Epigenome — V4 Target Opportunity Report (Public Edition)

Coverage/whitespace, trajectory & structural-decay signal, in one decision-first view

Target: Epigenome — programmable epigenome editing (non-DNA-cutting writers/erasers/recruiters) + small-molecule epigenetic drugs (HDAC, EZH2/EED/PRC2, DOT1L, menin-MLL, BET/BRD, DNMT, LSD1, TET). **Date:** 2026-07-23 · **Substrate: 95 on-target patent families**, real filing span **1989→2025**. Expanded this run from an innovator-core base by exhaustively re-ingesting the fields flagged as round-1 sampling gaps (programmable editing/recruiters, epigenetic-target degraders/PROTACs, BET/DOT1L/TET inhibitors). **What this is for:** to let an inventor or BD&L reader *see* where the foundational core is lapsing, where the field is heading, **and — now that the gap fields are sampled — which empty cells are genuine whitespace versus coverage holes we still cannot see into.**

> **Read-me (the guardrails, once):** This is **candidate generation + a coverage / trajectory / decay signal — NOT a freedom-to-operate (FTO) opinion and NOT a whitespace clearance.** Structural overlap is a **lower bound**. Whitespace = absence of *strong structural claim coverage in this corpus* (necessary-not-sufficient), never FTO. **The headline whitespace finding of this rerun is a negative one and it is deliberate** — see §2.
>
> **Public edition note.** This is the public edition of the Epigenome V4 Target Opportunity Report. It presents the coverage, whitespace, trajectory, and decay reads that accompany the article. The underlying patent-craftsmanship (KC) strength scoring, per-family score tables, machine-readable knowledge graph, and raw scored inputs are part of the internal analyst deliverable and are not reproduced here; the strength layer is relative-ordering, calibration-pending work and is deliberately not published as a ranking.

PART I — ACT ON THIS

1. Actionable summary

Ranked by how well-grounded each call is in the expanded 95-family corpus. Each is a *starting hypothesis for a dedicated search*, not a clearance.

#	Observation you can act on	Grounding	What to do / verify
1	**The foundational HDAC & EZH2 chemical genera are lapsing/expired** — romidepsin (expired), vorinostat/SAHA (lapsed), GSK126 (lapsed 2018),	**Strong** (real legal timelines, §4)	First-generation composition-of-matter is no longer defended → **non-obvious HDAC/EZH2 chemotypes have design-**

#	Observation you can act on	Grounding	What to do / verify
	oxadiazole HDACi (lapsed 2019); +several more granted-but-lapse-contested families flagged this run.		**around room around a decaying core.** Confirm each lapse + any live continuations.
2	**The strongest, freshest claims are menin-MLL and LSD1** (well-drafted, maintained, deep-Markush estates; menin-MLL is Orange-Book-linked, LSD1 granted to 2039).	**Strong** (native claim-strength + legal)	Entering these nodes means facing well-drafted, maintained, deep blockers — **an inhibitor me-too is not viable**; differentiate the modality.
3	**BET/BRD is the densest inhibitor estate in the class** — 43 families, the only cell COVERED on **multiple** broad live genera (3), spanning oncology, cardiometabolic and inflammation.	**Strong** (now fully sampled: 43 families vs 0 in round-1)	BET small-molecule inhibitor space is **crowded and well-defended** — treat as a mature, contested node; do not plan a generic bromodomain me-too.
4	**Direction of travel is inhibitor → degrader / recruiter-editor.** Degraders go 0→9 filings and recruiter-editors 0→11 in the 2019–2024 era; inhibitors fall from 100% (pre-2013) to 62% of recent filings.	**Moderate–Strong** (trajectory on real 1989–2025 dates, N=95)	Treat "degrade or recruit rather than inhibit the enzyme" as the **field's momentum vector.** The editing-recruiter estate is real and named.
5	**China/Korea follow-on filers are clustering** on EZH2, LSD1, BET and degraders.	**Moderate** (assignee origin)	Fast-followers on the mature nodes; freedom is narrowing where they cluster.
—	**Degrader-direction empties on DNMT / menin / HDAC / LSD1 / EED / TET**	**COVERAGE-LIMITED / NOT-SEARCHED — NOT whitespace**	See §2. Do **not** treat as open room; these are holes in what we could ingest, confirmed this run.

The one-sentence read: the defensible signal is that **the first-generation small-molecule *inhibitor* foundation (HDAC/EZH2) is decaying while the strongest, freshest claims sit on the newest nodes (menin-MLL, LSD1), BET is a crowded mature estate, and the field's filing momentum is turning from *inhibit* toward *degrade / recruit-and-edit*** — and, crucially, that **no cell in this class can honestly be called clean whitespace yet**, because the one modality where the empties cluster (degraders) is the one modality we could not finish ingesting.

2. The whitespace finding is a NEGATIVE one — and that is the point

A V4 report normally leads with a whitespace map naming open cells. **This rerun deliberately declines to name any.** That is the honest result, and it is the direct product of doing what round-1 could not:

- **Round-1 could not support whitespace at all** — empty cells were sampling gaps. The audit that run **retracted** a "recruiter/editing whitespace" call for exactly that reason.
- **This run closed the sampling gaps as far as the tooling allows:** exhaustive discovery found **215 families** across the flagged fields; **~100 new families were ingested** (BET 0→43, editing-recruiter 3→14, degrader 1→9, DOT1L 0→9). The previously-empty cells are now **populated and tested.**
- **With the gaps closed, the empties that remain resolve into three honest categories — none of them "whitespace":**

Empty-cell class	Count	Meaning	Example
COVERAGE-LIMITED	4	Families provably **exist** but could not be ingested (null publication numbers / discovery-leg errors). Not whitespace — a hole in our lens.	HDAC-degrader: 5 real families found this run **all with null pubnums → un-ingestable.** LSD1/EED/PRC2 degrader legs errored upstream.
NOT-SEARCHED	19	The exact query was never issued (e.g. "menin degrader"), OR the direction is mechanistically possible but unqueried. "We didn't look" ≠ opportunity.	DNMT-degrader, menin-MLL-degrader, recruiter-of-enzyme cells.
IMPLAUSIBLE-BY-MECHANISM	2	Incoherent combination.	A programmable editor cannot be a small-molecule inhibitor or a PROTAC.
GENUINE-OPEN (whitespace)	**0**	—	—

- The audit **REJECTED the first version of this rerun** because 5 of 6 "open" degrader cells coincided with a documented coverage hole — the round-1 error repeating. We then **attempted to close the hole** and, confirming it is unclosable with current tooling, **downgraded all 6 to coverage-limited/not-searched.** The audit then **ACCEPTED**: the read is legitimate *because it declines to assert whitespace on an incompletely-sampled modality.*

What the corpus *does* legitimately support is the coverage-strength grid (below), the lapse/decay map (§4) and the trajectory (§4).

Coverage grid (what EXISTS and how strong, node × direction):

mechanism cell	verdict	scored core	live grants	max genus recall (lower bound)
BET-BRD / inhibitor	**COVERED** (densest; 3 broad live genera)	18	9	0.773
menin-MLL / inhibitor	**COVERED** — strongest/newest	3	3	0.716
EZH2 / inhibitor	**COVERED** (GSK126 lapsed)	5	3	0.57
LSD1-KDM1A / inhibitor	**COVERED** — consolidating	4	4	0.515
DNMT / inhibitor	**COVERED (single family — weak)**	1	1	0.857
HDAC / inhibitor	**PARTIAL — decaying** (no broad live genus; lapses)	6	2	0.903
programmable-editor / recruiter-repressor	**PARTIAL** (13 families, no compiled genus)	9	2	—
DOT1L / inhibitor	**THIN** (the broad genus is an application)	2	0	0.852
BET / degrader · EZH2 / degrader · PRC2-EED / inhibitor · editor / recruiter-activator	**THIN** (1–4 families)	1–4	0–1	0.14–0.55
DOT1L / degrader · TET	**DEFERRAL_GAP / ADJACENCY_ONLY** (present but biblio-only / method-only)	0	0	—

3. The chemical core — where live structural coverage sits

Across the genus-bearing subset, **median family_best recall ≈ 0.39 (a lower bound)** — even COVERED cells leave in-scaffold analog room. This is a *coverage-breadth* read, not FTO. The broadest live genera sit in the HDAC (lapse-contested), DNMT, BET, and menin-MLL cells. The BET-inhibitor cell is the only one resting on **three** independent broad live genera — structurally the best-defended node in the class.

PART II — METHOD & LIMITS (public summary)

How the reads were built

- **Full-landscape discovery** (USPTO ODP + BigQuery index + EPO OPS, US/EP/WO/CN) → **gap-field expansion ingest** (215 discovered → ~100 new

families ingested) → Wikidata-grounded 4-axis scope ontology → claim-scope binning (95 families) → structural/genus + coverage grid + trajectory + knowledge graph → **three independent adversarial audits.** Candidate generation, NOT FTO; no prior heuristic score is an input.

- **Full-coverage node grid** (all 95 on-target): BET-BRD 43 · programmable-editor-effector 14 · DOT1L 9 · EZH2 8 · HDAC 8 · menin-MLL 4 · LSD1-KDM1A 4 · DNMT 3 · PRC2-EED 1 · TET 1. By direction: inhibitor 69 · recruiter-repressor 14 · degrader 9 · recruiter-activator 2 · diagnostic 1.
- **Trajectory (real filing dates 1989→2025, N=95):** pre-2013 was 100% inhibitor; 2013–2018 added the first recruiter-repressor editors; 2019–2024 saw inhibitors fall to 62% while degraders rose from 0→9 and recruiter-repressors to 11. The class is diversifying away from pure enzymatic inhibition toward targeted degradation and programmable recruitment. (Caveat, non-blocking: the recent-era magnitude is somewhat inflated because this run deliberately went looking for recent degrader/editor filings — the *direction* of the shift is robust; the absolute recent count is an upper estimate.)
- **Lapse / structural decay:** foundational estates confirmed lapsed/expired — romidepsin (expired), vorinostat/SAHA (lapsed), GSK126 EZH2i (lapsed 2018), oxadiazole HDACi (lapsed 2019); a legal-consistency validator additionally flagged several families whose `granted_active` verdict is contradicted by lapse/discontinuation timeline events (read as contested and excluded from "live-granted" coverage).
- **Independent adversarial review:** the whitespace audit **REJECTED → fixes → ACCEPTED** — it caught 6 "open" degrader cells coinciding with documented coverage holes; after an unsuccessful attempt to close them, all 6 were downgraded to coverage-limited/not-searched (0 genuine-open), and application-kind guards were added (an application no longer counts as a live grant).

Named coverage limits (absence-first — not hidden)

- **A third of the on-target families are bibliographic-only** (deferred; CN/EP/WO members with no English claims) — counted for **coverage** but not **strength**.
- **The degrader modality was NOT exhaustively ingested.** HDAC-degrader and other LSD1/EED/PRC2 degrader families are **null-publication-number applications** the ingest tool cannot accept, and several discovery legs returned upstream errors. Degrader empties are therefore coverage holes, not opportunities.
- **Commercial epigenome-editing estates** that are US applications with null publication numbers are **not ingestable** and remain under-counted — the editing-recruiter coverage is a lower bound.
- **TET is genuinely thin** (1 family, method-combination only) — a lower-bound signal, partly a real scarcity and partly a coverage artifact.

The honest bottom line. This rerun did what was asked — it exhaustively ingested the flagged fields and re-ran the whole playbook. The result is a *legitimate* whitespace analysis, and its legitimate answer is: **on the chemistry that could be fully sampled (the inhibitor estates), coverage is dense and maturing, with BET/menin/LSD1 strong and**

HDAC/EZH2 decaying; and no cell qualifies as clean whitespace — the empties cluster in the degrader modality, which is precisely the modality current tooling cannot finish ingesting. The correct next action is a **degrader-focused ingest once null-pubnum applications and coverage-limited jurisdictions are reachable**, not an opportunity claim on today's holes.

This public edition presents the Epigenome coverage, whitespace, trajectory, and decay reads. Every number is computed on the 95-family on-target corpus on real filing dates. Candidate generation, NOT freedom-to-operate.