AI Draft Policy Withdrawal, an uncomfortable moment but a useful one

Darren Olivier

The recent withdrawal of South Africa's draft national AI policy by Communications Minister Solly Malatsi has caused understandable embarrassment. Fictitious academic references found their way into a draft official policy document that advocates for proper AI governance and human oversight is not only acutely ironic but also damaging. Still, beyond the headlines and the inevitable commentary, this episode offers something more useful than outrage. It offers a lesson.

I spend a fair amount of time working at the intersection of artificial intelligence and intellectual property. One theme keeps returning, sometimes quietly and sometimes with a bang. Reputation is fragile and takes years to build. AI can damage it quickly, and once trust is shaken, recovery can take considerable time. In some professions or industries, it may never fully happen.

We have already seen lawyers disciplined for filing documents that contained invented case law or non-existent authorities. Courts have reacted firmly. Professional bodies have done the same. Clients tend to react even faster, because they rely on the assumption of complete integrity and oversight. So, as this episode has illustrated, the risk is not theoretical. It is already playing out in practice and damage to brands, be those personal in the case of a lawyer or their firm or for an organisation or government (as it is in this case), and the damage is tangible.

That is why this moment is important. The very instrument and body designed to create trust around AI now finds itself serving a case study in governance. This is not just a drafting error or a technical slip, it is a reminder that AI governance is of critical importance to get right. In that sense, it is a valuable lesson.

Governance is not about banning AI

At its simplest, governance means knowing when AI has been used, who reviewed the output, and how the final decision was made.

Most organisations now use AI in some form, and government is no exception. That is sensible. AI is very good at producing structure, summarising large bodies of information and as a sounding board. Increasingly, AI is integrated into work products in a myriad of ways and industries. In practice, the best results usually come from collaboration. AI accelerates the work. Humans apply judgement.

My own working rule is fairly simple. Do not outsource human thinking. Let AI do the heavy lifting as required. Let humans supply judgement as the work takes shape. Then, before anything leaves the building, let humans verify every reference, every citation and every assumption. The final check is where credibility is protected, and it is also where many organisations still underestimate the risk.

As with the draft AI Policy, a document produced with the help of AI can look polished and authoritative while still containing serious errors.

Credit where it is due

The decision to withdraw the draft policy quickly was the right one. It showed awareness of the problem and a willingness to act before the issue grew even larger. In a governance context, speed matters. The longer a flawed document remains in circulation, the harder it becomes to rebuild confidence.

Many organisations hesitate in similar situations. They hope the issue will fade or that the error will go unnoticed. Experience suggests that this strategy rarely works. Withdrawal, by contrast, signals respect for the process and for the public that relies on it.

So while the situation is clearly embarrassing, the response deserves recognition. It protected the integrity of the policy process and created space to fix the problem properly.

The next draft matters more than the first

South Africa still needs a credible national AI policy, and the demand for clarity is only going to grow. Businesses want guidance on compliance. Regulators want consistency. Citizens want reassurance that new technologies are being managed responsibly. That is the real context in which the next draft will be judged.

The next version of the policy does not need to be perfect, but it does need to be credible. That credibility will come from three visible elements.

First, the policy should be clearly rooted in South African realities. Our legal framework, infrastructure constraints and economic priorities are different from those of Europe, the United States or Asia. Borrowing ideas from other jurisdictions is sensible, but copying them wholesale rarely works.

Second, the drafting process should show clear human oversight. That does not require a long list of committees or complicated structures. It simply means that responsibility is visible and accountability is clear. People trust systems more readily when they know who stands behind them.

Third, and perhaps most importantly, the policy should demonstrate verification. Sources should be checked carefully. Evidence should be traceable. Assumptions should be tested against real-world conditions. These steps may feel routine, yet they are the foundation on which public trust is built.

A broader lesson for all of us

It is tempting to treat this episode as a government problem. That would be an easy conclusion and, in some circles, a popular one. It would also miss the point.

Every organisation is now experimenting with AI. Every professional is under pressure to work faster and deliver more. Every team is learning, often in real time, how to integrate new tools into existing workflows. The risk of error sits everywhere, not only in government departments.

What happened here simply brought the issue into the open, in a very public way. Visibility has value. It forces conversations that might otherwise be postponed. It encourages organisations to look more carefully at their own controls and processes.

In that sense, the incident performs a useful function. It shows how quickly credibility can be tested when governance falls behind technology. It also reminds us that verification remains a human responsibility, even in a world increasingly shaped by automation.

South Africa now has an opportunity to rebuild the policy on stronger foundations. If the next draft reflects careful oversight, disciplined verification and a clear understanding of local realities, this uncomfortable moment will have served a constructive purpose.

That would be a worthwhile outcome, and one that extends well beyond government.

Darren Olivier is a partner at Adams & Adams spearheading our AI governance initiative and leading the SmartAIIP team that has developed www.smartaiip.adams.africa. SmartAIIP is an Africa focussed platform designed to assist organisations prepare for and manage AI governance risks and opportunities.