



Multimodal & Foundation Models in Banking



REPRESENTATION
LEARNING



MULTIMODAL
FUSION



EVALUATION
UNDER
DEPLOYMENT



COUNTERFACTUAL
REASONING



EXPLANATION &
GOVERNANCE



SIMULATION &
SYNTHETIC DATA

FROM REPRESENTATION LEARNING TO RESPONSIBLE DEPLOYMENT

Call for Papers

We invite extended abstracts and short papers of at most four pages on any aspect of multimodal and foundation models applied to banking and financial decision-making. Review is double blind, and the workshop is non-archival — presenting here does not prevent you publishing the same work elsewhere.

Submissions are open. The deadlines below are final.

All deadlines are 23:59 AoE (Anywhere on Earth).

Milestone	Date
Submission site opens	13 August 2026
Paper submission deadline	2 October 2026
Author notification	14 October 2026
Camera-ready deadline	30 October 2026
Workshop day	14 or 15 November 2026

Application domains and model families

Deploying multimodal and foundation models in banking is an end-to-end, multi-stage problem, not only a model-architecture problem. The technical and methodological questions at each stage are research-worthy and are currently scattered across disparate venues.

In scope on the application side: consumer credit underwriting; fraud detection (transaction, identity, application, account takeover); collections and recovery; servicing decisions; identity verification and KYC; payment authorization; BNPL; insurance underwriting where data shapes are similar; and small-business credit.

On the model side: tabular foundation models, event-sequence transformers, multimodal fusion architectures, hybrid representations, and other foundation-model architectures applied to these domains.

Topics include, but are not limited to

- **Architectures & fusion strategies**
Novel neural architectures for integrating heterogeneous financial data — for example, aligning customer transaction histories with text from application forms, credit-bureau records, and identity documents.
- **Financial foundation models (FFMs)**
Pre-training objectives, scaling laws, and transfer-learning techniques specifically for tabular, transactional, credit, or market data.
- **Trustworthy AI & regulation**
Technical approaches to interpretability, algorithmic fairness, and robust governance in compliance with global financial regulations (EU AI Act, Basel, IFRS 9).
- **Efficient adaptation**
Parameter-efficient fine-tuning and distillation of massive models for low-latency financial environments.
- **Privacy & security**
Federated learning and differential privacy in the context of sensitive multi-party financial datasets.
- **Applications**
Innovations in credit scoring, high-frequency fraud detection, automated auditing, and generative macroeconomic modelling.
- **Evaluation under deployment conditions**
Out-of-time evaluation protocols, selection-bias-aware metrics, segment-conditional performance, calibration under operating thresholds, paired comparisons and statistical rigour for rare-label scoring, and the benchmark–deployment evidence gap.
- **Counterfactual reasoning & off-policy evaluation**
Reject inference, policy-change estimation, treatment-effect estimation for interventions (collection strategies, servicing offers, credit-line changes), adverse action explanations, and the interaction between foundation-model representations and causal identification.
- **Simulation & synthetic data**
Calibrated simulators and synthetic data for model improvement on under-represented segments and rare events; simulator-based stress testing under counterfactual regimes (recession scenarios, segment-mix shifts, novel attack patterns); validation methodology for simulator fidelity.

Submission instructions

Length: Extended abstracts and short papers, maximum 4 pages excluding references.

Format: ACM sigconf two-column template, consistent with the ICAIF '26 main track. Use the sigconf class with the anonymous parameter. [[download](#), [Overleaf](#)]

Review: Double blind. Every paper is reviewed by two expert reviewers drawn from the workshop **programme committee**.

Presentation: Accepted papers are presented as a 20-minute contributed talk with 10 minutes of Q&A, or in the animated poster session.

Anonymity: Review is double blind, following the ICAIF '26 policy. Submitted papers must not reveal the authors' identities, either directly or through obvious reference to previous

work. When citing your own work, use the third person — "Óskarsdóttir et al. [cite] found that..." rather than "in our previous work [cite], we found that...". In the ACM LaTeX template, use the sigconf class with the anonymous parameter.

Non-archival: MFMB is a non-archival venue. Accepted papers do not appear in the ACM proceedings, and presenting at the workshop does not preclude publishing the same work elsewhere — at a journal, at a conference, or at any archival venue, before or after the workshop. Work that has already appeared as a preprint or in another non-archival venue is welcome.

Publication of accepted papers: Accepted papers are posted on this website only with the explicit authorisation of the authors. Authors are asked for that permission after notification; consent is opt-in and may be withdrawn at any time. Where it is not given, only the paper's title and authors are listed.

Questions

Please direct questions about the call to the primary contact, [María Óskarsdóttir](#), or to any member of the organizing committee.