Every Proposal Is Well Written Now. That Is a Problem.

How artificial intelligence is changing grant applications and why funders must rethink what they measure

By Richard Chidiebere (Senior Officer, Grants & MERL, NYFF)

Somewhere around the fortieth proposal, I noticed something I could not unsee. The writing had stopped varying. In my role as Senior Grants & MERL Officer for the Nigeria Youth Futures Fund (NYFF), I coordinated the review process for the Youth Leadership Development Fund Applications against seven key criteria: alignment with the Fund's thematic focus, feasibility of the proposed intervention, coherence of the proposal, prudent budget management, inclusion, measurable outcomes, and sustainability. These criteria guided my assessment of whether each proposal was strategically aligned with the Fund's priorities, logically designed, realistically implementable, financially responsible, inclusive in its reach, capable of delivering measurable change, and likely to sustain its impact beyond the grant period.

The applicants were remarkably different. Some were informal groups of five young people with little experience managing grants. Others were established organisations with dedicated programme and finance teams. They came from different states, operated in different contexts and had varying levels of institutional capacity.

Historically, these differences were often apparent from the first page of an application. Some proposals arrived professionally written and carefully structured. Others were three urgent, grammatically imperfect paragraphs from people who understood their communities far better than they understood donor language.

That distinction was becoming less visible.

Sentences were polished. Implementation plans were properly structured. Risk matrices appeared where they had rarely appeared before. Monitoring and evaluation frameworks were increasingly sophisticated. Applicants seemed to have acquired, almost overnight, the language of professional grantmaking.

Generative artificial intelligence has changed proposal writing. But it has also introduced an uncomfortable question for funders: when everyone can produce a professionally written proposal, what exactly are we measuring when we assess writing quality?

For years, funders have relied, consciously or otherwise, on good writing as an indication of organisational competence. AI has weakened that relationship. Yet many grant assessment processes have not evolved to reflect this change.

The changing economics of proposal writing

To understand the implications, it is important to appreciate how grant applications function. A proposal is more than a request for funding. It is an attempt by an applicant to demonstrate that they understand a problem, have a credible solution, possess the capacity to implement it and can account for the resources entrusted to them. However, communicating these things effectively requires a particular set of skills.

Applicants must understand donor priorities, develop a theory of change, construct logical frameworks, identify indicators, prepare budgets and articulate anticipated outcomes.

For smaller organisations and informal youth-led groups, these requirements can become significant barriers to accessing funding. An organisation may be delivering valuable services in its community without having anyone capable of translating its work into a donor-compliant proposal. Another may have a strong understanding of its intervention but lack the technical expertise to develop an adequate monitoring and evaluation framework.

Generative AI is changing this equation.

Tools such as ChatGPT, Claude, and Gemini can help applicants structure their ideas, improve clarity, develop implementation plans and translate practical experiences into conventional proposal language. They can also provide guidance on developing budgets, indicators and results frameworks.

For organisations with limited institutional resources, this is potentially transformative.

A 2025 pilot study published in Innovative Surgical Trends examined ChatGPT's ability to independently develop a competitive grant application. Although reviewers considered the resulting proposal well written, their assessments differed considerably. One reviewer considered it good, while another identified significant concerns about feasibility. The researchers concluded that AI could complement grant writing but could not independently guarantee a competitive application.

This distinction is important. AI can improve how an applicant communicates an idea without necessarily improving the quality of the idea itself.

The opportunity, therefore, is not simply that more people can write better proposals. It is that people with valuable ideas and limited technical writing experience may now have a better opportunity to compete for funding.

When writing quality stops being a reliable indicator

Consider how grant applications are assessed in practice. The Nigeria Youth Futures Fund supports young Nigerians working on governance, accountability and civic participation. Its grantmaking approach recognises that applicants operate at different levels of institutional development. Consequently, funding is structured across three categories:

- The Individual and Small Grant category supports young people and informal groups, often at the earliest stages of developing an intervention.
- The Development category supports organisations with an existing track record that are ready to deepen or expand their work.
- The Catalyst category supports organisations seeking to deliver interventions at regional scale. This comes with a particularly demanding requirement: applicants must demonstrate genuine implementation capacity across every state in the geopolitical region they propose to cover.

In the most recent grant cycle discussed here, the fund received more than 3,000 applications. Every application was reviewed and scored by trained reviewers against established assessment criteria.

The initial review process narrowed the applications to 336 for further consideration: 200 Individual, 110 Development and 26 Catalyst applications.

Ultimately, 166 organisations and individuals received grants: 100 Individual awards, 60 Development awards and six Catalyst awards.

These figures tell an important story. Approximately one in eighteen applicants received funding. More importantly, only six of the 26 shortlisted Catalyst applications were successful.

In most cases, the challenge was not the quality of the writing. The proposals were generally well presented. The issue was whether the organisations could demonstrate the geographical presence and operational capacity required to deliver interventions across their proposed regions.

A well-written proposal cannot manufacture a network of implementing partners across states where an organisation has no operational presence.

This is where the distinction between what an applicant claims and what can actually be demonstrated becomes critical.

The final selection process also required considerations beyond application scores. For example, selecting 60 Development grantees from 110 shortlisted applications involved balancing geographical representation, gender, thematic coverage and the inclusion of applicants facing particular personal circumstances.

Scores informed these decisions, but they did not independently determine them.

The quality of assessment therefore remains essential. It is the mechanism through which funders distinguish between the strength of an idea, the credibility of its implementation arrangements and the evidence supporting an applicant's claims.

Good writing was Never What we Meant to Measure

A grant proposal is fundamentally an assessment instrument.

It is designed to help a funder establish four things:

1. Does the applicant understand the problem they are trying to solve?
2. Does the applicant have the capacity to implement the proposed intervention?
3. Is there a credible relationship between the proposed activities and the intended outcomes?
4. Are the requested resources reasonable, justified and accountable?

Writing quality was never supposed to be the outcome of interest. It was a proxy.

Historically, organisations with established systems, experienced programme staff and exposure to donor-funded projects were more likely to produce polished proposals. Reviewers consequently became accustomed to interpreting professional writing as an indication of institutional competence.

For a long time, that assumption appeared reasonable.

But measurement theory offers an important warning: a proxy is useful only when it maintains a meaningful relationship with the underlying construct it is supposed to represent.

When that relationship weakens, the proxy becomes unreliable.

This is precisely what generative AI has introduced into grant assessment. It has made professional writing more accessible without necessarily making the underlying organisational capacity equally accessible.

There is also a more uncomfortable consideration. The system was not entirely fair before AI arrived.

Grantmaking processes have historically rewarded applicants who understand the language, conventions and expectations of funders. Organisations with access to experienced consultants and grant writers often enjoyed an advantage over smaller, less resourced groups.

AI has not created this inequality. In some respects, it has reduced it by making technical writing support more accessible.

The challenge is ensuring that funders do not mistake improved access to professional writing for evidence of improved institutional capacity.

The Benefits are Real, And Funders Should Acknowledge Them

It would be a mistake to approach AI-assisted applications primarily as an integrity problem.

Consider a women's cooperative in a rural local government area that has developed a reliable referral system connecting pregnant women to health facilities. Its members understand the problem, have established community relationships and can demonstrate the value of their intervention.

However, nobody on the team has ever developed a formal results framework or written a donor-compliant proposal.

The nearest consultant with that expertise may charge more than the grant itself.

For such an organisation, AI is not necessarily a shortcut around the work. It can be a bridge between demonstrated community experience and the technical language required to access resources.

There are other benefits.

Basic proposal quality has improved. Applicants can more easily structure their interventions, distinguish activities from outcomes and articulate implementation risks. Reviewers may consequently spend less time interpreting poorly organised submissions and more time interrogating the substance of proposed interventions.

Research from MIT Solve, developed through a collaboration involving Harvard Business School, the University of Washington and ESSEC Business School, has also explored AI-assisted grant application assessment. Its initial findings suggest that AI can support objective screening, such as checking eligibility requirements and alignment with funding priorities, while human reviewers retain responsibility for more nuanced judgements about impact and organisational capacity.

This presents an opportunity for funders to improve their processes without surrendering the expertise and contextual understanding required for good grantmaking.

The objective should not be to discourage applicants from using technology. It should be to ensure that technology improves access without compromising the integrity of funding decisions.

The New Problem: Polished Vagueness

The challenges are equally significant.

First, writing quality is becoming less useful for differentiating between applicants.

When strong and mediocre proposals demonstrate similar levels of grammatical accuracy, structure and professional presentation, reviewers lose an important distinction. Scores may become inflated, particularly on criteria that reward logical presentation and technical completeness.

Second, weaknesses in applications have changed form.

Previously, weak proposals often failed visibly through poor organisation, incomplete sections or inadequate articulation of implementation strategies.

Today, some proposals are weak in ways that are considerably harder to identify.

They sound convincing but provide little context-specific information. They present sophisticated theories of change without explaining the practical relationships between activities and outcomes. They include comprehensive risk matrices without demonstrating an understanding of the risks that actually exist in their operating environments.

The tell is often the absence of specificity.

A proposal may describe an intervention that could be implemented in virtually any Nigerian state, community or sector, with only a few words changed.

Everything appears correct, but very little is verifiable.

Polished vagueness is becoming a significant warning sign in grant assessment. It is particularly difficult to identify when reviewers rely heavily on the quality of written submissions.

Third, there are questions of application integrity.

During the review of 110 shortlisted Development applications, eight were flagged for integrity concerns.

Interestingly, most of these concerns were not necessarily related to AI. Some applications contained blank or corrupted documents. Others included boilerplate content from shared drafts that had circulated between unrelated applicants, occasionally retaining another organisation's name. Some budgets were remarkably similar across organisations without any obvious connection. One application had been recycled almost word for word from a previous funding round.

These practices existed before generative AI became widely accessible. Templates, shared documents and recycled proposals have circulated within grantmaking ecosystems for years.

What AI changes is the speed, accessibility and sophistication with which generic content can be reproduced.

The implication is not that AI-assisted applications should automatically attract suspicion. Rather, funders need more deliberate processes for checking originality, consistency and the credibility of information supplied by applicants.

Why AI Detection is not the Answer

One tempting response is to introduce AI detection software into grant assessment.

This approach, however, raises serious questions about reliability and fairness.

A study by Liang and colleagues, published in Patterns in 2023, found that several GPT-detection tools disproportionately misclassified essays written by non-native English speakers as AI-generated. The

researchers demonstrated that linguistic patterns associated with second-language writing could trigger false accusations of AI use.

This finding has implications for grantmaking, particularly in diverse contexts where applicants have different levels of formal education, language proficiency and access to professional writing support.

A detection tool could end up penalising precisely the applicants whose access to funding AI has helped improve.

More fundamentally, detecting whether a proposal was written with AI does not establish whether its contents are accurate, feasible or original.

An entirely human-written proposal can contain fabricated statistics, unrealistic budgets and exaggerated claims. Equally, an AI-assisted proposal can describe a genuine intervention with verifiable evidence.

The relevant question should not simply be whether AI was used. It should be whether the information presented can withstand scrutiny.

Four Things Funders Should Change

If professional writing is becoming less reliable as an indicator of capacity, funders must reconsider how they design and implement their assessment processes.

There are four practical adjustments worth considering.

1. Prioritise verifiable detail over professional presentation.

Funders should deliberately assess the specificity and credibility of information provided by applicants.

Ask for named communities, wards, implementation partners and referral institutions. Where applicants present participant figures, ask for the basis of those numbers. Where budgets include unit costs, examine whether they reflect actual local market conditions.

AI can generate plausible information, but applicants should be able to substantiate their own operational realities.

Verifiable evidence should be treated as an assessment criterion in its own right.

2. Give greater weight to demonstrated delivery.

Intentions are relatively easy to describe convincingly. Evidence of implementation is considerably harder to manufacture.

Where applicants have previous experience, reviewers should examine what they have actually delivered, who benefited, what challenges emerged and how the organisation responded.

This does not mean that first-time applicants should be disadvantaged. Rather, assessment should distinguish between demonstrated experience, emerging capacity and proposed implementation arrangements.

For new organisations and informal groups, evidence may include community relationships, previous volunteer activities, documented initiatives and credible partnerships.

3. Introduce structured conversations with shortlisted applicants.

A well-facilitated conversation can reveal what a written application may conceal.

Fifteen minutes with an applicant can provide valuable insight into their understanding of the problem, the rationale behind their intervention, their knowledge of the operating environment and their ability to respond to implementation challenges.

This is particularly important for larger grants, where the financial and programmatic risks are greater.

The conversation should not become an informal test of presentation skills. It should be a structured opportunity to validate the applicant's understanding and clarify uncertainties in the proposal.

4. Interrogate the budget and strengthen due diligence.

Budgets remain one of the most revealing components of a grant application.

They demonstrate what applicants believe implementation will cost, how they intend to allocate resources and whether their proposed activities are financially realistic.

Reviewers should examine unit costs, personnel allocations, logistics, procurement assumptions and the relationship between expenditure and expected results.

Funders should also strengthen reference checks, verification of organisational records, validation of implementation partners and examination of previous project evidence.

Alongside these measures, funders need clear policies on responsible AI use, including disclosure expectations where appropriate and safeguards against uploading confidential applicant information into unapproved AI platforms.

The National Institutes of Health, for example, prohibits its peer reviewers from uploading confidential grant applications or original concepts into generative AI tools. Its policy highlights the importance of confidentiality in protecting applicants and maintaining the integrity of the review process.

The Future of Grantmaking Requires Better Measurement

The conversation about AI and proposal writing should ultimately lead to a more fundamental discussion about what funders consider evidence of capacity. For too long, grant assessment has sometimes conflated an organisation's ability to communicate its work with its ability to deliver that work. These are related, but they are not the same. A professionally written proposal does not guarantee implementation capacity. A sophisticated results framework does not guarantee that an organisation understands how to measure change. A detailed budget does not automatically establish financial management capacity.

Likewise, an applicant who struggles to articulate an excellent intervention should not be automatically considered less capable of delivering it.

The emergence of generative AI provides an opportunity to correct some of these longstanding weaknesses in grantmaking:

- **Score the intervention and the document separately**. Be explicit in the rubric about which criterion measures which.
- **Reward detail that can be checked**. Named partners, local unit costs, last year’s numbers. Software can produce a believable figure, not your figure.
- **Weigh what has been delivered above what is intended**. Intentions are now cheap to describe well. Delivery leaves a trail.

- **Talk to shortlisted applicant**s. Ask them to explain their own logic without the document. The gap between how the proposal reads and how well they understand it is the most honest signal available.
- **Interrogate the budget**. Still the hardest page to fake, still the most underweighted.
- **Track what actually predicted delivery**, and revise the rubric on evidence rather than assumption.

The principle is straightforward: stop asking applicants to prove that they can write like funders. Start asking them to demonstrate that they understand their communities, can deliver what they propose and have evidence to support their claims.

In the age of AI, better writing is becoming easier to access. Better grantmaking will depend on our ability to measure what actually matters.

References and further reading

1. Chhabra, B. R., et al. (2025). Assessing ChatGPT's potential as an autonomous grant writing tool: A pilot study. Innovative Surgical Trends. https://doi.org/10.1007/s44414-025-00007-x
2. Liang, W., Yuksekgonul, M., Mao, Y., Wu, E., & Zou, J. (2023). GPT detectors are biased against non-native English writers. Patterns, 4(7), 100779. https://doi.org/10.1016/j.patter.2023.100779
3. Wagh, P., & Hanna, H. (2024). The Future of Grant-Making: How AI is Reshaping Proposal Evaluation. MIT Solve. https://solve.mit.edu/articles/the-future-of-grant-making-how-ai-is-reshaping-proposal-evaluation
4. National Institutes of Health (2023). The Use of Generative Artificial Intelligence Technologies is Prohibited for the NIH Peer Review Process. https://grants.nih.gov/grants/guide/notice-files/NOT-OD-23-149.html
5. UNESCO (2023). Guidance for Generative AI in Education and Research. https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research
6. National Health and Medical Research Council (2026). Policy on Use of Generative Artificial Intelligence in Grant Applications and Peer Review. https://www.nhmrc.gov.au/about-us/resources/policy-use-generative-artificial-intelligence