can decide whether an idea makes it to market. And for FINMA, it frees up our experts' time for the genuinely difficult cases – where judgement is needed most.

To conclude: FINMA is aware that high demands are placed on public authorities when it comes to the use of AI – and rightly so. Our decisions must be reasoned, and based on facts and a correct interpretation of the law. In a principles-based legal system like ours, discretion and context are particularly important. Simply delegating them to a machine is clearly not our goal, even if it were technically possible. So let me be equally clear: no letter will leave FINMA without one of our experts reviewing and signing it. AI prepares, but people decide.

At the same time, FINMA is run according to the principles of sound economic management, and we are constantly working to increase the efficiency and effectiveness of our work. Wherever language is at the heart of the work, the enormous potential of our AI tools is already evident. But realising that potential is not a matter of flipping a switch. It happens iteratively, with everyone on board who wants to do the best possible work together with AI. Close cooperation across disciplines has proven a key success factor, and we will deepen that cooperation in the next stages of development.

I am convinced that this is not fundamentally different at your institutions, and I look forward to discussing it with you.

Thank you for your attention.