

A204 · Perturb-LM: Leakage-Aware Language Retrieval of Cell Painting Morphology

Presenter: Makenna Rodriguez — National Institutes of Health

Authors: Makenna A. Rodriguez, Adam Diaz

Session: Day 1 · Thu Oct 1, 2026 · 2:15–4:15 PM

High-content Cell Painting allows researchers to measure how cells respond to different treatments, but it remains difficult to search these large datasets. Using natural-language interfaces could help make phenotypic data more accessible. However, retrieval results might look better than they are if they rely on treatment names, metadata, or other shortcuts. To address this, we created Perturb-LM, a framework that assesses whether biomedical language models can identify relevant Cell Painting data without relying on simple text similarity.

We tested Perturb-LM on 4,524 quality-checked JUMP CPJUMP1 profiles, each described by 904 morphology features. For the retrieval analysis, we used 4,190 profiles that had perturbation labels. We built queries using descriptive biological metadata but omitted treatment and sample names, target sequences, chemical identifiers, plate, well, batch, and other features that could be used for direct lookup. We compared a TF-IDF method with identifiers removed, frozen BiomedBERT embeddings, and a ridge projection that mapped BiomedBERT outputs into morphology space.

For the main evaluation, we excluded candidates from the same plate and well, allowing us to assess 180 out of 1,079 queries. The TF-IDF method reached a mean average precision of 0.1574 (5,000-bootstrap 95% CI, 0.1367–0.1795). In comparison, unaligned BiomedBERT scored 0.0092, and projected BiomedBERT scored 0.0332. Projected BiomedBERT stayed below TF-IDF by -0.1243 mAP (paired 95% CI, -0.1519 to -0.0946), and this pattern was consistent across four different evaluation settings.

These results show that it is important to use strong lexical controls and leakage-aware evaluation when applying biomedical language models to complex cellular phenotyping data.