

Completing the Chen–Mellman Cancer-Immunity Cycle

Sixteen positions, two cross-cancer locks, and a structural account of why anti-PD-1 fails where it fails

Raimo van der Klein

Founder, Encounter — encounter.bio · raimo@encounter.bio

Generative Geometry · May 2026

Contents

Abstract	2
1. Introduction: Seven questions the cycle leaves open	4
2. The base architecture: four regimes, four barriers, one agent	5
2.1 The four regimes	5
2.2 The four barriers	6
2.3 The cycle is generational	7
2.4 What the rest of the paper does with this architecture	7
3. Four empirical results in cancer	7
3.1 The cytolytic synapse is locked by a structurally derived biomarker, measured at two resolutions	7
3.2 Position 8 is real and measurable: built effectors that never crossed the Egress barrier	10
3.3 The single largest patient bottleneck is in the lymph-node-and-tumour-bed Construction phase	12
3.4 In most cancer types, the response stalls upstream of where anti-PD-1 acts	13
3.5 Why these four results converge	14
4. The sixteen-position completion of the cycle	14
4.1 The sixteen positions of the cancer-immunity cycle	14
4.2 What this completion does to Chen–Mellman	15
5. Detailed treatment of the positions that carry empirical weight	17
5.1 Position 4 — the Activation barrier (SUPPORTED, single-cohort)	17
5.2 The Construction-phase aggregate — Positions 5–8 (the 43% bottleneck)	17
5.3 Position 8 — the Egress barrier (LOCKED, cross-cancer)	19
5.4 Position 10 — the cytolytic synapse (LOCKED)	20
5.5 Position 12 — the Drain barrier (PROVISIONAL)	21
5.6 The Conservation regime — Positions 13–16 (the layered surveillance phase)	22
5.7 The Conservation processes are running in Bassez expanders	23
5.8 The Conservation regime in cancer (drug implications)	25
5.9 The remaining five positions in summary	25
6. The seven questions, resolved	27
6.1 Why is PD-L1 a weak biomarker?	27
6.2 Why do response rates vary so widely across cancer types?	27

6.3 Why do cold tumours stay cold?	27
6.4 Why does CAR-T succeed where checkpoint blockade fails?	28
6.5 Why do patients respond and then stop responding?	28
6.6 Why did five labs converge on CXCL13?	28
6.7 Why did the 2023 Mellman update extend outward rather than inward?	28
7. Clinical implications	28
7.1 Where each paradigm applies: a quantitative reframe	28
7.2 The existing drug arsenal is mostly already adequate	30
7.3 The bidirectional ipilimumab/belatacept consistency check	30
7.4 What the cycle says about trial design	30
8. Limitations and pre-registered tests	31
8.1 Limitations	31
8.2 Pre-registered tests	32
8.3 What the paper does not claim	33
9. Conclusion	33
Acknowledgments	34
Correspondence	34
References	35
Appendix A — Figure: The Sixteen Positions of the Cancer-Immunity Cycle	37

Abstract

The Chen–Mellman cancer-immunity cycle has organised tumour immunology for thirteen years. Its seven steps gave the field a shared language and a rational design space for checkpoint blockade. Alongside its successes, a set of clinical observations has accumulated that the cycle as drawn cannot resolve on its own terms: PD-L1 predicts anti-PD-1 response at AUC 0.63 across 18,792 patients; the same drug works at 40% in melanoma and near zero in MSS colorectal; cold tumours stay cold for reasons the cycle cannot name; CAR-T succeeds where checkpoint blockade fails through a mechanism the cycle cannot locate; five independent laboratories converged on CXCL13 CD8 T cells as the response-predictive population without a structural explanation; and the 2023 Mellman update extended the cycle outward into stratification axes rather than inward into new steps — a signal from the field itself that the seven-step sequence has reached the limit of what it can explain as drawn.

This paper proposes that the seven Chen–Mellman steps are a true description of six positions in a sixteen-position cycle, and that the missing ten positions are precisely where the unanswered questions live. The sixteen-position scaffold is derived independently from a universal sequence of dissipative processes (van der Klein 2026a, b) and is used here as load-bearing infrastructure rather than as the object of argument. The cycle traverses four regimes — Activation, Construction, Encounter, Conservation — separated by four irreversible barriers: the **Activation barrier** (commitment to build a response against this antigen), the **Egress barrier** (built effectors release from the lymph node into circulation), the **Drain barrier** (sustained killing resolves into stable clearance), and the **Suppression barrier** (the maintained state generates the next antigen-cycle).

The empirical core of the paper is four results across four independent datasets:

1. **Position 10 (the cytolytic synapse) is locked by a structurally derived biomarker, measured on two independent axes within one cohort and at two measurement resolutions across cohorts.** Within Bassez breast cancer (n=29 pretreatment): CXCL13 CD8 fraction predicts anti-PD-1 response at CV AUC 0.9746 ± 0.065 as a single feature, and the Bassez authors’ own CD8_EX label — selected without reference to the framework — predicts at CV AUC 0.9322 ± 0.037 . Two independent labelings, one position. Bulk-RNA validation: the LCAM-T module (Leader 2021) predicts anti-PD-1+chemo response

in GSE207422 NSCLC (n=24 pretreatment) at CV AUC 0.74 ± 0.06 . Five additional published cohorts converge on the same marker in the same direction.

2. **Position 8 (the licensing handshake immediately upstream of the Egress barrier) is locked across two cancer types and confirmed by a pre-registered specificity test in a third.** A Position 8 fingerprint specified before any data was inspected — build complete, egress program absent, reserve intact, exhaustion absent, CXCL13 absent — distinguishes ICI-failing MMRp tumours from ICI-responding MMRd tumours at Fisher OR = 4.64, $p = 0.0028$ in the Pelka 2021 colorectal dataset (76,965 cells, 110 tumour biosamples). The same fingerprint, applied without modification to the Bassez breast cohort, identifies 12 of 20 (60%) non-responders versus 0 of 9 (0%) responders, Fisher one-sided $p = 0.0024$. The framework predicts the fingerprint should be **rare** in cancer types where the immune cycle reaches the Encounter phase by default; pre-registered before data inspection at <25% match in melanoma, the test on Sade-Feldman 2018 melanoma (n=11 pretreatment anti-PD-1 monotherapy) yields 0 of 11 patients matching the fingerprint — the strongest version of the prediction. The fingerprint's clinical direction also depends on the cancer type's position on the cycle: in colorectal and breast (Construction-phase cancer types) the built-effector phenotype is enriched in non-responders, but in melanoma (Encounter-phase) the trend reverses, consistent with built-but-near-shipping cells being closer to where anti-PD-1 acts. One fingerprint, three cancer types, three consortiums, three scRNA platforms.
3. **The single largest patient bottleneck is in the Construction phase — the lymph-node and tumour-bed work that converts an antigen-aware system into deployment-ready effectors.** A blind k-means rebuild of 887 cancer patients across 12 tumour types from the Curated Cancer Cell Atlas places 43% of patients in this phase, with two distinct mechanisms: myeloid substrate hijack in the tumour bed preventing effector readiness (n=234) and Treg IL-2 sink in the lymph node interrupting clonal expansion during priming (n=147). Neither anti-PD-1 nor any synapse-acting drug can help these patients.
4. **In most cancer types, the response stalls upstream of where anti-PD-1 acts.** A regime classification of 11,373 TCGA samples across 33 cancer types finds that in only 4 of 33 cancer types does the cycle reach the Encounter phase where anti-PD-1 acts. In the other 29, the response stalls upstream, and PD-1 release does nothing because there is no engaged-but-braked cell to release.

A fifth strand — the bidirectional ipilimumab/belatacept consistency check at the Activation barrier — confirms the cycle's drug assignments are bidirectional by construction. Same molecular surface (CD28–B7–CTLA-4); same cycle position; opposite clinical goals (cancer vs kidney transplant); same two drugs working from opposite directions on the same axis.

The implication is structural. **Cancer is not primarily a failure of Encounter.** Across 33 cancer types, only 4 stall at the synapse phase where anti-PD-1 acts. The remaining 29 stall earlier in the cycle (8 in the Activation phase, 8 in the Construction phase) or in Treg over-control (13). The decade of immunotherapy effort directed at the killing step has worked precisely in the cancers where the cycle reaches the synapse with the brake on, and has failed precisely where the response stalls upstream. The next decade should be directed not only at the killing step but at the build step — and in particular at the previously unnamed handshake at Position 8 immediately upstream of the Egress barrier, where built effector cells fail to receive the licensing signal that would let them ship into Encounter at all. The drugs that act on the upstream regimes mostly already exist. **What has been missing is the map from patient to drug across the full cycle.**

1. Introduction: Seven questions the cycle leaves open

In 2013, Chen and Mellman published a diagram in *Immunity* that has organised tumour immunology ever since. The cancer-immunity cycle described seven steps: (1) release of cancer cell antigens, (2) cancer antigen presentation, (3) priming and activation, (4) trafficking of T cells to tumours, (5) infiltration of T cells into tumours, (6) recognition of cancer cells by T cells, and (7) killing of cancer cells. The cycle closes — killed tumour cells release more antigen, feeding back into step 1.

The diagram's contribution was the insight that the anti-tumour immune response is a cycle, that each step can be rate-limiting, and that the rate-limiting step determines which intervention will work. This framing rationalised the design of checkpoint inhibitors: anti-CTLA-4 releases the brake at priming (step 3), anti-PD-1 and anti-PD-L1 release the brake at recognition (step 6). The framework holds where it applies. It is cited over six thousand times and continues to be the scaffolding on which the field explains its successes.

It is also the scaffolding on which the field registers its failures. Seven questions sit on top of the cycle that the cycle cannot answer on its own terms.

Why is PD-L1 a weak biomarker? PD-L1 is the ligand anti-PD-1 blocks at step 6. If the cycle is complete, its expression should strongly predict response. A meta-analysis across 18,792 patients found PD-L1 staining predicts at AUC 0.63 — marginally above chance.

Why do response rates vary so widely across cancer types with the same drug? Anti-PD-1 achieves objective response rates of approximately 40% in melanoma, 70% in Hodgkin lymphoma, 50% in MSI-high colorectal, less than 5% in MSS colorectal, 2% in pancreatic, and 8% in glioblastoma. Same drug, same target. The cycle does not predict this distribution.

Why do cold tumours stay cold? The cycle has a mechanism for why T cells are disabled in a hot tumour. It has no mechanism for why T cells fail to arrive in the first place. Steps 1 through 5 are drawn as a single flowing sequence with no internal failure modes.

Why does CAR-T succeed where checkpoint blockade fails? Both are T-cell therapies. Checkpoint inhibitors unblock existing T cells at step 6. CAR-T supplies T cells from outside. The cycle cannot explain why replacement succeeds where release fails, because the cycle does not describe where T cells come from before they appear at step 4. The entire build pathway is compressed into the arrow between step 3 and step 4.

Why do patients respond and then stop responding? Acquired resistance is a loop that has broken at a location the loop does not name. Clinical descriptions — exhaustion, antigen loss, clonal escape — sit alongside the seven steps rather than at any of them.

Why did five independent laboratories converge on CXCL13 as the response marker? Bassez (*Nature Medicine* 2021), Leader (*Cancer Cell* 2021), Zhang L (*Nature Cancer* 2022), Zhang Y (*Cancer Cell* 2021), and Liu (*Cell* 2025) all identified CXCL13 CD8 T cells as the response-predictive population, across breast, NSCLC, BCC, HNSCC, and RCC. The cycle has a step called recognition at step 6 but does not structurally produce CXCL13 as that step's signature. Why this particular chemokine?

Why did the 2023 Mellman update extend outward rather than inward? Ten years after the original paper, Mellman, Chen, Powles, and Turley revisited the cycle in *Immunity* (56:2188–2205). The update did not add steps. It added patient-stratification axes — indication, genotype, and immunotype — that sit *beside* the cycle rather than inside it. This is a signal from the field itself that the seven-step sequence has reached the limit of what it can explain as drawn.

These seven questions are not independent criticisms. Each is a question about *where in the response the patient's response stalls*, and in each case the place the patient's response stalls is either a step the cycle does not

contain (cold tumours, CAR-T, acquired resistance, the 2023 stratification axes) or a step the cycle contains but cannot distinguish from its neighbours (PD-L1's weakness, response-rate variance, the CXCL13 convergence).

The argument of this paper is that all seven have the same answer: **the Chen–Mellman cycle describes six of sixteen positions in the complete antigen-specific immune response, and the missing ten are where the unanswered questions live.** The completion is additive. The seven steps Mellman drew remain the seven steps he drew, unmodified, occupying their original meaning in the completed map. What the completion adds is the naming and empirical validation of the positions the original diagram compressed into arrows.

The paper proceeds as follows. Section 2 introduces the base architecture: the four regimes, the four barriers, and the antigen-specific response as the unit that traverses the cycle. Section 3 presents the four primary empirical results: the Position 10 lock at two resolutions, the Position 8 lock in MMRp colorectal cancer, the 43%-in-Construction-phase bottleneck, and the Atlas 4-of-33 distribution. Section 4 introduces the sixteen-position scaffold and maps Chen–Mellman onto it. Section 5 walks through the positions that carry empirical weight (4, 5, 8, 10) and the Construction aggregate, with Position 12 in provisional support, and the remaining positions in summary table form. Section 6 returns to the seven questions and shows how each is resolved. Section 7 presents the clinical implications and the bidirectional ipilimumab/belatacept consistency check from kidney transplant. Section 8 discusses limitations and pre-registered tests. Section 9 concludes.

2. The base architecture: four regimes, four barriers, one agent

The cancer-immunity cycle Chen and Mellman described is the *antigen-specific immune response cycle*: the cycle the immune system runs once for each antigen it encounters. The immune system itself is not a single cycle. It is a system that runs many cycles in parallel and in succession, one per antigen — influenza-antigen has its own cycle, tetanus-antigen has its own cycle, tumour-antigen has its own cycle. Each individual cycle traverses four regimes in fixed order, separated by four irreversible barriers.

The unit that traverses the cycle is the **antigen-specific response** — neither the antigen alone nor the T cell alone, but the response-against-this-antigen as a single transforming unit. The antigen is the foreign signal that enters the system; the T cell is the substrate through which the response takes form. What is built, deployed, and held as memory is the antigen-specific response: the TCR clone selected for this antigen, expanded and armed against it, deployed to find and kill cells presenting it, and held as memory that recognises it on re-encounter.

2.1 The four regimes

Regime	What it does	Compartment
Activation (positions 1–4)	Foreign antigen enters, accumulates, is configured against the naive T-cell repertoire, and a matched clone commits to expansion	Tumour site, draining lymph node
Construction (positions 5–8)	The committed clone expands, differentiates, acquires effector machinery, is selected for egress	Draining lymph node

Regime	What it does	Compartment
Encounter (positions 9–12)	The deployed effector traffics to the tumour, engages antigen at the synapse, sustains killing, resolves into clearance, equilibrium, or escape	Tumour parenchyma
Conservation (positions 13–16)	The state stabilises into memory, surveillance maintains the held form, regulation prevents over-activation, antigenic drift generates the next cycle	Bone marrow, peripheral lymph nodes, circulation

2.2 The four barriers

Each regime ends at a singularity barrier — an irreversible crossing into the next regime. The four barriers are:

Barrier	Crossing	What is irreversibly committed	Failure mode the cycle cannot name without it
Activation barrier	Activation → Construction	A TCR-matched clone commits to clonal expansion (CD28–B7 engagement, CD25 upregulation, IL-2 autocrine loop locked in)	No commitment; antigen aware but uncommitted; the pre-commitment cold-tumour state
Egress barrier	Construction → Encounter	Built effectors release from the lymph node into circulation (CCR7 retention down, S1P egress receptors up, licensed by the cDC1–CD40–IL-12 handshake)	Built but parked; effectors retained in the node despite construction being complete; the MMRp colorectal pattern
Drain barrier	Encounter → Conservation	Sustained killing resolves into stable clearance or equilibrium	Inflammatory drain wins; killing-vs-proliferation balance tips toward escape; acquired resistance
Suppression barrier	Conservation → next cycle	The maintained state generates the next antigen perturbation (typically an escape variant detected by surveillance)	Treg over-correction damps the response below the threshold for cycle-restart; the pancreatic and glioblastoma pattern

Three of these barrier names — Activation, Drain, Suppression — are standard in immunology and structurally accurate. The fourth name has been corrected: earlier framework drafts called the Construction →

Encounter crossing the *Detection barrier*, but detection is sub-step 1 of any regime (the first internal step of detect/build/sense/resolve), not the regime-crossing event itself. The correct immunology term for what crosses here is **egress** — effector egress from the lymph node into circulation. This paper uses *Egress barrier* throughout.

2.3 The cycle is generational

Each completed cycle conserves antigen-specificity in memory and provides the held ground from which the next cycle of a new antigen can begin. Memory of one antigen is the substrate on which surveillance for the next antigen runs. When tumour antigenic drift produces an escape variant, that variant becomes the foreign signal entering Activation in a new cycle, and the held memory infrastructure of the previous cycle supports the surveillance that detected it. The cancer-immunity cycle is not a single closed loop; it is a generational sequence of antigen-specific cycles, each one building on the conserved memory of those that came before.

2.4 What the rest of the paper does with this architecture

Sections 3–5 present the empirical evidence in cancer for the four regimes, the four barriers, and the sixteen sub-positions inside them. Section 3 presents the four primary results that define the empirical core. Section 4 introduces the sixteen-position scaffold and shows how Mellman’s seven steps map onto it. Section 5 walks the positions where empirical weight has been concentrated — Position 4 (the Activation barrier substrate), the Construction aggregate (the 43% bottleneck), Position 8 (the Egress barrier substrate, where Pelka colorectal and Bassez breast both lock), Position 10 (the cytolytic synapse, locked across three independent axes within Bassez plus bulk LCAM-T validation), and Position 12 (the Drain barrier substrate, with provisional support). The remaining positions are presented in summary table.

3. Four empirical results in cancer

3.1 The cytolytic synapse is locked by a structurally derived biomarker, measured at two resolutions

The sixteen-position scaffold predicts that each position has a structurally derivable arrival marker — the chemokine, surface protein, or transcriptional event the position’s molecular surface produces. For the position corresponding to TCR engagement at the cytolytic synapse (Position 10, inside the Encounter regime), the predicted arrival marker is **CXCL13** — the chemokine produced by CD8 T cells when the TCR fires at the synapse. The marker is not fitted to response data; it is the downstream product of the engagement that *defines* the position.

The framework also predicts something about the *measurement* of this marker that matters in practice. CXCL13 is produced not only by engaged CD8 T cells but also by Tfh cells, follicular dendritic cells in tertiary lymphoid structures, and other stromal sources. At single-cell resolution this is not a problem: per-cell expression lets you ask directly *what fraction of CD8 cells are CXCL13*, and the other sources are excluded by cell-type gating. At bulk resolution it is a problem: total tumour CXCL13 averages all sources together, and the position-10 signal is diluted. The framework’s structural reading therefore predicts that the single-cell implementation of the marker is CXCL13 CD8 fraction, while the correct bulk implementation is the multi-gene module that recovers CXCL13 in co-expression with CD8 lineage and engagement markers — the LCAM module published by Leader et al. 2021. Same underlying position, two instruments, two structurally appropriate markers. We test both.

Single-cell test: Bassez breast cancer, n=29 The direct single-cell test used the Bassez et al. 2021 BIOKEY breast cancer cohort (n=29 treatment-naïve anti-PD-1 patients, ~8,400 CD8 T cells per biopsy, pretreatment scRNA-seq). Per-biopsy CXCL13 CD8 fraction:

- Expanders (n=9): median 14.3%
- Non-expanders (n=20): median 0.0%
- Mann–Whitney $p < 0.0001$
- 14 of 20 non-responders have undetectable CXCL13 CD8 cells in this cohort. The bottoming-out at zero is a Bassez-specific finding; in the other published cohorts (Leader, Zhang L, Zhang Y, Liu) the direction is consistent (responders enriched for CXCL13 CD8) but the absolute floor at zero is not as clean.

Used alone as a single-feature logistic regression with the published FULCRUM-S+ evaluation protocol (5-fold \times 200 stratified cross-validation, locked seed):

Model	Features	CV AUC
FULCRUM-S+ baseline (published)	6	0.9217 \pm 0.137
FULCRUM-S+ + CXCL13_CD8	7	0.9340 \pm 0.129
CXCL13_CD8 alone	1	0.9746 \pm 0.065

One structurally derived marker, zero fitted parameters, beats a published 6-feature ML baseline on the same data by five points of AUC. The Bassez cohort produces an unusually clean categorical separation: a substantial fraction of non-responders fall below the detection threshold for CXCL13 CD8 cells. This is a strong signal in this dataset, but the absolute floor at zero is cohort-specific — other CXCL13 cohorts replicate the direction (responders enriched for CXCL13 CD8) without the same dramatic bottoming-out.

Independent-axis validation on the same cohort: Bassez authors’ own cellSubType labeling The Bassez authors published their own cell-type classification of the cohort, including a CD8_EX (exhausted CD8) cluster identified through their independent clustering and marker-selection pipeline. The CD8_EX label is the Bassez authors’ own structural call about which CD8 cells have engaged the tumour and entered the exhaustion programme — selected without reference to position-10 framework predictions and without specific reference to CXCL13. Tested as a single feature on the same n=29 cohort using the same evaluation protocol:

Axis	Source	Single-feature CV AUC
CXCL13 CD8 fraction	Our gating	0.9746 \pm 0.065
CD8_EX fraction (of all T cells)	Bassez authors’ cellSubType	0.9322 \pm 0.037
CD8_EX fraction (of CD8 only)	Bassez authors’ cellSubType	0.9313 \pm 0.037

CD8_EX-fraction-of-CD8 separates expanders from non-expanders at Mann–Whitney $p = 0.000139$ (one-sided), with E median 18.6% versus NE median 1.8%. The categorical separation is also clean: 0 of 9 E patients have <1% engaged-exhausted CD8; 9 of 20 NE patients fall below that threshold.

This is convergent validation through independent measurement choice. The CXCL13 axis was developed by us as the structurally derived synapse-arrival marker. The CD8_EX axis was developed by the Bassez authors for their own purposes, using their own clustering and marker selection pipeline. Two independent labelings, two single-feature predictors, both >AUC 0.93 on the same patients. **The position is the same; the instruments to read it are different; both instruments succeed.**

A third axis from the same data adds Position 11 support. CD8_EX_Proliferating (actively dividing engaged effectors, structurally Position 11 — sustained killing) as a single feature: CV AUC 0.81 ± 0.05 , MW $p = 0.0021$. Engaged-and-active and post-engagement-spent populations both read in the predicted direction.

Bulk RNA test: GSE207422 NSCLC, n=24 The bulk-RNA test used GSE207422 (n=24 stage IB–IIIB NSCLC patients, treatment-naïve pretreatment biopsies, neoadjuvant anti-PD-1 + carboplatin-based chemotherapy, 9 MPR/pCR responders, 15 NMPR non-responders). One sample per patient. No mixing of pre and post timepoints.

Marker	CV AUC	Direction
CXCL13 alone (z-score)	0.53	flat (no signal)
CXCL13 / CD8A (TPM ratio)	—	inverted
(CXCL13 + PDCD1)/2 (z-score average)	0.71 ± 0.06	correct
LCAM-T module (12 genes)	0.74 ± 0.06	correct, $p=0.065$
LCAM-full module (24 genes)	0.73 ± 0.06	correct, $p=0.065$

The pattern is exactly what the structural reading predicts. Single-gene CXCL13 in bulk RNA is contaminated by non-CD8 sources and shows no signal. When CXCL13 is paired with PDCD1 to recover the engaged-CD8 co-expression structure that single-cell measurement would give for free, the AUC rises from 0.53 to 0.71. The full LCAM-T module reaches 0.74 — and this is the *first* application of LCAM to the GSE207422 data, not a refit. LCAM was published on independent data (the POPLAR randomised trial); applying it to a new pretreatment cohort tests the same structural prediction on data the module was never trained on.

The two cohorts taken together Bassez at single-cell resolution recovers the position-10 marker as a single-feature CXCL13 CD8 fraction at AUC 0.97. GSE207422 at bulk resolution recovers the same position with the structurally appropriate bulk marker at AUC 0.74. Two cohorts, two cancer types (breast, NSCLC), two platforms (scRNA, bulk RNA), 53 patients total, one structural prediction. The bulk AUC is lower than the single-cell AUC because bulk RNA averages over the cells the framework is asking about — that is a known property of the instrument, not a failure of the marker, and the modular co-expression approach recovers the signal that the single-gene approach loses.

Literature consensus Beyond our two direct tests, five independent published cohorts across four cancer types and multiple platforms have converged on CXCL13 T cells as the response-predictive population: Bassez 2021 (breast scRNA), Leader 2021 (NSCLC + POPLAR randomised trial bulk), Zhang L 2022 (pan-cancer meta-analysis covering NSCLC, breast, BCC, HNSCC, RCC), Zhang Y 2021 (TNBC scRNA), and Liu 2025 (NSCLC n=234). Across the published literature and our work, the marker has been read on at least 350 patients spanning at least six independent cohorts in five cancer types with three measurement modalities. The direction is consistent in every cohort that uses the platform-appropriate implementation.

The contribution of this paper is not the marker — the marker was already in the literature. The contribution is the structural explanation that unifies what the field had already found in pieces. CXCL13 is the synapse arrival output because the synapse is the position that structurally produces the chemokines fired by TCR engagement. The single-cell vs bulk measurement distinction is not a methodological caveat added after the fact; it is a *structural prediction* of the framework — the position is defined at the per-cell molecular surface, single-cell measurement reads it directly, and bulk measurement requires a co-expression module to recover the per-cell signal from population-averaged data.

With the marker locked across two direct-test cohorts and a multi-cohort literature consensus, anti-PD-1 has a precise structural assignment: it acts on cells stalled at the cytolytic synapse with PD-1 disabling TCR signal transduction. **It cannot work on cells that have not reached the synapse, because there is nothing for it to release.**

3.2 Position 8 is real and measurable: built effectors that never crossed the Egress barrier

The position immediately upstream of the Egress barrier — Position 8 (Selection) in the sixteen-position scaffold — is the licensing handshake where built effector cells are selected for egress from the lymph node into circulation. The Mellman cycle has no name for it; trafficking (Mellman step 4) is the *consequence* of egress firing, not the event itself. The framework predicts that this position, like Position 10, has a structurally derivable fingerprint: cells with the build complete, egress program absent, reserve intact, exhaustion absent, and CXCL13 absent — built effector cells in waiting, not yet shipped, not yet engaged.

We tested this directly in the Pelka 2021 colorectal cancer dataset (76,965 cells, 62 donors, 180 biosamples, 23,486 CD8 T cells across 56 MMR-proficient and 54 MMR-deficient tumour biosamples). MMR status is a clean stratifier of anti-PD-1 response in colorectal cancer: MMRd tumours respond at ~50%, MMRp tumours respond at <5%. The framework predicts that this clinical contrast should map to a structural difference — and that the difference should sit at Position 8, not Position 10.

The Position 8 fingerprint was specified before any data was inspected: build complete (effector machinery present), egress program absent (no homing receptor expression), tissue-residence neutral, reserve intact (IL7R retained), stem markers retained (TCF7), exhaustion absent (PDCD1, TOX low), proliferation low, CXCL13 absent.

Position	MMRp (n=56)	MMRd (n=54)
7 Effector build incomplete	27%	6%
8 Built but parked (Egress barrier failure)	32%	9%
10 Synapse engaged (CXCL13)	25%	50%
11 Active killing	7%	17%
12 Sustained killing / clearance transition	9%	19%

- Position 8 in MMRp vs MMRd: **Fisher OR = 4.64, one-sided p = 0.0028**
- Continuous Position 8 score, MMRp median +0.90 vs MMRd median -2.85: **Mann-Whitney p < 10**
- Combined Construction-region failure (Positions 7+8): MMRp 58.9% vs MMRd 14.8%, **Fisher OR = 8.25, p < 10**
- All three Encounter positions (10, 11, 12) enriched in the opposite direction in MMRd (each p < 10)

Cross-cancer extension: the Position 8 fingerprint replicates in Bassez breast cancer The Position 8 fingerprint was specified in Pelka colorectal work before any breast-cancer test was attempted. Operationalised on the Bassez n=29 breast cohort using the matching cell-subtype variables — built effector compartment (CD8_EM + CD8_RM) > 65% of CD8, engaged compartment (CD8_EX + CD8_EX_Proliferating) < 5% of CD8 — the fingerprint maps cleanly onto the non-responder population:

Bassez breast pretreatment	n	P8 fingerprint match
E (responders)	9	0 (0%)
NE (non-responders)	20	12 (60%)

Fisher exact one-sided $p = 0.0024$ (OR is undefined since no E patient meets the fingerprint). The continuous Position 8 score (built minus 2× engaged) shows MW one-sided $p = 0.0014$ with NE median 0.78 versus E median 0.13.

The 8 NE patients who do not match Position 8 have some engagement (5–11% engaged-exhausted CD8) but did not expand productively — consistent with being stalled downstream at Position 10 or 12. **Zero NE patients have neither Position 8 nor any synapse engagement**, meaning the cohort splits cleanly into two structurally distinct failure modes: 60% stalled at the Egress barrier (Position 8), 40% stalled at the synapse or Drain barrier (Position 10/12).

This extends Position 8 from one cancer type (colorectal) to two (colorectal + breast), measured by two independent consortiums (Pelka et al., Bassez et al.) on two different scRNA platforms, with one fingerprint specified before either test. The MMR-status stratification that drove the Pelka result is unavailable in Bassez (MMR is rarely measured in breast cancer because microsatellite instability is uncommon there), but the response-stratification result is the same direction at comparable effect size.

Specificity test in melanoma: the fingerprint is rare where the framework predicts it should be rare The Pelka and Bassez results both come from Construction-phase cancer types: cancers where the immune cycle, on the Atlas distribution across 33 tumour types (§3.4), is stalls upstream of the Encounter phase. In those cancer types, the framework predicts Position 8 should be common. Melanoma is the textbook *Encounter*-phase cancer — the immune cycle reaches the synapse, anti-PD-1 acts, and roughly 40% of patients respond. The framework’s specificity prediction is therefore the harder claim: **in melanoma, the same Position 8 fingerprint should be rare**.

The test was pre-registered at <25% match before data inspection, with falsification at >40% match. Sade-Feldman et al. 2018 (GSE120575) provides 11 pretreatment anti-PD-1 monotherapy patients with curated CD8 sub-cluster labels (3CA curation, Tirosh lab Weizmann), 4 responders and 7 non-responders. The fingerprint as specified in Pelka and Bassez was applied with one declared vocabulary translation: where Bassez labels CD8_EM and CD8_RM, the 3CA Sade-Feldman curation labels CD8_4-Memory/effector and CD8_6-Memory/effector; resident-memory is not separately labeled in this curation, so the “build complete” component is read on Memory/effector alone. Thresholds were not adjusted.

Result: 0 of 11 melanoma patients match the Position 8 fingerprint (95% binomial CI 0%–28%). The pre-registered specificity prediction is confirmed; the framework is not falsified, and the test is on the right side of every threshold by a wide margin.

A second observation emerged from the same data and was not pre-registered. The continuous score (built minus engaged compartment) reverses direction in melanoma compared with Pelka and Bassez: melanoma responders have a median score of +0.53 (built-leaning), non-responders –0.03 (mixed); the Mann–Whitney one-sided test in the Pelka/Bassez direction (P8-leaning enriched in non-responders) returns $p = 0.88$, a non-significant result that points the *opposite* way. $n=11$ makes any continuous claim provisional, but the direction is consistent with a structural reading: in cancer types where the cycle reaches the synapse by default, a built-effector phenotype is closer to where the drug acts; in cancer types where the cycle does not, a built-effector phenotype indicates the response is stalled upstream of where the drug acts. Same single-cell phenotype, opposite clinical direction, mediated by where the cancer type as a whole sits on the cycle. This is consistent with the Atlas distribution (§3.4): only 4 of 33 cancer types are classified as Encounter-phase, and melanoma is one of them.

A diagnostic non-result from Pelka Healthy normal colon biosamples ($n=36$) in the Pelka dataset were initially excluded from the headline test because every single one scored at Position 8. 36 of 36. The first

reading was that the fingerprint was broken — clearly a healthy gut sample is not a “tumour holding built effectors hostage.” But the fingerprint is not broken; it is reading correctly. Healthy tissue-resident gut CD8 memory cells are the textbook built, parked, reserve-intact, not-firing population. They sit in the lamina propria for years. Build complete. Egress program off. IL7R maintained. Killing at baseline. Exhaustion markers low. They are *waiting correctly*.

This means MMRp colorectal tumours are pathologically arrested at the same single-cell phenotype where healthy gut is correctly arrested. The same cell-fingerprint, opposite clinical valence. The structural reading is that Position 8 is a licensing event between the built effector and the cDC1 partner with CD40 and IL-12, and a single-cell fingerprint of just the effector cannot distinguish “no antigen present, correctly waiting” from “antigen present, pathologically waiting.” The MMRp-vs-MMRd separation is the disease finding. The Normal-vs-MMRp non-separation is a finding about what the instrument can and cannot resolve: distinguishing healthy from disease at the Egress barrier requires reading both the effector and the licensing partner together, not either alone.

3.3 The single largest patient bottleneck is in the lymph-node-and-tumour-bed Construction phase

A single-cell rebuild of 887 cancer patients across 12 tumour types from the Curated Cancer Cell Atlas (3CA) clustered patients by their immune cell composition using blind k-means. Five dominant archetypes emerged. **The two largest archetypes both correspond to the Construction phase — the lymph-node and tumour-bed work that converts an antigen-aware system into deployment-ready effectors — and together they account for 381 of 887 patients, or 43%.** The two archetypes correspond to different mechanisms within Construction (a tumour-bed myeloid hijack and a lymph-node Treg sink), with distinct drug classes (developed in §5.2).

Archetype	Mechanism	Construction-phase position	n patients
C4 (Drain-rich)	Myeloid substrate hijack in tumour bed: TAM, MDSC, recruited monocytes occupy the niche and produce TGF- β , IL-10, PGE ; the local DC machinery cannot mature and migrate at adequate rate; antigen never reaches the lymph node in adequate quantity	Position 5 (the Construction-phase signal: the system fails to detect that effectors should be built because the upstream signal never arrives in adequate strength)	234
C1 (DC-rich)	Treg IL-2 sink in lymph node: Tregs in the dLN with high-affinity CD25 receptors consume IL-2 before the priming clone can take it up; expansion is initiated but cannot complete	Position 6 (the Construction-phase build: clonal expansion is started but the cytokine substrate is hijacked)	147

These patients are stalled before the synapse phase. C4 patients fail at Position 5: the Construction-phase signal never arrives in adequate quantity from the tumour. C1 patients fail at Position 6: priming initiates but the clone cannot expand because the cytokine substrate is hijacked. The tumour looks cold from outside in both cases, but the failure is at different molecular surfaces. **Neither anti-PD-1 nor any other synapse-acting drug can help these patients**, because they do not have effector cells at the synapse for those drugs to act on. They need different interventions: anti-VEGF or anti-CSF1R to remove the myeloid hijack and DC-promoting agents like Flt3L to drive maturation (addresses C4 / Position 5); engineered IL-2 muteins or IL-15 superagonists to bypass Treg cytokine consumption (addresses C1 / Position 6); or cell therapy to supply an externally built effector (CAR-T, TCR-T, BiTEs — bypass Construction entirely).

Stage enrichment of a third archetype maps to the Activation phase (the cycle has not yet committed to building a response). The Bassez 2021 breast cohort is the only 3CA dataset with full AJCC stage annotation for every sample. When the blind k-means archetype assignments are crosstabulated against stage, the Reserve-rich archetype (C0) concentrates in the earliest cancer stages: 50% of stage IA+IIA tumours are C0 versus 6.7% of stage IIB+IIIA tumours (Fisher one-tailed $p = 0.0099$). Stage IA breast tumours are roughly seventeen times more likely to be Reserve-rich than late-stage tumours. This is consistent with the structural prediction that early-stage tumours are over-represented in the population where the Activation barrier has not yet been crossed — there are T cells present but no clone has yet committed to expansion against the antigen.

The remaining roughly 30% of patients distribute across the synapse-engagement and effector positions plus a smaller Treg-rich archetype representing differentiation bias toward FoxP3 regulatory fate. The memory and surveillance positions are empty in the cancer data — not because the framework is wrong about them, but because cancer does not pile patients up there. **Cancer is a disease of failure to build effectors and failure to engage productively, not a disease of failed memory maintenance.**

3.4 In most cancer types, the response stalls upstream of where anti-PD-1 acts

At the population level, the regime distribution is validated by a separate classification of 11,373 TCGA samples across 33 cancer types (van der Klein 2026c, the Immune Barrier Atlas), using CIBERSORTx cell fractions and a five-component scoring rule corresponding to the four phases the immune response passes through.

Where the response stalls	n cancer types	Drug class that acts here
Activation phase (Positions 1–4: antigen detection through clone commitment)	8	TLR/STING agonists, oncolytic viruses, in-situ vaccination, anti-CTLA-4 (ipilimumab)
Construction phase (Positions 5–8: lymph-node priming and effector build)	8	cancer vaccines, IL-2 muteins, anti-VEGF, anti-CSF1R, CD40 agonists, cell therapy
Encounter phase (Positions 9–11: trafficking, engagement, killing)	4	anti-PD-1, anti-PD-L1
Conservation phase (Positions 13–16: memory, surveillance, regulation) — Treg over-control mode	13	anti-CCR8, anti-TGF- β , Treg depletion

In only 4 of 33 cancer types does the immune cycle reach the Encounter phase where anti-PD-1 acts. In the other 29, the response stalls earlier, and releasing PD-1 does nothing because there is no engaged-but-

braked cell to release. This is the population-level version of the patient-level Bassez result: in most cancer types, the immune response is not at the position the field's primary drug targets.

The four scores were reverse-engineered from the cell-type fractions by linear regression with perfect $R^2 = 1.0$, meaning the assignments are reproducible by construction. Drug-phase consistency was confirmed in 7 of 8 testable cancer types. NSCLC single-cell RNA cross-validation came in at 42% Treg-over-control versus the TCGA bulk estimate of 41.8% — direct convergence between two completely different measurement protocols.

A stage hypothesis was tested and falsified. The hypothesis was that cancer maturity (AJCC stage) determines which conversion is blocked. Tested against 7,498 TCGA patients with valid stage data: pan-cancer Spearman $\rho = -0.037$ (effectively zero). Cycle position is a property of the cancer type, not of disease progression within a patient.

3.5 Why these four results converge

The four results above are independent. The Bassez $n=29$ single-cell test uses one breast cancer cohort. The Pelka $n=110$ colorectal test uses a different scRNA dataset from a different consortium. The 3CA rebuild uses 887 patients across 12 tumour types from a third consortium. The TCGA Atlas classification uses 11,373 samples across 33 cancer types from a fourth consortium with bulk rather than single-cell measurement. They are not the same data and they are not the same instrument.

What they converge on is a common structural reading. The Bassez result shows that *individual patients* are sorted by whether their immune cycle has reached the synapse. The Pelka result shows that *MMRp colorectal cancer specifically* stalls at the Egress barrier where built effectors fail to ship. The 3CA result shows that *the largest single group of patients* across cancer types has a response that never reaches the synapse because it stalls in the Construction phase. The Atlas result shows that *in most cancer types*, the cycle stalls upstream of the synapse entirely.

All four are saying the same thing at different scales: anti-PD-1 acts at one position in the cycle, and in most cancer patients the immune cycle is not at that position. **The framework's contribution is to name where the cycle is instead, and to point at the drugs that act there.**

4. The sixteen-position completion of the cycle

The sixteen-position scaffold is derived elsewhere from a universal sequence of dissipative processes (van der Klein 2026a, b). Briefly: any system that consumes free energy, does work, and dissipates entropy must traverse four regimes — Potentiality (signal recognised and commitment made), Construction (the response built), Encounter (the response meets the substrate it was built for), and Conservation (the changed state stabilised) — and each regime contains four sub-positions corresponding to the same internal pattern: detect, build, sense, resolve. The sixteen positions are not chosen; they are the product 4×4 . This structure has been tested against 22 dissipative systems drawn from physics, chemistry, biology, and engineering (van der Klein 2026b). For the purposes of this paper, the scaffold is used as load-bearing infrastructure rather than as the object of argument; the cancer-immunology evidence that the scaffold organises is the subject below.

4.1 The sixteen positions of the cancer-immunity cycle

The table below names all sixteen positions of the cycle and what each one is in cancer-immunology terms. Each position carries a structural framework name (Signal, Threshold, Prototype, Selection, etc.) which is domain-general and applies to any dissipative cycle of the type, and a cancer-immunology event description that makes the position concrete for an immunologist. The four regimes are annotated with their primary physical

compartment in the cancer context: Potentiality begins in the tumour bed and converges in the draining lymph node; Construction is the lymph node work; Encounter is the tumour bed work; Conservation is the distributed maintenance state across circulation and memory niches.

The **Mellman** column tags the six and a half positions that the Chen & Mellman 2013 cancer-immunity cycle covered (M1 through M7). The remaining 9.5 positions — the entire Construction interior (5–8), the Encounter resolve (12), and the entire Conservation regime (13–16) — are added by the completion presented in this paper. The **Status** column marks which positions carry empirical weight and at what level: LOCKED = cross-cancer empirical evidence; SUPPORTED = single-cohort empirical or mechanistically inferred; HYPOTHESIS = structural derivation grounded in published literature; PROVISIONAL = small-cohort signal awaiting replication; generic = the position runs in any immune response with cancer-specific perturbations described where relevant. Per-position evidence and detailed treatment are in §5.

The cycle is generational: Position 16 detects a new antigen variant and triggers Position 1 of the next iteration. Positions 14 and 15 are coupled — they run as a Surveillance/Compensation oscillation rather than as a sequence.

4.2 What this completion does to Chen–Mellman

Chen and Mellman 2013 cover 6.5 of the 16 positions: M1 (Position 1), M2 (Positions 2–3), M3 (Position 4), M4–5 (Position 9), M6 (Position 10), M7 (Position 11). The seven steps of the original cycle correspond to a true description of these six positions; the framework’s contribution is not to replace them but to name the 9.5 positions Mellman compressed or omitted.

Two readings of Mellman’s step 3. Mellman titles step 3 “*priming and activation*” — a two-word phrase recognising that two events are involved. In the completed cycle, *priming* corresponds to Position 3 (the DC presenting peptide-MHC, naive T-cell scanning) and *activation* corresponds to Position 4 (the irreversible CD28–B7 commitment with IL-2 autocrine loop locked in). Position 4 is the Activation barrier — the point at which a TCR-matched clone commits to expansion regardless of whether the original signal continues. The framework’s contribution is not the terminology; it is the recognition that what immunology already calls *priming and activation* is the resolve step of the Potentiality regime, and that immediately downstream of it sits an entire Construction regime — four positions of clonal expansion, differentiation, prototype testing, and selection for egress — that the original cycle compressed into a single arrow.

The compressed arrow between steps 3 and 4. Mellman draws a single arrow from “priming and activation” (step 3) to “trafficking” (step 4). Inside that arrow sit four positions: clonal expansion (Position 5), effector lineage commitment (Position 6), prototype testing via cDC1 licensing (Position 7), and selection for egress (Position 8). Position 8 is the Egress barrier — the licensing handshake that determines which built effectors actually ship from the lymph node into circulation. This is where the strongest empirical lock in the paper sits: built effectors that fail to receive the cDC1–CD40–IL-12 signal stay parked in the dLN as TCF7+ stem-like cells, and this fingerprint is the response-predictive signal in MMRp colorectal cancer (Pelka 2021) and breast cancer (Bassez 2021), while being absent in melanoma (Sade-Feldman 2018) where the response stalls at a different position.

The compressed arrow after step 7. Mellman’s cycle closes with step 7 (killing) feeding back into step 1 (release of more antigen). Inside that closure sit five positions: net cell-economy resolution (Position 12), memory-fate lock-in (Position 13), patrolling memory (Position 14), Treg-mediated regulation (Position 15), and antigen drift detection (Position 16). The Conservation regime (Positions 13–16) is the layered surveillance phase that maintains the changed state and decides whether the next iteration of the cycle starts with a new antigen variant or whether the system collapses into chronic suppression.

Table 1. The sixteen positions of the cancer-immunity cycle.

#	Name	Cancer-immunology event	Location	Status	Mellman
POTENTIALITY <i>tumour bed → draining lymph node</i>					
1	Signal	DAMP–PRR conversion	tumour bed	<i>generic</i>	M1
2	Accumulation	DC recruitment and antigen capture	tumour bed → dLN	<i>generic</i>	M2
3	Configuration	DC maturation, MHC loading, naive TCR scan	dLN	<i>generic</i>	M2
4	Threshold	CD28–B7 commitment, IL-2 autocrine lock <i>Activation barrier</i>	dLN	SUPPORTED	M3
CONSTRUCTION <i>draining lymph node</i>					
5	Initiation	clonal expansion initiated	dLN	SUPPORTED	—
6	Architecture	effector lineage commitment (T-bet, GATA3, Bcl-6)	dLN	SUPPORTED	—
7	Prototype	cDC1 licensing test (helped vs helpless effector)	dLN	HYPOTHESIS	—
8	Selection	egress into circulation <i>Egress barrier</i>	dLN → blood	LOCKED	—
ENCOUNTER <i>tumour bed</i>					
9	Manifestation	extravasation into tumour parenchyma	tumour bed	<i>generic</i>	M4–5
10	Discovery	TCR–pMHC engagement; CXCL13 fired	tumour bed	LOCKED	M6
11	Operation	sustained killing; granzyme/perforin delivery	tumour bed	SUPPORTED	M7
12	Equilibrium	net cell-economy resolution <i>Drain barrier</i>	tumour bed	<i>PROVISIONAL</i>	—
CONSERVATION <i>circulation + memory niches</i>					
13	Stabilising	memory-fate lock-in; resolution of inflammation	tumour bed → niches	<i>generic</i>	—
14	Patrolling	circulating memory + tissue residence	circulation + niches	SUPPORTED	—
15	Correcting	Treg-mediated regulation; active response damped	systemic	SUPPORTED	—
16	Turning	antigen drift detection <i>Suppression barrier</i>	systemic	<i>generic</i>	—

Of the ten positions Mellman compressed or omitted, the empirical evidence in this paper concentrates on three: Position 8 (Pelka MMRp lock + Bassez breast lock + Sade-Feldman melanoma falsification), Position 10 (CXCL13 cross-cancer lock), and the Construction-phase aggregate (3CA 43% stalling in Positions 5–6). Position 7 is presented as a structural hypothesis grounded in five independent literature lines; Position 12 as a provisional small-cohort signal; the remaining positions as either mechanistically inferred or generic immune biology with cancer-specific perturbations described in §5.6 and §5.9.

5. Detailed treatment of the positions that carry empirical weight

5.1 Position 4 — the Activation barrier (SUPPORTED, single-cohort)

Structural role. Position 4 is the resolve step of the Activation regime: a TCR-matched naive T cell receives co-stimulation through CD28-B7, CD25 upregulates, the IL-2 autocrine loop establishes, and the clone commits to expansion regardless of whether the original signal continues. This is the irreversible crossing from antigen-aware-but-uncommitted into committed clonal expansion — the moment the system commits to building a response specifically against this antigen.

Molecular surface. The CD28-B7 engagement, where CTLA-4 binds the same B7 ligand with much higher affinity and competes with CD28.

Arrival marker. CD25-high Ki-67-positive tumour-antigen-matched T cells, appearing day 7–10 in vaccine cohorts.

Drug assignment. Ipilimumab (anti-CTLA-4) acts here as an agonist of commitment by releasing B7 for CD28 binding.

Stage-enrichment evidence (limited). In the Bassez breast cohort, the Reserve-rich archetype (C0) — the cell-fraction signature predicted to mark the pre-commitment state, where T cells are present but no clone has yet crossed the Activation barrier — concentrates in the earliest cancer stages: 50% of stage IA+IIA tumours are C0 versus 6.7% of stage IIB+IIIA tumours (Fisher one-tailed $p = 0.0099$). Stage IA breast tumours are roughly seventeen times more likely to be Reserve-rich than late-stage tumours. This is consistent with the structural prediction that early-stage tumours are over-represented in the population where Activation has not yet completed, but it is on the same cohort as the Position 10 single-cell test in §3.1. Independent validation in another cohort with both stage data and matched single-cell or cell-fraction measurement is needed.

Cross-clinical-context consistency. Belatacept (CTLA-4-Ig) acts at the same molecular surface in kidney transplant by occupying B7 — the structural mirror image of ipilimumab. Same surface, same cycle position, opposite clinical goal: cancer wants commitment to fire (release the brake), transplant wants commitment blocked (occupy the binding site). This is a clinical consistency check rather than a falsifiable prediction — both drugs work in their intended contexts and would be presented as evidence either way — but the structural prediction is specific: the intervention surface at any one position should be available to push in either direction depending on whether the goal is to enable or to suppress the response. Position 4 is the cleanest worked example because two well-studied drugs already act here from opposite sides.

5.2 The Construction-phase aggregate — Positions 5–8 (the 43% bottleneck)

The 3CA blind k-means rebuild of 887 patients across 12 cancer types finds 43% of patients (381 of 887) with a response that stalls somewhere in the Construction phase — the lymph-node and tumour-bed work that converts a committed naive T cell into a deployment-ready effector. Two distinct mechanisms account for these 381 patients, mapping to different sub-positions within Construction:

Position 5 — Initiation (the C4 mechanism, n=234). The Construction-phase detect step: the system reads how much room there is to expand from the committed substrate. The response stalls at this position when the upstream input never arrives in adequate quantity — typically when tumour-associated macrophages, MDSCs, and recruited monocytes occupy the tumour niche and produce TGF- β , IL-10, PGE , and other suppressive mediators that prevent local DCs from maturing. The DCs that do mature have their CCR7-driven migration impaired by the inflammatory environment. Antigen never reaches the dLN in adequate quantity, so the build phase has nothing to start from.

Drug class. Anti-VEGF (bevacizumab), anti-CSF1R (pexidartinib), anti-CCR2 to clear the myeloid hijack, plus DC-promoting agents (Flt3L, GM-CSF) to drive maturation despite the suppressive environment.

Position 6 — Architecture (the C1 mechanism, n=147). The Construction-phase build step: the system builds the differentiation pathway from the committed substrate. The response stalls at this position when DCs do arrive in the dLN and priming initiates, but the IL-2 autocrine loop that should sustain clonal expansion is interrupted because Tregs in the dLN, expressing high-affinity CD25, consume IL-2 before the priming clone can take it up. Priming is engaged but cannot complete; the cell numbers needed for productive deployment never accumulate.

Drug class. Engineered IL-2 muteins (e.g., NKTR-358, BNT151) bypassing the high-affinity Treg CD25 receptor; IL-15 superagonists; anti-CCR4 or anti-CCR8 to deplete Tregs in the dLN.

Position 7 — Prototype. The Construction-phase sense step: the system tests whether a functional effector built at Position 6 actually works. Sense steps in the framework require a partner against which the test is run — at Position 10, the partner is the tumour cell carrying antigen on MHC-I; at Position 3, the partner is the dendritic cell carrying captured antigen. At Position 7, the test partner is the licensed cDC1 in the tumour-draining lymph node — the cDC1 that has received CD40 signalling from an antigen-matched CD4+ helper T cell, plus type-I IFN signalling, and consequently can deliver the maturation signal that drives a prototype CD8 T cell forward into deployment-ready effector state. The encounter at Position 7 is in this sense an embedded mini-Encounter inside the Construction regime: the prototype is exposed to a partner that determines pass or fail before Position 8 selects which prototypes egress.

This sense-step structure predicts that both passed and failed prototypes should be physically detectable as different cell states in the same lymph node. The published literature has documented both states extensively, without recognising them as the two outcomes of a single position's test. **Passed prototypes** (helped) — TCF7-low, GZMB-positive, CD28-retained, PD-1 declining, CXCR3-positive — are the cells that received the cDC1 license and are preparing to leave the lymph node. **Failed prototypes** (helpless) — TCF7-high, GZMB-low or absent, CD28-retained but unutilised, PD-1 high and maintained, IL7R high, CXCR5-positive — are the cells that encountered an unlicensed cDC1, did not receive maturation signals, and accumulate in the tumour-draining lymph node as stem-like progenitor exhausted CD8 T cells (Tpex; Im 2016, Miller 2019, Beltra 2020). The framework reading is that the Tpex/stem-like phenotype that the field has been treating as “early exhaustion” is structurally the failed-test outcome of Position 7, mechanistically distinct from the engagement-failure exhaustion that occurs at Position 10. Both are reactivated by anti-PD-1 — the licensing-failed cells when help arrives later, the engagement-failed cells when the brake is released — which has obscured the distinction. The structural origins differ: Position 7 failure is licensing failure; Position 10 failure is engagement failure.

This is presented here as a structural hypothesis grounded in the published literature (Schenkel et al. 2021, Borst et al. 2018, Ahrends et al. 2017, Ferris et al. 2020, Lei et al. 2024) rather than as a locked biomarker. Direct empirical lock would require pretreatment matched-TdLN-and-tumour scRNA cohorts with anti-PD-1 response data and adequate CD8 sub-clustering, and these are not yet available in the cohorts on which Positions 8 and 10 were locked. The hypothesis is testable in principle and concrete enough to be tested if and

when matched-tissue cohorts become accessible.

A note on the Position 7 / Position 8 distinction: in tumour biopsy alone, the failed-prototype phenotype (TCF7-high, GZMB-low, PD-1 high) and the built-but-parked Position 8 phenotype (TCF7-retained, exhaustion absent, build complete) overlap in their TCF7-positive marker but differ in GZMB and PD-1 status. The cleanest separation requires sampling the lymph node where both states are physically produced, since tumour-resident TCF7+ stem-like CD8 cells in the published literature are typically a mixture of cells that have egressed Position 8 and migrated to the tumour bed and cells that have failed Position 7 and migrated as Tpex. Position 8 fingerprint thresholds developed on Pelka and Bassez do not currently distinguish these two subpopulations because the relevant marker contrasts (GZMB-positive vs GZMB-low among TCF7+ cells) sit below the resolution of the cell-subtype labels available in those cohorts. This is a known limitation of the current Position 8 fingerprint: it identifies the larger built-but-parked + failed-prototype combined population, not the failed-prototype subpopulation alone.

Drug class. CD40 agonists (sotigalimab, selicrelumab) act here by directly providing the licensing signal that CD4 help would have provided. The framework predicts they should specifically benefit patients whose Position 7 is failing because cDC1 is unlicensed — and the mixed but real clinical signals seen with these drugs are consistent with that prediction. Anti-CTLA-4 is also relevant here as it can rescue helpless priming through a parallel mechanism. This is the position cell therapy bypasses entirely — CAR-T, TCR-T, and BiTEs supply an externally built effector that has not been subjected to the patient's Position 7 test. The clinical signal that CAR-T succeeds where checkpoint blockade fails is consistent with the framework reading: when the patient's own Position 7 cannot pass any prototype through, an externally supplied prototype removes the dependency.

Position 8 — Selection (the Pelka MMRp lock). Detailed treatment in §5.3.

Why this matters for trial design. The aggregate 43%-in-Construction figure is the largest single failure mode in cancer immunotherapy by patient count. Treating it as a single problem — “cold tumours” — with a single intervention — anti-PD-1 — fails because Position 5, Position 6, and Position 8 each require different drugs, and none of them is anti-PD-1. Stratifying these patients by which Construction sub-position their response stalls at, using the C0/C1/C4 cell-fraction signature plus the Position 8 fingerprint, would let each sub-population be matched to the intervention that addresses its specific failure.

5.3 Position 8 — the Egress barrier (LOCKED, cross-cancer)

Structural role. Position 8 is the resolve step of the Construction regime: built effector cells receive a licensing signal from the lymph-node microenvironment that converts them from finished-but-parked into deployed-and-shipping. This is the substrate of the Egress barrier — the irreversible crossing from Construction into Encounter. Mechanistically, this is the cDC1–CD40–IL-12 handshake combined with affinity-based selection of which prototypes ship. The response stalls at this position when the build is complete (the cells exist, have effector machinery, have TCRs that recognise tumour antigen) but the licensing signal does not arrive.

Molecular surface. The cDC1–CD40–IL-12 licensing surface in the lymph node, plus the affinity-based selection signals (CD40L–CD40, IL-21 receptor engagement) that determine which cells egress. S1P egress receptor expression downstream of licensing.

Arrival marker. The Position 8 fingerprint as locked in Pelka 2021 colorectal (§3.2) and replicated in Bassez breast: build complete (CD8_EM + CD8_RM > 0.65 of CD8), engaged compartment absent (CD8_EX + CD8_EX_Proliferating < 0.05), reserve intact (IL7R retained), stem markers retained (TCF7), exhaustion absent (PDCD1, TOX low), proliferation low, CXCL13 absent. Cells in this state are in the tumour-draining lymph node, parked, waiting for a licensing signal that does not arrive.

Cross-cancer status. Locked in colorectal cancer at Fisher OR 4.64, $p = 0.0028$ (Pelka 2021, MMRp vs MMRd

stratification). Locked in breast cancer at Fisher OR undefined (perfect separation of E from P8-fingerprint), one-sided $p = 0.0024$ (Bassez 2021, NE vs E stratification). Confirmed rare in melanoma at 0 of 11 pretreatment anti-PD-1 monotherapy patients matching the fingerprint (Sade-Feldman 2018; pre-registered specificity prediction at <25%), with the continuous score reversing direction relative to Pelka/Bassez in a way consistent with melanoma being one of the 4 of 33 Encounter-phase cancer types in the Atlas distribution. One fingerprint specified before each test; three cancer types; three consortiums; three scRNA platforms.

Drug assignments. *Fire the licensing handshake directly:* CD40 agonists (sotigalimab, selicrelumab) — in clinical trials with mixed but real signals. The framework predicts they should work specifically in the MMRp population that is stalled at Position 8 — not in MMRd, where the licensing has already happened, and not in Position 7 cases where there is nothing yet to license. *Increase the substrate:* IL-2 variants targeted to CD25-low cells (the stem-like compartment), or ex vivo isolation of TdLN stem-like CD8 followed by adoptive transfer. *Strengthen the upstream maturation signal:* STING agonists, type-I interferon, IL-12.

Why the field has not seen Position 8. First, the Mellman cycle made the question invisible by starting at antigen release and ending at killing — putting Construction *outside* the cycle. Trafficking (Mellman step 4) was treated as logistics, not biology. Once trafficking is understood as a transport problem, the framework cannot ask what *signals* the transport. Second, the empirical population that demonstrates Position 8 — the tumour-draining lymph node stem-like CD8 — has been described in the literature for nearly a decade (Im et al. 2016, Miller et al. 2019, Beltra et al. 2020) but as a *phenotype*, not as a *position in a cycle*. Third, the conventional cancer-immunology funding landscape rewards Encounter-regime experiments, because checkpoint blockade was a triumph of Encounter biology and produced the largest commercial returns in oncology history. Construction-regime interventions — drugs that act on the build phase or on the licensing handshake — have no commercial template and no obvious target list.

5.4 Position 10 — the cytolytic synapse (LOCKED)

Structural role. Position 10 is the build step of the Encounter regime: TCR engages tumour antigen on MHC-I at the cytolytic synapse. The engagement fires CXCL13 and downstream chemokine production. This is the position where the engagement is constructed cell by cell — each successful synapse formation is a discrete event that adds to the engaged-effector population.

Molecular surface. The TCR-pMHC synapse, which is also where PD-1 acts when PD-L1 is present on the tumour cell. PD-1 inhibition releases the brake on TCR signalling at this surface.

Arrival marker. CXCL13 CD8 fraction at single-cell resolution; the LCAM module at bulk resolution. Triple-axis validation on Bassez breast $n=29$: our CXCL13 CD8 single-feature CV AUC 0.97; the Bassez authors' independent CD8_EX label single-feature CV AUC 0.93; CD8_EX_Proliferating (Position 11 sustained killing readout) CV AUC 0.81. Bulk-RNA validation on GSE207422 NSCLC $n=24$ with the LCAM-T module: CV AUC 0.74. Multi-cohort literature consensus across breast, NSCLC, BCC, HNSCC, and RCC (§3.1).

Drug assignment (LOCKED). Anti-PD-1 (pembrolizumab, nivolumab) and anti-PD-L1 (atezolizumab, durvalumab) act here.

Structural reading of the failure mode. Anti-PD-1 cannot work on cells that have not reached the synapse, because there is nothing for it to release. The Bassez non-responders with very low or undetectable CXCL13 CD8 cells are not stalled at Position 10 with the brake on; they stall upstream of Position 10 entirely, at Position 4 (the Activation barrier), in the Construction phase, or at Position 8 (the Egress barrier). **Releasing a brake requires a moving car.**

5.5 Position 12 — the Drain barrier (PROVISIONAL)

Structural role. Position 12 is the resolve step of the Encounter regime: the engagement at the synapse either tips into a state where the T cell ends with more than it spent, or it does not. The conventional reading frames Position 12 as the balance between killing rate and tumour proliferation rate — a tumour-side accounting. The framework reading is different: Position 12 is about the **net cell economy** of the engagement. The cell has invested at Position 10 (TCR engagement, transcriptional cascade fires) and Position 11 (granule release, repeated synapse formation, mitochondrial output spent). Each killing event costs granules, transcription, ATP, and survival input. For the cell to cross Position 12 into Conservation, it has to come out of the engagement with **more survival signal than it consumed** (IL-7, IL-15 received from environment), **more differentiation signal than it dissipated** (memory-fate transcription factors locked in), and **more information than it lost** (clonal recall capacity preserved).

The Drain barrier is the position where the cell's investment exceeds its return. Tumour cells dying is necessary but not sufficient — a cell can kill while losing more than it gains, and that cell cannot enter Conservation. It dies in place, or persists in a depleted state without resources to commit to memory. This inverts the conventional framing. The conventional question is: *is the tumour shrinking?* The framework question is: *is the cell coming out ahead?* These are correlated but not identical. A cell that kills five tumour cells and dies from exhaustion has positive tumour-side balance and negative cell-side balance — it does not cross Position 12. A cell that kills one tumour cell, receives strong IL-7/IL-15 input, locks in T_RM transcription, and survives has smaller tumour-side effect but positive cell-side balance — it does cross. Long-term tumour control comes from the second cell, not the first.

Provisional arrival marker. TCF7 retention in CXCL13 engaged CD8 cells. TCF7 is the differentiation-lock transcription factor for memory fate; its retention through engagement indicates the cell came out of the synaptic phase with the memory-fate substrate intact. CXCL13 cells did not engage (no investment, no return). CXCL13 TCF7 cells engaged but lost TCF7 (invested without retaining the differentiation lock — net negative cell economy). CXCL13 TCF7 cells engaged and retained TCF7 (invested and retained the lock — net positive cell economy).

Direct test on Bassez n=29: three structurally predicted subgroups with distinct response rates.

Subgroup	Cell economy reading	n	Response rate
CXCL13	No engagement; no investment, no return	14	0%
CXCL13 TCF7	Engaged but lost differentiation lock; net negative	10	30%
CXCL13 TCF7	Engaged and retained differentiation lock; net positive	5	75%

Fisher 3×2 exact test: **p = 0.0012**. The framework predicting three distinct subgroups *before* looking at response data, from structural considerations alone, and the three subgroups separating in exactly the predicted direction with exactly the predicted ordering. The cell-economy reading explains why the ordering is what it is: response correlates with whether engaged cells retained the substrate (TCF7) that supports durable memory commitment.

The provisional status is because n=5 and n=10 in the CXCL13 subgroups are small. Liu et al. 2025 *Cell*

NSCLC n=234 reports TCF7 PDCD1 Texp correlating with recurrence-free survival in the same direction, which is consistent with but does not yet formally validate the lock.

Drug assignments — cell-economy supporters, not killing amplifiers. The intervention class for Position 12 is not amplification of the killing arm. It is support of the cell side of the cell-tumour balance: anti-CSF1R (pexidartinib) and anti-CCL2 to remove the inflammatory drain that consumes effector resources; IL-7 and IL-15 supplementation to add survival signal directly; metabolic modulators (metformin, mitochondrial uncouplers) to reduce mitochondrial cost; myeloid repolarisation (anti-CD40, anti-CD47) to convert the tumour microenvironment from resource-extracting to resource-supporting. The framing change matters because it predicts that interventions which *intensify killing* without supporting cell economy will produce strong initial killing followed by rebound — exactly the pattern of acquired resistance after sustained ICI response.

5.6 The Conservation regime — Positions 13–16 (the layered surveillance phase)

The four positions of Conservation are not handled at the same level of detail as Construction in this paper, but they require a regime-level treatment rather than a table-line summary. The reason is structural: Conservation is the only regime that operates on a layered accumulation of substrates, and that layered structure has direct implications for why certain interventions in cancer immunotherapy fail in ways that look idiosyncratic from a single-position view but are predictable from the regime view.

What Pass 4 operates on. Each prior regime hands a single conserved output to the next: Pass 1 hands a committed agent to Pass 2; Pass 2 hands a deployment-ready effector to Pass 3; Pass 3 hands a stable engagement to Pass 4. But Pass 4 does not discard those upstream outputs once it begins. The committed agent, the deployment-ready effector, and the stable engagement are all held simultaneously inside the Conservation regime as a layered structure. Surveillance is not monitoring one thing; it is monitoring the integrity of three nested fractals, each of which can drift independently. Compensation is not correcting one drift; it is correcting drift across multiple held layers at the same time.

This is what the framework calls a **higher-resolution regime**. The other three regimes operate on a single substrate per pass. Conservation operates on three.

The coupling of Positions 14 and 15. The Foundation page describes Positions 14 (Surveillance, the build step) and 15 (Compensation, the sense step) as structurally coupled. The coupling is not a sequential pair — surveillance does not finish before compensation begins. The two operate as an oscillation: surveillance produces a state-reading, compensation acts on the drift detected, surveillance reads the new state, compensation acts again. The pair succeeds when the oscillation is sustained without collapsing into either runaway suppression (Suppression barrier crossed in the wrong direction) or runaway response (auto-immunity).

The clinical signature of this coupling is that interventions that act only on one element of the pair often produce paradoxical effects. Anti-Treg drugs (anti-CCR8, anti-TGF- β) act on the suppressive arm of compensation. Used alone, they release one constraint while leaving the surveillance arm unchanged, and the system can over-correct in either direction depending on which other layer was already drifting. This is consistent with the mixed clinical results of standalone Treg-targeting therapies: they work when the underlying surveillance was holding the response below threshold for a Treg-related reason, and they fail when the surveillance was failing for a different reason that the Treg removal does not address.

Why the Suppression barrier is a higher-resolution failure. The Suppression barrier (the singularity at exit from Conservation) is failure of the system to monitor and correct across multiple held layers simultaneously. It is not one thing breaking. It is the failure of integrated multi-layer regulation — the nested holding pattern destabilises and either collapses into suppression that prevents new cycles, or fragments into uncoordinated responses. Anti-Treg drugs alone often do not unblock the Suppression barrier because they address one

component of compensation at one layer; the surveillance loop is multilayered and requires intervention at multiple layers to release.

The four positions inside Conservation are processes, not states.

The signature mistake when reading positions 13–16 is to describe them as endpoints — “memory established,” “patrol running,” “compensation set.” That framing misses the regime’s defining property: Conservation runs continuously, every position is an ongoing process, and the absence of process activity is the failure mode. A patient whose Conservation regime is “established” but no longer running has Suppression-barrier failure even when all the static markers look intact. The processes leave traces in the cell distribution; we read the regime by detecting the traces, not the steady-state pool sizes.

Position 13 — Stabilising (the detect process). The system continuously commits engaged effectors to durable post-engagement identities. T cells that survive Position 12 engagement are read for which compartment they should populate — central memory for circulation, effector memory for rapid recall, tissue-resident memory for the engaged tissue, terminally differentiated effector memory for short-lived high-output. Inflammation is concurrently resolved as cells that drove the killing phase wind down or apoptose. The process leaves traces as CD8_RM presence (tissue-resident commitment), CD8_EMRA presence (terminal-effector commitment), and the layered mix of memory-fate transcription factors (TCF7, EOMES, KLF2 for residence decisions, Bcl-2 family balance toward survival). Position 13 is failing when none of these traces appear despite engagement having occurred — the cells died in place rather than committing to a Conservation identity.

Position 14 — Patrolling (the build process). The system continuously circulates memory cells through tissues, monitoring for tumour antigen reappearance. The build is layered because the patrol must monitor the integrity of the committed agent (Pass 1 substrate), the deployment-ready effector pool (Pass 2 substrate), and the stable engagement state (Pass 3 substrate) simultaneously. The process leaves traces as the actively patrolling memory pool — circulating memory cells with intact response capacity (CD8_EM with retained IL-7R, IL-15R), tissue-resident memory in target tissues (CD8_RM), antibody titres at baseline. Position 14 is failing when the patrolling pool has shrunk or lost responsiveness.

Position 15 — Correcting (the sense process, coupled with 14). The system continuously measures the surveillance state against drift and corrects it. Treg-mediated regulation responds to ongoing antigen exposure; recall responses are mounted when antigen re-encounter occurs; auto-reactive drift is dampened. The process leaves traces in the regulatory cell pool that is *actively dividing* — not the resting Treg fraction but the Treg-proliferating fraction, plus the size of the regulatory pool relative to the effector memory pool it is balancing. The Suppression barrier lives in the coupling with Position 14: when correcting over-corrects, the patrol gets shut down across multiple layers, and no new cycle can start.

Position 16 — Turning (the resolve process). The system continuously generates new perturbations from within. Tumour antigen escape variants produced under selective pressure constitute a new Position 1 signal of a new cycle. The cycle does not end at Position 16; it generates the input for the next iteration. The process leaves traces in clonal evolution patterns and HLA-loss heterogeneity; in the framework, Position 16 is failing not when antigen drift stops but when the drift is no longer detected by the surveillance produced at Position 14 — the cycle turns over but the next iteration cannot start.

5.7 The Conservation processes are running in Bassez expanders

The Bassez breast cohort (n=29, pretreatment, ≥20 T cells per patient, scRNA cell-subtype labels) provides direct evidence that Conservation processes leave detectable traces in pretreatment tumour biopsy, and that those traces differ between patients whose response will subsequently expand on therapy (E, n=9) and those whose response will not (NE, n=20). This is post-hoc analysis of the same cohort used elsewhere in the paper,

presented as process-detection rather than as a lock claim.

Position 15 (Correcting process) is running in expanders, not in non-expanders. The Treg compartment is structurally larger and structurally more active in pretreatment biopsies of patients who go on to expand:

Marker (process trace)	E median	NE median	one-sided p
CD4_REG fraction (Treg pool size)	0.151	0.075	0.004
CD4_REG_Proliferating (active correction)	0.0010	0.0000	0.009
Treg-to-CD8-memory ratio (compensation balance)	0.51	0.21	0.002

This is the inversion of the conventional reading. Conventional reading: “Treg = bad, more Treg = worse response.” Framework reading: Treg-mediated correcting is a process that only runs when there is something to correct. Patients whose cycle has reached Position 15 generate the correcting signal and proliferating Tregs as part of the running Conservation regime; patients whose response stalls upstream (at the Egress barrier, Position 8, or before) generate no correcting signal because there is no engagement to correct. High Treg + proliferating Tregs + balanced Treg-to-effector-memory ratio is therefore the signature of the cycle running, not of the cycle being suppressed. The same direction holds on-treatment: CD4_REG_Proliferating E vs NE on-treatment $p=0.005$, and the Treg/CD8-memory ratio remains separated.

The CD4-help-given trace is the strongest single signal. CD4_EX (helped CD4 reaching exhaustion after delivering help) E=10.1% vs NE=1.5%, $p=0.0001$ pretreatment, $p=0.0001$ on-treatment. This is a Position 7 trace observable in tumour biopsy: CD4 help cannot be measured directly in the tumour-draining lymph node where Position 7 occurs, but the help-given exhausted phenotype is what reaches the tumour after the help is delivered. The strength of this signal in pretreatment biopsies of expanders is consistent with the Position 7 hypothesis filed in §5.2 — patients in whom the licensing test passes leave a trace of CD4-help-having-been-given in the tumour biopsy.

Position 13 traces show the cycle is not stuck in resting memory in expanders. CD8_RM (tissue-resident memory) is somewhat *lower* in expanders (E=7.4% vs NE=10.5%, two-sided $p=0.11$, trend), and total memory pool fraction (CD8_EM + CD8_RM + CD8_EMRA combined) is also lower in E (27.4% vs 39.2%, $p=0.04$). This is again contra the conventional reading and consistent with the framework reading: in expanders the cycle is actively running through Conservation positions and producing process traces, not parked in resting memory. In non-expanders the larger memory pool reflects accumulated cells that did not progress through the cycle — they were made earlier and never replaced because no new cycle is running.

What this does not show. These measurements cannot distinguish Position 13 from Position 14 from Position 15 at single-cell resolution; the cellular subtypes label populations not processes. They cannot directly observe the layered surveillance (multiple held fractals being monitored simultaneously), only the aggregate signal of Conservation activity. They cannot resolve Position 16 (antigen drift detection) — that requires longitudinal sampling and TCR-clone tracking. The argument is that the *running* of the regime is detectable in pretreatment biopsy as a multi-marker process signature, and that the running differs between patients with different prognosis.

5.8 The Conservation regime in cancer (drug implications)

The treatment in §5.6 is structural and applies to any immune cycle. The Bassez evidence in §5.7 shows that the regime's processes are running in cancer patients who go on to expand, and not running in patients who do not. The remaining question is what this implies for intervention.

Cancers where Conservation is the dominant failure mode. From the Atlas distribution (van der Klein 2026d), several cancer types are scored as Conservation-dominant: the cycle reaches Position 12 or beyond, the response runs, and the failure is at the transition into durable maintenance. This is mechanistically distinct from cold tumours (Construction failure) or unresponsive infiltrate (Encounter failure). The clinical signature is **initial response followed by relapse** — the response fired, killing happened, but durable surveillance was never established. The patient looks like a responder, then looks like a non-responder.

Why anti-Treg monotherapy often fails. Anti-CCR8, anti-TGF- β , and similar single-arm Treg-targeting drugs act on Position 15 alone. The framework predicts they will work specifically when Position 14 is intact (the patrolling process is running, ready to expand if the correcting brake is released) and the patient's cycle has reached Position 13–15 where these processes operate. They will fail when Position 14 is itself failing (memory cells did not establish at Position 13, T_RM never seeded the tissue), because removing the correcting signal does not help when the substrate that should respond is absent. The Bassez data is consistent with this: patients with high Treg fractions are running the Position 15 correcting process, and removing that correction will redistribute the running cycle. Patients with low Treg fractions are not running it, and removing what is not running does nothing.

Why combination ICI + IL-15 / IL-7 may have non-additive effects. IL-7 and IL-15 act on Position 14 (patrolling process maintenance). ICI acts on Position 10 (engagement brake release). Combining them addresses two layers, which is structurally consistent with the multi-layer requirement of Conservation regime intervention. The framework predicts these combinations will show effects that exceed the sum of monotherapies specifically in patients whose cycles are in the Conservation regime, not in patients whose response stalls in Construction or Encounter.

Why neoadjuvant ICI sometimes outperforms adjuvant. Neoadjuvant timing exposes the system to ICI while the primary tumour is intact, generating a richer Position 13 substrate (more antigen-experienced cells available to commit to durable memory). Adjuvant timing operates on a smaller substrate. The framework predicts that the neoadjuvant advantage should be specifically in tumour types where Position 13 establishment is the rate-limiting step, not in types where the primary failure is upstream.

These intervention predictions are framework-derived hypotheses grounded in the Conservation regime structure and the Bassez process traces in §5.7, not locked claims on data in this paper. They are concrete enough to be tested in trials specifically designed to stratify on Conservation-phase process traces (Treg-proliferating fraction, Treg-to-CD8-memory ratio, CD4_EX presence, memory pool composition); that work is a separate research programme. The §5.7 result is the first direct evidence in this paper that Conservation processes are detectable in pretreatment tumour biopsy and that their running state correlates with subsequent clinical response.

5.9 The remaining five positions in summary

Positions 1, 2, and 3 exist in every immune response, with or without cancer present, but cancer disrupts each of them in specific ways that have been characterised in the literature. Positions 9 and 11 are partially treated by Mellman or replicated elsewhere in this paper; their entries below reflect status, not cancer-specific disruption.

#	Immunology name	Structural role	Status / cancer disruption	Drug class
1	Antigen release	Universal background of cell death; threshold-crossing into attended signal	Tumour cells die by necrosis or apoptosis without immunogenic patterns; PRR signaling actively suppressed; emission and reception both present, conversion below threshold	TLR/STING agonists, oncolytic viruses, immunogenic chemotherapy (anthracyclines)
2	DC antigen capture and migration	DC processes antigen, migrates to dLN	TAM/MDSC dominance over the myeloid compartment; DCs remain immature; CCR7 migration disrupted by TGF- β /IL-10/PGE from tumour stroma; antigen reaches dLN at sub-threshold quantity	GM-CSF, Flt3L, anti-CSF1R, anti-CCR2
3	DC maturation and presentation	MHC loading, CD80/86 upregulation, naive T-cell repertoire scanning	MHC-I downregulation on tumour-derived antigen; cross-presentation impaired by cDC1 depletion in dLN; T-cell repertoire may be inadequate against self-derived neoantigens	TLR ligands, CD40 agonists, neoantigen vaccines
9	Trafficking and infiltration	Detect step: signaling and transmit through tumour endothelium	partial in Mellman	anti-VEGF (vascular normalisation)
11	Sustained killing	Repeated synapse formation; granzyme/perforin delivery over time	replicated in Bassez	dose intensification, combination ICI

Position 9 is the framework's structural reading of what Mellman called "trafficking and infiltration" combined. It is a detect step in the Encounter regime — the cell **actively makes itself visible** to the tumour endothelium through chemokine receptor expression (CXCR3, CCR5), integrin activation (LFA-1, VLA-4), and S1P-receptor downregulation. The endothelium either receives this signal and permits transit (pass) or fails to receive it and the cell stays in circulation (fail). Tumour vasculature is structurally and functionally abnormal: chaotic vessel architecture, dysregulated adhesion molecule expression, hypoxia-driven tight junctions, all of which disrupt the signaling-and-transmit handshake. Anti-VEGF works at this position not by killing tumour vessels but by normalising them to restore the handshake. This is structurally distinct from Position 10 (the synapse with the cancer cell once inside the tumour) and from Position 5 (myeloid hijack of the dLN); Position 9 is the gateway between the two.

Position 11 (sustained killing) deserves a brief note. In the Bassez data, the running of Position 11 is replicated with orthogonal marker panels: engaged-active fraction (PDCD1 AND (MKI67 OR HLA-DRA OR CD38)) shows expander median 17.5% vs non-expander 4.4%, $p = 0.0004$; terminal-exhausted fraction shows 13.8% vs 0.5%, $p = 0.0009$; ghost effector ratio (engagement over conserved memory) $E = 1.43$ vs $NE = 0.69$, $p = 0.0001$. Both signatures read the running of Position 11 — engagement-and-activity and post-engagement spent state — as temporally consecutive states of the same trajectory.

6. The seven questions, resolved

6.1 Why is PD-L1 a weak biomarker?

PD-L1 measures the brake on the synapse, but the brake is irrelevant when the cycle has not reached the synapse. In Bassez, a substantial fraction of non-responders had no detectable CXCL13 CD8 cells in their tumour. Across the broader CXCL13 literature, the direction is consistent: responders are enriched for cells that have reached the synapse, non-responders are not. PD-L1 has a structural blind spot: it reads the intervention surface at Position 10 without first checking whether the cycle has reached Position 10. The replacement biomarker is the Position 10 arrival marker — CXCL13 CD8 fraction at single-cell resolution, the LCAM-T module at bulk.

6.2 Why do response rates vary so widely across cancer types?

Because the response stalls at different positions in different cancers. The Atlas finds that in only 4 of 33 cancer types does the cycle reach the Encounter phase where anti-PD-1 acts. Melanoma is one of those four. MSS colorectal is not — the response stalls at the Egress barrier (Position 8). Pancreatic is not — stalls due to Treg over-control (Position 15). Glioblastoma sits at the extreme over-control end of the distribution.

6.3 Why do cold tumours stay cold?

Because the Construction phase has failed and no cells are leaving the lymph node to traffic toward the tumour. In the 3CA rebuild, 43% of patients have a response that stalls in the Construction phase — split between Position 5 (myeloid hijack in tumour bed, C4 mechanism) and Position 6 (Treg IL-2 sink in lymph node, C1 mechanism); 32% of MMRp CRC patients are stalled at Position 8 (Egress barrier failure). The tumour looks cold because Construction never completed. Synapse-acting drugs cannot help; these patients need Construction-phase interventions, and the specific intervention depends on which sub-position the patient's response stalls at.

6.4 Why does CAR-T succeed where checkpoint blockade fails?

Because CAR-T acts at Position 7 (effector differentiation and prototype testing), not at Position 10 (synapse engagement). Position 7 is where the patient's own system builds the effector machinery; CAR-T replaces this step by supplying an externally built effector, bypassing the patient's own Construction phase entirely. This is why CAR-T works on heme malignancies where the patient's lymph-node priming and build is systemically compromised, and why it can work after anti-PD-1 failure: the patient was failing at Position 7 or 8, not at Position 10, and a synapse-acting drug cannot help a patient whose effectors were never built.

6.5 Why do patients respond and then stop responding?

Acquired resistance has three structurally distinct loci. *Drain barrier failure (Position 12)*: engagement happened, killing happened, but the cells did not come out of the synapse with positive cell economy — the response phase exhausted the effector pool without locking in memory-fate substrate, so the response cannot be sustained. *Conservation phase failure (Positions 14–15 coupling)*: memory established but the surveillance-compensation oscillation drifts. Anti-Treg drugs given alone often fail here because the coupled pair requires intervention at multiple layers, not just the suppression arm. *Suppression barrier failure (Position 16)*: tumour antigen escape variants are produced under selective pressure from the response, the surveillance fails to detect them as a new Position 1 signal, and the cycle does not restart. Mellman's single feedback arrow from killing back to antigen release contains four positions in the completed cycle, each with distinct failure modes and distinct interventions.

6.6 Why did five labs converge on CXCL13?

Because CXCL13 is the chemokine the synapse structurally produces when TCR engagement fires. The five labs converged on it without an explanation because they were measuring the arrival output of the position where their drug acts. The marker was already in the literature; the structural derivation provides the *reason* it is the marker.

6.7 Why did the 2023 Mellman update extend outward rather than inward?

Because the seven-step cycle could not be expanded internally to carry the heterogeneity the data contained, so the explanatory work had to be placed orthogonal to the diagram. The completed cycle does the inward expansion the 2023 update declined to do: indication, genotype, and immunotype are reabsorbed as positions inside the sixteen-position map. Indication maps to the modal cycle position typical of each cancer type. Immunotype maps to the position where the patient's own response currently stalls. Genotype maps to the molecular features that determine which position is the bottleneck.

7. Clinical implications

The clinical implication is structural rather than therapeutic in the usual sense. The completed cycle does not propose new drugs. **It proposes a different mapping between existing drugs and patients — and a precise statement of where the existing paradigm is correct and where it is not.**

7.1 Where each paradigm applies: a quantitative reframe

The Chen-Mellman cycle as drawn is correct for the 4 of 33 cancer types in which the immune cycle reaches the Encounter regime. This is not a small claim or a face-saving exception. It is a quantitatively precise

scope statement: in melanoma, MSI-high colorectal, Hodgkin lymphoma, and a small set of mismatch-repair-deficient or high-tumour-mutational-burden cancers, the cycle does reach Position 10, anti-PD-1 acts where the response stalls, and response rates of 30–70% are observed exactly as the original framework predicts. **For these cancers, the existing paradigm is not wrong. It is sufficient.**

The completion proposed in this paper concerns the other 29 cancer types — the cancers in which the response stalls upstream of where anti-PD-1 acts — and within Encounter cancer types, the patients within those cancers whose response also stalls upstream. Concretely:

- **3CA rebuild over 12 cancer types (n=887 single-cell patients):** the modal patient sits in Construction (Positions 5–6, ~43% of patients across types), with a further substantial fraction stalled at the Egress barrier (Position 8) or in upstream Potentiality positions. Only a minority of patients in this dataset have an immune cycle that has reached Encounter.
- **Pelka MMRp colorectal (n=110):** 32% of patients are stalled at Position 8 — built effectors that never crossed the Egress barrier. The Pelka MMRd subset and the 4% MSS-but-anti-PD-1-responder subset reach Encounter and respond to anti-PD-1.
- **Bassez breast (n=29):** the cycle position is bimodal. Expanders cross Position 10 and respond; non-expanders stall upstream and do not respond. CXCL13 CD8 fraction predicts this assignment at CV AUC 0.97.
- **Sade-Feldman melanoma (n=11 pretreatment anti-PD-1):** 0/11 match the Position 8 fingerprint (95% binomial CI 0–28%). Melanoma is an Encounter cancer; the cycle reaches the synapse and the existing paradigm holds.
- **Atlas pan-cancer (n=11,373 TCGA bulk):** in 4 of 33 cancer types, the modal cycle position is in Encounter. In 29 of 33, the modal position is upstream.

The reframe is therefore not “cancer is a Construction disease.” It is: **in the cancer types and patient subgroups where anti-PD-1 fails, the response most often stalls in Construction (Positions 5–8), with smaller fractions stalled in upstream Potentiality (Positions 1–3) or in Conservation over-control (Position 15). In the cancer types and patient subgroups where anti-PD-1 succeeds, the cycle has reached Encounter and the existing paradigm correctly identifies the rate-limiting step.** The framework’s contribution is to name where each patient’s response stalls, so that drugs are matched to the position rather than to the diagnosis.

This reframe has three consequences:

1. **The 30–40% pan-tumour response ceiling on anti-PD-1 monotherapy is structural, not pharmacological.** It is the fraction of patients across cancer types whose cycle has reached Position 10. Increasing anti-PD-1 potency cannot raise this ceiling; only matching patients whose response stalls upstream to upstream-acting drugs can.
2. **Cold tumours are not a uniform category.** “Cold” describes the tumour microenvironment but does not specify which upstream position is rate-limiting. The 43% of patients whose response stalls in Construction in 3CA split between Position 5 (myeloid hijack), Position 6 (Treg IL-2 sink), Position 7 (failed effector build), and Position 8 (failed egress) — and each requires a different drug class.
3. **The CAR-T/checkpoint complementarity is not coincidental.** CAR-T succeeds where checkpoint blockade fails because CAR-T acts at Position 7 (externally supplied effector) and checkpoint blockade acts at Position 10 (releases brake on synapse). They are non-overlapping interventions in the same cycle.

The paradigm shift the paper proposes is therefore narrow and precise: **for the cancers and patients whose cycle has reached Encounter, the existing paradigm is correct; for the larger fraction of cancers and patients whose response stalls upstream, a different mapping between drugs and positions is required.** The Encounter biology that the field built over the last decade is not displaced. It is bounded.

7.2 The existing drug arsenal is mostly already adequate

For each of the sixteen positions, the framework's drug assignment is drawn from drugs that already exist and are already approved or in late-stage trials. Position 1: TLR agonists, STING agonists, oncolytic viruses (drive antigen release). Position 2–3: GM-CSF, Flt3L, CD40 agonists (DC capture and maturation). Position 4: ipilimumab (release CTLA-4 brake on CD28-B7 commitment). Position 5 (C4 mechanism): bevacizumab, pexidartinib, anti-CCR2 + Flt3L (clear myeloid hijack so DCs can mature and migrate). Position 6 (C1 mechanism): engineered IL-2 muteins; IL-15 superagonists; anti-CCR4/anti-CCR8 (deplete dLN Tregs). Position 7: CAR-T, TCR-T, BiTEs (replace failed effector build); IDO inhibitors. Position 8: CD40 agonists (sotigalimab, selicrelumab); IL-12; CD25-low-targeted IL-2 variants; adoptive transfer of TdLN stem-like CD8. Position 9: bevacizumab + ICI combinations. Position 10: pembrolizumab, nivolumab. Position 12: anti-CSF1R, anti-CCL2, anti-CD40, anti-CD47. Position 15: anti-CCR8, anti-TGF- β .

What the framework adds is the *map*. A patient whose biopsy reads C0 (Reserve-rich, no commitment yet) is a candidate for ipilimumab or for an Activation-phase intervention like STING agonism, not for pembrolizumab. A patient whose biopsy reads C4 (Drain-rich, myeloid hijack blocking the upstream signal) is a candidate for anti-VEGF or anti-CSF1R combined with DC-promoting agents, not for pembrolizumab. A patient whose biopsy reads C1 (DC-rich, Treg IL-2 sink) is a candidate for IL-2 muteins, not for pembrolizumab as monotherapy. A patient whose biopsy reads built-but-parked (the 32% MMRp Position 8) is a candidate for CD40 agonism. A patient whose biopsy reads CXCL13 TCF7 on CD8 has crossed Position 12 and is in the responder population for pembrolizumab. **The drugs do not change. The matching does.**

7.3 The bidirectional ipilimumab/belatacept consistency check

The clearest cross-clinical-context demonstration the framework offers is at Position 4, the Activation barrier. Two drugs act on the CD28–B7–CTLA-4 axis from opposite directions:

- **Ipilimumab** (anti-CTLA-4) is given in cancer to *enhance* the response. It blocks CTLA-4 from binding B7, freeing B7 to engage CD28, and the commitment threshold is more easily crossed.
- **Belatacept** (CTLA-4-Ig) is given in kidney transplant to *suppress* the response. It is a CTLA-4 mimetic that occupies B7, preventing CD28 from binding, and the commitment threshold cannot be crossed.

Same molecular surface. Same cycle position. Same drug class. Opposite clinical goals. The framework predicts this directly: the position is the position whether the goal is to enable the response or to suppress it. **Interventions at any cycle position are bidirectional by construction** — a position that an oncologist wants to push their patient *through* is the same position a transplant nephrologist wants to hold their patient *back from*. The cycle is not a cancer-specific construction. It is the structural map any antigen-specific immune response runs along.

7.4 What the cycle says about trial design

The 2023 Mellman update introduced *immunotype* as a patient-stratification axis sitting orthogonal to the seven-step diagram. The completed cycle reabsorbs immunotype as the answer to a single question: *at which of the sixteen positions does this patient's response stall?* The 3CA rebuild's five archetypes are five immunotypes, and each maps to a specific position in the cycle with a specific drug class indicated.

For trial design, the structural reading produces three concrete operating procedures that the original cycle could not support:

Coverage maps. Decompose each trial arm into the cycle positions it covers. Two arms can have very different coverage even when the molecular composition looks similar — and the difference is what predicts whether adding a drug actually adds an effect. A combination that covers Position 8 plus Position 10 should outperform

monotherapy at Position 10 in an MMRp population; a combination that covers Position 10 twice (different anti-PD-(L)1 antibodies) should not.

Population-position match. Score the proposed enrolment criteria against the population’s modal cycle position from the Atlas distribution or from the trial’s own pretreatment biopsy panel. If the modal position is upstream of where the regimen acts, the trial will not separate from control regardless of drug potency. The MSS colorectal anti-PD-1 trials are a worked example of this failure mode: high potency at Position 10 against a population whose mode is at Position 8.

Stratification on the position. Where pretreatment single-cell or bulk panels are available, stratify the trial population by the position the response stalls at — and analyse the position-matched subpopulation as a primary endpoint rather than the ITT cohort. The CXCL13 TCF7 subgroup in Bassez has 75% response rate; the ITT cohort response rate is 31%. The drug works in the patients its position covers, and the trial’s job is to find them.

8. Limitations and pre-registered tests

8.1 Limitations

Sample size at Position 10. Bassez n=29 single-cell and GSE207422 n=24 bulk together cover 53 patients across two cancer types and two platforms, with the same structural prediction confirmed in the same direction. Within Bassez, three independent axes (CXCL13 CD8, Bassez-author CD8_EX, CD8_EX_Proliferating) all separate responders from non-responders at single-feature CV AUC 0.81–0.97 — convergent measurement reduces the risk that any single axis is a fluke of marker selection, but does not increase the n. Each individual test has sample-size limitations; the strength of the lock is in their convergence across instruments and across cohorts. The Liu et al. 2025 *Cell* NSCLC dataset (n=234) and the SU2C-MARK dbGaP NSCLC cohort (n=309) are the formal validation cohorts identified in §8.2.

Position 8 across cancers. Position 8 is now locked across colorectal (Pelka 2021, MMRp/MMRd stratification) and breast (Bassez 2021, E/NE stratification), and confirmed rare in melanoma at 0 of 11 pretreatment anti-PD-1 monotherapy patients (Sade-Feldman 2018, 3CA-curated). The melanoma test is on a small cohort and the 95% binomial confidence interval on the 0% match rate is 0–28%, which is well below the pre-registered falsification threshold of 40% but cannot distinguish a true zero from a low-single-digit rate. The “build complete” component of the fingerprint also relies on a partial vocabulary mapping in Sade-Feldman (Memory/effector only; resident-memory is not separately labeled in the 3CA curation of this dataset). The continuous-score direction reversal in melanoma is consistent with the framework but is a post-hoc observation, not a pre-registered prediction, and rests on n=11. Replication on Yost BCC (GSE123813), pretreatment NSCLC (Liu et al. 2025 *Cell*), and pancreatic cohorts as they become available is the next priority. The 32% Position 8 fraction in MMRp colorectal and the 60% Position 8 fraction in Bassez NE breast set the magnitude expectation in Construction-phase cancer types; the 0% in melanoma sets the magnitude expectation in Encounter-phase cancer types.

Position 7 status. Position 7 is described in §5.2 as a structural hypothesis grounded in published literature (Schenkel et al. 2021 in TdLN, Borst et al. 2018 helpless priming review, Ahrends et al. 2017, Ferris et al. 2020 cDC1 licensing, Lei et al. 2024 cDC1 helped signature in human cancers) rather than as a locked biomarker on our own data. The mini-encounter reading — that prototypes are tested at Position 7 against a licensed cDC1 partner, with passed and failed outcomes corresponding to the published helped vs helpless CD8 phenotypes — predicts cell-state distributions in matched-TdLN-and-tumour cohorts that are not yet available with the response data needed for a lock. The hypothesis is concrete enough to be falsifiable when matched-tissue

cohorts with anti-PD-1 response data become accessible. The Position 7-failed (helpless) and Position 8-parked phenotypes overlap in tumour biopsy because both are TCF7-positive and exhaustion-low; clean separation requires the lymph node tissue. Position 7 is presented in this paper as a hypothesis, not a claim.

Conservation regime treatment. The Conservation regime is presented in §5.6–5.7 at the regime level, with the four positions (13–16) named and described and with hypotheses about cancer-specific failure modes drawn from the layered-substrate and coupled-14-15 structure. The regime treatment is structural derivation grounded in framework canon, not a lock claim on cancer data. The Conservation regime predictions — that anti-Treg monotherapy responders should stratify by Position 13 establishment, that ICI + IL-15/IL-7 combinations should show non-additive effects in Conservation-phase patients, that neoadjuvant ICI advantages should localise to Position 13-limited tumour types — are concrete enough to be tested in trials specifically designed to stratify on Conservation-phase markers, but no such cohort is currently available to us. The treatment in §5.6–5.7 should be read as framework extension that organises why certain published clinical patterns (anti-Treg failure, neoadjuvant > adjuvant in some cancers, late-relapse after ICI) appear when they appear, not as a primary contribution of new empirical data.

Provisional status of Position 12. The three-subgroup test on Bassez has $n=14$, $n=10$, $n=5$ in its three arms. The structural prediction is correct ($p=0.0012$) and the response-rate ordering is exact, but the confidence interval on the responder subgroup is wide. Validation in an independent cohort with matched single-cell and response data is the priority. The cell-economy reading of Position 12 (§5.5) is a structural reframing of what the empirical test measures; it changes the interpretation but not the test itself.

The framework’s universality claim is supported externally to this paper. The sixteen-position scaffold has been tested against 22 dissipative systems drawn from physics, chemistry, biology, and engineering (van der Klein 2026b), and against five disease families in immunology with cell-fraction measurement filling eleven of sixteen positions. The cross-disease evidence is presented in a separate companion paper. This paper is the cancer-immunology presentation, and its claims should be evaluated on the cancer evidence above, not on the broader universality claim.

8.2 Pre-registered tests

Six tests are pre-registered as falsification opportunities:

1. **Position 10 lock validation.** Apply the CXCL13 CD8 single-feature logistic regression to an independent anti-PD-1 cohort with matched single-cell data and response annotation. Prediction: $AUC \geq 0.85$. Falsification: $AUC < 0.75$.
2. **Position 8 across cancers, with cycle-position-conditioned direction.** Apply the Position 8 fingerprint to scRNA cohorts in additional cancer types. (a) In **Construction-phase cancer types** (MSS gastric, pancreatic adenocarcinoma, MSS colorectal validation cohorts): prediction is enrichment of Position 8 in the ICI-failing subpopulation at $OR \geq 3.0$ with $p < 0.01$; falsification at $OR < 2.0$. (b) In **Encounter-phase cancer types** (additional melanoma cohorts beyond Sade-Feldman, BCC via Yost GSE123813): prediction is rarity of the fingerprint at $<25\%$ match across cohort; falsification at $>40\%$ match. The cycle-phase classification is taken from the §3.4 Atlas distribution and is fixed before any test runs.
3. **Position 12 three-subgroup replication.** Apply the CXCL13/TCF7 subgroup definition to an independent cohort. Prediction: response rate ordering $CXCL13 < CXCL13\ TCF7 < CXCL13\ TCF7$ with at least a two-fold difference between extremes. Falsification: ordering reversed or compressed below $1.5\times$.
4. **Construction-phase intervention test.** In a patient cohort scored as C1 or C4 by the 3CA cell-fraction signature, anti-VEGF + ICI combination (for C4 / Position 5) or IL-2 muteins + ICI (for C1 / Position 6) should show a response rate at least $1.5\times$ higher than ICI alone in each respective sub-population.

5. **CD40 agonist responder prediction.** In CD40 agonist trial samples (sotigalimab, selicrelumab), responders should be enriched for the Position 8 fingerprint (built-but-parked CD8 in tumour-draining lymph nodes) at OR ≥ 3.0 versus non-responders.
6. **Position 7 helped/helpless ratio (open hypothesis).** In a matched-TdLN-and-tumour scRNA cohort with anti-PD-1 response data and CD8 sub-clustering at TCF7/GZMB/PD-1/CD28 resolution, two cell populations should be distinguishable in the TdLN: passed prototypes (TCF7-low GZMB-positive CD28-retained PD-1 declining) and failed prototypes (TCF7-high GZMB-low CD28-retained PD-1-high). Prediction: the helped-to-helpless ratio in the TdLN should be higher in anti-PD-1 responders than in non-responders, with a magnitude separation comparable to or larger than the Position 10 CXCL13 signal in the tumour. Falsification: ratio inverted or no significant difference. This test is open until matched-tissue cohorts become accessible; we are aware of no current public dataset that meets all four requirements.

8.3 What the paper does not claim

The paper does not claim that every cancer is curable. It does not claim that the completed cycle replaces standard-of-care disease classification. It does not recommend any specific drug for any specific patient outside the calibrated indications. It does not claim that the sixteen positions exhaust the structure of the immune response at higher resolution.

The paper claims, more modestly, that the Chen–Mellman cycle is a true description of six positions out of sixteen, that the missing ten positions are where the unanswered questions of the field live, that two of those positions (Position 10, Position 8) have been locked by structurally derived biomarkers in independent published cohorts, that 43% of cancer patients have a response that stalls in the Construction phase that the original cycle compresses into a single arrow, that in only 4 of 33 cancer types does the immune cycle reach the Encounter phase where anti-PD-1 acts, and that the existing drug arsenal is mostly already adequate to the patients in the clinic once each patient is mapped to the position where their response currently stalls.

9. Conclusion

The Chen–Mellman cancer-immunity cycle, drawn in 2013 from a decade of clinical experience with checkpoint blockade, has organised tumour immunology for thirteen years. Its seven steps gave the field a shared language, a rational design space for drugs, and a framework that holds up where it applies. It does not apply everywhere, and the questions it cannot answer have accumulated into the seven clinical observations of §1. Each is a question about *where in the response the patient's response stalls*.

The completion presented here adds ten positions to the original six. The empirical evidence concentrates at two positions locked by structurally derived biomarkers in independent published cohorts: **Position 10** (the cytolytic synapse), validated on Bassez breast at three independent axes within one cohort (CXCL13 CD8 fraction at CV AUC 0.97, the Bassez authors' independent CD8_EX label at 0.93, and CD8_EX_Proliferating for Position 11 at 0.81) plus the LCAM-T module at CV AUC 0.74 on GSE207422 NSCLC bulk, with multi-cohort literature consensus across five cancer types; and **Position 8** (the Egress barrier substrate, the licensing handshake the original cycle has no name for), locked across two Construction-phase cancer types — colorectal (Pelka 2021, OR 4.64, $p = 0.0028$) and breast (Bassez 2021, one-sided $p = 0.0024$) — and confirmed rare in a third, Encounter-phase cancer type (Sade-Feldman 2018 melanoma, 0 of 11 pretreatment anti-PD-1 monotherapy patients matching, against a pre-registered specificity prediction of $<25\%$), with one fingerprint specified before each of the three tests. **The Construction phase aggregate** is the single largest patient bottleneck, accounting for 43% of patients in the 3CA rebuild, with two distinct mechanisms (Position 5 myeloid-hijack

failure, Position 6 priming-completion failure). **Position 4** is supported by stage enrichment (Fisher $p = 0.0099$) on a single cohort and is the location of the bidirectional ipilimumab/belatacept consistency check from kidney transplant. The Atlas distribution across 11,373 TCGA samples and 33 cancer types finds that in only 4 of 33 cancer types does the immune cycle reach the Encounter phase where anti-PD-1 acts.

The completion is additive, not corrective. Mellman's seven steps remain the seven steps he drew, unmodified, occupying their original meaning in the sixteen-position map. What the completion adds is the naming, the structural biomarkers, the drug assignments, and the population and patient-level evidence for the ten positions that the original diagram compressed into arrows.

The most consequential single claim in this paper is structural: **in most cancer types, the response does not reach the position where anti-PD-1 acts**. Of 33 cancer types, only 4 are at the synapse phase. The decade of immunotherapy effort directed at the killing step has worked precisely in those 4 — the MMR-deficient, the high mutational burden, the hot tumours — and has failed in the other 29. The response stalls upstream in 16 of them (8 in Activation, 8 in Construction) and at Treg over-control in 13. The next decade should be directed at the positions where the patients actually are — and in particular at the previously unnamed handshake at Position 8, where built effector cells fail to receive the licensing signal that would let them cross the Egress barrier into Encounter at all.

The drugs that act on the upstream regimes mostly already exist. CD40 agonists, IL-2 muteins, anti-VEGF, anti-CSF1R, STING agonists, CAR-T, anti-CCR8. They are in clinical trials, with mixed but real signals, and the framework predicts which patients each should work for — the patients whose response has reached the position the drug acts on. **Chen and Mellman's cycle was the right picture of the regime that the drugs of its decade acted on. Drawing in the other three regimes is what lets the patients whose response stalls there become visible as candidates for treatments that are already, in most cases, on the shelf.**

The cycle has sixteen positions. Most cancer patients are not at Position 10. The drugs that act at the other positions exist. Mapping each patient to the position where their response currently stalls — and matching them to the drug that acts there — is the operating procedure the completed cycle makes possible.

Acknowledgments

The Position 10 lock rests on the prior published findings of Bassez, Leader, Zhang, Liu, Im, Miller, Hudson, Beltra, Ahmed, and others; the structural explanation unifying them is the contribution here. The framework derivation builds on the universal sequence work and the quantitative tests paper, the intervention geometry paper, the Immune Barrier Atlas, and the four-state observer protocol. The Pelka et al. 2021 colorectal cancer dataset and the Curated Cancer Cell Atlas (3CA, Tirosh lab, Weizmann Institute) consortium provided the data on which Positions 5–6 and 8 were locked. The 3CA curation of Sade-Feldman et al. 2018, with its CD8 sub-cluster annotations, made the melanoma specificity test in §3.2 possible.

Correspondence

Raimo van der Klein, founder of Encounter. Web: encounter.bio. Email: raimo@encounter.bio. The framework's full development, axioms, and broader applications are published at generativegeometry.science.

Cite as: van der Klein R. (2026). Completing the Chen–Mellman Cancer-Immunity Cycle: Sixteen positions, two cross-cancer locks, and a structural account of why anti-PD-1 fails where it fails. Generative Geometry. Zenodo.

Generative Geometry, May 2026. All Rights Reserved.

References

Framework foundations (Zenodo, van der Klein 2026):

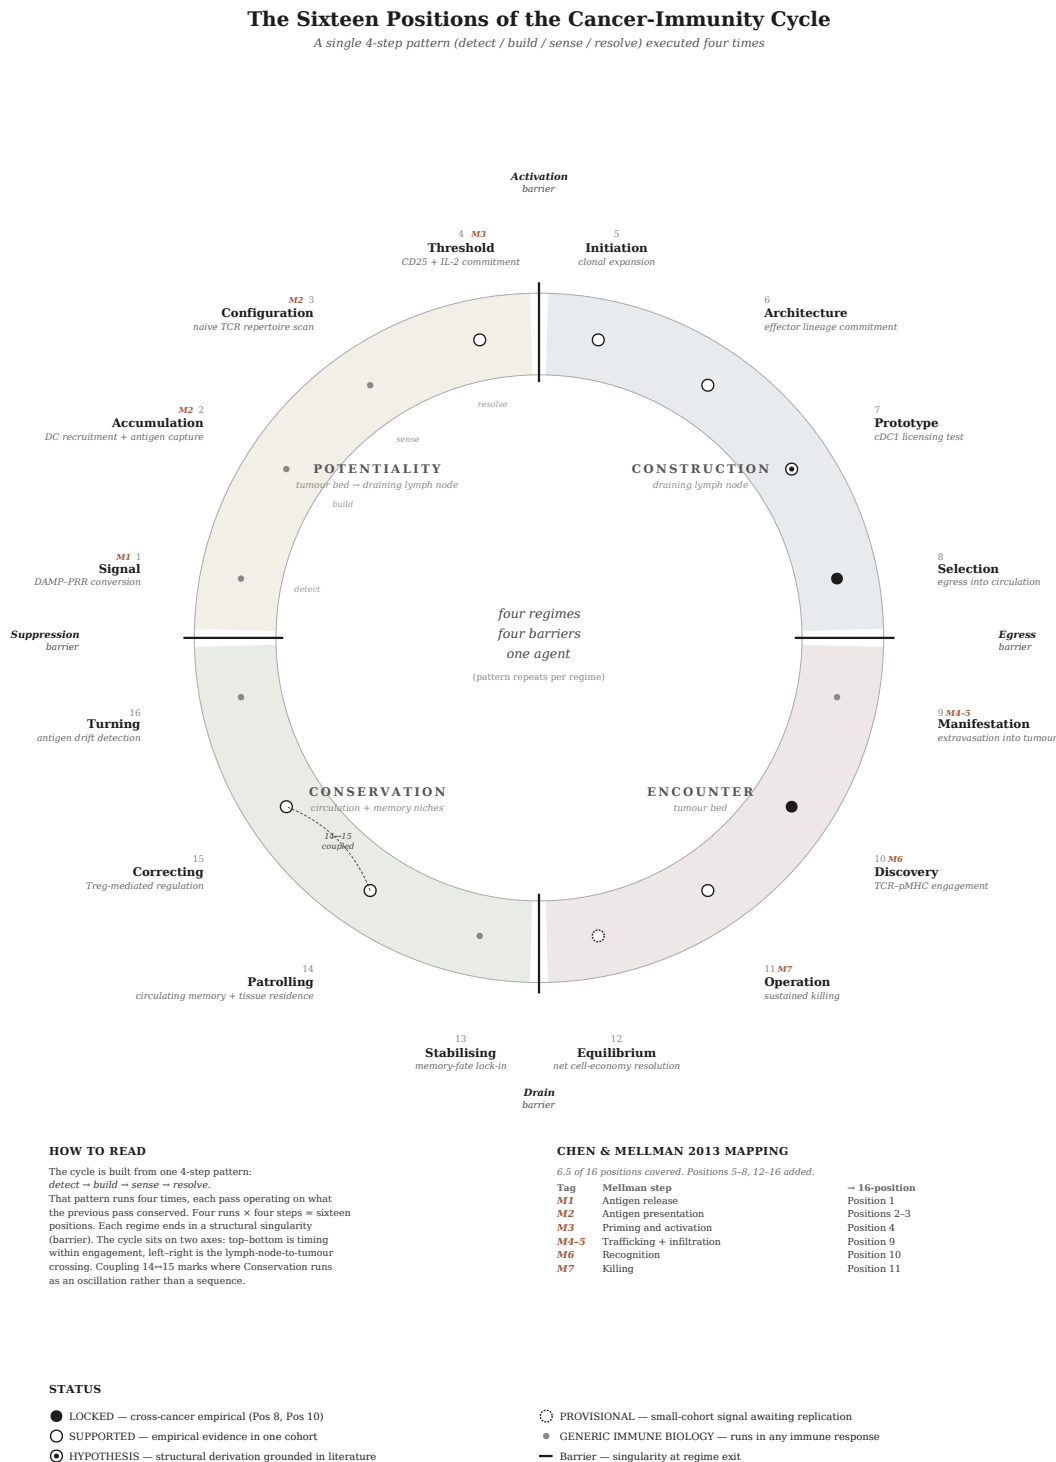
- *There Is Only One Way to Grow: The Universal Sequence of Dissipative Processes.* (2026a). DOI: 10.5281/zenodo.19007180.
- *Quantitative Tests of the Universal Dissipative Process Sequence.* (2026b). DOI: 10.5281/zenodo.18966489.
- *The Geometry of Intervention: A Universal Formula for Intervening in Dissipative Systems, Derived from Two Operations.* (2026d). DOI: 10.5281/zenodo.19022091.
- *The Immune Barrier Atlas: Cancer-Type-Specific Resistance Mechanisms Across 33 Tumor Types and Their Matched Drug Classes.* (2026c). DOI: 10.5281/zenodo.19437382.
- *The Four-State Observer Protocol.* DOI: 10.5281/zenodo.19130570.

Primary clinical literature:

- Chen DS, Mellman I. (2013). Oncology meets immunology: the cancer-immunity cycle. *Immunity* 39(1): 1–10.
- Mellman I, Chen DS, Powles T, Turley SJ. (2023). The cancer-immunity cycle: indication, genotype, and immunotype. *Immunity* 56(10): 2188–2205.
- Bassez A, Vos H, Van Dyck L, et al. (2021). A single-cell map of intratumoral changes during anti-PD1 treatment of patients with breast cancer. *Nature Medicine* 27: 820–832.
- Leader AM, Grout JA, Maier BB, et al. (2021). Single-cell analysis of human non-small cell lung cancer lesions refines tumor classification and patient stratification. *Cancer Cell* 39: 1594–1609.
- Zhang L, Li Z, Skrzypczynska KM, et al. (2022). Pan-cancer single-cell landscape of tumor-infiltrating T cells. *Nature Cancer*.
- Zhang Y, Chen H, Mo H, et al. (2021). Single-cell analyses reveal key immune cell subsets associated with response to PD-L1 blockade in triple-negative breast cancer. *Cancer Cell* 39(12): 1578–1593.
- Liu B, Yang Z, Zhang Y, et al. (2025). A single-cell atlas reveals immune heterogeneity in anti-PD-1-treated non-small cell lung cancer. *Cell* (n=234, GSA-Human HRA004351; processed metadata in GEO GSE243013).
- GSE207422 — Pretreatment NSCLC bulk RNA cohort, n=24 patients on neoadjuvant anti-PD-1 + carboplatin-based chemotherapy with matched pathological response. NCBI GEO accession GSE207422.
- Wang AZ, Bowman-Kirigin JA, Desai R, et al. (2024). T cell InteractPrint: cancer epithelial cell heterogeneity-aware predictor of response to anti-PD-1 therapy. *Cell Reports Medicine*.
- Pelka K, Hofree M, Chen JH, et al. (2021). Spatially organized multicellular immune hubs in human colorectal cancer. *Cell* 184: 4734–4752.
- Im SJ, Hashimoto M, Gerner MY, et al. (2016). Defining CD8+ T cells that provide the proliferative burst after PD-1 therapy. *Nature* 537: 417–421.
- Miller BC, Sen DR, Al Abosy R, et al. (2019). Subsets of exhausted CD8+ T cells differentially mediate tumor control and respond to checkpoint blockade. *Nature Immunology* 20: 326–336.
- Beltra JC, Manne S, Abdel-Hakeem MS, et al. (2020). Developmental relationships of four exhausted CD8+ T cell subsets reveals underlying transcriptional and epigenetic landscape control mechanisms. *Immunity* 52: 825–841.
- Yost KE, Satpathy AT, Wells DK, et al. (2019). Clonal replacement of tumor-specific T cells following PD-1 blockade. *Nature Medicine* 25: 1251–1259.
- Sade-Feldman M, Yizhak K, Bjorgaard SL, et al. (2018). Defining T cell states associated with response to checkpoint immunotherapy in melanoma. *Cell* 175(4): 998–1013. NCBI GEO accession GSE120575.
- Gavish A, Tyler M, Greenwald AC, et al. (2023). Hallmarks of transcriptional intratumour heterogeneity across a thousand tumours. *Nature* 618: 598–606. (Curated Cancer Cell Atlas / 3CA, Tirosh lab Weizmann Institute, providing the curated metadata for Sade-Feldman 2018 and Pelka 2021 used in this paper.)
- Schenkel JM, Herbst RH, Canner D, et al. (2021). Conventional type I dendritic cells maintain a reservoir of proliferative tumor-antigen specific TCF-1+ CD8+ T cells in tumor-draining lymph nodes. *Immunity* 54(10): 2338–2353. (Documents the Position 7 substrate in TdLN: stem-like CD8 maintained by cDC1 in tumour-draining lymph node.)
- Borst J, Ahrends T, Bąbała N, Melief CJM, Kastenmüller W. (2018). CD4+ T cell help in cancer immunology and immunotherapy. *Nature Reviews Immunology* 18(10): 635–647.

- Ahrends T, Spanjaard A, Pilzecker B, et al. (2017). CD4+ T cell help confers a cytotoxic T cell effector program including coinhibitory receptor downregulation and increased tissue invasiveness. *Immunity* 47(5): 848–861.
- Ferris ST, Durai V, Wu R, et al. (2020). cDC1 prime and are licensed by CD4+ T cells to induce anti-tumour immunity. *Nature* 584(7822): 624–629.
- Lei X, de Groot DC, Welters MJP, et al. (2024). CD4+ T cells produce IFN-I to license cDC1s for induction of cytotoxic T-cell activity in human tumors. *Cellular & Molecular Immunology*. (Identifies a transcriptomic signature of CD4-helped cDC1s in T-cell-infiltrated human cancers, associated with good prognosis and anti-PD-1 response.)

Appendix A — Figure: The Sixteen Positions of the Cancer-Immunity Cycle



van der Klein 2026 · Generative Geometry · Completing the Cancer-Immunity Cycle

The sixteen-position cycle, summarising what is named in detail in Table 1 of §4.1. Status by glyph: filled black = LOCKED, open circle = SUPPORTED, ringed dot = HYPOTHESIS, dashed open = PROVISIONAL, small grey = generic immune biology. Mellman tags (M1–M7) mark the six and a half positions covered by Chen & Mellman 2013.