MARLUC

CASE STUDY · MARLUC SYNOPSIS

la Caixa Foundation AI pre-screening of biomedical grant proposals

IMPLEMENTED 2022–2024 · SCIENCE AND RESEARCH · SPAIN · BAND C · CLAIMED ONLY

VERSION 1.0 · PUBLISHED 25 NOVEMBER 2026 · UPDATED 25 NOVEMBER 2026 · FACTS CHECKED 25 NOVEMBER 2026

A Spanish research funder uses three language models to flag weak proposals in its flagship biomedical call. Two human reviewers check every flagged proposal before anything is removed. Published figures come from the funder and its suppliers.

ABOUT THIS SYNOPSIS

This is MarLuc's summary and grading of publicly available sources, for information only. It is not investment, legal, financial or professional advice, and should not be relied on for any decision. Readers must validate it against the original sources and, where appropriate, challenge its conclusions.

Facts and figures are as reported by the sources listed below, which may be incomplete, self-reported or out of date. Check the original source before relying on any figure.

The impact band, proof marker and SEAM view are MarLuc's opinion. The organisation is not a MarLuc client and has not reviewed or endorsed this synopsis. Names and trademarks belong to their owners.

NOTICE WORDING LAST UPDATED 16 NOVEMBER 2026

AT A GLANCE

AI used: Language-model classifiers that flag weak grant proposals, with human review of each flag.

Implemented: 2022–2024. Grade: band C, Process change, Claimed only.

Headline result, as reported: Models flag weak grant proposals and two human reviewers decide, with 35 of 98 flagged proposals put back in one call; figures are the funder's own

1 · The case

The facts, as reported by the 3 sources listed at the end. Where a source is silent, the field says "Not reported".

Organisation

la Caixa Foundation (Fundacion la Caixa). la Caixa Foundation is a Spanish philanthropic foundation. Its CaixaResearch Health programme, launched in 2017, funds biomedical research in Spain and Portugal. Canvas block: 4. Operations and process.

Business challenge

The challenge. The funder wants to reduce the administrative burden of its evaluation pipeline by detecting immature proposals before they reach peer review. Its 2023 call guidelines say that in all previous editions around 30% of proposals failed to meet the call's standards, according to remote reviewers.

Who implemented it

In-house (In-house): Runs the call and the human review.
SIRIS Academic (SIRIS Lab) (Consultancy): Co-author and technical provider, per the paper.
IThinkUPC (Software vendor): Analytics and AI provider and co-author, per the paper.

Technology (SEAM pillar E)

SEAM types: T5 Custom / fine-tuned models. Deployment: Not reported. Vendors and products: SIRIS Academic, IThinkUPC, BioLinkBERT-Base, BioELECTRA-Base.

Where it was built in: Three NLP models trained on full texts of past proposals and their peer-review scores, applied after eligibility checks and before peer review.

What changed

Before: all eligible proposals went to peer review (546 in HR22). After: in HR23 the model suggested removing 98 of 493; reviewers confirmed 63 and put 35 back. Time or cost saved: Not reported. [1] [2] [3]

What it cost

Not reported

What it saved

Not reported. The paper frames the aim as reducing reviewers' workload; it gives counts of proposals removed from peer review, not hours or euros. [1] [2] [3]

Governance and control

The funder ran a bias check on historical reviews before the trial, uses two human reviewers for every flagged proposal, and plans annual retraining. Its 2023 call guidelines state that no redress or appeal procedure is admitted for dismissed proposals; the paper's authors say redress mechanisms must be put in place.

S2 · Accountable owner · Claimed
Foundation staff run the programme and the pilot; the paper is written by them with their suppliers. [2]

A1 · Safe use · Claimed
The funder analysed historical reviews for bias before the trial and says it found no evidence of systematic bias. [2]

A2 · Human in the loop · Claimed
Two human eligibility reviewers examine every proposal the models flag; either can return it to peer review, per the paper and the 2023 call guidelines. [2] [3]

A3 · Monitoring and assurance · Claimed
A hidden parallel run in HR22 compared model flags with final outcomes; the models are retrained each year. [2]

A4 · Vendor AI · Not reported
Suppliers SIRIS Academic and IThinkUPC are named as co-authors; contract and oversight terms are not reported. [2]

A5 · Agent controls · Not reported
Not reported.

A6 · Standards and regulation · Claimed
The authors assess the system against EU AI Act principles and say it is not a high-risk system under Article 6 and Annex III. [2]

What worked, what did not, what is unproven

WORKED

– In the HR22 hidden run, the models unanimously flagged 116 of 546 proposals; reviewers put 30 back and none of those reached the panel stage.

– In HR23, 63 of 98 flagged proposals were confirmed for removal before peer review.

– The funder says a pre-trial analysis found no evidence of systematic bias in past selection.

DID NOT

– In HR22, two proposals that two independent reviewers judged should be discarded advanced to panel and were funded in the normal track; one had been flagged by the models.

– Of the 35 proposals put back in HR23, one reached the panel phase but was not funded, so reviewers caught some wrong flags.

– The authors say evaluators worry that being told a proposal was pre-discarded may bias them, and that redress mechanisms for removed applicants are still needed; the 2023 call guidelines admit no redress or appeal.

UNPROVEN

– Whether the step saves reviewer time or money, as no figures are given.

– Whether the models disadvantage unconventional or non-native-English proposals.

– Whether results hold in later calls or other programmes.

Timeline

2017 CaixaResearch Health programme launches as the funder's flagship biomedical call [2]

2022 HR22 call: hidden parallel run of the pre-screening models beside the normal process [2]

2023 HR23 call: model used in the live workflow; 98 of 493 proposals flagged [2]

2023 The funder's call guidelines add a content eligibility phase: an AI tool flags proposals with a very low probability of selection, two experts validate each flag, and no redress or appeal is admitted [3]

Oct 2024 Paper published; the funder decides to apply AI pre-screening systematically with mitigating measures [1]

2 · Grade and SEAM assessment

MarLuc's opinion, made on the evidence above. How cases are graded: marluc.co.uk/case-studies/how-we-rate

Level of change and impact

Claimed: Not applicable. Grade: C, Process change. Proof: Claimed only. Route: Process-led.

Process change (Level 2, band C). The selection process has a new model-assisted pre-screening step with human review of every removal. Effects on cost and fairness are not independently shown.

The SEAM view and lessons learnt

The drivers are a bottom-class vote needed from all three models and a human check on every removal, so the step is a filter with an override. The paper's own data are the evidence, and benefits to cost or funding quality are not measured. Before copying it, a leadership team should ask how often reviewers overturn the model, and what applicants who are removed can do about it.

1. Run the model in parallel before it affects any applicant.
2. Require a human check on every proposal the model removes.
3. Give removed applicants an explanation and a route to appeal.

Evidence

Score 2 of 4. Self-reported with a parallel comparison: a peer-reviewed paper by the funder and its suppliers (published in two places) and the funder's own call guidelines; no independent audit.

Evaluation method. Self-reported by the funder and suppliers in a peer-reviewed conference paper, with a hidden parallel run compared with actual outcomes; no independent audit.

Sources disagree. The paper names the models as BioLinkBERT-Base, BioELECTRA-Base and BioLinkBERT-Base with Adapter; the funder's 2023 call guidelines name BioBERT, BioELECTRA and BioBERT with Adapter blocks.

Status. Ongoing: the paper says the funder has decided to apply AI pre-screening systematically.

3 · Original sources

Read these in full. To challenge a fact or a grade, use the contact form at marluc.co.uk or the link on https://marluc.co.uk/case-studies/la-caixa-foundation-ai-prescreening-grants

[1] Cambridge Core (Data & Policy): AI-assisted prescreening of biomedical research proposals: ethical considerations and the pilot case of la Caixa Foundation
https://www.cambridge.org/core/journals/data-and-policy/article/aiassisted-prescreening-of-biomedical-research-proposals-ethical-considerations-and-the-pilot-case-of-la-caixa-foundation/D17CC86C884D4F9E3A6DD9ABF3F4605E

[2] Zenodo: Carbonell Cortes et al., full text of the la Caixa pre-screening paper (Data for Policy 2024)
https://zenodo.org/records/13736709/files/AI-assisted%20pre-screening%20of%20biomedical%20research%20propos-%20als;%20ethical%20considerations%20and%20the%20pilot%20case%20of%20%E2%80%9Cla%20Caixa%E2%80%9D%20Foundation.pdf

[3] la Caixa Foundation: CaixaResearch Health Call for Proposals 2023, eligibility and evaluation guidelines (copy hosted by irsjd.org)
https://irsjd.org/media/upload/domain_3/pdf/evaluation%20guidelines_1663845898.pdf

la Caixa Foundation (Fundacion la Caixa) website: https://fundacionlacaixa.org

VERSION HISTORY

v1.0 25 November 2026 First published