

ICCT2026 · Oral presentation · Paper #1786

A Multimodal Deep Learning Framework for Default Prediction in Peer-to-Peer Lending

Borrower text + structured features, deployed as a mobile app

Khadija, Muhammad Abid Hussain, Aiman Nadeem, Palwasha, Muhammad Azfar Yaqub, Mohsin Raza Jafri

NUST Balochistan Campus · Free University of Bozen-Bolzano

Why borrower text matters

Structured data dominates

Credit scores, loan amounts and income drive most default models.

Borrower text is discarded

Free-text descriptions may carry intent and financial discipline, but are rarely used.

A gap in the literature

Few studies compare text models on short narratives across fusion settings.

Question: can borrower descriptions improve default prediction when combined with tabular features?

How P2P lending works

Listing and screening


Borrowers request loans; the platform lists and screens them.

Credit risk assessment

The decision point where our default prediction applies.

Funding and repayment

Lenders fund approved loans and receive principal and interest.



Framework overview

1 • **Data:** 44,902 Lending Club loans


2 • **Text:** five deep learning models

3 • **Fusion:** Logistic Regression meta-classifier


4 • **App:** Flutter front end, FastAPI backend

Three-stage validation per text model

Hyperparameter search, imbalance strategy, threshold calibration.



Two-phase experimental design



Dataset and preprocessing

44,902

loans, issued 2007–2014, with
borrower descriptions

14.51%

default rate (6,515 defaulted,
38,387 fully paid)

70

engineered structured features
from 19 raw columns

Filtering: 2,260,701 → 466,345 → 451,059 → 123,470 → 44,902 loans.

Text: at least 40 words, mean 68.4 words, padded or truncated to 115 tokens.

Split: 20,000 train, 10,000 validation, 14,902 test (stratified, seed 42).

Structured-data baselines

Table 5. UnimodalStructured Data Results

Model	Validation ROC-AUC	Test ROC-AUC	Benchmark ROC-AUC
FICO-Only	0.597	0.598	0.675
Structured-NoFICO	0.703	0.685	0.698
All-Structured	0.703	0.684	0.702


0.598

FICO-only test ROC-AUC: a weak predictor

0.685

Structured-NoFICO test ROC-AUC

Text-only models



Test ROC-AUC. Bars start at zero.

A simple Average Embedding model matches DistilBERT on short borrower text.

Multimodal late fusion

0.713

best multimodal test ROC-AUC
(AE + all structured features)

0.684

structured-only baseline
(All-Structured)

0.615

best text-only model (Average
Embedding)

Without FICO: the same 0.713, so descriptions may partly compensate for missing credit-bureau data.

Operating point: recall 0.651, precision 0.236 for AE + all structured features.

Benchmark: Kriebel et al. report 0.712 for the same fusion setting.

Deployment and takeaways

Flutter app

Borrowers and lenders sign in with OTP, enter loan data and see an AI Risk Score.

FastAPI backend

Tokenizes the description to 115 tokens, then fuses the text probability with form features.

Risk output

Risk score, loan grade and default probability, returned in real time.

Threshold calibration is essential under class imbalance. Limits: 2007–2014 loans from one platform; next, larger datasets and non-linear fusion.

Thank you. Questions?