Silicon Valley Bank, read as a topology

An OptaCat v0.9.1 analysis set beside the official reviews

Corpus: 36 Scenes, 91 Lines, 115 events, 104 callbacks, 72 canonical entities

What OptaCat is

OptaCat is a topological cartography engine for corpus-based processes. It is not a summary system and not a character-frequency counter. It converts a document or a transcript into three strictly separated layers — a corpus container, a canonical identity hierarchy, and an event topology of atomic directed events with callback weights — and then computes a small set of derived readings from those frozen layers.

Three properties matter for reading what follows.

Events are atomic and directed. Every event is one source, one target, one label, anchored to one Line of the source text. Nothing is an event because two entities appear in the same paragraph.

Callbacks carry the structure. When a later event materially depends on an earlier one, the annotator records a callback, and only when removing the earlier event would cost the later one a specific part of its meaning. The event's weight is then mechanical: one plus the number of callbacks it receives. Weight is local. Load, the ninth reading, is the transitive closure of the same graph: how many events in the whole corpus depend on this one, directly or through a chain.

The policy is frozen and hashed. Every tunable coefficient lives in one declared file. This run used policy `optacat_svb_v0_9_1`, ledger mode `source_weight_target_flat`, and declared the four acts of the closure as the outcome events.

The corpus was built from three public-domain sources: the California DFPI review of its oversight of SVB (May 8, 2023), the Federal Reserve Office of Inspector General material loss review (September 25, 2023), and FDIC testimony to the Senate (March 27, 2023). Every Line is a verbatim excerpt with a page reference. Scene order is event time, from the 2018 threshold change to the bridge bank of March 12, 2023.

Official findings beside the OptaCat readings

The two reviews state seven findings between them. The table sets each against what the topology says, with the figure that carries the claim.

Official finding	OptaCat reading	Verdict
Fed OIG 1. The regional-bank supervisory approach did not evolve as SVB grew.	The growth itself is load-bearing: crossing $100 billion carries load 30, tier 1. The specific supervisory lapses the finding names carry nothing: the pandemic pause and the 2021 "satisfactory" rating both carry load 0.	Split
Fed OIG 2. The transition from the regional to the large-bank portfolio was ineffective.	Entering the large-bank portfolio carries load 28, tier 1, and the understaffed supervisory team hangs from it.	Agrees
Fed OIG 3. Examiners should have scrutinised interest-rate risk on the held-to-maturity portfolio.	The three highest-load Scenes of the whole corpus are this episode: removing the hedges (45), the counterintuitive deposit-duration assumptions (35), the year-end losses they produced (33). But the supervisory side of the same finding is weightless: the November 2022 interest-rate-risk letter carries load 2.	Agrees on the risk, contradicts on the response
DFPI 1. SVB was slow to remediate, and regulators did not ensure speed.	Invisible to load: the liquidity letter carries 3, the management action plan 0, the open MRIAs 0. Visible to recurrence, where the supervisory letter and the CAMELS rating are the two most recurrent objects after deposits.	Appears in another reading
DFPI 2. Rapid growth was not sufficiently accounted for.	The same two Scenes as OIG 2: load 30 and 28, both tier 1.	Agrees
DFPI 3. A high level of uninsured deposits contributed to the run.	Not testable here. One single event in the corpus touches uninsured deposits, and its Scene carries load 0.	Silent
DFPI 4. Digital banking and social media accelerated the run.	Partly testable. The two run Scenes carry load 23 and 19, so the panic is structurally real, but OptaCat has no measure of speed at all: an event that took eight hours and one that took eight months are the same event.	Partly; speed not measurable

Three readings OptaCat produces that neither review states.

The hidden attractor is the instrument, not the money. The strongest attractors are the CAMELS rating (14.3) and the supervisory letter (13.1), ahead of the liquidity framework (8.9) and the interest-rate simulations (6.3). There is no structural attractor at all: no object in this corpus is both highly visible and strongly attracting. What drew the most action in this story were the instruments of supervision themselves.

Nothing relational survives. Four ordered relations reach the minimum support of three transitions, and all four score zero. Not one "who acts on whom" pair persists across a single Scene transition. A supervisory system is meant to be a standing relationship; in the event topology it is a sequence of one-off contacts.

The dark horse is the Board's surveillance team. It issued the June 2022 report naming SVBFG among the banks with the highest unrealized losses and placed it on the watch list; the dedicated supervisory team did not act on it. Both reviews record the facts. Neither treats the team as a principal. The metric does, because its few acts all point the same way and it is almost never on stage.

The nine readings in full

Every reading below comes from the same frozen run. Each names the file it is written to and the score family it belongs to; the specification forbids mixing families into a single ranking, so they are reported side by side rather than combined.

1. What is the centre? 2. What is the bridge?

`node_metrics.yaml` (`center_score`), `bridge_cluster_analysis.yaml` (`bridge_score`)

Rank	Centre, actors		Centre, artifacts		Bridge	
1	Silicon Valley Bank	0.97	CAMELS rating	0.18	Silicon Valley Bank	1.00
2	bank examiners	0.69	supervisory letter	0.17	bank examiners	0.64
3	the holding company	0.34	withdrawal requests	0.14	the holding company	0.32

Both answers are structurally forced: in a corpus built around one institution, that institution is in nearly every Scene. The informative part is the artifact column, where the two supervisory instruments lead the bank's own money.

3. Where is the rupture?

`rupture_deformation_analysis.yaml` (`rupture_score`)

Fourteen of 35 transitions score exactly 1.000, so no single rupture can be named. The two lowest-scoring transitions are more interesting than the highest: 0.55 into the hedge removal of March 2022 and 0.58 into the January 2023 ratings meeting. Those are the two moments where the cast barely changes — the same parties simply continue.

4. Where is the hidden or structural attractor?

`hidden_attractor.yaml` (`hidden_attractor_score`, `structural_attractor_score`)

Hidden attractors	score	Actor attractors	score
CAMELS rating	14.3	Silicon Valley Bank	58.6
supervisory letter	13.1	the holding company	21.3
SVB liquidity risk framework	8.9	SVB senior management	12.5
SVB interest-rate simulations	6.3		

The structural list is empty: no object here is both highly visible and strongly attracting.

5. Which object connects different Scenes?

`motif_recurrence.yaml` (`recurrence_score`)

Object	score	Scenes present	bridge Scenes	reactivation Scenes
SVB deposits	0.160	9	5	0
supervisory letter	0.118	7	3	0
CAMELS rating	0.111	7	2	0
liquidity risk framework	0.062	4	1	0

The last column is zero for every object in the corpus: not one artifact is carried by a callback from one Scene to an event in another. Recurrence here is presence and movement between spaces, never causal reactivation. That is the same absence question 7 reports from the other side.

6. Which space deforms Actors?

`space_deformation_analysis.yaml` (`deformation_score`)

Space	score	Scenes	affected Actors
FRB San Francisco	0.722	14	16
Congress	0.637	1	2
Silicon Valley Bank	0.592	11	14
DFPI	0.512	1	4

The Reserve Bank leads on breadth; Congress reaches second place on intensity alone, out of a single Scene — the 2018 threshold change.

7. What remains invariant across Scene transitions?

`invariant_transition_analysis.yaml` (`invariant_score`)

Node	score	opportunities	survivals
Silicon Valley Bank	0.55	20	11
bank examiners	0.20	12	4
FRB San Francisco	0.15	10	3
SVB board of directors	0.10	5	2

Four ordered relations reach the minimum support of three opportunities; all four survive none of them and score zero.

8. Whose appearance perturbs the field, and who is the dark horse?

`perturbation_actor_analysis.yaml`, `dark_horse_analysis.yaml`

Perturbators	score	shifted	Dark horses	score	vis.	acts
FRB San Francisco	0.32	12	Board surveillance team	0.225	0.06	3
DFPI	0.29	15	FDIC	0.216	0.11	9
SVB senior management	0.28	14	Board of Governors	0.211	0.17	8

The visibility gate excluded the examiners (ungated 0.30) and the Reserve Bank (0.22): both are too present in the text to be dark horses, which is the point of the gate.

9. What is everything standing on?

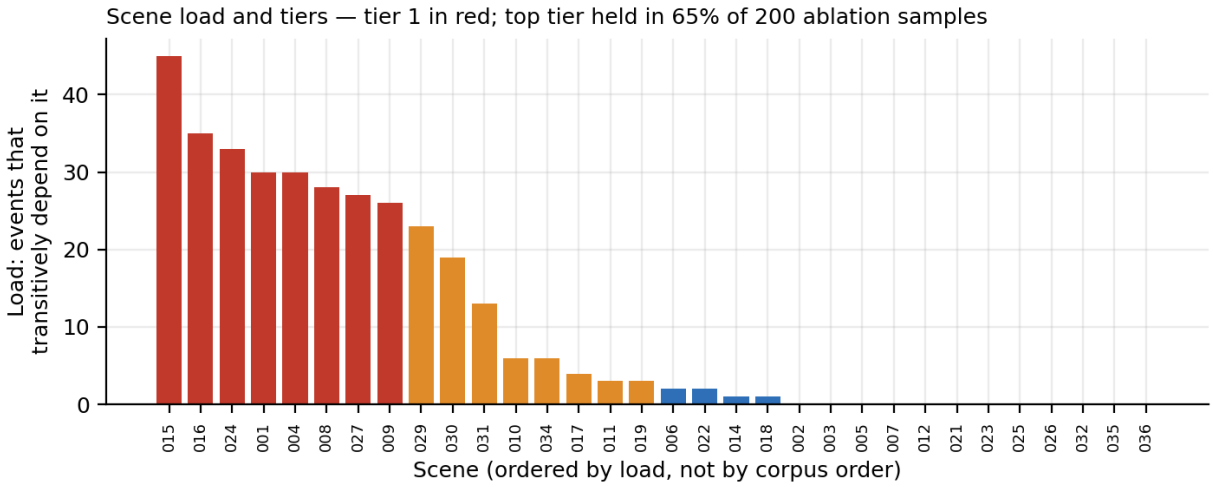
`load_analysis.yaml` (`load`, tier membership)

Tier-1 Scene	load	Actor load	
015 The hedges come off	45	Silicon Valley Bank	61
016 Counterintuitive assumptions	35	the holding company	34
024 Year end: the losses	33	Board of Governors	30
001 The threshold is raised	30	bank examiners	27
004 Past one hundred billion	30	depositors and customers	23
008 Into the large-bank portfolio	28		
027 The announcement	27		
009 The liquidity examination	26		

Outcome support, against the four acts of the closure declared as outcome events: 20 of 115 events, with set stability 0.58 mean Jaccard under ablation.

The figures

Scene load and tiers

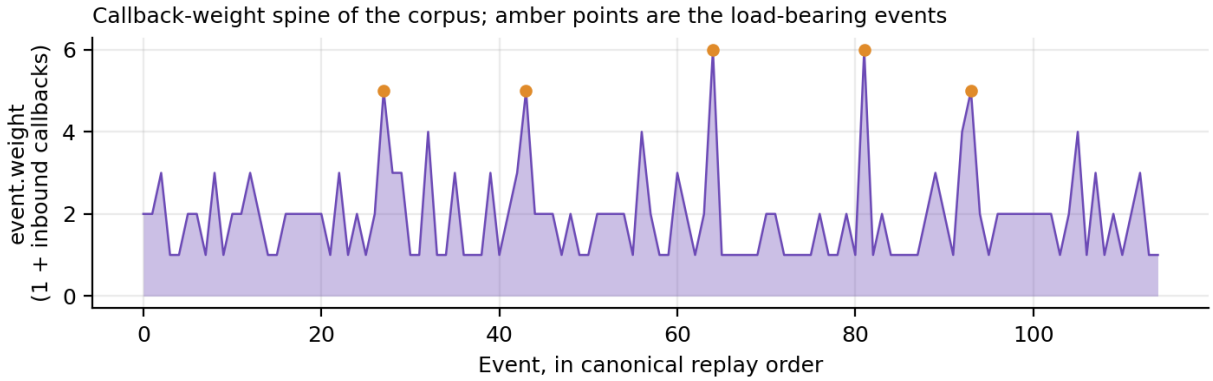


Load is the ninth reading: for each Scene, the number of events elsewhere in the corpus that depend on it, directly or through a chain of callbacks. Bars are ordered by load and coloured by tier; tier 1 is red. Tiers, not ranks, are the finding — the specification forbids presenting "the third most load-bearing Scene" as a result, because the ordering within a tier is not stable. Stability is measured rather than assumed: 200 deterministic ablation samples each delete about a fifth of the callbacks and recompute everything, and the top tier survived with 65 percent mean overlap.

The three tallest bars are Scenes 015, 016 and 024: the hedges coming off in March 2022, the deposit-duration assumptions of April 2022, and the losses at year-end. The March 8 announcement, which triggered the run, is seventh.

One check matters for reading this chart. Load is not a measure of how much the reports wrote about a Scene: across the 32 eligible Scenes, load correlates with the number of events in a Scene at +0.05 and with the number of Lines at +0.21. Scene 005 contains six events and carries load 0; Scene 015 contains three events and carries load 45.

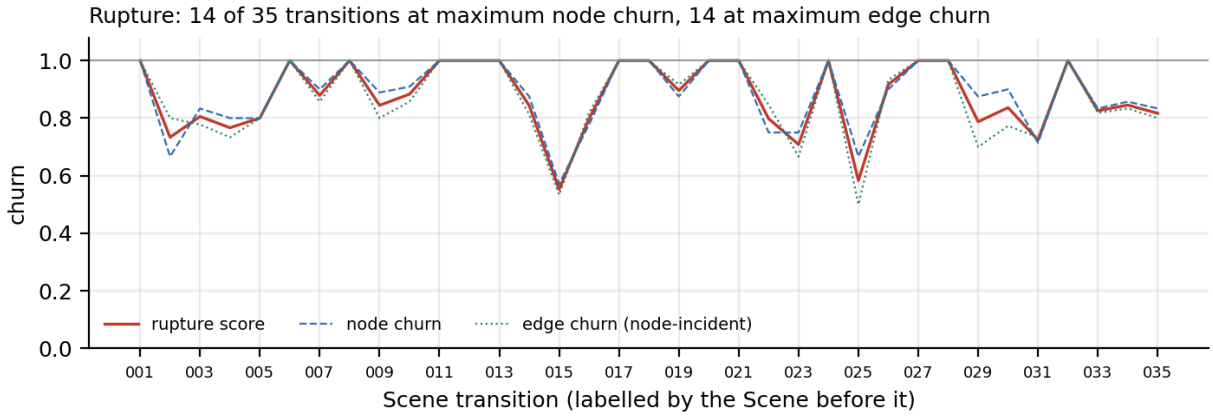
The callback-weight spine



Each point is one event in replay order; the height is `event.weight`, which is one plus the number of later events that call back to it. This is the local reading of the same graph the load chart reads globally. A tall

point is an event that many later events refer back to directly; the amber points are the five tallest. Flat stretches are passages where the narrative moves forward without reactivating anything earlier.

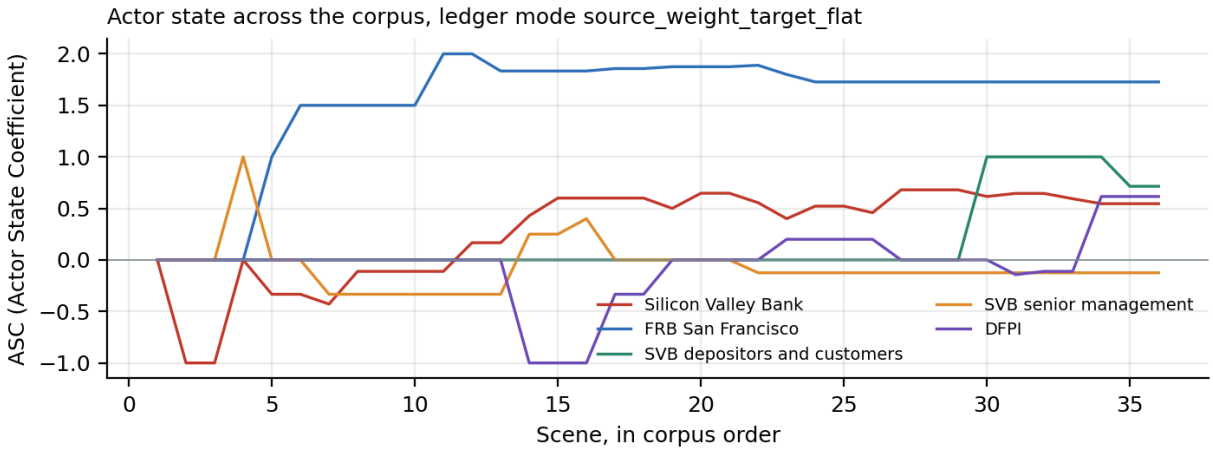
Rupture



Rupture asks where the field reconfigures between adjacent Scenes. Two components are measured across a window before and after each transition: **node churn**, the Jaccard distance between the sets of participants, and **edge churn**, the weighted difference between the event weight carried by each participant. A value of 1.0 means the two sides share nothing at all.

This corpus saturates. Fourteen of 35 transitions sit at maximum node churn, because each Scene happens in a different institutional room and adjacent Scenes routinely share no participant. The v0.9.1 node-incident edge metric brought edge saturation down from 100 percent to 40, which is why the advisory now passes, but the reading remains weak here: fourteen of the 35 transitions still tie at exactly 1.000. The honest conclusion is that this question does not separate transitions on corpora of this shape, and the specification requires saying so rather than naming a winner.

Actor state



ASC, the Actor State Coefficient, is the accumulated signed value of an actor's participation divided by the number of times they participated. Acting on someone adds the event weight; being acted upon subtracts a flat one under this ledger mode. It is a positional measure, not a moral one: a high ASC means an actor more often initiates than receives.

The Reserve Bank climbs early and stays high — it is the body that issues letters, ratings and requirements. The bank's own line sits near zero throughout: it receives as much as it initiates. The depositors' line is flat at zero until Scene 29 and then rises sharply, which is the run entering the ledger. That step is also the clearest illustration of this corpus's declared limitation: thousands of depositors withdrawing $42 billion in eight hours enter the topology as a collective actor performing a handful of ordinary events.

Terms used in the figures

Term	Meaning
Scene	A container of the corpus in event time. Here, one episode of the supervisory record or of the March window.
Line	The smallest segment of source text used as evidence. Every Line in this corpus is a verbatim excerpt with a page reference.
Event	One atomic directed relation: one source, one target, one label, anchored to one Line.
Callback	A recorded dependency from a later event to an earlier one, admitted only when removing the earlier event would cost the later one a specific part of its meaning.
`event.weight`	One plus the number of callbacks an event receives. Local: it counts direct references only.
Load	The transitive version: how many events depend on this one directly or through a chain. Reported by tier, never as an ordinal rank.
Tier	A band of the load distribution. Tier 1 is the highest. Equal loads always share a tier.
Tier stability	The share of ablation samples in which the top tier survives. Each sample deletes about a fifth of the callbacks at random and recomputes the whole reading.
Outcome support	The set of events upstream of the declared outcome. Here the outcome is the closure of the bank on March 10; the support set holds 20 of 115 events.
Node churn	Jaccard distance between the participant sets of two adjacent Scene windows. 1.0 means no shared participant.
Edge churn	The weighted difference in event weight carried by each participant across the same two windows.
ASC	Actor State Coefficient: accumulated signed participation value divided by the number of participations.
Attractor	A node with strong incoming event weight. Hidden if it stays below the visibility threshold; structural if above it.
Dark horse	An actor with concentrated positive state and low visibility, under a declared visibility gate.
Visibility	The share of Scenes in which an entity is named in the text — distinct from how often it acts.

How much of this is independent of the reviews

Three qualifications belong with any comparison above.

The corpus comes from the reviews. Events exist here because the reviews narrated them. If a fact went unrecorded by the investigators, it cannot be load-bearing in this corpus. What was checked is narrower and does hold: load does not simply track how much each Scene was written about, as the correlations of +0.05 and +0.21 show. The weighting inside the selected material is structural; the selection is not ours.

One agreement followed a correction. The first run of this corpus produced an outcome-support set of 13 events out of 113, and inspection showed no callbacks at all between the 79 events of the supervisory history and the 34 events of March 2023: the closure rested, in the annotation, on the March cascade alone. The sources state the missing link — unrealized losses grew after the hedges were removed, and the bank "aimed to cap its losses by selling bonds" — but the corresponding Lines were not in the corpus. They were added through the correction loop, two events and three callbacks followed, and the top of the load ranking moved from the 2018 threshold change to the March 2022 hedge removal.

That correction was prompted by a result that looked wrong, and "wrong" partly meant "unlike what the reviews describe". A correction loop can import the reviews' conclusions into the corpus. The protection is procedural rather than technical: the added Lines, the new events, and the counterfactual reason for each callback are recorded in the audit files, and anyone can remove them and see what the reading becomes.

Load is dependency, not causation. A Scene can carry a large load without being a cause, and a cause can carry little load if the corpus records no later event that depends on it. The three highest-load Scenes of this corpus are the interest-rate episode of 2022. That is a statement about what the rest of the record hangs from, and it is consistent with the OIG's third finding — but it was reached by counting dependants, not by judging what supervisors owed.